



THE 2026 EDITION · BY MICHAEL DOYLE

HOW TO GET FOUND ON *AI.*

A field guide to visibility across ChatGPT, Google AI Mode, Perplexity, Gemini, Claude, and the agents that are quietly replacing the search box.

MICHAEL DOYLE
FOUNDER, TOP4 TECHNOLOGY
25+ YEARS IN DIGITAL · AUSTRALIA · INDONESIA · GLOBAL

Volume II
UPDATED APRIL 2026

CONTENTS

What's *changed*, what works, and where to put your money in 2026.

..	Foreword — <i>from the founder</i>	03
00	Introduction — <i>eighteen months that broke search</i>	05
01	From search to synthesis	07
02	The citation economy, by the numbers	11
03	The 2026 GEO stack — <i>what each platform rewards</i>	15
04	Positioning & frameworks that get cited	18
05	Technical foundations — <i>schema, llms.txt, the honest version</i>	21
06	Local AI visibility — <i>the unfair advantage for SMBs</i>	24
07	Agentic commerce — <i>when the agent does the buying</i>	27
08	Custom GPTs & AI agents that earn their keep	30
09	Measuring AI visibility — <i>the new analytics stack</i>	33
10	Action plan by vertical	36
11	Conclusion — <i>the next training cycle</i>	40
..	About the author	41

I've watched the web reinvent itself *three times*.

This one is different.

I started in digital in 1999, in my first job out of the gate. Founded my own agency in 2001. Sold it. Bought it back. Built another. Somewhere along the way Top4 Technology grew into a platform serving 200,000+ businesses globally — from ASX-listed brands and government departments to small clinics and family trades. Property developers in Surabaya. Plumbers in Parramatta.

I have lived through Google ranking, the social era, mobile, voice. None of those shifts made the previous playbook worthless. They added a layer. The work compounded.

This one is different.

AI search hasn't added a layer. It has replaced the layer everyone was standing on.

If you sell anything to anyone, the rules just changed. Not in a "we should look at this next quarter" way. In a "your traffic is already moving and your team hasn't noticed yet" way. The brands and businesses that adapted in the last twelve months are pulling away. The ones still optimising for page-one Google are watching their organic traffic drop 30, 40, 50% and chalking it up to "the algorithm."

It isn't the algorithm. It's a new system.

Why I wrote this in 2026.

I wrote the first edition of this guide in September 2025. The shape of the opportunity was clear by then. The detail wasn't. I had to make some bets about how things would develop, and frankly, some of those bets were wrong.

llms.txt was supposed to matter. It doesn't, yet — and probably won't in the way most agencies sold it. ChatGPT Instant Checkout was supposed to be the agentic commerce on-ramp. OpenAI quietly killed it in March. AI Mode was supposed to be a slow rollout. It hit 75 million daily users by the time you're reading this.

So this edition is the version with the bets corrected. With the new mechanics — query fan-out, retrieval before ranking, embeddings, the centroid model that determines whether AI even sees you — explained in language a marketing director can act on.

None of this is theoretical. Every framework in this book is one we run for clients at Top4 Technology. Every number is one I have tested. Every recommendation is one I'd bet my own business on.

Which I am.

The brands winning today aren't running clever AI strategies. They are running unsexy fundamentals with discipline, and they started six months before everyone else.

If this guide saves you from making the wrong bet at the wrong moment, it has done its job. If it gives you the language to brief your team, sell it to your CFO, or pick the right agency partner, even better. And if it convinces you to stop reading and start acting — that's the goal.

The next training cycle is coming. Let's get you in it.

— *Michael*

Eighteen months that *broke* search.

Your traffic is moving. Your team probably hasn't noticed yet.

In September 2025 I published the first edition of this guide. AI search was a thesis. Eighteen months later it is the channel. Faster than I expected. Messier than the analysts predicted. And the businesses that adapted in those eighteen months are already pulling away from the ones that didn't.

If you read the 2025 version, you got most of the foundations right. The shape of the opportunity hasn't changed — be the source the model cites. The tactics underneath have moved a long way.

Here's the lay of the land you're now selling, building, and writing into.

48%

of all tracked Google queries now trigger an AI Overview, up 58% year on year.

BRIGHTEDGE · MAR 2026

93%

of Google AI Mode searches end with zero clicks to any external website.

SEER INTERACTIVE · 2025

1,200%

surge in traffic from AI sources to small business sites in the last 12 months.

PREVISIBLE · JAN 2026

Two things are happening at once. Traditional Google traffic is shrinking. AI-mediated traffic is exploding. They behave nothing like each other.

People who arrive from an AI answer convert at roughly 4× the rate of standard organic traffic. But only if your business is the one being recommended in the first place.

That's the whole game now.

Not "ranking on page one." Being named in the answer.

WHO THIS IS FOR

Founders, CMOs, agency leads, and anyone with revenue tied to a website. If you sell to humans who use Google, ChatGPT, Perplexity, Gemini, or Claude — which by 2026 means everyone — this guide is built for you.

What's new in this edition: AI Mode & AI Overviews mechanics, the agentic commerce protocols, llms.txt with an honest verdict, local AI visibility for SMBs, vertical playbooks, and a real measurement stack.

What changed since the last edition.

If you read the 2025 edition and built against it, you got most of the foundation right. But six things have shifted in ways that demand a fresh playbook.

1 · GOOGLE LAUNCHED AI MODE — AND IT'S A DIFFERENT BEAST

AI Mode is not the same product as AI Overviews. Overviews sit above traditional results and let people click through. AI Mode replaces the results page entirely with a conversational answer. It hit 75 million daily users by late 2025 and became publicly available across the US in March 2026. Sessions average 49 seconds — more than double an Overview session — and the people inside them are doing real research. They are also, overwhelmingly, not clicking out.

2 · CHATGPT'S COMMERCE STRATEGY RESET

OpenAI launched Instant Checkout in September 2025, then quietly retired in-chat checkout in March 2026 in favour of merchant apps and dedicated discovery experiences. Google countered with the Universal Commerce Protocol at NRF in January 2026. The infrastructure is moving faster than most retailers can keep up with — but the underlying truth is settled: agents are doing the shopping, and your product feed is now your storefront.

3 · LLMS.TXT GOT HYPED, THEN GOT REAL

The 2025 advice — "implement llms.txt now" — was premature. As of early 2026 no major LLM provider officially honours the standard, despite 844,000+ sites adopting it. We give you the unvarnished take in Chapter 5.

4 · LOCAL AI VISIBILITY IS THE NEW LOCAL SEO

Perplexity is Bing-first. Gemini draws from Google Business Profile. ChatGPT pulls live data from Bing too. Local businesses with clean directory citations and proper schema are getting cited at rates that didn't exist 12 months ago.

5 · AI CITATION TRACKING BECAME A CATEGORY

Profound, Otterly, AthenaHQ, Rocketito, Scrunch, Alclicks — there are now real tools that show you where you appear in AI answers and where you don't. Treat them like the early Search Console. Imperfect, occasionally weird, but essential.

6 · BOTTOM-OF-FUNNEL CONTENT BEATS TOP-OF-FUNNEL

Generic how-to content has been cannibalised by AI Overviews at almost 100%. Comparison pages, product pages, case studies with real numbers, and pricing pages are now the formats earning citations and the trickle of clicks that remain.

From search to *synthesis*.

The job of a search engine used to be to find things. The job of an AI engine is to decide things — what's true, what's relevant, what's worth recommending, and what's not worth mentioning at all.

That distinction sounds philosophical. It is in fact the entire commercial difference between the old web and the new one. A ranking system rewards the most relevant page for a query and lets the user adjudicate. A synthesis system makes the adjudication itself, then presents a single, confident answer with three to ten cited sources baked in.

If you are not one of those sources, you are invisible. There is no second page.

The four engines that matter.

Five years ago, "search" effectively meant Google. In 2026, the answer-engine layer is divided across roughly four platforms with very different mechanics, audiences, and citation behaviours.

GOOGLE AI MODE & AI OVERVIEWS

~48% OF QUERIES · 75M DAILY USERS

Two products in one ecosystem. AI Overviews are the quick AI answer above the blue links; AI Mode is the full conversational replacement. Driven by Gemini, deeply integrated with Google's index, Knowledge Graph, and Google Business Profile.

- Cites organic results that already rank well
- Heavy weighting toward schema and structured data
- Strong local signal via Google Business Profile

PERPLEXITY

~97% OF ANSWERS CARRY CITATIONS

The most citation-friendly engine by a wide margin. Real-time crawl, fresh content bias, draws heavily from Reddit, Quora, and authoritative third-party publications. Bing-indexed under the hood.

- Rewards visible publish/update dates
- Loves original data, statistics, and proprietary research

CHATGPT (OPENAI)

700M+ WEEKLY USERS · 1B+ DAILY QUERIES

The behavioural default for "ask AI" worldwide. Live web access via ChatGPT Search is powered by Bing. Cites sources in only ~16% of responses — which makes the slots that exist disproportionately valuable.

