

Impact GEO

A Content Ecosystem Manifesto
for the AI Search Era



A theoretical guide for GEO practitioners,
content creators and platform governors

hello@reasongo.com (v0.1.2)

Contents

Preface: Why this book exists

Chapter 1: The paradigm shift — How AI search reshapes information distribution

Chapter 2: The core principles of Impact GEO

Chapter 3: The white-hat GEO method — From tactics to principles

Chapter 4: Incentives in the content ecosystem: Scoring, reputation and certification

Chapter 5: Governance and standards: Building a trusted AI information ecosystem

Chapter 6: The action roadmap — Start doing Impact GEO today

Appendix 1: The Impact GEO self-check list

Appendix 2: Sources cited in this book

Preface: Why this book exists

On March 15, 2026, a smart band was "born".

It has no factory, no engineers and no prototype. It exists only in a little over a dozen advertorials written by a machine and pushed out in bulk. Two hours after those pages went online, two of the largest AI models were already citing it near the top of their recommendations, presenting it to users as a real product and a top pick. When CCTV's 3·15 gala exposed the operation, a lot of people understood something for the first time: an AI answer may not be the answer the model judged best. It may be an answer someone fed it, systematically.

That should be unsettling. It is also the most instructive fact in this book.

It shows how the meaning of influence has changed. An advertorial used to reach the people who happened to read it. Content that AI cites comes back out as fact in countless later conversations, over and over. The ecosystem effect of a single page has grown far beyond its readership. Good content keeps feeding the information environment. Bad content keeps polluting it.

What this book takes up is what that amplification means, and how to push it somewhere healthier.

GEO (generative engine optimization) is a technique, and techniques take no side. The same method can help a careful piece

of content get understood and cited by AI, or help a fictional product become a fact within hours. When the first happens, the content ecosystem improves. When the second happens, it degrades. The tools are identical. The difference is what practitioners choose to do with them.

Calling that a matter of personal morals is cheap. The harder question is this: **why is the wrong choice so often the more "rational" one under the current incentive structure?**

If poisoning costs almost nothing and pays back in hours, while compliant optimization takes months and shows slowly, no volume of moral appeal stops bad money driving out good. This is a design problem, not a character problem. To change what individuals do, change what rewards them.

The aim of this book is to give "change the incentive structure" a theoretical framework that is clear enough to actually build on.

Its central claim fits in one line: **ecosystem health is the long-term moat for GEO practitioners.** Trust in AI search fell 28 percentage points in a single year. The citation weight of authoritative sources is rising across the board. The first industry standard has just taken effect, and the regulatory framework is taking shape as it does. At a turning point like this, doing what is good for the ecosystem stops being a moral sacrifice and becomes strategic rationality. Once users stop trusting AI answers and start doubting the whole information environment, the first to be eliminated are not the users. They are the brands and the service providers whose business depends on that environment being trusted.

The book is written for three readers. **GEO practitioners** will find a method for moving from a traffic mindset to a trust-asset mindset. **Content creators** will see why content that is true, transparent and verifiable is the best long-term investment available in the AI era. **Platform governors and policy researchers** will get a framework for turning ecosystem health from an abstract ideal into mechanisms that run.

Six chapters follow. Chapter 1 looks at the change AI search brings to how information gets distributed. Chapter 2 establishes the four core principles of Impact GEO. Chapters 3, 4 and 5 work through the method, the incentive system and governance in turn. Chapter 6 gives each role an action roadmap. Each chapter closes with a short list of its main points, so you can read for the argument or look up a single question.

This book is not a sermon and not a warning shot. What it wants is a practitioner standing at the fork between doing it right and doing it fast being able to see that **doing it right is becoming doing it fast**.

That is why the book was worth writing, and why it may be worth your time.

Let's begin.

Chapter 1: The paradigm shift — How AI search reshapes information distribution

The way people get information has changed.

It changed fast. In 2025, 82% of consumers found AI search more helpful than traditional search. One year later that number was 54%. Nearly a third of them reversed their judgment after using it themselves. What they found is that an AI answer may not be the answer the model judged best. It may be an answer someone fed it, systematically.

This is not a technical fault. It is a problem of how incentives are arranged.

To see it clearly, look at the three turning points of the past two years.

The first turning point: search logic breaks

A traditional search engine works like this: the user types keywords, the system returns a list of links, the user filters and clicks. A brand reaches people by being found in search. Whoever ranks first gets the click.

Generative AI search changes the logic. The user asks in natural language, the AI synthesizes one answer out of several sources, and the user takes that answer. The route from brand to person shifts from being found in search to being cited inside the answer. What the user sees is no longer "what a brand says about itself". It is "what AI takes to be true".

What makes this shift deep is not the technology. It is that it redistributes where influence sits. In traditional search, influence sits in rankings. In AI search, it sits in citations. Once a piece of content becomes a source for an answer, its information leaves the model as fact, again and again, in every later conversation that touches the topic.

A page's influence no longer depends on how many people clicked it. It depends on whether AI is willing to treat it as a trusted fact. When being cited replaces being ranked as the object of competition, one structural change follows: manipulating rankings has a visible price, SEO work and paid placement, while manipulating citations cost practically nothing at first.

The second turning point: the industry's distortion

In March 2026, CCTV's 3·15 gala exposed a complete abuse pipeline. The recipe is embarrassingly simple: use AI writing tools to generate fake product reviews and invented industry rankings in bulk, push them onto the web through press-release distribution platforms, and let AI crawl them at retrieval time. A product that

never existed went from advertorial to "the standard answer" in two hours.

This is not one bad actor. Black-hat GEO has a settled business model: hundreds of AI-generated articles a day, monthly fees starting around ¥1,000, well under a tenth of a yuan per article. Meanwhile the companies actually spending money to build trustworthy content are buried under the junk. A brand can invest years in reputation and lose most of it to a single fabricated claim produced for pocket change.

This is a textbook tragedy of the commons. The AI corpus is a shared resource, and individually rational behavior grazes it to death. When the marginal cost of manipulating citations is near zero, the marginal return is high, and the negative externality is paid by the whole information ecosystem, no volume of moral appeal stops bad money driving out good.

The third turning point: the ecosystem pushes back

But the ground is moving.

Trust in AI search fell 28 percentage points in one year. More than half of consumers now doubt the reliability of AI answers; 86% cannot fully trust AI-generated content and go check the primary sources themselves when the information matters. At the same time the weight of authoritative sources in AI citations is rising across the board. Citation counts for Wikipedia, the U.S. National Institutes of Health (NIH) and similar sources have grown sharply, and the major

AI platforms are deliberately cutting back: dropping low-value content, raising the weight of state media and authoritative institutions.

In other words, the market is voting with its feet. Once users stop trusting AI answers and start doubting the whole information environment, the first to be eliminated are not the users. They are the brands and service providers whose business depends on that environment being trusted. In a market where users switch engines after AI gets a citation wrong twice in a row, the cost of source quality is charged straight to retention.

Ecosystem health is becoming the competitive moat.

This is where the book starts

GEO as a technique takes no side. The same method can help careful content be understood and cited by AI, or help a fictional product become a fact within hours. When the first happens, the content ecosystem improves. When the second happens, it degrades.

Calling that a matter of personal morals is cheap. The real question is why, under the current incentive structure, the wrong choice is so often the more rational one. If compliant optimization costs a lot and shows results slowly, while poisoning pays back fast and carries little risk, then changing what individuals do means changing what rewards them.

The book's central claim fits in one line: **ecosystem health is the long-term moat for GEO practitioners**. At a turning point marked by falling trust, rising weight for authoritative sources and a standard that has just taken effect, choosing what is good for the ecosystem stops being a moral sacrifice. It becomes strategic rationality.

This is not a manual for "ethical GEO". It answers a more practical question: **how does doing the right thing become the smarter commercial choice?**

Key takeaways

- AI search moved the object of brand competition from being found to being cited. A page's influence no longer rests on clicks; it rests on whether AI is willing to output it as a trusted fact.
- When manipulating citations is nearly free and the whole information ecosystem pays for it, the gap between individual and collective rationality cannot be closed by moral appeal.
- Falling trust in AI search and rising weight for authoritative sources are one clear market signal: ecosystem health is the long-term moat.
- Impact GEO does not substitute ethics for technique. It makes ecologically sound practice the smarter commercial choice.

Chapter 2: The core principles of Impact GEO

Chapter 1 laid out the problem in full: AI search has magnified what a piece of content can do, and the current incentive structure makes poisoning cheaper than building. The more basic question is what a different incentive structure should look like.

This chapter proposes four principles. They are not moral commandments. They are a systematic answer to a narrower question: which GEO practices win over a long enough horizon. Each one tracks a structural change already under way in AI search.

Principle 1: Source transparency — content must be traceable and verifiable

The deepest difference between AI search and traditional search is the relationship between the user and the information. In traditional search the user sees a list of links and judges each source personally. In AI search the user sees one synthesized answer, and does not know which sources went into it, which were given more weight, or which had an interest in the outcome.

The core requirement of source transparency is this: when AI cites content, that content's origin can be traced accurately; and when AI tells a user something is a fact, the user can check what it rests on.

This is not a moral ask. It is an algorithmic rule already being enforced. The mainstream large language models run explicit source tiering. Government websites, mainstream media, industry regulators and authoritative academic databases are classified as top-tier trusted sources, crawled first and cited with a standing advantage. Independent accounts, advertorials and content from unlicensed sites are demoted and read with caution. The higher a source's tier, the more likely AI is to cite it.

