STATE OF THE ART DIGITAL
Methodology Document — OPTIMUM Framework Companion

The Three-Layer AI Visibility Framework

A measurement methodology for AI citation, search behaviour and demand

Steve Coulter, State of the Art Digital
July 2026

---

1. Purpose

This document sets out a working methodology for establishing whether AI search systems, including ChatGPT, Perplexity, Gemini, Google AI Overview and Claude, are citing a business as a source, and whether that citation is producing a measurable change in demand.

It is a companion piece to the OPTIMUM framework and sits alongside the three intelligence reports published free at stevecoulter.co.uk: Your Business Is Becoming Invisible, Search Doomsday Is Q3 2027, and The Chaff Effect. Those reports set out the industry-level argument. This document converts that argument into a repeatable measurement practice that can be applied to a single client's own data.

2. The Attribution Problem

No analytics platform currently reports that an AI assistant has cited a page. Google Search Console shows position, impressions, clicks and CTR for Google Search only. Google Analytics 4 shows sessions and referral sources, but many AI applications either pass no referrer at all or pass one that is indistinguishable from a generic external link.

This means a business can see that a visit has occurred without being able to confirm what caused it. The visit can look identical whether the user typed the URL directly, opened a bookmark, or clicked a citation inside an AI answer.

Because no single tool closes this gap, the framework below does not rely on one. It builds a case from three separate sources of evidence and treats agreement between them as the signal, rather than treating any one metric as proof on its own.

3. Connection to the Three-Part Thesis

The framework exists to test, in a client's own data, whether the following three positions are actually playing out.

Your Business Is Becoming Invisible
Established that AI-driven answers are replacing clicks for a growing share of queries. The consequence is that visibility inside the answer itself now matters as much as position on the results page, and in some cases more.

Search Doomsday Is Q3 2027
Argued that Q3 2027 marks the point at which AI search overtakes traditional SERPs as the primary business case, the date at which enquiries, leads and sales generated through AI citation are expected to

outweigh those generated through conventional search ranking. This is a statement about where commercial outcomes shift, not about competitor readiness.

The Chaff Effect

Describes a separate, independently timed mechanism governing competition for citation itself, distinct from the Q3 2027 channel shift, which moves through two epochs.

In the first epoch, the field is open. Most competing businesses have not yet met Tier 1 technical compliance, let alone built Tier 2 citation-cluster architecture or AI-optimised content. A business that establishes both early is one of a small pool of qualifying sources in its category, and captures a citation share disproportionate to the effort involved. This is the citation windfall available to early movers.

The second epoch begins at point X, whenever mass Tier 1 adoption is reached across a given competitor set. This point is not fixed to Q3 2027 and will vary by sector. At that point, Tier 1 compliance stops separating competitors and becomes merely the entry ticket required to compete at all, while the real contest moves to Tier 2. Citation share can no longer be won simply by qualifying. It has to be won by being better structured than a field of other qualified competitors, a materially harder proposition.

The report also showed that the pool of clicks AI leaves behind is shrinking, and that what remains is skewing toward lower-intent, later-stage searchers. The headline figures are 74% zero-click for informational queries against 31% for transactional, evidence that the value of any single citation is itself changing as the second epoch approaches.

The three-layer framework is the mechanism for checking these three positions against a client's own numbers, rather than relying on industry-wide statistics alone.

4. The Three Layers

Layer	Question	Measured by
1. Visibility	Am I being cited?	Direct prompt testing (Citation Share)
2. Behaviour	Are more people coming to my site?	GA4: sessions, direct and referral traffic, landing pages
3. Search demand	Has my brand become more sought after?	GSC: branded query volume, clicks, position

No single layer is conclusive. Layer 1 confirms citation is happening but not that it is driving anything. Layer 2 confirms behaviour is changing but not why. Layer 3 confirms demand is moving but not what caused it. Read together, the three layers form a chain of evidence.

5. Method

Step 1: Baseline

Before any AI SEO or GEO work begins, record the current position across both platforms.

Google Search Console: total clicks, total impressions, average position, CTR, branded queries, non-branded queries, top landing pages.

Google Analytics 4: users, sessions, direct traffic, organic traffic, referral traffic, enquiries and conversions, landing pages.

Step 2: Citation Share testing

Compile 50 to 100 prompts that a realistic prospective customer would put to an AI assistant. Run each prompt through ChatGPT, Google AI Overview, Gemini, Perplexity and Claude, and record whether the business was cited, which page was cited, whether the citation was a direct link, the position within the answer, and which competitors appeared instead.

This produces a Citation Share figure per platform: the proportion of tested prompts in which the business was cited.

Step 3: Monitor Search Console

Over the following month, compare branded search volume, organic clicks and average ranking against the baseline.

Step 4: Monitor GA4

Ask specifically whether traffic to the pages identified as cited in Step 2 is increasing, independently of whether overall rankings have moved.

Step 5: Build the case

Combine the three layers. If Citation Share is rising, branded search is rising, and traffic to the cited pages is rising, while ranking position is largely unchanged, that combination is strong evidence that AI citation is generating demand that traditional ranking movement cannot explain on its own.

