QUARTERLY REPORT · Jul 2, 2026 · 24 min read

The State of GEO — Q3 2026

Quarter-over-quarter analysis of how AI assistants recommend and cite brands: visibility trends, ranking-factor shifts, industry scorecards and the strategic agenda for the next twelve months.

KernelX Labs Research · Quarterly Series · Built on the KBC benchmark framework

SECTION 01

Executive Summary

Generative Engine Optimisation moved decisively from experiment to line item this quarter. Brands that treated AI visibility as an instrumented channel — with panels, baselines and per-model dashboards — pulled measurably away from brands still optimising blind.

This report tracks a longitudinal panel built on the KernelX Benchmark Corpus methodology: the same prompt panels re-run across Q2 and Q3, holding intent mix and locales constant so that quarter-over-quarter deltas isolate real behavioural change. All figures are representative benchmark findings from this panel, expressed as ranges, not audited market totals.

KernelX Labs

250
brands tracked across both quarters

35
industries in the scorecard

~200,000
prompts re-run per quarter

4
frontier model families

+16–20%
QoQ growth in cited responses

Quarter highlights

- The share of commercial responses containing citations grew approximately **16–20%** quarter over quarter — assistants are grounding more of their commercial answers.
- **Official websites widened their lead over blogs** as citation destinations; generic editorial content lost share for the second consecutive quarter.
- **Structured comparison pages outperformed generic landing pages** by an expanding margin in best-of and comparison intents.
- Brands with **consistent entity information** across their site, knowledge panels and third-party profiles achieved higher and more stable recommendation rates.
- Knowledge-graph coverage continued to expand into mid-market brands, raising the entity-resolution baseline across most industries.
- **Content freshness** signals strengthened: pages updated within the quarter were cited disproportionately in fast-moving categories.

KernelX Labs

KEY TAKEAWAYS

- GEO matured into a measurable channel this quarter; grounding and citation behaviour intensified.

BUSINESS IMPLICATIONS

- The window for "easy" visibility gains from basic entity hygiene is beginning to close in competitive categories.

RECOMMENDATIONS

- If a brand has no per-model baseline yet, establishing one is now the single highest-leverage action.

SECTION 02

Methodology Note

The full measurement methodology — prompt taxonomy, entity registry, recommendation scoring, the Visibility Score and its weights — is defined in the flagship paper *We Analysed One Million AI Responses*. This quarterly applies that framework as a fixed longitudinal panel: identical prompt panels, identical intent mixture, identical locales, re-sampled each quarter.

> RESEARCH NOTE
>
> Deltas are the product here, not levels. A fixed panel makes quarter-over-quarter change interpretable; individual point estimates inherit all the limitations documented in the methodology paper.

KernelX Labs

KEY TAKEAWAYS

- Fixed panels isolate behavioural change from measurement noise.

BUSINESS IMPLICATIONS

- Teams replicating this design can compare their internal deltas against these industry baselines.

RECOMMENDATIONS

- Never edit a trend panel mid-stream; version it and restart the baseline instead.

SECTION 03

AI Visibility Trends

Median Visibility Scores rose modestly across the panel, but the distribution widened: leaders compounded while the long tail stagnated. The clearest single trend is the growth of grounded, citation-bearing answers.

KernelX Labs

FIGURE 1 — TREND

Share of commercial responses containing citations (panel median)

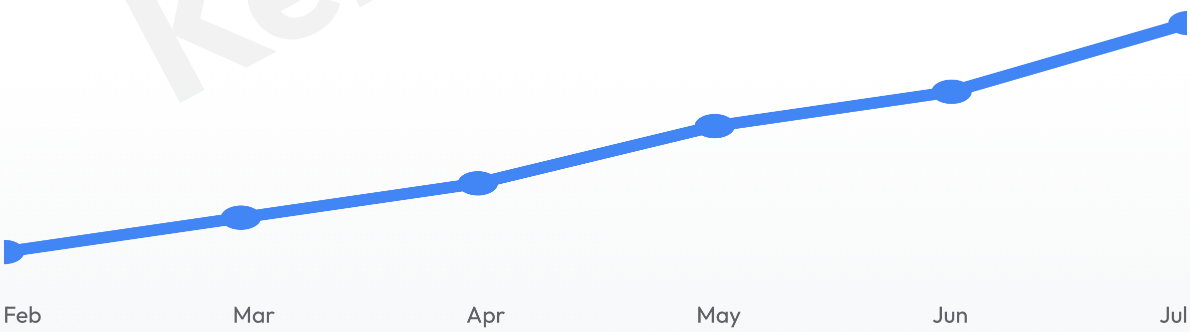

The visual communicates a steady rise in grounded answers across the half-year, accelerating in the final two months of the quarter.