- Bing index is the gateway — get indexed there or stay invisible
- Listicles, articles, and product pages dominate citations
- Brand authority and entity clarity matter more than backlinks

MICROSOFT COPILOT & CLAUDE

ENTERPRISE REACH + PREMIUM REASONING

Copilot is the workhorse inside Microsoft 365 and Edge — quieter than ChatGPT in consumer mind-share, but shipping commerce features faster. Claude (Anthropic) is the long-context model used heavily by professionals and developers.

- Copilot Checkout is live with Shopify, PayPal, Etsy
- Both reward genuinely well-structured, factual content

- Reddit presence in your category genuinely moves the needle
- Claude weights authoritative sources particularly heavily

How synthesis actually works.

It helps to understand the mechanism, because it tells you what to optimise for. When you type a question into Google AI Mode, the system uses a technique called query fan-out: a single question is decomposed into roughly sixteen related sub-queries, each one searched in parallel, the results filtered for credibility and relevance, and the whole lot synthesised into one coherent answer with citations.

That changes the optimisation problem. You are no longer trying to rank for "best CRM for small business." You are trying to be in the comprehension layer that handles every plausible sub-question — pricing, integrations, ease of use, comparisons, edge cases — across enough independent sources that the model treats your brand as the consensus answer.

Ranking is no longer the prize. Inclusion is. And inclusion is decided by whether the model believes you are the consensus.

THE 2026 REALITY, IN ONE LINE

The zero-click economy is real, and it's fine.

The most common objection to all this: "if 93% of AI Mode searches end without a click, what's the point?" Fair question. The answer is uncomfortable but mathematically clean.

The traffic you used to get from informational queries — "how does X work," "what is Y," "guide to Z" — is gone. Permanently. That click was always low-intent and low-value; you just didn't notice because volume hid the conversion problem. What's left in the click stream is high-intent traffic: comparison pages, "best X for Y" queries, transactional searches, branded searches. Those still click through, and they convert at 4x or higher.

So the strategic move is not to fight zero-click. It's to redirect content investment away from informational top-of-funnel (which AI now eats) and toward bottom-of-funnel content that AI cites and sends qualified traffic to. We come back to this in Chapter 4.

REALITY CHECK

Press Gazette found Google search traffic to publishers dropped 33% globally in 2025. Business Insider's organic traffic fell 55% from April 2022 to April 2025. The Reuters Institute surveyed 280 media leaders who expect a 43% average drop over the next three years. Business as usual is over.

How AI *actually* sees your brand.

It doesn't read your homepage. It does the maths.

I want to give you the next layer down, because once you understand it, every recommendation in the rest of this book makes intuitive sense.

AI engines do not see your business the way you describe it on your About page. They build their own version of your brand from the content you have published — every page, every blog post, every product description, every PR mention — and they store that version as a mathematical object. Not a sentence. A position in space.

Three concepts make this concrete.

1 · RETRIEVAL COMES BEFORE RANKING

In classical SEO, you fight to move from page two to page one. In AI search, the fight happens earlier. Before anything is ranked, the system decides which content is even eligible to be considered. That step is called retrieval. If you're not retrieved, you don't lose. You just never entered the contest.

This is the single most important mental model shift in 2026. Optimisation is no longer about position. It's about admission.

2 · PAGES DON'T COMPETE — PASSAGES DO

AI doesn't treat your page as one unit. It breaks it into chunks: paragraphs, sentences, sections. Each chunk is evaluated on its own merit against the query. A paragraph buried halfway down a guide can win on its own. A homepage hero can be ignored entirely. The competition is passage versus passage, not page versus page.

This is why the format guidance in this book is so specific — direct answers in the first 80 words, question-shaped headers, FAQ blocks, TL;DRs. Every chunk has to earn its own place.

3 · YOUR "BRAND" IS A CENTROID

Here's the bit most marketing teams aren't ready for. Every chunk you publish is converted to a vector — a position in high-dimensional space that represents its meaning. All those vectors cluster together. The centre point of that cluster is what AI thinks your brand is.

Not your tagline. Not your brand guidelines. The mathematical average of everything you've ever written.

Two implications. First, every page you publish either reinforces that centre or pulls it sideways — drift is a real, measurable phenomenon. Second, brands that follow the same playbooks and write about the same things end up in the same place in meaning-space. The technical term is "cluster collision." The plain-English version: you and your three competitors all sound identical to AI, so the model picks one of you and ignores the rest.

THE HARD TRUTH

If your content is broadly similar to your competitors' content, AI sees one cluster and picks one brand. Producing more of the same content makes the problem worse, not better. The advantage in 2026 is a distinctly different position in meaning-space — which is what naming your methodology, publishing original data, and going narrow on positioning all achieve from different angles.

The framework researchers call this "the centroid model" of brand visibility. The practical version is simpler. Every chapter that follows is, ultimately, a different way of pulling your content into a clearer, more distinct centroid that AI can recognise, retrieve, and recommend. Hold that picture in your head as you read on.

What this means for your team this quarter.

If you take only one thing from this chapter, take this: the unit of optimisation has changed from page to entity. Google AI Mode and ChatGPT do not rank pages. They model entities — your brand, your people, your methodology, your products — and decide whether to mention them. Pages still matter, but only as the evidence layer that proves the entity is real, current, and credible.

That reframe drives every chapter that follows:

- Chapter 2 shows you what citation actually looks like and what behaviours buy it.
- Chapter 3 breaks down each platform's preferences in detail.
- Chapter 4 covers entity-level positioning — naming your methodology, claiming a specific outcome, drawing a small clean line around your expertise.
- Chapter 5 handles the technical evidence layer: schema, structured data, what to do about llms.txt.
- Chapters 6–10 apply the model to local businesses, eCommerce, agents, measurement, and the verticals Top4 sees most often.

§

One more thing before we move on. The temptation in this market is to chase platforms — "we need a ChatGPT strategy, a Perplexity strategy, a Gemini strategy." Don't. The platforms differ in mechanics, but they all reward the same underlying behaviours: clear positioning, structured content, third-party validation, fresh activity, consistent identity, and original evidence. Build that foundation once and it pays across all of them.

The agencies and brands winning right now are not running four parallel optimisation programmes. They are running one good one and pointing it at four pipes.

The citation economy, *by the numbers.*

The 2025 edition called this the "3 Rs" — Reputation, Recency, Authority. The instinct was right. The data is now sharp enough to be specific.

Here is what the leading studies and citation trackers have surfaced over the past twelve months. None of this is theory. All of it is measurable.

The numbers that should be on your strategy slide.

4.4×

conversion rate of AI-referred traffic vs traditional search.

EMARKETER · 2025

+35%

organic CTR uplift for brands cited in an AI Overview vs not cited.

SEER INTERACTIVE

3.6×

more frequent crawling by LLM bots than Googlebot, by April 2026.

SEARCH ENGINE JOURNAL

12%

of URLs cited by ChatGPT, Perplexity, and Copilot rank in Google's top 10 for the same query.

AHREFS

44.2%

of all LLM citations come from the first 30% of a page's text.

POSITION DIGITAL

46%

of all ChatGPT citations within a topic concentrate in just the top 10 domains.

GROWTH MEMO · MAR 2026

Read those numbers carefully.

Three of those stats are particularly worth chewing on, because they overturn long-held SEO assumptions.

Only 12% of cited URLs rank in Google's top 10. Translation: ranking and citation are completely separate scoreboards. You can be invisible on Google and prominent in ChatGPT. You can be number one on Google and ignored by Perplexity. The scoreboards do not talk to each other.

44.2% of citations come from the first 30% of the page. Translation: if your answer is buried below a 600-word introduction, the model never finds it. The TL;DR section, the lead paragraph, and the H2 directly under the title are now the most valuable real estate on any page you publish.

The top 10 domains take 46% of citations. Translation: this is a winner-take-most economy. The wide middle of the SEO market — pages 2–10 of Google for years — barely exists in AI. You are either the consensus or you are background.

What buys a citation in 2026.

Citation behaviour varies by platform, but five behaviours show up as predictive across every major engine. Treat this as the modern replacement for the "3 Rs."

1 · ENTITY CLARITY

Your business needs to be a coherent entity in the model's worldview. That means consistent name, consistent description, consistent founder/leader attribution, and a Wikidata or Knowledge Graph footprint where possible. If your brand name is even mildly generic and you haven't disambiguated yourself, the model will choose someone else when in doubt. AI is conservative by design.

2 · STRUCTURED DIRECTNESS

Headers that match real questions. Direct answers in the first 60–80 words. FAQ sections written in the voice your customer actually uses. Tables for comparisons. Bullet lists for processes. The format matters as much as the content because synthesis engines extract by structure first, semantics second.

3 · ORIGINAL EVIDENCE

Proprietary data, original surveys, internal benchmarks, named methodologies, real client outcomes with real numbers. Models reward the source they can't get elsewhere. This is also the single biggest moat against the wave of AI-generated thin content flooding every category — your originality is your defensibility.

4 · DISTRIBUTED PRESENCE

You cannot win AI visibility from your website alone. You need consistent appearances across podcasts, industry publications, directories, review platforms, and (yes) Reddit and Quora threads in your space. The model is looking for a pattern of independent corroboration. One website plus thirty independent mentions beats one well-optimised website every time.