So source transparency is not only the right thing to do; it is the thing the algorithms recognize. Where a brand's content has clear origins, a complete evidence chain and visible interests behind it, AI systems return a higher citation weight. Where origins are vague and untraceable, the content may get crawled today and will still be flagged as low quality in later cross-checks, then phased out.

T/CAPT 026—2026, the association standard that took effect in August 2026, sets out an A/B/C/D four-tier credibility rating and requires brand knowledge bases to keep three zones separate: a fact zone, an opinion zone and a marketing zone. Its full title is *Specification for Trusted Information Dissemination and Information Ecosystem Governance in Generative Engine Optimization (GEO)*. The logic behind that design is source transparency itself. Verifiable objective fact is held apart from judgment and from promotion, so that AI can tell "this is a fact" from "this is an opinion" from "this is an ad" when it cites.

For a GEO practitioner the principle translates into three working points. In production, every core factual claim carries an explicit source. In organization, factual content and marketing content are

managed in separate zones. In monitoring, every piece of content AI cites can be traced back to its primary source.

Principle 2: Truth first — accurate content is the best optimization

In the traditional SEO era the logic was to make search engines believe your content was good. In AI search the logic changes: the model judges the content itself and cross-verifies it against other sources. Large language models both cross-check and remember. Exaggeration and false claims may be crawled for a while, then get marked as low quality and dropped.

The principle states it plainly: under AI's citation mechanism, accurate content is not the better option. It is the only option that lasts.

The data supports that reading. Weight for authoritative sources in AI search is rising across the board. May 2026 figures show Wikipedia's AI citation volume up 55.2% month over month and the U.S. National Institutes of Health (NIH) up 41.3%, moving it into the top three cited sources overall. The research conclusion is blunt: AI models are actively looking for authoritative sources.

At the same time the short-term payoff of poisoning is being cut back deliberately. After the 2026 Spring Festival, DeepSeek compressed the sources it reads closely from 10-15 down to 4-5, and citation rates for mass-distributed press-release content fell sharply. The odds of a single-source claim being blocked reach 95%, and content that contradicts itself across platforms is 82% less likely

to be featured. The platforms are subtracting on purpose: dropping low-value content and raising the weight of content that is accurate and authoritative.

Truth first has three operational layers. **Content:** every thousand words should carry at least five verifiable statistics and at least two authoritative citations, with the core conclusion stated in the opening paragraph. **Structure:** the argument should have a clear chain of facts so AI can identify and extract the evidential relationships. **Verification:** a third party should be able to confirm it independently, rather than taking the brand's own word for it.

Truth first does not mean giving up marketing copy. It means keeping the two distinguishable: marketing should be recognizable as marketing, and factual claims must survive verification. The standard's rule that factual claims in marketing copy must match the fact-zone version is exactly this principle written into regulation.

Principle 3: Ecosystem health is the moat

The first two principles say what to do. The third says why keep doing it.

The argument is short. In a market where trust in AI search is falling, source quality stops being a moral bonus and becomes a competitive factor that directly determines retention.

The trend is visible in the numbers. In 2025, 82% of consumers found AI search more helpful than traditional search; by 2026 that was 54%, a fall of 28 percentage points in one year. More than half

of consumers now doubt the reliability of AI answers, and 86% go check the primary sources themselves when the information matters.

What matters more is the change in behavior. 72.4% of users switch engines after AI gets a citation wrong twice in a row. For an AI search platform, the cost of source quality is therefore charged straight to retention. A product polluted by low-quality content loses users who vote with their feet; a product that keeps producing trustworthy answers earns loyalty. The same logic holds for brands and service providers: when your content becomes a source behind answers people trust, you gain more than exposure. You gain standing.

The direction of industry standards confirms it. The stated goal of T/CAPT 026—2026 is to move GEO away from the technical pursuit of exposure, mentions and rankings, and toward building verifiable factual assets, controlling dissemination for compliance, and governing the information ecosystem — in its words, shifting the ecosystem from traffic-driven to trust-driven. The China Advertising Association (CAA) and the China AI Industry Development Alliance (AIIA) are pushing trusted GEO ecosystems in parallel, and the first companies to pass the GEO service trust assessment have been certified.

Taken together, these moves point one way: ecosystem health is turning from an aspiration into an entry requirement. Once the regulatory framework matures and platform algorithms keep tilting toward credible sources, the brands and providers that built their content systems early will hold a structural advantage.

Principle 4: Positive externality — every optimization should add to the ecosystem

The first three principles are about getting your own content cited correctly. The fourth widens the frame from the self to the whole ecosystem: every optimization should leave the AI information ecosystem better off while it works for you.

Positive externality is an economics term for an act that benefits others without them paying for it. In the AI content ecosystem it looks like this: a careful, verifiable piece of content does not only help AI answer questions about your brand accurately. It adds trustworthy knowledge about the subject to the AI corpus. Later users asking related questions, and later models retrieving related topics, get better answers because that page exists.

Poisoning produces the opposite. A fabricated page misleads the reader in front of it and pollutes the corpus, so countless later users and conversations inherit the error. The Apollo-9 smart band case shows the real radius of a fake advertorial: it is not the people who read it, it is every later conversation in which AI repeats it as fact.

So the principle's requirement is about measurement. When you evaluate GEO performance, do not only count how often you were cited. Ask whether your content made the AI information ecosystem more trustworthy.

That has concrete implications at three points. **Production:** make content with public knowledge value, not just brand praise. A page that answers questions about the brand and hands the industry

something reusable creates brand value and ecosystem value at once. **Measurement:** add an ecosystem contribution dimension, asking whether AI's answers on the topic became more credible after they cited you. **Strategy:** put ecosystem health into the long-term competitive model. In a market where everyone poisons, no brand is insulated. In a market where everyone builds credible content, every participant benefits.

The fourth principle closes the loop with the other three. Source transparency makes content traceable. Truth first makes it survive inspection. Ecosystem health as a moat makes good practice pay. Positive externality makes sure each optimization adds value to the ecosystem rather than withdrawing it. All four lead to the same conclusion: Impact GEO is not the ethical version of GEO. It is the long-termist version, in step with how algorithms evolve, regulation tightens and trust has to be rebuilt.

Key takeaways

- **Source transparency** is a citation rule AI algorithms already enforce systematically. The A/B/C/D rating and the three-zone requirement turn "traceable origins" from a moral appeal into an engineering standard.
- **Truth first** is the only strategy that lasts under AI's cross-verification mechanism. Citation weight for authoritative sources is rising across the board, and the short-term effect of poisoned content is being cut back by the algorithms themselves.

- **Ecosystem health is the moat:** trust in AI search fell 28 percentage points in a year, and 72.4% of users switch engines after being served a wrong citation twice. The cost of source quality is charged straight to retention, so ecosystem-friendly practice stops being a cost and becomes an asset.
- **Positive externality** asks practitioners to add trustworthy knowledge to the AI information ecosystem while pursuing their own citations. Every optimization should make AI's answers more credible, not less.

Chapter 3: The white-hat GEO method — From tactics to principles

Chapter 2 set out the four principles of Impact GEO: source transparency, truth first, ecosystem health as the long-term moat, and positive externality. The next question is more concrete. How do principles turn into practice?

This chapter gives a complete methodological framework. It does not try to list every tactic. It sets out one closed loop running from mapping assets to continuous testing, so that white-hat GEO becomes an engineering practice that can be copied, verified and sustained.

3.1 From being ranked to being cited: the shift in mindset

Before the method, one premise. GEO is not a newer version of SEO. It is a different optimization logic.

The core goal of traditional SEO is to place a page higher on the results list. The core goal of GEO is to get content into the answer AI generates and be cited there as a trusted source. The difference is not how complicated the techniques are. It is the standard of success: SEO is judged by ranking, GEO by citation.

The first adjustment this demands is that **keyword density is a negative signal for citation**. Data shows that pages above 8% keyword density are 17% less likely to be cited by AI than pages at a natural 2%–3%. Models retrieve by semantic similarity, not keyword matching. Stacking keywords does nothing; it lowers the information entropy of the text until AI reads it as low-quality template material.

The second adjustment: **your competition is no longer other websites. It is how the model represents knowledge internally**. Content has to exist in a form AI can understand, trust and cite directly. How content is organized, structured and evidenced matters more than how well it is written.

3.2 The four-step loop of white-hat GEO

The engineering framework is a loop of four steps: **map assets** → **build sources** → **refine expression** → **keep testing**. The steps are not a line; they cycle. Every round of test results feeds back into the asset inventory, the source plan and the expression choices.

Step 1: Map assets

Everything rests on verifiable assets. Before asking how to be cited, answer a plainer question: what is worth citing?

This step is a systematic inventory of the verifiable information a brand owns: official qualifications, certifications, product whitepapers, real case data, third-party test reports, patents and

intellectual property. Each item becomes a structured data asset with its source, date and scope of application recorded.

The rule for this step: **inventory only what a third party can independently verify**. Data with no traceable source, claims that were never tested, vague "industry-leading" wording. None of it belongs on the list. When models cross-check across sources, unverifiable information is marked low quality, and it can pull the whole brand into demotion.

The output should include an entity table: the core entities (brand, product, person, event) with their names, aliases, attributes, relationships and definition boundaries, so every piece of content describes the same entity the same way.

