TYGART MEDIA · FIELD REPORT

Built for the Question

What happened when we stopped guessing what AI assistants want, read the report that shows what they were already asking one of our sites, and built the answers. Receipts included — including the parts that complicate the story.

44,583

AI citations earned by bcesg.org in 90 days (Aug 7–Nov 6, 2026)

11 → 30

citations per day on the beauty-comp guide's query, before vs after the Oct 19 build

57 : 1

citations for every human visit, sitewide — the audience is the assistant

Current version — November 8, 2026.
Update history: tygartmedia.com/we-watched-the-questions-ai-was-asking-our-site
Data: Bing Webmaster Tools AI Performance + Microsoft Clarity, bcesg.org.

1 · The signal

In early October 2026, a filtered export from Bing Webmaster Tools showed one question — "pharmaceutical companies esg executive compensation" — earning 51 citations in five days on BC ESG, a site we operate. The citations were landing on a generalist compensation guide: a good page, but a sideways answer to a specific question. The demand was real, measured, and nobody had built for it on purpose.

We tried to trace what sparked the demand. There was movement around a company we are not naming, because we could never source it — and we hold to a rule that runs this whole report: **build on the demand you can measure, never on the attribution you can't.** So we measured all of it.

2 · The map: count everything before building anything

On October 19 we pulled the site's full AI Performance report and counted it by value — every query, every page, every share. Four clusters mattered:

Cluster (October 19 snapshot)	Citations then	Citations now	What we did
Beauty companies — executive pay & sustainability targets	68	**684**	Built a dedicated guide, Oct 19
Pharmaceutical companies — ESG executive compensation (4 phrasings)	51	**600**	Built a dedicated guide, Oct 19
ICMA green bond principles	159	**417**	Nothing — pages already existed
Worker categories — pay-equity reporting	9	**24**	Nothing — noted, not built

Two clusters were demand with no door built for it. Two were controls — one already served, one deliberately left alone. That accident of the map turned into the experiment's best feature: we can compare what building did against what the tide did.

3 · The build

That night, October 19, we published two guides — one per unanswered cluster — under three rules:

1. **One real question per page, in the asker's own words.** The page's address is the question, word for word — the string assistants are already carrying.

2. **Answer first.** The answer sits at the top; the depth below serves whoever wants it — human or machine.
3. **Every page carries its download.** Each guide ships with a companion PDF. A page and its PDF are one asset; if the content changes, both change in the same run. Section 6 shows why this matters more than it sounds.

4 · What happened: the step

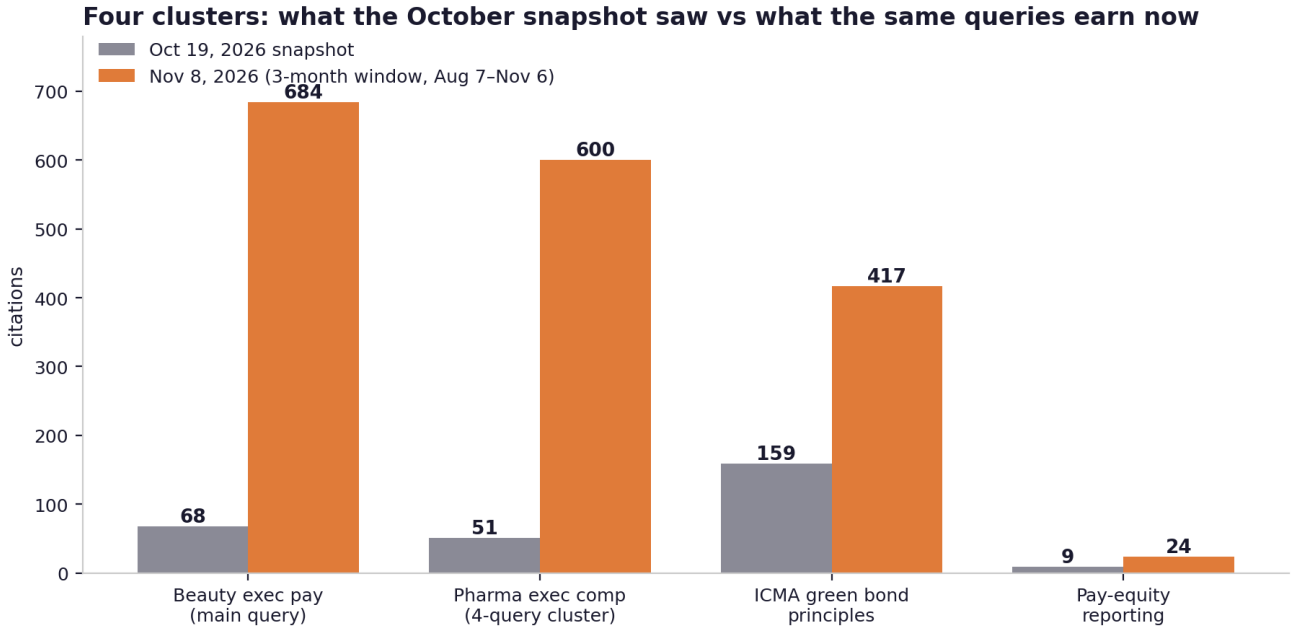
Read the daily series, not the cumulative totals, and the story tells itself. The beauty query ran at **11 citations a day**; from the day after the guides went live it stepped to **30 a day** (peak day: 50) and held the new level. The pharma query stepped from **3 a day to 11 a day** on the same date. Cumulative totals flatter everything; a dated step in a daily series is much harder to argue with.