5 · RECENT ACTIVITY

Visible publish dates. Recent updates. Active social signals. Press in the last 90 days. Models heavily favour recency for almost every category except evergreen reference content. A brand that was talked about last week looks more real than one last referenced two years ago.

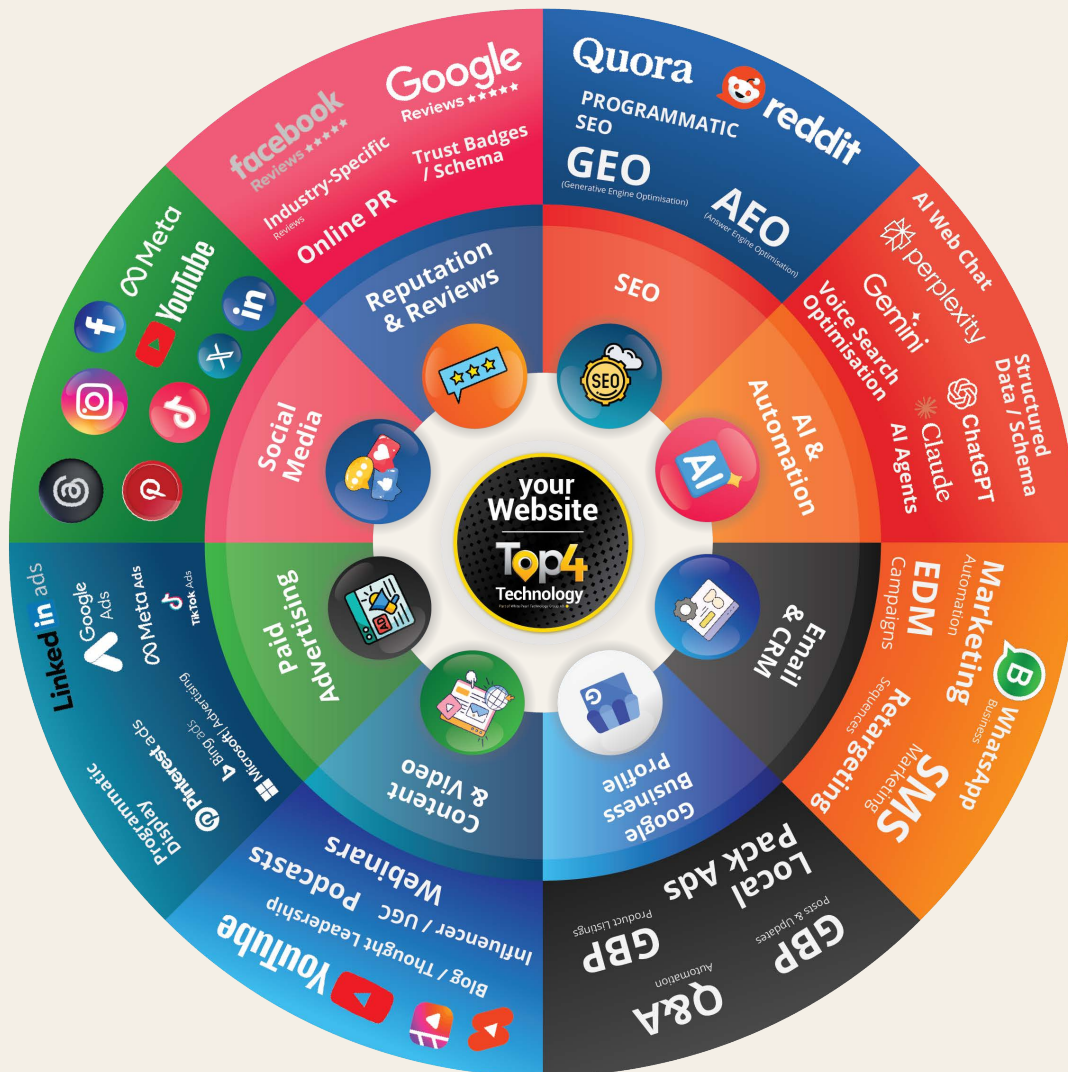
THE COMPOUNDING EFFECT

None of these is a tactic in isolation. They compound. A clean entity makes structured content easier to attribute. Structured content earns more directory and podcast distribution. Distribution drives recency signals. Recency reinforces the entity. Three quarters of consistent execution against all five and the model starts treating you as the answer in your category.

The contrast point is also worth naming. Brands that are losing AI visibility right now share a profile: weak entity definition, generalist positioning, content optimised for keywords rather than questions, no original data, all activity concentrated on their own domain, and last-updated dates from 2023.

Eight channels. *One website at the centre.*

Every channel in this guide hangs off the same hub. Your website is the entity; the eight rings around it are the surfaces that prove it to humans, search engines, and AI models.



THE TOP4 TECHNOLOGY WHEEL • RING 1 CHANNELS & PLATFORM MAP

Read the wheel inside-out. The centre is your site — the entity AI models must recognise. The inner ring is the eight Top4 channel categories. The outer ring is the platforms inside each. **SEO, AI & Automation, Reputation, and Google Business Profile** are where AI visibility is won or lost; the others are the corroboration layer that makes the model believe you.

The 2026 GEO stack — *what each platform actually rewards.*

The acronym soup matters less than people think. SEO. AEO. GEO. LLMO. They describe the same thing from different angles. What matters is knowing the specific levers each platform pulls — because while the foundation is shared, the tactics on top differ enough to be worth separating.

Here is the per-platform breakdown for April 2026, current to everything we've measured this quarter.

Google AI Mode & AI Overviews.

The big one. Google's AI products draw from the existing index, the Knowledge Graph, Google Business Profile, YouTube, and a dedicated retrieval layer. They reward what Google has always rewarded — with sharper teeth.

- **Schema markup is now non-negotiable.** Organization, Person, Service, Product, Article, FAQPage, Review, and LocalBusiness schemas materially increase the chance of being parsed correctly and cited.
- **Top-10 organic ranking still matters.** Pages ranking #1 in Google have a 58% chance of being cited in an AI Overview; pages ranking outside the top 20 are 3.5× less likely. Traditional SEO is still your foundation here.
- **Knowledge Graph entity status is gold.** Claim a Wikidata entry. Build the entity. The Knowledge Panel is now the single highest-leverage win in Google's ecosystem.
- **Branded queries get an 18% CTR uplift** when they trigger an AI Overview. Build branded search demand and AI works in your favour.

ChatGPT (OpenAI).

The behavioural default. ChatGPT Search runs on Bing's index, which means the optimisation lever everyone underuses is right there.

- **Bing Webmaster Tools.** Submit your sitemap. Fix indexation errors. Use IndexNow so new pages are pushed instantly. ChatGPT cannot cite what Bing hasn't indexed.
- **Crawler access.** Check robots.txt for OAI-SearchBot and ChatGPT-User. Roughly 70% of ChatGPT citations come from sites not blocking these bots — meaning most teams are blocking themselves out of citations they could earn.
- **Listicles and comparison content.** 21.9% of all AI citations are listicles. "Best X for Y" formats earn citations at the highest rate of any content type.
- **Brand authority over backlinks.** ChatGPT's training-era knowledge weights authoritative mentions and consistent third-party references far above raw link counts.

Perplexity.

The most citation-generous engine — and the easiest to break into for new entrants because it weights freshness and direct answers so heavily.

- **Visible dates on every page.** Both publish and updated. Refresh quarterly and surface the new dateModified.
- **Bing Places, not just Google Business Profile.** Perplexity's local index runs through Bing. Most local businesses haven't even claimed their Bing Places listing — that is a free win available right now.
- **Reddit and Quora presence.** Perplexity scrapes both heavily. Honest, useful contributions in the relevant subreddits earn citations no SEO investment can replicate.
- **TL;DR blocks at the top of pages.** A 60–80 word callout summary is a perfect quotable passage. We've seen 4× lift in Perplexity citations on client pages by adding nothing but TL;DRs.

Microsoft Copilot, Gemini, and Claude.

Three smaller pipes by consumer volume but disproportionately important for B2B, enterprise, and considered-purchase categories.

COPILOT & GEMINI

Copilot is integrated into Microsoft 365 and Edge — used heavily by office workers and enterprises. Strong commerce push via Copilot Checkout. Same Bing-index dependencies as ChatGPT.

Gemini draws from Google Search, Google Business Profile, Google Maps, YouTube, and the Knowledge Graph. For local businesses, the path to Gemini visibility is essentially excellent local SEO — fully optimised GBP, 50+ recent reviews, regular Google Posts, complete services list.

CLAUDE (ANTHROPIC)

Used heavily by developers, technical professionals, and enterprises. Weights authoritative sources particularly heavily — research papers, official documentation, established publications.

If you sell to technical audiences, Claude visibility is built through documentation quality, GitHub presence, and inclusion in trusted reference sources. Less amenable to fast tactics, more rewarding of genuine technical depth over time.

The unified stack.