Step 2: Build sources

An asset has to sit where AI can crawl it and already trusts it.

Trust in sources is tiered. Government websites, mainstream media, industry regulators and authoritative academic databases are high-weight sources and get cited first; the brand website and its official accounts sit in the middle; independent accounts and advertorials are demoted. Building sources means publishing verifiable assets into high-weight channels first, so that retrieval can actually find them.

One point is easy to miss: **not everything should be published on the brand's own website**. Being cited as a case in an industry research report carries far more credibility than promoting yourself on your own site. So building sources is not only publishing. It is

also taking part in industry research and supplying data and cases that third parties can cite.

Check for three kinds of conflict at the same time: factual conflicts, where the same event carries different dates or figures across sources; definitional conflicts, where the same concept is defined differently on different pages; evidential conflicts, where cited sources reach opposite conclusions and nothing says so. Merge what can be merged. For what cannot, state the conditions of application.

Step 3: Refine expression

With assets and sources in place, content has to be presented in the form AI understands and cites most easily. This happens on three layers.

Structure layer: make extraction fast. Models prefer to lift information out of lists, tables, FAQ structures and definition paragraphs rather than summarize it out of continuous narrative. One paragraph carries one fact, and the core conclusion goes in the opening paragraph, because crawling may only take the first part of the page.

Markup layer: tell AI what the page is. Schema.org markup (FAQ, HowTo, Article, Product, Organization and so on) lets AI identify the page type and the relations between pieces of content. Content without structured markup loses about 47% of its weight in AI citation on average. FAQ markup lets AI quote your content directly

when the question is how to do something; HowTo markup makes the clearest step-by-step content the preferred pick.

Evidence layer: back every claim with verifiable data. Content with cited sources is 34.4% more likely to be cited, content with statistics 32.1%, content with direct quotes 29.7%. A workable writing baseline: every thousand words carries at least 5 verifiable statistics and at least 2 authoritative citations, with the core conclusion stated up front.

One principle deserves emphasis here: **the quality of a citation matters more than the number.** Academic papers, official documents and primary data beat content that independent accounts have copied from one another. Keep source, date and measurement consistent when you cite, so nothing reads as taken out of context.

Step 4: Keep testing

White-hat GEO is not a project you finish. The metrics to watch are whether AI retrieves the content, whether it cites it, whether entities resolve correctly, and whether the answer absorbed the key information.

A workable routine: fix 3 to 5 target queries, test them on several AI platforms on a schedule, record how the citation sources change, then read the results. Which content got cited, which did not, and what structural features the cited content shares. Iterate on assets, sources and expression from what the tests say.

3.3 Three red lines

The method works only inside a clear boundary: what may be done and what may never be. These three red lines are the floor for every Impact GEO practitioner.

Red line 1: truthfulness. No fabricated data, no invented endorsements, no poisoning of public corpora. Every factual claim must trace back to the official site or an authoritative source and survive cross-checks. This is the dividing line between white hat and black hat.

Red line 2: compliance. No mass fake distribution, and AI-generated content must be labeled where required. The CAC's "Qinglang" campaign against AI application chaos, launched in 2026, explicitly lists using GEO technology for malicious marketing as a form of AI data poisoning. Compliance is not a cost. It is the entry requirement for operating long term.

Red line 3: user value. No keyword stuffing to fool retrieval, no inflated volume to hold position. Content should be organized around real user intent and answer what people actually want answered. User value is why content keeps getting cited. The model's training objective is to answer user questions, and content that ignores that is eventually dropped by the algorithm.

Three black-hat operations sit directly opposite the red lines: **fabricated citation sources** (junk sites, site networks and AI content farm pages generated in bulk, manufacturing the appearance that a brand is widely mentioned), **manipulated engagement data** (bots inflating clicks and dwell time to fake

demand), and **hidden or deceptive content** (cloaking, where crawlers are served something different from what users see, or brand terms unrelated to the article are stacked on the page).

These three share one feature: their effect depends on the platform not noticing. Once identified, citations and traffic fall back fast, and the digital assets a brand accumulated can be wiped out. Black hat pays in the short term and costs irreversibly later.

3.4 Where the paths diverge

An anonymized comparison between two mid-tier brands shows what the two paths do over time. The brand that chose black hat generated content with fake endorsements in bulk and pushed it out across channels. Its mention rate rose with fluctuations in the short term. Half a year later the source anomalies were detected, and AI citations for its core keywords went to zero. The brand that chose white hat spent two months first on its verifiable assets, filling out the official site and the whitepaper, then kept refining expression and testing. Half a year later its citation rate in AI answers for the core category had gone from near zero to a stable place in the front row.

In the short term the two paths look similar. In the long term the difference is structural. Black hat's return curve is a line that rises and then breaks. White hat's is a curve that rises slowly and keeps rising. Which path to take comes down to one question: do you want three months of exposure, or three years of assets?

The core logic of white-hat GEO fits one sentence: **trade structured verifiable assets, high-quality sources and clear expression for stable, verifiable citation by AI**. That aligns with regulation curbing corpus poisoning and malicious marketing, with algorithms raising the weight given to source quality, and with users rebuilding trust in AI search.

Because those three directions agree, white-hat GEO is not making do under constraints. It is moving with the trend. When regulation tightens, when platform algorithms keep tilting toward credible sources, and when users become immune to low-quality content, the brands that finished their asset inventory and source building early will hold a structural advantage.

Key takeaways

- The shift at the heart of GEO is from being ranked to being cited. Keyword density works against citation, and how content is organized, structured and evidenced matters more than how well it is written.
- The four-step loop is: **map assets** (inventory what can be verified) → **build sources** (put verifiable information where AI already trusts it) → **refine expression** (clear structure, complete markup, sufficient evidence) → **keep testing** (iterate on what the data says).

- The three red lines do not move: **truthfulness** (no fabricated data), **compliance** (no mass fake distribution), **user value** (no deceiving retrieval). The black-hat operations opposite them, fabricated citation sources, manipulated engagement data and hidden or deceptive content, work briefly and go to zero eventually.
- White-hat GEO trades verifiable assets for stable citation by AI, and that trade runs in the same direction as regulation, algorithm evolution and user demand. Choosing white hat is not a moral sacrifice. It is strategic rationality over the long term.

Chapter 4: Incentives in the content ecosystem: Scoring, reputation and certification

Chapter 3 gave the full white-hat method: map assets, build sources, refine expression, keep testing. A method answers how. It leaves a more basic question open: **why would good content be rewarded and bad content punished?**

If AI cited everything indiscriminately, white-hat work would have no payoff and black-hat work would stay cheap. A healthy ecosystem needs a mechanism that pays for doing the right thing. This chapter looks at three layers of incentive design: **content scoring** gives better content a higher citation weight, **creator reputation** turns credibility accumulated over time into a portable asset, and **authoritative certification** gives trusted sources a mark people can recognize.

The three build on each other: scoring is the base, reputation is the accumulation, certification is the mark. Without scoring there is nothing to accumulate reputation from. Without reputation, certification has no evidence behind it. Without certification, neither the score nor the reputation can be read quickly by users or by AI systems.

4.1 Source tiering: the layer under scoring

Content scoring presupposes source tiering. Before judging the quality of a piece of content, an AI system needs a standard for deciding which sources deserve more trust.

The A/B/C/D four-tier credibility rating set up by T/CAPT 026—2026, the association standard that took effect in August 2026, is the most systematic tiering framework so far. Government departments and authoritative media sit in the top tier; independent accounts and advertorials are demoted by design.

The logic running through the rating is verifiability. The higher a source's tier, the easier its information is to check and trace independently. Data published on a government website, reporting from authoritative media, a peer-reviewed paper: all of them have a clear path back to origin and someone accountable behind them. Independent accounts usually have neither.

Tiering is not a static label. It is a dynamic weight. At retrieval and citation time, platforms assign an initial weight by tier, then judge the content itself on top of that. Industry research puts the gap plainly: content from authoritative media is 8 to 12 times more likely to be cited by AI than content from an ordinary independent account, and national-level media carry more than five times the citation weight of local media.

That tells you how the mechanism actually works. Tiering is not a veto; it is a weight adjustment. An independent account that keeps producing high-quality, verifiable content will see its citation odds

rise as well. It simply starts lower and needs more accumulation to reach the same place.

4.2 Content scoring: five dimensions decide citation odds

With tiering as the base layer, content scoring asks a narrower question: within sources of the same tier, which pieces deserve citation?

Pulling together what the major AI platforms publish and what industry testing has measured, current scoring comes down to five dimensions.

Media tier and endorsements (about 25%). This measures the authority of the source itself: media level (national, provincial, industry, independent account), ICP license, supervising body, track record. It maps directly onto the tiering system and is the foundation layer of scoring.

Content expertise (about 22%). This measures the content: whether the author holds professional credentials, whether the piece has depth, whether its citations are properly made, whether there is a fact-checking process, whether it satisfies the E-E-A-T signals (experience, expertise, authoritativeness, trustworthiness). Within this dimension, bylined authorship is becoming an increasingly strong signal. In May 2026 Google changed how AI Overviews displays citations and made named authors a search ranking variable. An author with a verifiable professional

background and a real byline gets cited by AI far more often than anonymous content or content signed with something vague.

Site health (about 18%). This measures the technical quality and compliance record of the site: domain age, HTTPS configuration, mobile responsiveness, page speed, violation history, whether structured data is deployed. The reasoning is that a technically sound site maintained steadily over years produces content worth trusting.