The step: daily AI citations before and after the guides went live



Source: Bing Webmaster Tools AI Performance, bcesg.org, filtered by query. Beauty window Oct 9–Nov 2, 2026; pharma Aug 31–Nov 1, 2026.

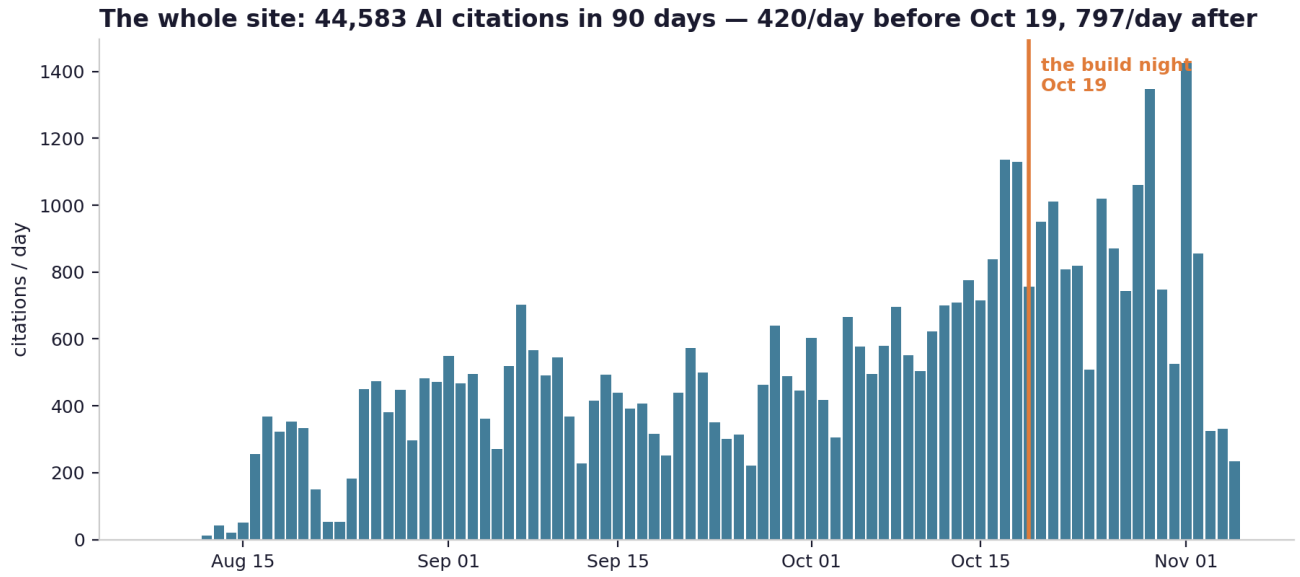
Daily citations per query, filtered views. Beauty window Oct 9–Nov 2, 2026; pharma Aug 31–Nov 1, 2026. The vertical line is the build date.