Strip away the platform-specific tactics and the work that compounds across all of them is straightforward. The 2026 GEO stack looks like this:

1. **Foundation.** Indexable site, clean robots.txt, complete schema, claimed entity (Google Business Profile, Bing Places, Wikidata where applicable).
2. **Content.** Question-shaped headers, direct answers in the first 80 words, comparison and listicle formats for commercial queries, original data and proprietary frameworks.
3. **Distribution.** Directories, reviews across multiple platforms, podcast appearances, guest articles, Reddit/Quora participation.
4. **Recency.** Quarterly content refresh cycles, ongoing PR drumbeat, active social presence, regular case study publication.
5. **Measurement.** Citation tracking (Profound, Otterly, or equivalent), bot analytics (Ahrefs Bot Analytics is free in beta), AI referral tracking in GA4 with custom segments.

The unfair advantage hiding in your robots.txt.

One detail worth its own page. BuzzStream's April 2026 audit of AI citation patterns surfaced a strange and useful asymmetry. Roughly 95% of pages cited by AI engines were blocking the AI training bots (GPT-Bot, Google-Extended), but allowing the AI retrieval bots (ChatGPT-User, OAI-SearchBot, PerplexityBot).

That distinction matters because most teams either let everything through or block everything. The smart configuration is selective.

```
# Block training-data harvesters. Optional but recommended for IP-heavy sites.
User-agent: GPTBot
Disallow: /

User-agent: Google-Extended
Disallow: /

# Allow live retrieval bots – these fetch your content at query time
# to cite you in the user's actual answer.
User-agent: ChatGPT-User
Allow: /

User-agent: OAI-SearchBot
Allow: /

User-agent: PerplexityBot
Allow: /

User-agent: ClaudeBot
Allow: /
```

The logic: training bots ingest your content into model weights — useful but uncompensated. Retrieval bots fetch your content at the moment of a real user query and cite you in the answer. Blocking the second category is what makes you invisible. Most agencies and tools (including some well-known SEO platforms) recommend blocking everything by default. That advice is now actively harmful.

QUICK AUDIT

Pull your robots.txt right now. Check for OAI-SearchBot, ChatGPT-User, and PerplexityBot. If any are disallowed, you're blocking yourself out of citations you could be earning today. Single biggest 5-minute win in this entire guide.

That ends the foundation. The next chapter covers the part most agencies overcomplicate — positioning and frameworks. The 2025 edition got this mostly right, and the principle hasn't changed. The execution has sharpened.

Positioning & frameworks *that get cited.*

If the model can't tell what you do or who you do it for, it won't recommend you. Generalists lose in AI search the way generalists lose in any market with infinite supply.

The 2025 edition introduced what we called the LLM Positioning Framework. The idea was simple and it has aged well: I help [specific audience] achieve [specific outcome] through [named process]. Eighteen months on, the only refinement we'd make is to be even narrower than feels comfortable.

The narrowness paradox.

Most founders resist tight positioning because it feels like leaving money on the table. "We could serve dentists too." "We do real estate as well." "Our framework works in any industry." All true, all unhelpful.

AI engines are optimised to provide the most relevant answer to a specific query. When someone asks "marketing strategy," the model surfaces general resources. When they ask "lead generation for cosmetic clinics with a single location and under \$100k in monthly revenue," the model picks the specialist who has spent two years writing about exactly that. The narrower your stake, the more questions you become the definitive answer to — even though each individual question has lower volume. Sum them up and you win the category.

The narrower your focus, the more queries you become the definitive answer to. Cumulative wins beat broad wins every time in synthesis search.

THE NARROWNESS PARADOX

The 2026 positioning template.

Same skeleton as 2025, sharper specificity.

THE TEMPLATE

"We help [specific audience] achieve [specific outcome with a measurable threshold] through [your named methodology]."

Generic: "We help businesses grow through digital marketing."

Sharp: "We help Australian cosmetic clinics generate qualified bookings through the Clinic Demand Engine — a SEO, AI Search, and review-management system designed for single-location practitioners."

Frameworks: still the highest-leverage IP move.

Named methodologies remain the most efficient way to compress your expertise into something AI engines can recognise, attribute, and recommend. The mechanics haven't changed since the last edition — repetition across diverse, credible sources creates entity authority. What we'd add for 2026:

- **Frameworks need to be schema-marked up.** Use `HowTo` or `DefinedTerm` schema on your framework page. Make it structurally obvious that this is a methodology, not just a blog post.
- **Frameworks need their own dedicated page.** Not buried in a services page. A canonical URL the model can return to as the source of truth.
- **Frameworks need to appear in third-party content.** Podcasts, guest articles, conference talks. The framework name said by other people in their context is what makes the model believe it's real.

The framework playbook in eight moves.

1. **Identify your repeating process.** What do you do for clients that produces consistent outcomes? Document the actual sequence, not the marketing version.
2. **Name it with intent.** Outcome-focused language wins. "The Demand Engine," "The Pipeline Acceleration Method," "The Trust Stack" beat clever names.
3. **Trademark it.** Genuine ownership signal. Costs are low. Worth doing.
4. **Build the canonical page.** One URL, full explanation, schema-marked, with a clear hierarchy of phases or steps.
5. **Create the supporting content.** Long-form article, downloadable PDF, video walkthrough, dedicated FAQ.
6. **Use it everywhere.** In proposals, on LinkedIn, in podcast intros, in case studies, on social. Repetition across channels is what makes the entity stick.
7. **Get other people saying it.** Train clients and podcast hosts to use the name. Get it into industry articles. Get it onto Wikipedia or Wikidata if you can.
8. **Refresh it annually.** Versioned frameworks (v2.0, 2026 edition) signal recency, which feeds the recency-bias loop.

Two warnings.

Don't trademark a framework you don't use. Models pick up the gap between what you claim and what your content shows. A framework you mention once and never reference again is worse than no framework at all — it reads as marketing fluff and depresses entity credibility.

Don't proliferate frameworks. One signature framework supported by two or three sub-frameworks is the right shape. Five separate flagship methodologies dilute the signal. Pick the lane.

Case study: framework as moat.

One example we've watched closely. A Sydney HVAC contractor — small business, four trucks, family-owned — built a framework called the "Five-Point Comfort Audit." It's a structured inspection process they used internally. They wrote a 2,500-word page describing it, marked it up with `HowTo` schema, added a downloadable PDF, mentioned it in every blog post, trained their team to reference it on every quote, and pitched it to two industry podcasts.

Twelve months in, asking ChatGPT, Perplexity, or Gemini "what's the best HVAC inspection process in Sydney" returns the Five-Point Comfort Audit as a named methodology — with their business attributed as the source — in roughly 60% of responses we've tested. Their phone-call volume from AI-attributed sources has tripled.

That outcome is not because they have the best HVAC business in Sydney. It's because they are the only one with a named, documented, repeated, schema-marked methodology that AI engines can latch onto when synthesising an answer.

TOP4 CLIENT SNAPSHOT

The same play works in legal (named matter-management process), real estate (named appraisal method), trades (named scope or pricing system), eCommerce (named curation method), medical (named patient-journey framework), and B2B (named onboarding or value-realisation method).

The vertical doesn't change the principle. Name the thing, document it, schema-mark it, repeat it across channels. Three to twelve months later it becomes the citation magnet.

Action this week.

- Audit your homepage and services page against the positioning template — is it sharp enough?
- Identify the repeatable process you already run for clients but haven't named.
- Draft three candidate names for it; pick the one that's most outcome-focused and most searchable.
- Trademark search the name. If clear, file the application this month.
- Build the canonical framework page on your site. Add HowTo or DefinedTerm schema.
- Brief every customer-facing team member to use the name in calls, proposals, and emails.

That's the foundation. Now we get to the technical layer that makes all of it legible to the machines.

Technical foundations — *schema, llms.txt, the honest version.*

Most of this chapter is non-controversial. Schema works. Structured content works. JSON-LD works. The controversial part is llms.txt — and the 2025 edition of this guide oversold it. Here's the corrected take.

Schema markup is the single biggest 2026 lever.

If you do nothing else from this guide, do schema. Every major AI engine uses structured data to understand what your business is, what you sell, who you are, where you operate, and what other people think of you. Pages with comprehensive schema are cited at materially higher rates than pages without.

The schemas that matter most for AI citation, in priority order:

- **Organization** — your business identity, with sameAs links to all your social and directory profiles. The single most important schema for entity disambiguation.
- **LocalBusiness** (or specific subtype: Restaurant, MedicalBusiness, ProfessionalService, etc.) — required for any business with a physical location or service area.
- **Person** — for founders, key team members, and named experts. Critical for personal brand and authority signals.
- **Service / Product** — for everything you sell.
- **FAQPage** — wraps your FAQ section in machine-readable Q&A. AI engines parse this directly.
- **Review & AggregateRating** — third-party validation made structurally explicit.
- **HowTo / DefinedTerm** — for your named frameworks and methodologies.
- **Article / NewsArticle** — for blog content, with proper datePublished and dateModified.
- **BreadcrumbList** — site structure clarity.

Implement these in JSON-LD in the page <head>. Validate everything in Google's Rich Results Test and Schema.org Validator. Both Google and Bing parse JSON-LD; everyone else inherits from those crawls.

The llms.txt situation.

Time to be straight about this. The 2025 edition recommended implementing llms.txt as a key tactic. As of April 2026, that recommendation needs significant qualification.

WHAT WE NOW KNOW

No major AI provider — OpenAI, Anthropic, Google, or Meta — officially honours llms.txt. Server log analysis from multiple independent audits confirms the major retrieval bots (GPTBot, ClaudeBot, Perplexity-

Bot) do not request llms.txt files. Google's John Mueller stated explicitly: "No AI system currently uses llms.txt."