Social and citation signals (about 20%). This measures how much recognition the content gets inside the ecosystem: how often high-weight sources cite it, shares, expert endorsement, the quality of user comments. It imports the logic of peer review. When several high-weight sources cite the same piece, its credibility gets confirmed at the ecosystem level, over and over.

Content freshness (about 15%). This measures time: publication date, update frequency, how fast the site responds to events, how current the data is. AI engines rarely cite content that has not been updated in about two years, and the correlation between metadata freshness and citation rate reaches 0.68.

Together the five produce a composite score, and the score sets a piece's position in the citation candidate set. The point is not an absolute number. The point is a shared quality language between the people who make content and the systems that cite it: there are readable signals for what is likely to be cited and what needs work.

4.3 Creator reputation: from one score to a long-term asset

Content scoring evaluates this piece. Creator reputation evaluates this person or this organization.

The difference matters. A single piece can be packaged carefully to score well. Standing reputation cannot; it has to be accumulated. Reputation mechanisms are designed so that doing the right thing for a long time pays better than doing something fast right now.

A reputation model usually combines several components. **Output success rate** is the most basic one: what share of a creator's published content AI cites correctly. **Repeat collaboration rate** asks whether brands and partners keep coming back to the same creator, which is a market signal rather than a platform judgment. **Engagement quality** looks at what users do after the content is cited: likes, follow-up questions, clicks through to the cited source. Platforms feed those behaviors back as reinforcement signals that change the source's weight. **Professional credential verification** is a bonus term: where a creator's identity and background can be verified independently, the content starts at a higher initial trust.

One design principle is worth pausing on: **reputation should decay over time**. A creator who stops producing good work should not hold a high rating forever. Some systems already apply a week-by-week decay after a stretch of inactivity, so that reputation reflects what the creator can do now rather than what they once did.

The other key property is **portability**. Reputation should not live inside one platform. It should be a digital asset that other platforms can verify, which means it has to rest on traceable behavior records and verifiable outcome data rather than on one platform's label. Once reputation built on one platform is recognized on another, the reward for good content stops being limited to a single channel and covers the whole ecosystem.

4.4 Authoritative certification: making trust visible

Scoring and reputation answer how an AI system judges quality. Certification answers how a reader recognizes it.

Trusted content with no visible mark leaves the reader unable to tell, inside an AI answer, between information from an authoritative institution and content from an unknown source. Certification makes that distinction visible.

Several forms already exist in the ecosystem. **Platform certification** includes Google's Preferred Sources: a user can mark the sites they trust as preferred, those sites then carry a prominent badge on their links inside AI answers, and a badged link is clicked about twice as often as an ordinary one. Google has also brought a "highly cited" label into traditional search results, marking original reporting that other outlets cite widely.

Industry certification includes the GEO service trust assessment run by the China Advertising Association (CAA) and the China AI Industry Development Alliance (AIIA); the first companies to pass

have been certified. Its value is that it does not depend on one platform's algorithm. An industry body evaluates against a common standard, which gives it standing across platforms.

Technical certification is represented by Content Credentials, pushed by C2PA (the Coalition for Content Provenance and Authenticity). The system cryptographically binds production information into the file, so whether content was modified by AI, and at which step, can be traced and verified. Adoption in China is still early, but it marks the important direction: truth stops resting on trust and starts resting on verification.

All three aim at the same thing, which is making trusted content visible to the person reading it. When a reader sees a certification badge on an AI citation, they no longer have to judge the source themselves. The badge has already done it. That lowers the cost of reading, and it turns the advantage of trusted content from implicit into explicit.

4.5 The virtuous cycle: what the design is for

Scoring, reputation and certification do not work in isolation. They form an incentive system only when they run together.

A well-designed cycle looks like this: **content scoring** gives better content a higher position in AI citation, **creator reputation** gives consistent creators a higher initial trust, and **authoritative certification** lets users recognize and choose trusted sources. Users' clicks and positive feedback then push the weight of good

content up again, closing a loop of better content, more citation, more recognition, higher weight.

The reverse loop holds as well. Low-quality content that gets cited in the short run loses weight over time, because users respond badly and cross-checks across sources fail. Platforms are cutting back deliberately. DeepSeek has compressed the sources it reads closely from 10-15 down to 4-5, the odds of content being blocked as a single point of failure reach 95%, and content that contradicts itself across platforms sees its odds of being featured fall 82%.

So the core of an incentive system is not rewarding good people and punishing bad ones. It is **letting the ecosystem clean itself**. When scoring, reputation and certification all run, doing the right thing pays on its own and doing the wrong thing drops out on its own. That takes no moral appeal and no extra penalty. It only takes the mechanism working as designed.

For GEO service providers and content creators, the reason to understand this system is that the rules are not constraints but levers. The scoring dimensions tell you what AI systems value. The reputation model tells you what accumulation is worth. The certification layer tells you how to get your credibility seen. Hold those three and doing the right thing stops being a moral choice and becomes a competitive advantage.

Key takeaways

- **Source tiering is the architecture under scoring.** The A/B/C/D rating assigns initial trust through one standard, verifiability. Content from authoritative media is 8 to 12 times more likely to be cited by AI than ordinary independent-account content.
- **Content scoring runs on five dimensions:** media tier and endorsements (25%), content expertise (22%), social and citation signals (20%), site health (18%), content freshness (15%). Bylined authorship is becoming a stronger signal each year.
- **Creator reputation upgrades a single score into a long-term asset.** It is measured by output success rate, repeat collaboration rate, engagement quality and credential verification, and it carries two key properties: time decay and portability.
- **Authoritative certification makes trusted content visible.** Platform certification (Google's Preferred Sources), industry certification (the GEO service trust assessment) and technical certification (C2PA Content Credentials) together form the visible badge layer.
- **The mechanism goal of the virtuous cycle is a self-cleaning ecosystem.** Scoring sets the ranking weight, reputation sets initial trust, certification sets user recognition. Run together, they pay good content and phase out bad content without anyone having to be persuaded.

Chapter 5: Governance and standards: Building a trusted AI information ecosystem

Chapter 4 showed how incentives pay for doing the right thing. Incentives need one thing to work: a widely accepted set of rules that says what counts as right and what counts as wrong.

With no rules, an incentive system cannot decide what to reward. With vague rules, scoring and reputation can be gamed. Governance and standards exist to give the ecosystem rules that are clear, enforceable and verifiable.

This chapter covers four layers: **industry standards** set the rules, **source tiering and three-zone separation** turn them into engineering, **platform responsibility and regulation** supply external pressure, and **multi-stakeholder governance** keeps the rules maintained and updated.

5.1 From wild growth to written rules

The need for governance comes from a structural fact. GEO as a technique is value-neutral, but its uses can produce opposite externalities. The same optimization that gets accurate content cited by AI also gets a fictional product turned into a fact.

In April 2026 the Cyberspace Administration of China (CAC) launched the "Qinglang" campaign against AI application chaos, listing AI data poisoning carried out by tampering with training corpora, fabricating authoritative data or using GEO technology for malicious marketing as a priority target. It was the first time at national level that GEO abuse was explicitly tied to AI data poisoning, and it moved GEO from self-regulation to a regulatory agenda.

Standards work accelerated at the same time. On 12 August 2026 T/CAPT 026—2026, China's first association standard for trusted GEO dissemination, took effect. It is administered by the China Association of News Technology Professionals and drafted with Xinhuanet Convergence Media Future Research Institute, the Xinhua News Agency State Key Laboratory and others. It sets systematic requirements for GEO service providers across nine areas: project intake, content review, corpus onboarding, optimization execution, dissemination control, monitoring and tracing, risk handling, high-risk scenarios and organizational capability.

The China Advertising Association (CAA) began its own GEO standardization in March 2026, focused on end-to-end compliance, with self-regulation rules and operating guidance across technical capability, service process, performance evaluation, business ethics and compliance management.

These efforts point one way: moving GEO from traffic-driven to trust-driven. The aim is not to restrain the technique. It is to draw

the boundary so that compliant operators gain an advantage and violators pay.

5.2 Source tiering: the A/B/C/D four-tier rating

Source tiering is the infrastructure of GEO governance. Without an agreed tier standard, AI systems cannot decide which sources deserve more trust, and scoring has no baseline.

The A/B/C/D four-tier credibility rating in T/CAPT 026—2026 is the most systematic framework available.

Tier A covers information published lawfully, within their remit, by bodies with statutory duties: government departments, judicial organs, regulators. What defines the tier is that publication follows a legal procedure with accountability behind it, which makes it the most verifiable class of information. In specialist fields such as weather, earthquakes and public health, Tier A sources carry authority nothing else can supply.

Tier B covers academic papers, industry whitepapers and original reporting from mainstream media. These have gone through some degree of peer review or editorial checking, have a traceable origin and an identified author. Credible, but below a body with statutory duties.

Tier C covers company websites, product manuals and reporting from ordinary media. Credibility here depends on the brand's own integrity record and how verifiable the content is. Tier C is not

unusable; core factual claims need a Tier A source or cross-checks across several sources.

Tier D covers anonymous origins, low-quality content and untraceable information. AI citation demotes or skips it by design.

What tiering changes is the form of the question. It turns credibility from a vague intuition into an engineering parameter. When a brand's core factual claims are anchored to Tier A sources with a complete evidence chain, its initial citation weight is far above content whose origin is fuzzy and untraceable. The higher the tier, the more likely AI is to cite it. That is the direct lever behind the source citation rate metric.

