

A black and white line drawing of a baseball glove is in the top left corner. A red baseball is in the top right corner. Another black and white line drawing of a baseball glove is in the middle right. A red baseball is in the bottom right corner. The title text is centered in the middle of the page.

# **Why ChatGPT Doesn't Cite Your Press Releases**



Published July 2026  
St. Nick Media Services  
& Pitch'd

## Executive summary

**Zero.** That is the number of press releases, out of 758 across fourteen years of B2B work, that hit the threshold this audit calls “AI wire ready.” The highest score in the corpus reached 74 out of 100. Nothing in the dataset cleared 80.

This matters because a new audience has entered the chat, literally. ChatGPT reports more than 900 million weekly active users. Google’s AI Overviews became the default surface for informational queries in 2024. Buyers, journalists, recruits, and analysts now do an increasing share of category research inside an AI chat. Companies that AI engines cannot cite immediately drop out of consideration and customer awareness.

PR teams have always written for two audiences: Editors and search crawlers. A third audience emerged in late 2022, AI or answer engines, and this audience grades releases on different criteria. Most teams aren’t writing for it.

This audit measures the gap between current press release writing, and the criteria AI engines use to decide what to cite. It does so at scale, across fourteen years, and with the first within-release measurement of what an editorial pass actually moves.

### **Three findings dominate.**

**First: Not one release in fourteen years scored “AI wire ready.” The dataset for this audit topped out at 74 out of 100. The gap is real, and most businesses have not yet adapted.**

**Second: The “good for AI” share fell by more than half** as the AI-search era emerged. In the Pre-AI era (2012–2019), 21.5 percent of final releases scored “good.” In the AI era (2023–2026), 10.8 percent did. The bar for AI visibility moved upward when the channel emerged. Writing practice has not caught up.

**Third: Editorial input substantially improves releases. Releases scored at both draft and final stages showed the editorial pass delivered a median lift of +22 on Timeliness, +22 on Authority, and +14 on Specificity. The editorial pass reliably moves the needle on AI readiness.**

#### BY THE NUMBERS

**0 / 758**

Releases that reached AI wire-ready (score  $\geq 80$ ) across 14 years

**21.5%  $\rightarrow$  10.8%**

Share of FINAL releases scoring “good for AI” ( $\geq 65$ ), Pre-AI era vs. AI era

**+22 / +22 / +14**

Median within-release Timeliness, Authority, and Specificity lift from the editorial pass (n=156 paired)

**24 / 156**

Releases moved from below the “good for AI” threshold to above it through editing alone

**92 / 100**

Mean Originality across every era and both statuses. Cliché discipline held.

#### A NOTE ON CALIBRATION

*The Scout, an AI optimization tool built by St. Nick Media Services’ Pitch’d, grades releases for AI visibility against general-audience answer engines: ChatGPT, Gemini, Claude, and Perplexity. It does not grade them for trade-press readability. The two audiences are different. Trade press in complex, highly regulated industries parses specialized vocabulary, regulatory shorthand, and technical terminology natively. That is the working language of those audiences. Releases that score lower on Readability, Originality, or Citation Potential for AI-citation purposes can still serve their intended trade-press audiences effectively, and many in this corpus did. The audit measures how well these releases serve a new audience that emerged in 2022 to 2024, separately from the audiences for which they were originally written. Both audiences now matter. The findings here are about whether the writing has adapted to serve both.*

## Press releases now answer to three audiences

ChatGPT launched in November 2022 and now reports more than 900 million weekly active users. That count more than doubled in twelve months. Google’s Gemini reports roughly 400 million monthly users. Perplexity reports approximately 45 million and doubled year over year. Bing’s index now powers ChatGPT’s search functionality. Google’s AI Overviews became the default surface for informational queries in 2024.

PR teams once wrote for two audiences: Editors and search crawlers. The list now runs to three: Editors, search crawlers, and AI answer engines. The third audience grades on different criteria, and most PR teams are still optimizing only for the first two.