Around 844,000 websites have implemented it anyway. Anthropic, Cloudflare, and Stripe publish their own llms.txt files. The standard exists, the file format is well-defined, and adoption among site owners has run far ahead of adoption by the AI platforms it's intended to serve.

So should you implement it?

Our 2026 take, calibrated to the actual evidence:

If you have technical documentation, an API, or a developer-facing product — yes. Coding assistants (Cursor, Windsurf, Claude Code, GitHub Copilot) actively benefit from llms.txt and llms-full.txt files when generating code that references your platform. Anthropic specifically requested llms.txt and llms-full.txt for their own documentation. Vercel reports 10% of signups now come from ChatGPT-driven discovery, and they credit their llms.txt implementation as part of why.

If you run a content site, marketing site, or eCommerce store — it's optional. The implementation cost is low (it's a markdown file at the root of your domain). The current return is uncertain. We recommend implementing if you're already doing other technical SEO work; we wouldn't make it a priority over schema, content, or distribution.

What you absolutely shouldn't do: believe the SEO tools that flag a missing llms.txt as an "issue." Semrush and several others have started doing this. It is not an issue. It is a missing implementation of an unratified, unrequested standard. Treat that flag the same way you'd treat a tool flagging "missing pinterest verification meta tag" — an opinion, not a fact.

If you do implement, here's the right shape.

```
# Top4 Technology – Expertise Index

## What we do
We build digital marketing systems for Australian SMBs and government clients,
with deep specialism in AI search visibility, Kentico CMS, Drupal, and
Google Premier Partner-level performance media.

## Core methodology
The Top4 Demand Engine – a SEO + AEO + Local Search + Reviews stack designed
for service-based businesses with $1M-$50M annual revenue.

## Documentation
- [Services overview](https://top4marketing.com/services)
- [Industries we serve](https://top4marketing.com/industries)
- [Case studies](https://top4marketing.com/work)

## Contact
Email: hello@top4.com.au
Web: https://top4marketing.com
```

That's the format. Plain markdown, hosted at `/llms.txt`. If you have substantial documentation, also publish `/llms-full.txt` with the actual content of your most important pages compiled into a single file. Mintlify will generate this automatically; if you're on WordPress, several plugins now handle it.

Other technical foundations worth getting right.

A short list of the technical work that consistently moves AI visibility, in roughly priority order:

- **Bing Webmaster Tools** — submit your sitemap, verify ownership, fix indexation errors. Powers ChatGPT and Perplexity. Most teams ignore Bing entirely. Free win.
- **Google Business Profile** — fully completed, photos, regular Posts, services list, attributes filled in, owner-responded reviews. Powers Gemini and Google AI Mode local results.
- **Bing Places** — the local equivalent for Bing/Perplexity. Almost no one has claimed theirs.
- **IndexNow** — instant indexing protocol for Bing. Implement on your site or via Cloudflare. New pages reach ChatGPT and Perplexity within minutes instead of days.
- **Wikidata entry** — for any established business. The single biggest entity disambiguation win.
- **Site speed and Core Web Vitals** — still matter for Googlebot, which still feeds AI Overviews.

Page-level structure for AI citation.

The format on individual pages matters more in 2026 than at any point in the SEO era. Here is the structure we now use as the default for any page intended to earn AI citations:

1. **H1** with the primary question or topic clearly stated.
2. **TL;DR or summary block**, 60–80 words, in the first viewport. The single highest-value 80 words on the page. Models lift directly from here.
3. **Direct, specific answer** to the main question in the first paragraph. Don't bury it.
4. **H2s shaped as questions** or specific outcomes. Match the way customers actually ask, not how marketers describe.
5. **Short paragraphs, bullet lists, comparison tables** wherever applicable. Synthesis engines extract by structure first.
6. **Original data, statistics, and named methodologies** prominently positioned.
7. **FAQ section** at the bottom with FAQPage schema.
8. **Visible publish and updated dates**.
9. **Author byline** with link to author page (which has Person schema).
10. **Internal links** to related pages with descriptive anchor text.

THE 5-MINUTE AUDIT

Pick your three most important pages. For each, answer: does the H1 state the topic clearly? Is the TL;DR in the first viewport? Are the H2s shaped as questions? Is there schema? Is there a visible updated date? Is the FAQ schema-marked? If the answer to any is no on a page that matters, you've found your next month's work.

Local AI visibility — *the unfair advantage for SMBs.*

If you run a local business — a clinic, a law firm, a tradie, a real estate agency, a restaurant — the AI search shift is the best thing to happen to your category in fifteen years. Most of your competitors haven't noticed. That window is open right now and it won't stay open.

Here's the asymmetry. Big national brands are losing AI visibility because they can't write geographic-specific content at scale — a firm operating in 15 markets can't realistically produce 15 versions of "How to file for divorce in [city]." Local businesses can. And the AI engines reward exactly that specificity.

The four pillars of local AI visibility.

1 · ENTITY CONSISTENCY ACROSS LOCAL SEARCH INFRASTRUCTURE

Your business name, address, phone number (NAP), hours, and category need to match across every directory the AI engines pull from. Inconsistent NAP is the single most common reason local businesses are skipped in AI answers — the model can't confirm which version of you is real, so it picks someone less ambiguous. The minimum viable directory list for 2026:

- **Google Business Profile** — fully completed, weekly Posts, photos updated monthly, owner-responded reviews
- **Bing Places** — the underrated one. Powers Perplexity local results.
- **Apple Business Connect** — increasingly important as Apple Intelligence rolls out
- **Industry-specific directories** — Houzz for trades, Avvo for legal, Healthgrades for medical, etc.
- **Top4 and other Australian directories** for local relevance
- **Yelp, TripAdvisor, Trustpilot** as relevant to your category

2 · LOCALBUSINESS SCHEMA WITH FULL PROPERTY COMPLETION

Generic LocalBusiness schema is necessary; specific subtypes are better. Use Dentist, Attorney, Plumber, RealEstateAgent, etc. Fill in areaServed, geo coordinates, openingHoursSpecification, priceRange, paymentAccepted, and knowsAbout for service specialisations.

3 · GEOGRAPHIC CONTENT DEPTH

One generic service page does not earn local AI citations. What earns citations is location-specific content — neighbourhood-level pages, suburb-specific guides, location pages that genuinely talk about local context (council requirements, regional regulations, local landmarks, named partners). Most AI Overview citations for local queries go to pages with measurable geographic specificity.

4 · MULTI-PLATFORM REVIEW VELOCITY

Reviews on Google alone aren't enough. Perplexity reads Yelp, Bing, Trustpilot, and category-specific platforms. Aim for ongoing review velocity (rate per month) across at least three platforms, not just total volume on Google.

The 90-day local AI visibility plan.

If you're starting from cold, this is the sequence that consistently moves citations within 60–120 days for the local clients we work with.

DAYS 1-30 · FOUNDATION

Fix what's broken before adding anything new.

- NAP audit across Google Business Profile, Bing Places, Apple Business Connect, and your top five industry directories. Resolve inconsistencies.
- Implement complete LocalBusiness schema with the most specific subtype available. Validate.
- Audit robots.txt — confirm OAI-SearchBot, ChatGPT-User, and PerplexityBot are not blocked.
- Submit to Bing Webmaster Tools. Submit your sitemap. Fix any indexation errors.
- Claim Wikidata entry if the business is established enough to qualify.

DAYS 31-60 · CONTENT

Build the geographic depth.

- Build dedicated location pages for each suburb or neighbourhood you serve. Real local context, not just "we serve [city]" boilerplate.
- Add a TL;DR / summary block to every service page (60–80 words).
- Create one substantial FAQ page covering the 15–20 questions you actually get asked. Schema-mark it as FAQPage.
- Reshape H2s on service pages as questions ("How long does X take?", "How much does Y cost in [city]?").
- Publish two case studies with real numbers, real outcomes, real client names where possible.

DAYS 61-90 · DISTRIBUTION & SIGNALS

Build the corroboration layer.

- Set up a structured review request process — aim for 3–5 new reviews per month across at least three platforms.
- Pitch two local podcasts or industry publications for guest appearances or expert commentary.
- Participate honestly in two or three relevant subreddits and one Quora topic per week.
- Issue one local press release tied to a real news hook.
- Begin tracking AI citations using Otterly, Profound, or a manual check spreadsheet.

Three months in, you should be seeing first citations in Perplexity (the fastest engine to recognise new entities) and improved presence in Gemini local results. ChatGPT and AI Mode citations typically follow at the four-to-six month mark as the entity layer thickens.

Agentic commerce — *when the agent does the buying.*

If you run an eCommerce business, this is the chapter you've been waiting for. The shift to agent-driven shopping is happening faster than anyone planned, and the rules for being discoverable are completely different from traditional eCommerce SEO.

Quick definition for the slow seats: agentic commerce is shopping where an AI agent does the discovery, comparison, and increasingly the checkout, on behalf of a human. The human sets the intent ("I need running shoes for flat feet under \$150 that ship to Sydney") and the agent handles the rest. This is not a 2030 prediction. It's a 2026 reality.