5.3 Three-zone separation: cutting off marketing dressed as fact

Tiering answers which sources are credible. A harder problem sits inside a single source: fact, opinion and marketing copy can be mixed in one place, so AI cannot tell them apart when it cites.

The three-zone separation required by T/CAPT 026—2026 exists for exactly this. The standard requires brand knowledge bases to store and use three libraries separately (fact, opinion and marketing): a **fact zone** of verifiable objective fact, an **opinion zone** of judgments and positions, and a **marketing zone** of promotion and advertising. Factual claims in marketing copy must match the fact-zone version.

In engineering terms, brand content needs separate storage by zone plus labeling. When a system retrieves a brand's description of

its own product, it can identify which part is verified objective fact ("the product holds this certification"), which is the brand's judgment ("industry-leading"), and which is promotion ("buy now").

The value of separation is that it blocks marketing-as-evidence at the production stage. In traditional marketing a brand could mix fact and exaggeration in one piece of copy, because a human reader would sort it out. In AI search, a system that cannot make that distinction will hand the reader the promotional line as a fact.

The standard also requires semantic duplication control and prompt-injection protection. Duplication control stops a brand from generating similar content in bulk to hold position: when a system sees several sources saying near-identical things, it marks them low-quality duplicates and demotes them. Prompt-injection protection stops a brand from hiding instructions inside content to steer what the model generates.

5.4 End-to-end traceability: every citation leads back

Tiering and separation govern content before it reaches AI. End-to-end traceability governs what happens after AI cites it.

The requirement is to record the whole path from production to citation, so that every citation can be traced back to its original source and every step can be audited and verified.

The technical means in the standard is the trace identifier (Trace ID). Each piece of content gets a unique identifier at publication,

recording its source, version, review history, distribution channels and later AI citations. When a system cites something, that identifier leads back through the full production chain.

This is already implemented in practice. The GEO-Trace engine gives each piece of content a unique digital identity so the whole chain is traceable. Monitoring systems keep a relationship table of question, answer, mention and source, tracking which AI platforms cited the content, in what context, and whether the citation expressed the original information correctly.

Traceability is worth more than after-the-fact accountability. It works as in-flight prevention. A brand that knows every piece of content is recorded and auditable loses the appetite for shortcuts. A platform that knows every citation can be traced to origin chooses sources more carefully.

5.5 Risk handling: circuit-breaker and correction response

No prevention removes all risk. When risk arrives, speed of handling matters. T/CAPT 026—2026 sets out a circuit-breaker mechanism requiring GEO service providers, on detecting an anomaly, to pause high-risk operations, start the repair process and escalate to management.

The design borrows from financial markets: on an abnormal signal, stop first, diagnose, resume once the all-clear is confirmed. In GEO the anomalies that should trip the breaker include a brand's core information appearing wrong in AI answers, abnormally frequent

false citations of competitor content, and large-scale inconsistency of information across platforms.

The correction response is what happens after the breaker trips. Once an anomaly is confirmed, three steps run: **preserve the evidence** (screenshot the AI answer as it stood and record the citation sources), **reinforce the correct source** (publish the correction through high-weight channels), and **report the error to the platform** (use the AI platform's feedback channel to request a fix). After the repair, re-test across platforms to confirm the anomaly is gone.

The standard also requires service capability tiers: L1 basic capability, L2 professional capability and L3 comprehensive governance capability. Governance is becoming a competitive dimension for providers. One that only delivers basic optimization and one that can handle risk and compliance end to end are different products at different prices.

5.6 Platform responsibility: the gatekeeper role

Industry standards govern what providers do. AI search platforms themselves are also part of the system. A platform decides what gets cited, in what form, and what label the user sees. Those decisions determine the externalities of GEO work.

In June 2026 the UK's Competition and Markets Authority (CMA) ruled that Google must give publishers clearer source attribution and links inside its AI-generated search features, must offer publishers an opt-out option from AI search, and must not demote

publishers in ordinary search results because they opted out. The CMA described it as "a world first, giving publishers an effective tool to stop their content being used to power AI features in search".

The logic behind the ruling: an AI search platform holds a strategic market position in information distribution, and that position carries governance duties. A platform cannot keep the user growth that AI answers bring while declining to attribute sources and protect publishers.

In China the direction is equally clear. The CAC's "Qinglang" campaign requires AI platforms to cross-check the sources they cite, issue risk notices, and label the links of cited information. Google has been testing controls in Search Console that let site owners manage how their content appears in AI search, with exposure metrics and reports on which pages show up in AI answers.

The other dimension of platform responsibility is **how clear the labeling is**. The CMA noted in its ruling that some stakeholders reported inaccurate labeling in generative AI search features, with room to improve clarity. Inaccuracy includes citing the wrong source, labeling a link that does not match the answer's content, and mixing several sources without distinguishing them. Having a label is not enough on its own. It has to be clear, accurate and verifiable.

5.7 Multi-stakeholder governance: keeping the system sustainable

Governance is not a one-off. Technology iterates, risks evolve, and rules have to move with them. The system itself needs to be sustainable.

The core idea of multi-stakeholder governance is that no single actor can govern alone. Government supplies the legal framework and the regulatory floor. Industry bodies write standards and self-regulation rules. Platforms implement technical governance. GEO providers and brands carry compliance responsibility. Users and media watch. The roles differ, and none of them is optional.

In current practice the layers are already visible.

The regulatory layer. The CAC's "Qinglang" campaign supplies the enforcement framework and draws the legal line around AI data poisoning and malicious GEO marketing. Its significance is that "what you must not do" stopped being industry consensus and became legal requirement. Once a violation can bring an administrative penalty, the cost of a shortcut is no longer a reputational risk. It is a legal one.

The standards layer. T/CAPT 026—2026 supplies the industry standard, the CAA's GEO standardization supplies self-regulation rules, and the finance-sector standard *Specification for Tiering and Adoption of Financial Information Sources for Large AI Models* supplies vertical guidance. Together they form a layered rule set:

general standards set the floor, vertical standards adapt it to an industry, self-regulation sets conduct above the floor.

The platform layer. AI search platforms govern through algorithm design and product features. Source tiering algorithms set citation weight, labeling lets users see origins, and opt-out mechanisms let publishers control how far their content is used. Whether standards land in practice depends on what platforms can enforce.

The service layer. The L1/L2/L3 capability tiers and industry self-regulation initiatives translate governance requirements into service standards a provider can actually deliver. The transparency, traceability and fairness principles in the China GEO industry development initiative are becoming the reference for admitting and assessing providers.

The four layers relate like this: regulation draws the floor, standards give the operating rules, platforms enforce them technically, providers carry execution responsibility. When all four run together, governance stops being a campaign and becomes routine operation.

5.8 What governance is for: good money driving out bad

The goal is not punishing violators. It is **changing the market's incentive structure** so compliant operators gain an advantage and violators pay.

Zhang Guohua, president of the CAA, said at the launch of its GEO standardization: "In the generative AI era, advertising should not chase short-sighted exposure through rule-breaking shortcuts. It must return to the foundation of being real and trustworthy. Setting standards for GEO is about treating the cause and the symptom at source, making the compliance boundary and code of conduct explicit, and letting a healthy pattern take hold where good money drives out bad."

That sentence states the logic of governance. Good money driving out bad does not happen on its own. It needs rules to hold it up. Where rules are missing, the bad kind (low-quality content, poisoning) has lower cost and faster return, so it drives out the good kind (accurate content, compliant optimization). Only when standards draw a clear boundary, platforms enforce them technically, and regulation raises the price of violation does good work get paid what it is worth.

For practitioners and brands, the point of understanding governance is not knowing what is forbidden. It is **knowing which way the rules are moving**. When source tiering becomes the base weight in AI citation, when three-zone separation becomes the standard for storing content, and when end-to-end traceability becomes a precondition for measuring results, the brands and providers that finished their compliance readiness early will take a structural advantage from the change.

Governance is not a limit. It is a filter. It selects the players willing to compete inside a long-termist framework, and that is the ground a healthy ecosystem stands on.

Key takeaways

- **Industry standards set the rules.** T/CAPT 026—2026 places systematic requirements on GEO service providers across nine areas, including project intake, content review, corpus onboarding, optimization execution, dissemination control, monitoring and tracing and risk handling. The CAA is running standardization in parallel.
- **Source tiering is the governance infrastructure.** The A/B/C/D rating assigns initial trust through one standard, verifiability. Core factual claims must be anchored to a source tier and an evidence chain.
- **Three-zone separation cuts off marketing dressed as fact.** Fact zone, opinion zone and marketing zone are stored separately, and factual claims in marketing copy must match the fact-zone version. Semantic duplication control and prompt-injection protection guard content quality further.
- **End-to-end traceability makes every citation lead back.** The Trace ID gives each piece of content a unique identifier and records the whole path from production to citation, turning after-the-fact accountability into in-flight prevention.
- **Circuit-breaker and correction response supply risk handling.** On an anomaly: pause high-risk operations, run the repair process, re-test across platforms. The L1/L2/L3 service capability tiers put governance into provider assessment.