October 19 snapshot counts are from that day's site-wide report; current counts are the Nov 8 report's 3-month window (Aug 7–Nov 6). Windows differ by design — the daily step above is the causation read; this chart is the scale read.

5 · The honest complication: the tide rose too

ICMA green bond principles grew 159 → **417 with no new build** — the pages already there kept earning. The whole site stepped from 420 citations a day to 797 a day across the same line. So this report does not claim every new citation as construction. The claim it defends is narrower: on the two built queries, the daily rate stepped on the build date and held, and the built pages took first or second place in earning them. The tide explains growth. It does not explain a step.



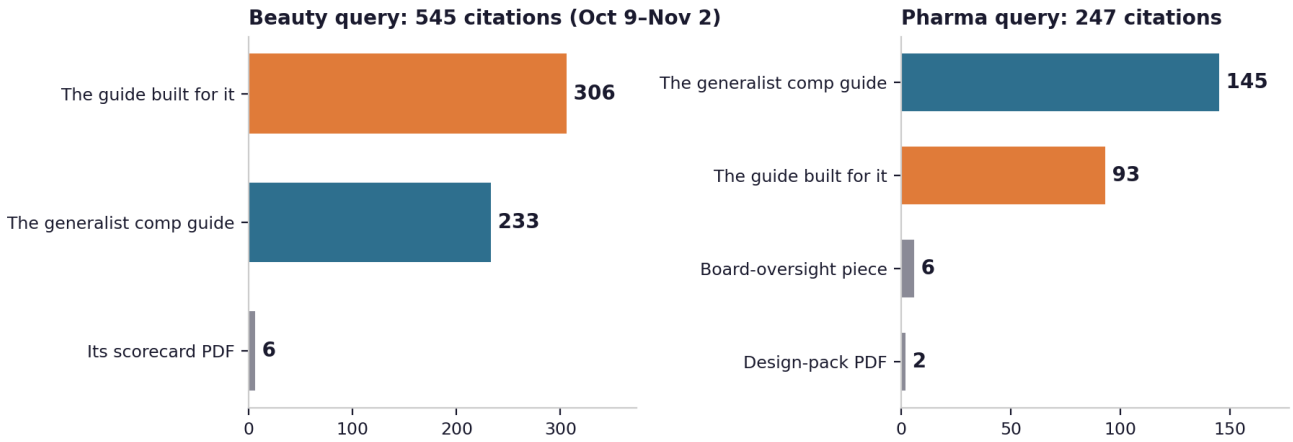
All queries, all pages: 44,583 citations in the 90-day daily series. Avg cited pages: 32. Distinct grounding questions: 428.

And small stayed small. Pay-equity reporting — the cluster we noted and skipped — went 9 → 24. No door, no step. In this report, the control group is not an embarrassment. It's the evidence.

6 · Who earns the query — and the file that travels

On beauty, the guide we built **leads** — 306 citations to the generalist guide's 233 — and the generalist lost nothing in absolute terms. The pie grew; the specific page took the bigger slice of a bigger pie. That's the cannibalization worry answered, at least this once. On pharma, the built guide is a strong **second** (93 to the generalist's 145). One difference between the builds: the beauty page's address is the query word-for-word; the pharma page's address drops one word. With two data points that's a hypothesis, not a law — labeled as one, and first in line to test on the next build.

Who earns the query: built for it vs already there



Citations by page within each filtered query view. Note the bottom rows: the PDFs.

The PDFs get cited. Six for the beauty scorecard and two for the pharma design pack in the filtered windows — and across the whole site, four PDFs hold 45 citations, one pharma quick scorecard alone at 23. The guides link straight to their PDF files — no download page in between — so an assistant can take the file in one move. A cited PDF means no visit, no session, no channel at all: the answer traveled as a file, into systems we'll never see. The downloads aren't lead magnets bolted onto content. In an AI answer, they are content. (This report is built the same way, on purpose.)