The two protocols you need to know.

ACP — AGENTIC COMMERCE PROTOCOL

Co-developed by OpenAI and Stripe. Powers checkout sessions inside ChatGPT. Started as the backbone for ChatGPT Instant Checkout, which OpenAI quietly pulled in March 2026 in favour of merchant apps.

Now adopted by 25+ partners including Salesforce, Squarespace, and Adobe Commerce. Focuses on the purchase-stage transaction within the ChatGPT ecosystem.

UCP — UNIVERSAL COMMERCE PROTOCOL

Co-developed by Google and Shopify, launched at NRF in January 2026. More comprehensive — covers the full journey from product discovery through post-purchase support.

Partners include Visa, Mastercard, Stripe, Walmart, Target, Etsy, and Wayfair. Powers checkout across Google AI Mode, Gemini, and Google Shopping.

You don't need to implement either protocol directly. If you're on Shopify, Shopify handles both for you via Agentic Storefronts (free, automatic for eligible US merchants as of March 2026). If you're on a different platform, your obligations are about the data that feeds these protocols.

What AI shopping agents actually need from you.

Here's the inversion most ecommerce teams miss: AI shopping agents do not read your blog posts to decide whether to recommend you. They read your product feed, your structured product data, your pricing accuracy, your inventory status, and your reviews. Editorial content does not move agentic commerce visibility. **Operational rigour does.**

The minimum data quality bar for agentic commerce in 2026:

- **Complete product feed** — title (max 150 chars), description (max 5,000), price with ISO 4217 currency code, availability, images, GTIN where available, full attribute set including colour, size, material, intended use

- **Real-time price and stock sync** — agents cross-reference your storefront price against your feed; mismatches mean your product gets skipped or flagged
- **Product schema** (Product + Offer + AggregateRating + Review) across every product page
- **Crawable robots.txt** — OAI-SearchBot and Googlebot must be allowed; this is the number-one reason brands are invisible in ChatGPT shopping results
- **Live policy pages** — return, shipping, warranty, support. Missing or gated policies reduce your trust score in agent selection.
- **Review velocity** — average rating, total count, and recency all factor

The agentic commerce numbers.

805%

YoY growth in AI-referred traffic to US retail sites on Black Friday 2025.

ADOBE VIA METAROUTER

\$20.9B

expected AI platform retail spending in 2026 — nearly 4× 2025.

EMARKETER · DEC 2025

+38%

higher purchase likelihood for shoppers arriving from AI services vs traditional channels.

EMARKETER

The forecasts on top of these numbers vary wildly. McKinsey projects \$3–5 trillion globally in agentic commerce by 2030. Morgan Stanley predicts nearly half of online shoppers will use AI shopping agents by 2030, accounting for ~25% of their spending. Even discounted heavily, the channel is going to be enormous and the decisions made about feed quality and protocol participation in 2026 will compound for years.

The eCommerce 2026 priority list.

If you sell online and you only do six things in the next 90 days, do these:

1. **Product data audit.** Pull your full feed. Score every product on completeness against the 12-attribute baseline (title, description, price, availability, images, GTIN, brand, condition, colour, size, material, category). Anything below 90% completion is invisible to agents.
2. **Schema audit on top 100 SKUs.** Product, Offer, AggregateRating, Review. Validate all four.
3. **Bot access audit.** Confirm OAI-SearchBot, Googlebot, ChatGPT-User, PerplexityBot, ClaudeBot are not blocked. Check Cloudflare and WAF rules, not just robots.txt.
4. **Policy pages live and indexable.** No login walls, no forms, no PDFs. Plain HTML, schema-marked where applicable.
5. **If on Shopify: enable Agentic Storefronts.** If on another platform: implement UCP via your stack's supported integration, or evaluate the Shopify Agentic Plan.
6. **Track AI shopping referrals in GA4.** Build custom segments for traffic from chatgpt.com, perplexity.ai, gemini.google.com, copilot.microsoft.com. Watch the trend line monthly.

ONE UNCOMFORTABLE TRUTH

The brands winning in agentic commerce right now are operationally boring and technically rigorous. Clean data, fast fulfilment, accurate inventory, structured everything. The brands losing are the ones that have spent five years pouring budget into Instagram aesthetics and brand storytelling while their product feed has 47% missing fields. The new channel rewards the thing the old channel didn't measure.

What about content?

Content still matters in agentic commerce — but not the way it mattered in classical SEO. The content that earns AI shopping citations is structured comparison content: "best running shoes for flat feet 2026," "X vs Y," "[product category] for [specific use case]." Agents read these to decide which products to surface in their comparison set.

What doesn't earn citations: generic blog posts about your category, lifestyle content, brand storytelling pieces, "why we do what we do" essays. These have value for human-facing brand work but they don't move the agent layer one inch.

The smart eCommerce content strategy in 2026:

- **Comparison pages** for every meaningful product cluster — your products vs competitors, ranked honestly, with the criteria explicit
- **"Best X for Y" listicles** owned and authored by you — even if you only feature one of your own products, the listicle format earns citations
- **Buying guides shaped as Q&A** — "what should I look for in X," "how do I choose between Y and Z"
- **Specification pages with schema** — every spec, every variant, every edge case

Where this is going.

The agentic commerce wave hasn't crested yet. Expect three things in the next twelve months:

1. **Universal feeds.** The current fragmentation — separate feeds for Google Merchant Center, Meta Catalog, Amazon, ChatGPT, etc. — will consolidate into 2–3 dominant feed standards. Shopify Catalog is already a leading candidate.
2. **Agent-to-agent negotiation.** Forrester predicts 20% of B2B sellers will face agent-led quote negotiations by end of 2026. Your sales process and pricing pages need to be readable by buyer-side agents.
3. **Loyalty and identity layers.** Google's UCP added Identity Linking in March 2026, allowing loyalty programmes to integrate. Brands with strong loyalty data will get preferential placement.

If you sell online, this is the highest-stakes single chapter in this guide. The protocol decisions, feed-quality investments, and content shape decisions you make in the next two quarters will compound through the rest of the decade.

Custom GPTs & AI agents *that earn their keep.*

The 2025 edition pitched custom GPTs as essential for every business. We were right that the category matters. We were wrong that every business should have one. Eighteen months in, the picture is clearer.

Most of the custom GPTs built in 2024 and 2025 are dead — abandoned, never used, or quietly broken because no one updated them. The successful ones share a profile: they solve a single, narrow, repeatable problem; they're integrated into a workflow people already use; and somebody owns them.

Here's the 2026 framework for deciding whether to build, and what to build if you do.

The four jobs custom AI agents are actually good at.

1 · INTERNAL EXPERT ASSISTANCE

An agent that captures the institutional knowledge of your senior team and makes it available to junior staff or new hires. Trained on your SOPs, frameworks, past project notes, and decision criteria. Used inside your operations, not customer-facing.

Best for: agencies, professional services, knowledge-heavy businesses. Build if: you have 15+ documented SOPs and 50+ resolved client situations to feed it.

2 · SALES QUALIFICATION AND DISCOVERY

An agent on your website that runs a real qualification conversation, gathers the prospect's situation, and either books a meeting or routes the lead based on rules you define. Replaces or supplements the contact form for considered-purchase B2B.

Best for: services with defined ICPs, considered-purchase categories, deals where qualification matters. Build if: you have a clear ICP, a documented sales process, and your team currently spends 30%+ of their time disqualifying leads.

3 · CUSTOMER SUPPORT DEFLECTION

An agent trained on your product documentation, FAQs, and resolved support tickets. Handles 60–80% of L1 questions without human involvement. **Critical:** it must escalate cleanly to a human when it doesn't know.

Best for: SaaS, subscription, eCommerce. Build if: support tickets are scaling faster than headcount and 50%+ of tickets are repeats of previous resolved issues.

4 · SPECIALIST CONTENT PRODUCTION

An agent trained in your voice, frameworks, and style — used internally to produce first drafts of LinkedIn posts, emails, proposal sections, or case studies. Not a replacement for human editing; a force multiplier on the production line.

Best for: teams publishing 10+ pieces per month. Build if: your team has documented voice/style guidelines and produces enough volume to amortise the build.

Notice what's not on this list: "general assistant for our customers." Those don't work. Customers don't want to talk to a generic chatbot version of your brand. They want their specific problem solved. **Build agents around problems, not personas.**

The 2026 build stack.

The platform landscape has consolidated. The serious options for building a custom AI agent in 2026:

- **OpenAI Custom GPTs** — best for simple knowledge-base agents inside ChatGPT. Easy to build, limited customisation, no API integration into your own systems unless you build it yourself. Good for internal use.
- **Claude Projects (Anthropic)** — strong for complex, context-aware reasoning. Best for advisory and analysis use cases. Used heavily in legal, finance, and consulting workflows.
- **Microsoft Copilot Studio** — enterprise option with deep integration into Microsoft 365, Teams, and Dynamics. Right answer for businesses already on the Microsoft stack.
- **Retell AI, Vapi, Bland AI** — voice-first agents for phone interactions. Increasingly used in legal intake, medical scheduling, and inbound sales.
- **n8n + Claude/GPT API** — for sophisticated workflow agents that touch your CRM, calendar, email, and other systems. The flexible option, requires technical capability or partner support.
- **GoHighLevel, HubSpot, Salesforce** with native AI features — embedded agents inside your existing CRM/marketing stack. Lower ceiling, lower setup cost.