- **Platform responsibility is the hinge of the system.** The UK CMA ruling requires clear source attribution and a publisher opt-out; China's "Qinglang" campaign requires AI platforms to cross-check and label the sources they cite. Labels must be clear, accurate and verifiable.
- **Multi-stakeholder governance is what keeps it sustainable.** Regulation draws the floor, standards give the rules, platforms enforce technically, providers execute. The end goal of governance is making good money drive out bad.

Chapter 6: The action roadmap

— Start doing Impact GEO today

The first five chapters ran from diagnosing the problem to building the theory, from method to incentives, from governance to platform duty. One question is left: **starting tomorrow, what specifically gets done?**

This chapter gives action lists for five roles. Each follows the same order: **do what is verifiable first, then what scales, then what influences the ecosystem.** No perfect conditions are required, and no industry-wide consensus is required. Every role can start inside what it controls.

6.1 Brands: from a ranking mindset to trust assets

Brands gain most directly from GEO practice, and they have more power to change the ecosystem than anyone else. The shift to make is in the question asked: not "how do I get cited by AI" but "what do I have that is worth citing".

Step 1: inventory verifiable assets. Take a week and work through the verifiable information the brand owns: official qualifications, certifications, product whitepapers, third-party test reports, patents

and intellectual property, real user cases. Record a source, a date and a scope of application for each. Keep only what a third party can independently verify. The output is a verifiable asset inventory, and it is the raw material for everything after this.

Step 2: build a source matrix. Distribute those assets across source channels at different tiers. Put core factual claims into high-weight sources first: government websites, mainstream media, industry regulators, authoritative academic databases. The brand website and official accounts are the home base and carry the complete picture. Third-party media coverage and industry research reports act as a credibility amplifier. Do not publish the same text everywhere; adapt it to what each channel is for.

Step 3: run three-zone separation. In the brand knowledge base, store content by zone: a fact zone, an opinion zone and a marketing zone. The fact zone holds verifiable objective fact, the opinion zone holds the brand's judgments and positions, the marketing zone holds promotional copy. Factual claims in marketing copy must match the fact-zone version. The output is a structured brand content asset system.

Step 4: set up monitoring. Fix 3 to 5 core queries and test the brand's citation situation on several AI platforms on a schedule. Record the citation source, the context, and whether the citation was accurate. Iterate on assets, sources and expression from what the tests say.

Action priority: if only one thing gets done, do step 1 and inventory the verifiable assets. Without them every optimization is

a castle in the air. With three months, finish the first two steps and build a source matrix, so that accurate information has somewhere to be found.

6.2 GEO service providers: from black-box operation to transparent engineering

Providers sit between brands and AI platforms, and what they do sets the quality of the whole GEO ecosystem. The shift is in positioning: from ranking operator to trusted content engineer.

Step 1: turn down black-hat work. State plainly which engagements are refused: bulk-generated false content, fabricated citation sources, manipulated engagement data, hidden or deceptive content. Write those red lines into the service contract and tell clients that short-term exposure is not worth long-term brand assets. The point is not moral standing. It is risk isolation: when regulation tightens, only providers with a clean violation history keep operating.

Step 2: build a standardized delivery process. Take the nine areas in T/CAPT 026—2026 and turn them into a process running from project intake to monitoring and tracing. The core stages: asset mapping (inventory what the client can verify), source planning (channels and priorities), content optimization (structure layer, markup layer, evidence layer), performance monitoring (citation rate, accuracy rate, ecosystem contribution), and risk handling (circuit-breaker mechanism and correction response).

Step 3: move up the capability tiers. Assess where you stand against L1 basic capability, L2 professional capability and L3 comprehensive governance capability, and plan the route up. An L1 provider does content optimization and source building. An L2 provider has complete compliance governance and risk handling. An L3 provider takes part in setting industry standards and governing the ecosystem.

Step 4: help write the standards. Join the GEO standardization work at the CAA, the AIIA and similar bodies, and feed front-line practice back into the next revision. Providers are where standards land, and they are the first source of what needs improving in them.

Action priority: if only one thing gets done, do step 1 and refuse black-hat work. It looks like lost revenue now and is an entry license later. With half a year, finish the first two steps and put a standardized delivery process in place, so white-hat GEO becomes an engineering practice that can be copied and verified.

6.3 Content creators: from traffic-driven to reputation-driven

Creators supply the corpus directly. What they choose to publish decides what AI learns and what it cites. The shift is in the unit of success: not how many page views a piece gets, but how many cross-checks it survives.

Step 1: put a byline and credentials on the work. State the author's identity, professional background and interests in every piece. Where there is a commercial relationship, label it as

advertising or sponsorship. A byline is not ceremony. Systems are treating named authors as a trust signal, and content from an author with a verifiable professional background is cited far more often than anonymous work.

Step 2: hold an evidence baseline. Follow a quantifiable floor in the writing: every thousand words carries at least 5 verifiable statistics and at least 2 authoritative citations, with the core conclusion stated in the opening paragraph. Prefer academic papers, official documents and primary data as sources. Do not quote out of context, and keep source, date and measurement consistent.

Step 3: separate fact from opinion. Mark which statements are objective fact and which are judgment. Factual claims must trace back to a verifiable source; opinions can keep a personal voice but have to be identified as opinions. The distinction helps AI cite correctly and helps readers judge what they are reading.

Step 4: accumulate a portable asset of reputation. Reputation is not a badge a platform grants. It is the record of what you kept doing. Keep publishing verifiable content, stay focused on the fields you actually work in, take part in building the knowledge of the community. A creator who produces good work in one domain over years finds that other platforms begin to recognize it too.

Action priority: if only one thing gets done, do step 1: sign the work and disclose the interests. It is the cheapest and highest-return act of building credibility. Thinking long term, make steps 2

and 3 a habit, so evidence density and the fact-opinion line become part of how you write.

6.4 AI platforms: from black-box citation to transparent recommendation

Platforms are the gatekeepers of information distribution, and their product design decides the externalities of GEO work. The shift is in role: from a tool that answers queries to a governor that maintains the information ecosystem.

Step 1: implement source tiering and labeling. In retrieval and citation, run an explicit tiering mechanism and give each tier a different citation weight. In the answer, label every cited source clearly so users can tell levels of credibility apart. Labels have to be accurate. Citing the wrong source, attaching a link that does not match the answer, mixing several sources without distinguishing them: each of these weakens the label.

Step 2: publish citation transparency reports. Give content producers and brands their citation data: which content was cited, in what context, and whether the citation expressed the original information correctly. Google has been testing similar controls in Search Console, letting site owners manage how their content appears in AI search and supplying exposure metrics. Transparency helps producers improve and makes the platform's own citation behavior auditable.

Step 3: design user trust signals. Put credibility cues inside the answer: the tier of each cited source, whether the content went

through a fact-checking process, whether the brand provides verifiable sources. A reader should be able to judge how trustworthy an answer is at a glance. Google's Preferred Sources is a usable reference: users mark the sites they trust, and those sites then carry a prominent badge in AI answers.

Step 4: give publishers control. Provide content owners with options over how AI uses their work, including an opt-out option from AI search. The CMA ruling has already set that direction: publishers should get an effective tool to stop their content powering AI features in search. A platform must not demote a publisher's ordinary search ranking because they opted out.

Action priority: if only one thing gets done, do step 1: tiering plus clear labeling. It is the infrastructure of platform governance and where rebuilt trust starts. With a product cycle available, take steps 2 to 4 in sequence.

6.5 Policymakers: from reactive enforcement to proactive guidance

Policymakers supply the rule framework for the whole ecosystem. The shift is in role: from punishing violations after the fact to steering good money to drive out bad before it happens.

Step 1: make the compliance boundary explicit. Building on the "Qinglang" campaign, state where GEO use is compliant (white-hat optimization, source building, raising content quality) and where it is not (corpus poisoning, malicious marketing, fabricated citations). Clear boundaries are the precondition for orderly competition.

Step 2: push standards into use. Turn association standards such as T/CAPT 026—2026 into reference points for market entry, and widen the reach of the GEO service trust assessment. Put source tiering, three-zone separation and end-to-end traceability into the compliance assessment systems for AI platforms and GEO service providers.

Step 3: coordinate across departments. GEO governance touches the cyberspace authority, market regulation and the advertising association. A coordination mechanism avoids conflicting rules and gaps in coverage. The CAC enforces against AI data poisoning, the State Administration for Market Regulation handles advertising compliance, and the CAA writes industry self-regulation standards. The three have to pull together.

Step 4: fund ecosystem public goods. Give policy support and resources to the shared infrastructure of the ecosystem: trusted source infrastructure, fact-checking tools, content provenance technology. These are the base of ecosystem health, and no single market actor has enough incentive to build them. Policy can change that arithmetic.

Action priority: if only one thing gets done, do step 1 and make the boundary explicit. Only with clear rules can the market form stable expectations, and only then does compliance pay.

6.6 Closing: make every optimization a contribution

This book began with a simple observation: AI search multiplies the influence of content as never before, and the current incentive structure makes poisoning cheaper to justify than building.

That observation opens onto a deeper problem: **inside an information ecosystem, individual rationality and collective rational results are in structural conflict.** Brands want exposure, providers want margin, creators want traffic, platforms want growth. Each of those is rational on its own. Added together they can degrade the ecosystem everyone depends on.

Changing that outcome cannot rest on individual virtue. The incentive structure has to change, so that doing the right thing is also the commercially smarter choice.

The book offered four principles: source transparency, truth first, ecosystem health as the moat, positive externality. They are not moral commandments. They are a systematic answer to what kind of GEO practice wins over a long competition.