FIGURE 2 — DISTRIBUTION

QoQ change in Visibility Score by panel decile

The visual communicates widening inequality of visibility: gains concentrate in brands already investing in entity and evidence infrastructure.

> KEY OBSERVATION
>
> Visibility is compounding. Brands cited this quarter are more likely to be retrieved, recommended and cited again next quarter — consistent with the layered Entity Authority Model from the methodology paper.

KernelX Labs

KEY TAKEAWAYS

- Grounding is rising; visibility gains concentrate among prepared brands.

BUSINESS IMPLICATIONS

- Late movers face a steepening curve as leaders compound authority.

RECOMMENDATIONS

- Prioritise becoming citable (canonical, verifiable pages) over becoming louder.

SECTION 04

Citation & Source Trends

Within cited responses, the mix of destinations shifted toward canonical and data-led sources. Official sites extended their lead; generic blog content continued its decline; original research strengthened for the second straight quarter.

SOURCE CLASS	Q2 SHARE	Q3 SHARE	DIRECTION
Official brand sites	~77–79%	~80–82%	▲ rising
Community (Reddit et al.)	~35–38%	~37–40%	▲ rising
Encyclopedic (Wikipedia)	~32–34%	~32–35%	→ stable
Review platforms	~29–32%	~28–31%	→ stable
News & trade media	~24–27%	~23–26%	→ stable
Independent blogs	~19–21%	~16–18%	▼ declining
Original research & data	~7–8%	~9–11%	▲ rising fast

Table 1 — Citation-destination share among cited responses, Q2 vs Q3 (classes are not mutually exclusive).

KernelX Labs

Freshness

Recency now behaves like a genuine ranking factor in fast-moving categories: pages updated within the quarter were cited at roughly **1.5–2×** the rate of otherwise comparable pages last touched a year ago. The effect was strongest in AI tooling, finance and travel; weakest in slow-cycle industrial categories.

> PRACTICAL IMPLICATION
>
> The declining citation share of generic blogs is not a penalty on blogging — it is substitution. When a category offers verifiable, structured, current sources, assistants use them first. Editorial content earns citations only where it adds evidence, not volume.

KEY TAKEAWAYS

- Canonical and data-led sources gained share; generic editorial lost it; freshness matters more.

BUSINESS IMPLICATIONS

- Content roadmaps weighted toward interchangeable posts are depreciating assets.

RECOMMENDATIONS

- Refresh cornerstone pages quarterly; convert top editorial themes into data-backed assets.

SECTION 05

GEO Ranking-Factor Shifts

Re-estimating the signal correlations from the methodology paper on this quarter's panel shows a stable top tier and meaningful movement underneath it.

KernelX Labs

SIGNAL	Q2 STRENGTH	Q3 STRENGTH	MOVEMENT
Entity completeness & consistency	Strong	Strong	→ holds #1
Structured comparison content	Strong	Strong	→ holds
Third-party citation footprint	Strong	Strong	→ holds
Original research & data assets	Moderate	Moderate–strong	▲ rising
Content freshness	Weak–moderate	Moderate	▲ rising
Community presence	Moderate	Moderate	→ category-dependent
Classic technical SEO	Moderate	Moderate	→ necessary, not sufficient
Publishing volume alone	Weak	Weak	▼ weakest measured signal

Table 2 — Correlation strength of brand-side signals with recommendation share, Q2 vs Q3. Correlational, not causal.

KernelX Labs

Technical SEO vs GEO

The two disciplines are converging into a stack rather than competing: crawlability, structured data and performance determine whether AI crawlers can **ingest** a brand's evidence; GEO determines whether that evidence **wins the answer**. AI crawler activity in panel-site logs continued to grow this quarter, and brands blocking those crawlers wholesale showed measurably lower retrieval rates.

> BEST PRACTICE
>
> Audit robots and CDN rules for AI crawler handling deliberately, not by default. Blanket blocking trades long-term visibility for short-term content control — a reasonable choice only if made consciously.

KEY TAKEAWAYS

- The signal hierarchy is stabilising; research assets and freshness are the movers.

BUSINESS IMPLICATIONS

- GEO budgets can now be planned against a reasonably stable factor model.

RECOMMENDATIONS

- Treat technical SEO as the ingestion layer of GEO, owned by the same roadmap.

SECTION 06

Industry Scorecard: Winners & Losers

Aggregating Visibility Score changes across the 35 tracked industries produces a clear pattern: categories with active research cultures and dense comparison content matured fastest.