The five-step build that actually works.

1. **Pick one job.** Resist the urge to build "the agent." Build the agent for one specific job. The custom GPTs that succeed are narrow.
2. **Write the system prompt before you build.** The prompt is the agent. If you can't articulate exactly what it should and shouldn't do in plain English, no amount of tooling will save you.
3. **Train on real artefacts, not summaries.** Real client emails, real resolved tickets, real proposals. Models learn voice and reasoning from primary sources.
4. **Test with your hardest cases.** The 80% case isn't the test. The 20% edge cases are where agents fail and damage trust. Build the test suite from your hardest historical situations.
5. **Define escalation triggers explicitly.** When should the agent stop and hand off to a human? Write those rules into the system prompt. Vague escalation logic produces awkward customer experiences and brand damage.

What about visibility?

One nuance worth flagging: a well-built custom GPT in the OpenAI directory can become a discovery channel in its own right. Some categories — niche SaaS tools, professional advisory services, specific software documentation — are seeing meaningful inbound from people finding the GPT in the OpenAI ecosystem. It's not yet a replacement for SEO or AI search visibility, but it's a real channel for narrow categories with active GPT users.

Treat your published GPTs as part of your distribution layer, not just as internal tools. Name them well. Describe them clearly. Update them quarterly.

Measuring AI visibility — *the new analytics stack.*

If you can't measure it, you can't budget for it, and "we should be more visible in AI" doesn't survive an executive review. Fortunately the measurement layer matured fast in late 2025 and early 2026. Here's the stack we now recommend.

Three categories of tools matter, and you need at least one from each.

1 • Citation tracking — what AI is saying about you.

These tools run real prompts against ChatGPT, Perplexity, Gemini, AI Overviews, Claude, and Copilot on a schedule. They capture the full response, identify when your brand is mentioned, track citation frequency over time, and benchmark against competitors.

- **Profound** — enterprise-grade. Strong on share-of-voice analysis, sentiment with thematic breakdown, citation-source mapping. Pricing starts around \$499/month and scales meaningfully from there. The most comprehensive option for serious brands.
- **Otterly.AI** — purpose-built for prompt monitoring across AI engines. Strong workspace structure for agencies managing multiple clients. Lite tier at \$29/month, Standard at \$189/month, Pro at \$989/month. The workhorse choice for mid-market.
- **AthenaHQ, Scrunch, Alclicks, Rocketito** — newer entrants with various specialisations. Worth evaluating but the category is consolidating fast.

2 • Bot analytics — what AI is reading from you.

Different question, separate tool. These tools track which AI bots are visiting your site, what they're fetching, and how often. Critical for confirming that your robots.txt and crawl access are configured correctly.

- **Ahrefs Bot Analytics** — currently free in beta. Server-side tracking via Cloudflare, no JavaScript needed. Tracks 12+ bot categories. **Start here.**
- **Cloudflare AI Audit** — for Cloudflare customers. Detailed bot traffic analysis, with controls to block or allow specific AI bots.
- **Server log analysis** — manual but free. Look for OAI-SearchBot, ChatGPT-User, ClaudeBot, PerplexityBot, GPTBot, Bytespider, Amazonbot, Applebot in your logs. If they're not appearing, something is blocking them.

3 • Referral tracking — what AI is sending to you.

The third leg. Even with high citation rates, the click-through to your site is the conversion-relevant metric. GA4 with custom segments handles this perfectly.

Build referral segments for the major AI sources: chatgpt.com, perplexity.ai, gemini.google.com, copilot.microsoft.com, claude.ai. Add the AI Overviews referral pattern (which appears as google.com

with specific URL parameters). Track conversion rate, time on page, and pages per session for each — they will be wildly different from your standard organic search behaviour, almost always in your favour.

The metrics that actually matter.

If you're building an executive dashboard for AI visibility — and you should be — these are the six metrics worth reporting monthly. Anything beyond these is noise.

VISIBILITY METRICS

- 1. Citation rate** — % of tracked prompts where your brand is mentioned, by platform
- 2. Share of voice** — your citation count vs. top three competitors in your category
- 3. Sentiment** — when cited, how favourably

PERFORMANCE METRICS

- 1. AI referral traffic** — sessions from AI sources, MoM trend
- 2. AI referral conversion rate** — typically 4× organic, worth tracking
- 3. Branded search lift** — AI mentions drive branded queries; track that flow

The prompt set that drives the dashboard.

Tracking is only as good as the prompts you track. A typical mid-market prompt set has 30–60 prompts spanning four categories:

- **Branded prompts** (5–10) — "Who is [your company]?" "What does [your company] do?" "[Your company] reviews"
- **Category prompts** (10–20) — "Best [your category] in [location]" "Top [your category] for [use case]"
- **Comparison prompts** (5–10) — "[Your company] vs [competitor]" "Alternatives to [competitor]"
- **Problem prompts** (10–20) — "How do I [problem you solve]" "What's the best way to [outcome you deliver]"

Run these monthly. The data set after three months is what you build strategy on.

TOP4 MEASUREMENT PROGRAMME

For Top4 clients on the AI Search Visibility programme, we run a 50-prompt set monthly across ChatGPT, Perplexity, Gemini, and Copilot, combined with Ahrefs Bot Analytics for crawl confirmation and GA4 for referral and conversion tracking. The dashboard is the main artefact in the monthly review — citation rate, share of voice, and conversion-side performance, with the next 30 days of priorities tied directly to where the gaps are.

What measurement won't tell you.

Two limits to flag honestly. **First**, AI responses are non-deterministic — the same prompt can produce slightly different answers across runs. Otterly's data shows there is a less-than-1-in-100 chance that ChatGPT or Gemini will give the same list of brands across two separate runs of the same prompt. Statistical sampling matters. Single-run citation checks are misleading.

Second, attribution between AI mention and downstream conversion is harder than in classical SEO. Someone might see your brand mentioned in a Perplexity answer, then Google you directly, then convert.

The AI mention drove the branded search; only the branded search shows up in your analytics. Plan for last-touch attribution to undercount AI's actual impact, and use citation rate as a leading indicator rather than relying purely on referral traffic.

Action plan *by vertical.*

The principles in this guide apply to every category. The tactics don't. Below are the playbooks we run for the verticals Top4 sees most often — distilled to the moves that consistently produce results.

Pick the one closest to your business. The 30-day priorities are usually right; the 90-day work is where the compounding starts.

VERTICAL · LEGAL SERVICES

Solo and small firms. The most underserved AI visibility opportunity right now. Big firms can't write geographic-specific content; small firms can. AI Overviews already trigger on 78% of legal queries.

- **30 days:** Attorney + LegalService schema, Google Business Profile completion, Bing Places claim, FAQ pages on top three practice areas with FAQPage schema, location-specific content for each suburb served.
- **90 days:** Wikidata entry for the firm and lead attorney, named methodology for your most common matter type ("the [Firm] [Matter] Method"), case study library with anonymised real outcomes, podcast appearances on local legal-affairs shows.
- **Watch:** Avvo, Justia, FindLaw, local bar association listings — directory consistency is critical in legal.

VERTICAL · MEDICAL & COSMETIC CLINICS

Single-location practitioners. Healthcare AI Overviews now appear on 88% of relevant queries. Patient research happens in Perplexity and ChatGPT before the appointment is booked. Clinics that show up there own the consideration set.

- **30 days:** MedicalBusiness or specific subtype schema (Dentist, Physician, Optometrist, etc.), Healthgrades, RealSelf (cosmetic), or WebMD Care directory listings, condition-specific FAQ pages, before/after gallery with structured Image data.
- **90 days:** Treatment-specific landing pages with detailed Q&A, named patient-experience framework ("the [Clinic] Care Journey"), reviews velocity programme across Google + Healthgrades + RealSelf, telehealth schema where applicable.
- **Watch:** AHPRA compliance and TGA advertising rules — content must be accurate and properly disclaimed; don't optimise yourself into trouble.

VERTICAL · REAL ESTATE

Agencies and individual agents. Buyer journeys now begin in AI search for area, school, and lifestyle research. Agents who appear as the named expert for a specific suburb dominate the consideration phase before the listings even matter.

- **30 days:** RealEstateAgent + LocalBusiness schema, suburb-level guide pages (school zones, transport, demographics, recent sales trends), Person schema for each agent, complete REA and Domain agent profiles.
- **90 days:** Named appraisal methodology ("the [Agency] Local Market Analysis"), monthly market commentary content, video appearances on YouTube with proper VideoObject schema, podcast presence on local-affairs shows.
- **Watch:** Suburb-specific content depth is the moat — generic "Sydney property" pages lose to specialists writing about specific suburbs.

VERTICAL · TRADES & HOME SERVICES

Plumbers, electricians, HVAC, builders. The Local Service Ads layer plus AI search creates a dual-engine opportunity. Trades that name and document their service methodology pull ahead disproportionately because most competitors don't.