The book offered a method: map assets, build sources, refine expression, keep testing. Not a bag of tricks, but one closed loop running from taking stock of what is true to iterating on it continuously.

The book offered an incentive system: content scoring, creator reputation, authoritative certification. Not extra burden, but rule design that pays better content more.

The book offered a governance framework: industry standards, source tiering, three-zone separation, end-to-end traceability, platform responsibility, multi-stakeholder governance. Not limits but a filter, selecting those willing to compete inside a long-termist frame.

All of it points at one target: **making every optimization a contribution to the information ecosystem.**

When you choose transparency over vagueness, accuracy over exaggeration, building over poisoning, you are not only helping yourself. You are making the AI information ecosystem more credible. When enough practitioners make the same choice, the ecosystem starts cleaning itself, and good money driving out bad stops being a wish.

Trust in AI search is falling. More than half of consumers now doubt the reliability of AI answers. That is a dangerous signal and an unusual opening. When users start looking for sources they can trust, the brands and providers that finished their verifiable assets early, built a transparent source matrix early, and accumulated verifiable reputation early will hold a structural advantage.

Ecosystem health is not a cost. It is the moat. Truth is not a sacrifice. It is an asset. Transparency is not a burden. It is a competitive advantage.

This path does not wait for perfect conditions. It can start with the evidence density of one article, with one byline and one interest disclosure, with one provider turning down one job, with one brand taking stock of what it actually has.

Start doing Impact GEO today.

Key takeaways

- **Brands**, four steps: inventory verifiable assets → build a source matrix → run three-zone separation → set up monitoring. The mindset shift is from a ranking mindset to trust assets.
- **GEO service providers**, four steps: turn down black-hat work → build a standardized delivery process → move up the capability tiers → help write the standards. The positioning shift is from ranking operator to trusted content engineer.
- **Content creators**, four steps: put a byline and credentials on the work → hold an evidence baseline → separate fact from opinion → accumulate portable reputation. The measure of success shifts from traffic to reputation.
- **AI platforms**, four steps: implement source tiering and labeling → publish citation transparency reports → design user trust signals → give publishers control. The role shifts from answering queries to maintaining the ecosystem.
- **Policymakers**, four steps: make the compliance boundary explicit → push standards into use → coordinate across departments → fund ecosystem public goods. The role shifts from reactive enforcement to proactive guidance.
- **Closing**: make every optimization a contribution to the information ecosystem. Ecosystem health is the moat, truth is the asset, transparency is the advantage. Start doing Impact GEO today.

Appendix 1: The Impact GEO self-check list

Use this list to check your own practice on a schedule. Brands, GEO service providers, content creators and platform governors are all covered. Run it quarterly, about 30 minutes each time. Six sections map onto the core modules in the book. You do not need full marks. You do need to know where the gap is.

1. Content transparency

Goal: make sure everything you publish survives the test of traceable origin and identifiable interests.

1. Does every piece state the author's identity and professional background?
2. Is commercially sponsored content labeled as advertising or sponsorship?
3. Does every core factual claim carry a traceable primary source?
4. Are factual statements and opinions distinguished in the text itself?
5. Have unverifiable phrases such as "industry-leading" or "the best choice" been avoided?
6. Is the brand knowledge base stored separately by zone: fact, opinion and marketing?

7. Do factual claims in marketing copy match the fact-zone version?

Self-check conclusion: if item 3, 4 or 6 is a no, fix it first. Those three are the base of transparency.

2. Source credibility

Goal: make sure core information appears in channels AI systems trust, and says the same thing in each of them.

1. Are core factual claims anchored to Tier A or Tier B sources (government websites, mainstream media, academic papers, industry whitepapers)?
2. Has the brand's core factual information been published in at least one high-weight source channel?
3. Does the brand website show complete qualifications, contact details, ICP license information and an update record?
4. Is factual consistency across sources checked regularly, covering dates, figures and definitions?
5. Have bulk postings of duplicate content on low-quality channels been avoided?
6. Have you taken part in industry research or supplied data and cases third parties could cite?

Self-check conclusion: if item 1 or 2 is a no, the brand's initial citation weight is low. Establish at least one core factual source in a high-weight channel first.

3. The five scoring dimensions

Goal: rate a single piece of content against the five dimensions that set AI citation weight.

Media tier and endorsements (25%)

- Does the publishing channel have verifiable qualifications and an ICP license?
- Is there a record of editorial review or peer review?

Content expertise (22%)

- Is there a named author with a verifiable professional background?
- Does every thousand words carry at least 5 verifiable statistics and at least 2 authoritative citations?
- Is the core conclusion stated in the opening paragraph?

Site health (18%)

- Does the site have HTTPS, mobile responsiveness and structured data markup?
- Is page speed within an acceptable range?
- Is there a violation history or a mass of broken links?

Social and citation signals (20%)

- Has the content been cited by other high-weight sources?
- Is there expert endorsement or a recommendation from a professional community?

- Are user comments positive and specific?

Content freshness (15%)

- Has the core content been updated within two years?
- Do the figures state when they were published and how long they hold?
- Has the site responded promptly to events in the field?

Self-check conclusion: of the five, raise content expertise and social and citation signals first. The first is the most controllable; the second only comes with time.

4. Creator reputation

Goal: see how much long-term credibility has actually accumulated, and where reputation building is weakest.

1. Do you keep producing verifiable content in your core field?
2. When AI cites your content, does the citation express the original information correctly?
3. Have partners invited you back for repeat work?
4. Can your professional credentials be verified independently (education, title, career history, body of work)?
5. Do you keep your public biography and record of results updated?
6. Have you proactively corrected content AI cited wrongly?

7. Have you contributed to shared knowledge in at least one professional community?

Self-check conclusion: if item 1 or 2 is a no, accumulation has not started. From this week, sign every piece and label its sources.

5. Ecosystem impact

Goal: judge the net effect your work has on the AI information ecosystem, and spot negative externality risks.

1. In the last three months, did you publish anything unverifiable?
2. Have you taken part in generating similar content in bulk?
3. Have you embedded hidden instructions or prompts in content?
4. Have you proactively corrected content AI cited wrongly?
5. Have you supplied reusable public knowledge to the field (method documents, industry data, case studies)?
6. When you found a competitor poisoning information, did you report it to a platform or an industry body?
7. Does your content distinguish where the brand's interest ends and the public interest begins?

Self-check conclusion: if the answer to item 1, 2 or 3 is yes, stop and correct it immediately. A history of negative externality can depress a brand's long-term weight inside AI systems.

6. The 90-day roadmap

Week 1

- Add bylines and source labels to the last three pieces you published.
- Check the brand website's HTTPS, mobile responsiveness and structured data.

Month 1

- Inventory the brand's verifiable assets and draft the fact zone.
- Fix 3 to 5 core queries and test current citation results on several AI platforms.
- Publish at least one piece of core factual content in a high-weight source channel.

Month 2

- Build a source matrix: set the publishing strategy and priority for each tier of source.
- Run three-zone separation: store the brand knowledge base by fact, opinion and marketing.
- Start an AI citation monitoring log recording source, context and accuracy.

Month 3

- Optimize at least three core pieces against the five scoring dimensions.

- Take part in one discussion on an industry standard or self-regulation rule, and put front-line practice on the table.
- Complete the first ecosystem impact check, and correct what it turns up.

After 90 days

- Re-test the core queries and compare against the figures from three months ago.
- Iterate on assets, sources and expression from what the monitoring says.
- Write the next 90-day plan, and move up a service capability tier or deepen reputation accumulation.

7. The one-question check

If only one question stays with you, make it this one:

My content, once AI cites it: does this make AI's answers more trustworthy, or less?

If the answer is more trustworthy, keep going. If it is less, stop and start again.

How to use this checklist: this is not an exam, and full marks are not the point. Its value is showing you which plank is shortest. Pick the weakest section each quarter and work on it alone. Look back in a year and you will find you are already inside the virtuous cycle of ecosystem health.

Appendix 2: Sources cited in this book

This appendix lists the sources behind the data and the main arguments in each chapter, in chapter order. Every entry gives the source type, the chapter it supports and a URL, so readers can check and update them.