6. Interpreting the Outcomes

The most important, and most commonly misread, output of this framework is what a stable ranking means once it is read alongside search volume and clicks rather than in isolation. The same Search Console position can indicate two opposite situations depending on the direction those figures move.

Outcome A: Stable ranking, increased branded searches and clicks

Average position in Search Console has not moved, yet branded query volume and clicks are rising. Because position is unchanged, this increase cannot be explained by conventional on-page ranking gains.

When this pattern appears alongside a rising Citation Share, the likely explanation is that AI systems are naming or recommending the business directly. Users are searching for the brand by name, or arriving from an AI answer, generating demand that page position alone did not produce.

This is confirmation that the Tier 2 AEO and GEO layer is working. It is the outcome the OPTIMUM three-tier commercial structure is designed to produce, and it is direct evidence in support of the Tier 1 gate argument in the Chaff Effect thesis: once technical compliance is in place, further gains come from citation work rather than further technical work.

Outcome B: Stable ranking, decreased searches and clicks

Average position in Search Console has again not moved. In isolation, an unchanged position is usually read as meaning nothing has gone wrong. Search volume and clicks, however, are falling.

This is the more dangerous pattern, precisely because it is the one most likely to be missed. A stable ranking measures relative competitiveness within a pool of available clicks. It says nothing about the size of that pool. This is the exact signature predicted by The Chaff Effect: the available pool of clicks is contracting as AI answers absorb queries before a click ever takes place. A business can hold its competitive position within that pool while the pool itself continues to shrink beneath it.

If Citation Share is flat or falling at the same time, the business is losing on both fronts: it is not capturing the citations that are absorbing the demand, and the demand itself is contracting. This combination is the leading indicator that the Chaff Effect has entered its second epoch for that competitor set, the point at which citation has become a genuinely contested arena rather than an open field, and it is precisely the pattern that conventional SEO reporting will not flag as a problem, because the ranking line on the report looks unchanged.

Why the distinction matters

A stable ranking is not, by itself, informative. It becomes informative only once it is read against the direction of branded search volume and clicks, and against Citation Share. Rising figures alongside a stable ranking indicate that citation is doing the work that ranking used to do. Falling figures alongside a stable ranking indicate that the pool is shrinking and that ranking has become a measure of position within a shrinking market rather than a measure of overall visibility.

This is why ranking alone is an unreliable KPI in an AI search environment, and why Layer 1 testing cannot be skipped. Without it, the two outcomes above are indistinguishable from a Search Console screen showing the same flat line.

7. Using the Framework in Practice

The framework is designed to run on a quarterly cycle, aligned with the Monthly Citation Watch product, so that movement in Citation Share can be checked against Search Console and GA4 data from the same period rather than compared across mismatched timeframes.

Local versus national and international businesses

The standard prompt set and citation criteria are built around location-based signals. This suits businesses whose customers search with a place in mind: an estate agency covering Skipton, Grassington, Settle and Ilkley, or a car dealership such as Fenwick Motors. For these businesses, geography is not incidental, it is the criterion. A prospective customer's prompt naturally contains a town, a county, or the phrase "near me", and a citation that ignores this is answering a different question to the one being asked.

A national or international business does not work this way. A pension provider operating across the UK, a wealth management firm used by clients nationwide, or a financial platform with no physical branch network is not competing to be recommended in a place. It is competing to be recommended as a category leader, a comparison winner, or the provider a particular type of saver or employer should choose. Applying the local template to this kind of business does not produce a low score because the business is weak. It produces a low score because the test is measuring something the business was never trying to do.

Before Step 1, the prompt set and citation criteria need to be rebuilt around entity and topic signals rather than location, so that the baseline reflects the actual competitive question the business is being judged against.

Example: adjusting a prompt set for a national business

The table below shows five prompts from a standard local template alongside the equivalent prompts rebuilt for a national pension provider.

Before (local template, unsuitable)	After (rebuilt for national scope)
Best pension provider near me	Best UK workplace pension provider
Pension advisor in [town]	Which providers offer the best workplace pension schemes
Pension consultant near [postcode]	Compare UK workplace pension providers
Local pension advice services	Leading workplace pension provider for UK employers
Which pension providers operate in [county]	Pension providers used by large UK employers

The left-hand column asks where a provider operates. The right-hand column asks which provider leads a category. This is the distinction that matters when adjusting the baseline: local prompts test geographic reach, national and international prompts need to test category authority and comparison standing instead.

8. Conclusion

No single metric currently proves that an AI assistant has cited a business. What can be built is a chain of evidence across three layers: whether citation is happening, whether behaviour is changing, and whether demand is moving, checked against a stable baseline over time.

This framework turns the argument set out across Your Business Is Becoming Invisible, Search Doomsday Is Q3 2027 and The Chaff Effect from an industry-level thesis into a measurement practice that can be run against a single client's own numbers, and it is this practice that distinguishes a business genuinely benefiting from AI citation from one whose stable ranking is masking a shrinking pool beneath it.