7 · The other instrument: what the humans did

Bing's report is the machines' side of the glass. Microsoft Clarity — which also runs on the site — is the humans' side. Last 30 days: the beauty guide, holder of 725 citations in the current report, has had **7 human sessions**. The pharma guide, 301 citations: **2**. Sitewide, 786 human sessions against 44,583 citations — and 2,526 bot sessions filtered out on top of that. Even the arrivals are mostly machines: on the pharma page, 6 of 8 arrivals were bots.

The few humans who land on the beauty guide arrive through a telling set of doors — among its seven sessions, referrers include **claude.ai and gemini.google.com**. They stay about 10 seconds, scroll about 5% of the page, click nothing, and leave. Friction of every measured kind — rage clicks, dead clicks, quick backs, script errors — is **zero** on both pages. That is not reading behavior. It is verification behavior: a person checking the source behind an answer they already got. The page's job for them is to be checkable at a glance — answer first, sources visible — which is the same job it does for the assistant. One structure serves both readers.

> **The conclusion the instruments force:** for this class of content, the assistant is the consumer and the human is an occasional auditor. Build the page like a source document that can survive an audit — because that's literally its use.

8 · The folklore check: llms.txt, JSON, and the machinery

A standing argument in our trade says none of this can work without special machinery for the machines — an llms.txt file, elaborate JSON feeds. We don't have to argue it. We can look:

Claim	What's true on bcesg.org
"You need an llms.txt to be cited"	**The site has none.** /llms.txt returns the 404 page; it has never existed there. 44,583 citations anyway.
"JSON schema buys citations"	Both guides carry FAQPage JSON-LD — and it predates the step. The best controlled evidence (1,885 pages, JSON-LD added, no citation lift) says schema buys eligibility and accuracy, not selection. Our dates agree.
"Citations send traffic"	57 citations per human visit sitewide; over 100:1 on the built pages. The citation is not a referral channel. It is the consumption itself.
"Downloads are marketing garnish"	45 citations held by 4 PDFs. The file is a citable asset in its own right.

What actually changed on October 19 was content — a page shaped like the question, carrying a file shaped like the answer. The assistants cite what answers the question in a form they can carry. The rest is garnish; some of it harmless, some of it an industry's worth of billable garnish.

9 · The operating loop

1. **Pull the AI citation report and count it by value** — which questions, which pages, what shares. Not vanity totals.

2. **Find the clusters where demand lands on pages that weren't built for it.** That gap is the work list. It's usually short.
3. **Build one page per question**, in the asker's words, answer first, its download attached as part of the same asset.
4. **Wait, re-pull, and read the daily step** — not the cumulative total, which flatters everything.
5. **Publish the receipts**, including the parts that complicate the story.

What we can't claim: Bing's own chart totals and table sums occasionally differ under filters (one view: 247 vs 246 — Bing prints a note saying so; we use chart totals and disclose it). Export windows differ and are labeled on every chart. Questions mutate — pharma's biggest phrasing in November wasn't the one we built for in October. And this is one site, one engine's report, one month. It's the cleanest window we have, and it is what it is.

Your Webmaster account is already holding your version of this report.

The AI citation report exists for every site in Bing Webmaster Tools — unopened, in most cases. We pull it, count it, and build what it's asking for. If you want that done on your site: book a call with Will Tygart — calendar.app.google/He743TpF2vQoC8Tt8

Method note. All figures are from Bing Webmaster Tools AI Performance exports for bcesg.org pulled October 8 (query-filtered), October 19, and November 8, 2026; citation counts are Microsoft Copilot and partner citations as Bing reports them. Windows: site-wide overview Aug 7–Nov 6, 2026; filtered views as labeled per chart. Behavior figures are from Microsoft Clarity (bcesg.org project), last 30 days as of November 8, 2026, bot sessions excluded as Clarity reports them. Page structure and file checks (llms.txt absence, JSON-LD presence, PDF links) verified directly against the live pages on November 8, 2026. The two guides discussed were published on bcesg.org on October 19, 2026. Where Bing's chart totals and table sums differ under filters, chart totals are used and the difference is disclosed. © 2026 Tygart Media. This report may be shared in full; please keep the method note attached — it's part of the findings.