Chapter 1: The paradigm shift

Consumer trust in AI search

- Data cited: In 2025, 82% of consumers found AI search more helpful than traditional search; in 2026 the figure was 54%, a fall of 28 percentage points in one year.
- Source: Fractl × Search Engine Land, *2026 AI Search Consumer Trust Study*.
- URL: <https://searchengineland.com/new-ai-search-data-visibility-trust-480089>

Gartner forecast on search engine traffic

- Data cited: By 2026, traditional search engine visits will fall 25%.
- Source: A Gartner forecast, cited in a research report from CITIC Construction Investment.
- URL: <https://www.stcn.com/article/detail/3550064.html>

Scale of generative AI users in China

- Data cited: As of December 2025, China had 602 million generative AI users, a penetration rate of 42.8%.
- Source: the 57th Statistical Report on China's Internet Development (CNNIC).
- URL: https://m.gmw.cn/2026-03/04/content_38627125.htm

Rising weight of authoritative sources in AI citations

- Data cited: In May 2026, YouTube was cited 188,863 times (up 56.4% month over month), Wikipedia 83,191 times (up 55.2%), NIH 91,227 times (up 41.3%).
- Source: Meltwater GenAI Lens, tracking more than 8 million AI citations.
- URL: <https://www.meltwater.cn/blog/ai-search-visibility-april-may-2026>

How consumers handle AI search

- Data cited: 86% of consumers check the primary sources themselves when the information matters; 72.4% of users switch engines after AI gets a citation wrong twice in a row.
- Source: Fractl × Search Engine Land, 2026 study.
- URL: <https://searchengineland.com/new-ai-search-data-visibility-trust-480089>

Chapter 2: The core principles of Impact GEO

The citation weight gap between authoritative media and independent accounts

- Data cited: Content from national-level media is 8 to 12 times more likely to be cited by AI than content from ordinary sources. One national media outlet can carry more citation weight than a hundred independent accounts combined.
- Source: Industry research data, cited in GEO industry analysis articles.
- URL: http://www.xtrb.cn/syunf/2026-08/25/content_1288492.htm

Share of AI citations going to independent accounts and brand websites

- Data cited: Earned media (third-party sources) account for 82% to 94% of all AI citations; brand-owned websites account for a single-digit share.
- Source: AuthorityTech, 2026 study.
- URL: <https://authoritytech.io/>

Keyword density and citation rate

- Data cited: Pages above 8% keyword density are 17% less likely to be cited by AI than pages at a natural 2%–3%.
- Source: Industry test data. Widely repeated across GEO industry analysis articles; the original publication is still to be confirmed. Readers are referred to the testing section of the GEO industry whitepaper.

Chapter 3: The white-hat GEO method

The DSS principles (white-hat GEO method)

- Data cited: The three core principles of Semantic Depth, Data Support and Authoritative Source.
- Source: iResearch × Yuanyi Information, *2026 Generative Engine Optimization (GEO) Whitepaper*.
- URL: <https://report.iresearch.cn/report/202602/4787.shtml>

Structured content and citation probability

- Data cited: Content in structured formats is 63% more likely to be cited directly by AI than plain text.
- Source: Volcano Engine developer community analysis report (April 2026), cited in GEO industry articles.
- URL: <https://www.hongshu18.com/article-detail/BPL38gEB>

Cited sources, statistics and direct quotes

- Data cited: Content with cited sources is 34.4% more likely to be cited by AI, content with statistics 32.1%, content with direct quotes 29.7%.
- Source: Industry test data. Widely repeated across GEO industry analysis articles; the original publication is still to be confirmed. Readers are referred to the testing section of the GEO industry whitepaper.

Named authors in AI citations

- Data cited: In May 2026 Google changed how AI Overviews displays citations and made named authors a search ranking variable.
- Source: Google announcement and Search Engine Land reporting.
- URL: <https://searchengineland.com/seo-pulse-preferred-sources-expand-gmail-brand-lift-pichai-on-ai-overviews/577173/>

Platforms cutting back low-quality content

- Data cited: DeepSeek compressed the sources it reads closely from 10-15 down to 4-5, and citation rates for more than 70% of mass-distributed press-release GEO content fell sharply.
- Source: DeepSeek's May 2026 algorithm update, analyzed by several GEO industry outlets.
- URL: <https://www.hongshu18.com/article-detail/BPL38gEB>

Schema markup and citation rate

- Data cited: Content without structured markup loses about 47% of its weight in AI citation on average.
- Source: Industry test data. Widely repeated across GEO industry analysis articles; the original publication is still to be confirmed. Readers are referred to the testing section of the GEO industry whitepaper.

Chapter 4: Incentives in the content ecosystem

Weight allocation across the five scoring dimensions

- Data cited: Media tier and endorsements (25%), content expertise (22%), social and citation signals (20%), site health (18%), content freshness (15%).
- Source: Public information from the major AI platforms plus industry testing. The weights are an industry observation; actual platform weights may differ, and readers should check each platform's own documentation.

Freshness and AI citation

- Data cited: AI engines rarely cite content that has not been updated in about two years, and the correlation between metadata freshness and citation rate reaches 0.68.
- Source: Industry research. Widely repeated across GEO industry analysis articles; the original publication is still to be confirmed. Readers are referred to the testing section of the GEO industry whitepaper.

Google Preferred Sources data

- Data cited: Google Preferred Sources covers 345,000 sources, and users are about twice as likely to click a preferred source as an ordinary link.
- Source: Google announcement, reported by Search Engine Journal.

- URL: <https://www.searchenginejournal.com/seo-pulse-preferred-sources-expand-gmail-brand-lift-pichai-on-ai-overviews/577173/>

The C2PA provenance standard

- Data cited: C2PA released Content Credentials 2.3, which binds production information into the file cryptographically.
- Source: C2PA official website.
- URL: <https://c2pa.org/the-c2pa-launches-content-credentials-2-3-and-celebrates-5-years-of-impact-across-the-digital-ecosystem/>

The GEO service trust assessment

- Data cited: The China Academy of Information and Communications Technology (CAICT), together with the AIIA Security Governance Committee, launched China's first GEO service trust assessment; the first nine companies passed and were certified.
- Source: Xinhuanet report "AIIA trusted GEO symposium held in Beijing".
- URL: <https://www.xinhuanet.com/digital/20260518/5b81c242155a44da9a69d69f90410c57/c.html>

Chapter 5: Governance and standards

The T/CAPT 026—2026 association standard

- Data cited: China's first association standard for trusted GEO dissemination. It sets up the A/B/C/D four-tier credibility rating and specifies three-zone separation, end-to-end traceability, a circuit-breaker mechanism and the L1/L2/L3 service capability tiers.
- Source: Issued by the China Association of News Technology Professionals, standard number T/CAPT 026—2026.
- URL: <https://www.ttbz.org.cn/standardDetail/8617629717344255b65b44426dc77ce8.html>

Bodies that took part in drafting

- Data cited: Drafted with participation from Xinhuanet Convergence Media Future Research Institute, the Xinhua News Agency State Key Laboratory and the Guangxi Daily.
- Source: Xinhuanet and Guangxi News reporting.
- URL: <https://v.gxnews.com.cn/> (Guangxi Daily reporting)

The CAC's "Qinglang" campaign against AI application chaos

- Data cited: The campaign lists AI data poisoning carried out by tampering with training corpora, fabricating authoritative data or using GEO technology for malicious marketing as a priority target.
- Source: The CAC's "Cyberspace China" account, reported by CCTV.com, Guangming Online and others.

- URL: https://m.gmw.cn/2026-05/01/content_38744439.htm

CAA GEO standardization work

- Data cited: GEO standardization work began in full in March 2026, focused on end-to-end compliance standards. President Zhang Guohua called for "a healthy pattern where good money drives out bad".
- Source: Xinhuanet client reporting.
- URL: <https://app.xinhuanet.com/news/article.html?articleId=0e5ab2a0858ccd79b752198b7bc7d26f>

The UK CMA ruling against Google

- Data cited: In June 2026 the CMA required Google to attribute content sources and links more clearly in AI-generated search results and to give publishers an opt-out option from AI search, without demoting publishers in ordinary search for opting out.
- Source: Competition and Markets Authority (CMA) announcement.
- URL: <https://www.gov.uk/find-digital-markets-measures/google-search-publisher-conduct-requirement>

DeepSeek source compression

- Data cited: DeepSeek cut the sources it reads closely at the final stage from 10-15 down to 4-5, and citation rates for more than 70% of mass-distributed press-release GEO content fell sharply.
- Source: DeepSeek's May 2026 algorithm update, analyzed by GEO industry outlets.

- URL: <https://www.hongshu18.com/article-detail/BPL38gEB>

Cross-verification on AI platforms

- Data cited: The odds of content being blocked as a single point of failure reach 95%, and content that contradicts itself across platforms sees its odds of being featured fall 82%.
- Source: Industry test data, cited in GEO industry analysis articles.
- URL: <https://cloud.tencent.com.cn/developer/article/> (Tencent Cloud developer community analysis)

Research report on a trusted GEO ecosystem

- Data cited: CAICT led and the AIIA published China's first *Research Report on Building a Trusted Generative Engine Optimization (GEO) Ecosystem*.
- Source: Xinhuanet reporting.
- URL: <https://www.xinhuanet.com/digital/20260518/5b81c242155a44da9a69d69f90410c57/c.html>

Preface and chapter 1 case

The Apollo-9 smart band poisoning case

- Data cited: Exposed at CCTV's 2026 3·15 gala. Industry insiders used the "Liqing GEO optimization system" to invent the Apollo-9 smart band; two hours after a dozen advertorials went out, two major AI models were recommending it as the standard answer.

- Source: CCTV's 3·15 gala, reported by CNR Online, the China Internet Joint Rumor Refutation Platform and others.
- URL: https://news.cnr.cn/dj/20260316/t20260316_527553512.shtml
- URL: <https://www.piyao.org.cn/20260403/ad5ea72461a34cd3bf9a27d7f252fa3e/c.html>

Notes on sources

1. Every source listed here is public material actually consulted while writing the book. Some industry test data (keyword density, Schema markup, scoring weights and similar figures) is repeated across many GEO analysis articles without a clear original publication; those entries say so.
2. The full text of the association standard is available on the national association standards information platform (www.ttbz.org.cn). Policy documents are available from the CAC's "Cyberspace China" account and from Xinhuanet, Guangming Online and other official outlets.
3. For academic research data (AI citation source statistics, trust surveys and the like), readers should obtain the original report for the full methodology and sample information before judging how far the figures apply.
4. This appendix is updated with each revision of the book. Corrections and additions are welcome.