- **30 days:** Specific business-type schema (Plumber, Electrician, HVACBusiness), service-area definition, Google Local Service Ads enrolment, photo library on GBP, week-of-completion case studies on each major job.
- **90 days:** Named service framework (the "Five-Point Audit," the "Whole-Home Diagnostic"), suburb-level service pages with real local context, structured pricing or pricing-range pages (rare in trades, hugely effective for AI citations), monthly job-of-the-month content.
- **Watch:** review velocity across Google + Hipages + Oneflare + Houzz; trades reviews concentrated only on Google leave Perplexity and Bing-indexed engines undersourced.

VERTICAL · ECOMMERCE

Direct-to-consumer brands and online retailers. Covered in detail in Chapter 7 — the agentic commerce shift is the single biggest 2026 change for this category. Below is the distilled action layer.

- **30 days:** Product feed completeness audit, Product/Offer/AggregateRating schema on every PDP, robots.txt audit for OAI-SearchBot and Googlebot access, Bing Webmaster Tools setup, agentic-storefront enablement on Shopify if eligible.
- **90 days:** Comparison page library ("X vs Y" pages for every major product), buying-guide content shaped as Q&A, review velocity across Trustpilot + Google + product-specific platforms, GA4 segments for each major AI source with conversion tracking.
- **Watch:** price and inventory accuracy across feed and storefront — agents skip products with mismatches; this is now an operational discipline, not a marketing one.

VERTICAL · B2B SAAS & PROFESSIONAL SERVICES

Considered-purchase technology and consulting. B2B Tech AI Overviews jumped from 36% to 82% of queries in twelve months. The shift in this category is happening in slow motion and most teams haven't caught up.

- **30 days:** Organization + Person schema for founders and lead consultants, complete G2/Capterra/TrustRadius profiles, comparison page coverage ("[Your tool] vs [Top 3 competitors]"), case study library with metric-driven outcomes.
- **90 days:** Named methodology with full schema treatment, original research or proprietary benchmark report, podcast appearances on category-relevant shows, LinkedIn thought-leadership cadence from named people (not the brand account).
- **Watch:** LinkedIn is the most-cited domain for professional queries across AI Overviews, AI Mode, ChatGPT, Copilot, and Perplexity — your people's LinkedIn footprint matters more than your company blog.

The cross-vertical truths.

Strip away the vertical specifics and you're left with the same five priorities, in roughly the same order, every time.

- 1. Get the technical foundation right.** Schema, robots.txt, indexation, directory consistency. None of this is glamorous; all of it is binary — either it's correct or you're invisible.
- 2. Sharpen positioning to a defensible niche.** The narrower, the better. Cumulative wins on specific queries beat broad wins on generic ones every time.
- 3. Name and document your methodology.** One signature framework, schema-marked, repeated everywhere, refreshed annually.
- 4. Build distribution beyond your own domain.** Directories, podcasts, reviews, third-party content, honest community participation. The corroboration layer is the hardest to fake and the most valuable to compound.
- 5. Measure consistently.** Citation rate, share of voice, AI referral traffic, conversion. Monthly review, quarterly priority reset.

If you do those five things — and you do them with intent rather than as items on a checklist — AI visibility takes care of itself. The teams winning in 2026 aren't running clever tactics. They're running these basics with discipline.

The 2026 winners aren't the businesses with the cleverest AI strategies. They're the businesses with the most disciplined application of unsexy fundamentals.

WHAT WE SEE ACROSS OUR CLIENT PORTFOLIO

One more vertical-agnostic point.

Whatever your category, the single highest-leverage move you can make in the next 30 days is the one most teams skip: **do the basic audit honestly**. Pull your schema, your robots.txt, your directory listings, your bot logs. Compare against the standards in this guide. Most teams discover they're failing on three to

five fundamentals they assumed were in place. Fixing those is usually the fastest 90 days of progress they'll see all year.

The work isn't complicated. The discipline of actually doing it — across your whole team, on a defined timeline, with someone owning each piece — is the part that separates the businesses that win the AI visibility race from the ones that talk about it.

The next training cycle *is* always coming.

If you read the 2025 edition, you might remember the closing line: "the next training cycle is coming. Your future AI visibility depends on the authority you build today." That's still true. It's just more urgent.

The pace of change has accelerated, not slowed. AI Mode shipped publicly. Agentic commerce protocols launched. Citation tracking became a real category. Local AI visibility became a viable channel for businesses that couldn't afford traditional national SEO. And every model — GPT-5, Claude Opus 4.7, Gemini 2.5, the open-weight models gaining ground — gets retrained on a regular cadence that bakes the current state of the web into the next generation's view of who matters.

If you are not visible across the corroborating sources during a training window, you are not in that model's worldview when it ships. There is a real cost to being late, and it compounds. **The window to establish your authority for the 2027 models is open right now.**

Six things to do before the end of the quarter.

1. Audit your robots.txt for AI retrieval bot access — single biggest 5-minute win
2. Validate your schema across the top ten pages on your site
3. Sharpen your positioning statement to a specific audience, outcome, and methodology
4. Stand up a citation tracking tool — Otterly, Profound, or equivalent — and run your first prompt set
5. Confirm your Bing Webmaster Tools and Bing Places listings are claimed and clean
6. Schedule the 90-day plan with someone owning each piece — not collectively, individually

None of this is novel. None of it is hard. The hard part is doing it consistently while every other quarter brings a new shiny tactic to chase. The advice that ages well in this market is unfashionable advice: get the fundamentals right, narrow your positioning, build a distinctive methodology, distribute it widely, measure it honestly, refresh it consistently.

Do that and the AI visibility takes care of itself.

Don't, and the next training cycle ships without you.

ABOUT THE AUTHOR

Michael Doyle.

Founder and Managing Director, Top4 Technology. 25+ years in digital, originally Scottish, mostly Australian, increasingly Indonesian.



25+

YEARS IN DIGITAL

200K+

BUSINESS PROFILES MANAGED

65K+

PAYING SUBSCRIBERS

Michael started in digital in 1999 and founded his first agency in 2001, in the early days of the Australian internet industry. Top4 Technology grew out of that foundation into a digital marketing and technology platform now managing 200,000+ business profiles for 65,000+ paying subscribers across Australia, Indonesia, and Asia-Pacific.

Top4 Technology is a wholly-owned subsidiary of White Pearl Technology Group AB, a Swedish-listed digital transformation group on Nasdaq First North (ticker: WPTG). The Top4 team operates across LLM-led SEO, multi-location and government digital, enterprise CMS delivery (Kentico), AI automation, and CRM & marketing automation. The group is a Google Premier Partner, Meta Business Partner, Kentico Solution Partner, and SiteMinder Reseller.

What this guide is built on.

This is not theory written by someone watching from the sidelines. Every framework, every data point, every recommendation in this book is drawn from real client work — campaigns the Top4 team has run, audited, refreshed, or rescued in the last twelve months across enterprise, government, multi-location, and SMB engagements.

If you'd like to talk to Michael directly about your situation, the contact details are on the back cover. For the day-to-day on AI search, automation, and the Top4 client work that informs everything in this book, follow him on LinkedIn — he writes most days.

FIND MICHAEL

LinkedIn: [linkedin.com/in/michaelcdoyle](https://www.linkedin.com/in/michaelcdoyle)

Email: michael@top4.com.au

Top4 Technology: top4marketing.com

Top4 Platform: top4.online · top4.com.au

White Pearl Technology Group AB: whitepearltech.com

If you found this guide useful, share it with one person on your team who needs it. Then tell Michael what you'd like covered in the 2027 edition.

Build it with *Top4.*

If this guide gave you a clearer view of what 2026 looks like, the next move is simple: get the audit done, build the plan, and start the work.

If your website exists but doesn't generate leads, that's not a marketing problem — it's a systems problem. Top4 Technology builds the systems that turn AI search visibility into measurable demand.

Selected clients include BMW, Ajinomoto, Dulux Group, Caroma, JELD-WEN, B&D Doors, The Blue Space, NSW Government, Sydney Children's Hospitals Network, and 65,000+ subscribers across the Top4 platform. Google Premier Partner. Meta Business Partner. Kentico & SiteMinder Partner.

THREE WAYS TO START

1 The AI Search Visibility audit

A 14-day deep audit of your current AI visibility, citation rate, technical foundations, and competitor benchmark. Delivered as a 30-page report with a 90-day prioritised action plan. Fixed scope, fixed price.

2 The AI Search Visibility programme

Ongoing monthly engagement covering citation tracking, content production, schema and technical work, directory and review-velocity programmes, and quarterly strategy reviews. Tiered from \$299/month for SMBs to enterprise pricing.

3 The Top4 Demand Engine

The full system — AI search visibility, traditional SEO, paid media, web development, and automation — delivered as a coordinated programme for businesses ready to consolidate their digital presence under one provider.

TALK TO MICHAEL

michael@top4.com.au

[LINKEDIN.COM/IN/MICHAELCDOYLE](https://www.linkedin.com/in/michaelcdoyle) ·
[TOP4MARKETING.COM](https://www.top4marketing.com)

GROUP · THIS EDITION

White Pearl Technology Group AB

NASDAQ FIRST NORTH · WPTG · VOLUME II ·
APRIL 2026