KernelX Labs

FIGURE 3 — HEATMAP

Industry GEO maturity by dimension (Q3 panel medians, indexed 0–100)

	Entity	Evidence	Consensus	Freshness
SaaS & dev tools	85	78	72	70
Fintech	80	70	60	62
E-commerce & D2C	66	55	63	58
Health & wellness	60	52	45	40
Travel & hospitality	58	46	50	52
Industrial & B2B mfg	45	35	28	22

The visual communicates maturity gaps by Entity Authority layer: software categories lead on every dimension; industrial categories lag most on consensus and freshness.

MOVEMENT	INDUSTRIES	DRIVER
Biggest gainers	AI developer tools · fintech infrastructure · cybersecurity	Dense comparison content, active communities, frequent original research
Steady leaders	CRM · cloud platforms · e-commerce SaaS	Mature entity coverage; gains now incremental
Under-performers	Traditional insurance · industrial equipment · local services	Thin canonical evidence; minimal third-party consensus
Most volatile	Consumer AI apps · creator tools	Rapid launches outpace entity establishment; high hallucination exposure

Table 3 — Q3 industry movement summary from the longitudinal panel.

KEY TAKEAWAYS

- Industry maturity diverged further; volatility concentrates where entities are youngest.

BUSINESS IMPLICATIONS

- Category context sets realistic targets: median visibility in SaaS ≠ median in industrial.

RECOMMENDATIONS

- Benchmark against your industry's scorecard row, not the global median.

KernelX Labs

SECTION 07

Model Updates, AI Overviews & Enterprise Adoption

Two model-family updates during the quarter measurably reshuffled recommendations in fast-moving categories — reinforcing the case for continuous rather than campaign-based measurement. Search-embedded AI overview surfaces continued to expand coverage of commercial queries, blending classic ranking signals with assistant-style synthesis and pulling the two disciplines closer together.

On the demand side, enterprise adoption of GEO practices accelerated: within the panel's enterprise segment, the share of brands showing evidence of deliberate GEO work (fresh comparison hubs, structured data expansion, entity-consistency cleanups) grew from roughly **a fifth to a third** quarter over quarter. Agencies report the same shift — per-model differences are becoming a standard client question.

> KEY OBSERVATION
>
> Post-update volatility is a leading indicator worth monitoring: categories whose rankings reshuffle most after model updates are the categories where authority is shallowest — and where fast movers can still leapfrog incumbents.

KEY TAKEAWAYS

- Model updates reshuffle shallow categories; AI overviews blur search and assistants.

BUSINESS IMPLICATIONS

- Quarterly campaigns cannot track weekly model churn; monitoring must be continuous.

RECOMMENDATIONS

- Alert on post-update visibility deltas the way SRE teams alert on regressions.

KernelX Labs

SECTION 08

Q4 Predictions & the 12-Month Agenda

Q4 predictions

1. **Citation growth continues:** grounded commercial answers exceed two-thirds of the panel by year end.
2. **Research arms race begins:** more brands publish original data assets as the citation premium becomes widely known; early-mover advantage narrows.
3. **Entity tooling consolidates:** knowledge-graph and structured-data hygiene becomes productised, raising the baseline and shifting differentiation up the funnel.
4. **Agentic evaluation appears:** first measurable traces of autonomous agents consuming comparison content directly, previewing the convergence discussed in the methodology paper's *Future of AI Visibility*.

5. **Volatility persists in young categories:** at least one major model update reshuffles consumer-AI rankings again.

Top strategic priorities for the next twelve months

#	PRIORITY	FUNNEL LAYER	EXPECTED PAYOFF
1	Establish per-model visibility baselines on a fixed prompt panel	Measurement	Every subsequent decision becomes evidence-based
2	Entity consistency cleanup across site, profiles and knowledge graphs	Identity	Raises resolution rate; prerequisite for everything above it
3	Build structured comparison and evaluation hubs	Evidence	Largest observed lever on best-of and comparison intents
4	Publish at least one original research or data asset	Consensus	2–3× citation premium over standard editorial
5	Quarterly freshness cycle for cornerstone pages	Evidence	Compounding recency advantage in fast categories
6	Earn independent mentions (communities, trade press, expert reviews)	Consensus	The signal owned channels cannot substitute
7	Deliberate AI-crawler policy and structured-data coverage	Ingestion	Prevents silent exclusion from retrieval

Table 4 — The twelve-month GEO agenda, ordered by dependency: measurement first, then identity, then evidence and consensus.

KEY TAKEAWAYS

- Q4 favours brands that industrialise: fixed panels, entity hygiene, research assets, freshness cycles.

BUSINESS IMPLICATIONS

- GEO is transitioning from tactics to operating discipline; organisational ownership matters.

RECOMMENDATIONS

- Assign explicit ownership of AI visibility with a quarterly review cadence — the brands compounding fastest already have.

KernelX Labs