The industry has begun adapting. Wire services shipped AI-optimization tools in early 2026, but access has its limits. Notified’s AI Press Release Optimizer, for example, grades on significantly fewer criteria AND only works for contracted Notified customers. Until now, there has been no

equivalent AI optimizer for 99 percent of releases distributed by agencies and in-house teams without enterprise wire contracts.

Independent research has begun mapping how AI engines source their answers. 5W Public Relations' AI Platform Citation Source Index, the most prominent example, analyzed more than 680 million individual citations across the major answer engines. The infrastructure to write for AI exists. The question this audit asks is: **How far has actual press release writing moved with it?**




Chart 1 illustrates the score distribution across 758 B2B press releases scored against The Scout. The structural ceiling sits visibly below the 80-point AI wire-ready threshold.

The audit looked at fourteen years of B2B press release work for the answer. The pages that follow trace what the data shows.

*“Press release writing has become corporate or executive vanity, not news announcement. Brand awareness, market share, and pipeline building left the room years ago. The executives who admit the problem and refocus PR will win the AI era.”*

— Jennifer Overhulse, Founder of St. Nick Media Services and Pitch'd

## AI invisibility has a price

These findings reveal a structural gap between current press release writing and the criteria AI engines use to decide what to cite. The reasonable response: So what? Press releases that AI does not cite still do real work. Brand awareness. Stakeholder communication. SEO baseline. Credibility transfer through trade press pickup. None of those mechanisms have disappeared.

## The case for press releases predates AI

The case for press releases predates AI by decades and the original metrics still hold up.

**Credibility transfer through earned media.** Edelman's annual Trust Barometer has tracked the gap between paid and earned media trust for more than twenty years. The 2024 edition reports earned media trust running in the low 60s and paid advertising in the high 30s. Press releases are the most reliable input into earned media. A release that gets picked up by trade press carries that trust transfer. A release that does not stays invisible to the channel.

**Stakeholder reach beyond the journalist.** Press releases do get read by employees, partners, customers, investors, recruiters, and analysts. The Muck Rack 2024 State of Journalism survey found that 67 percent of journalists prefer learning about news through press releases over other formats. The same survey notes that PR teams report internal stakeholders consuming releases at higher rates than the external audience they were ostensibly written for.

**Operational infrastructure that compounds.** Backlinks from wire distribution still build domain authority. Branded-search volume reliably lifts in the 30 days following a wire release. Press releases get repurposed internally as one-pagers, sales decks, and outbound email proof points. A track record of regular news visibility makes a company easier to find when something goes wrong, easier to correct the record about, and easier to defend when the alternative is silence. This is operational infrastructure that compounds. It cannot get built retroactively.

## Six business costs accrue when AI cannot cite you

When AI engines cannot cite a release, six specific business costs accrue:

**1. Consideration-set exclusion in B2B research.** A buyer asks ChatGPT, "What are the top platforms in [category]?" or "What fits a company like ours?" The AI returns a list. Companies absent from that list do not get evaluated. Forrester research has documented for years that 60 to 70 percent of the B2B buying journey completes before a vendor talks to sales. Absence from AI citations now effectively functions just like absence from analyst rankings.

**2. Trust-signal absence on direct-brand queries. A prospect, journalist, or candidate asks an AI, "Tell me about [company]." The response: "I don't have**

**current information about them.” The implicit signal: Fringe, dormant, or unverified. Worse than no information. A negative signal generated by absence.**

**3. Narrative-control loss.** AI engines still answer questions about companies they cannot cite. They use whatever they can find. Older news. Competitor framings. Outdated product descriptions. LinkedIn mentions. Sometimes hallucinations. Press releases are the primary mechanism companies have used to influence the canonical version of their own story. That mechanism still works, provided the release gets written so AI can use it.

**4. Recruitment and partnership discoverability.** Today, your potential talent pool researches employers via AI before applying. Partnership prospects ask AI what tools fit their stack. Companies not surfaced in those answers typically don’t make the shortlist. The cost shows up in HR pipeline quality and partnership deal flow.

**5. SEO-PR convergence.** Google’s AI Overviews now intercept many informational queries that used to deliver organic traffic. Press releases written for the older search funnel fall out of the AI surface and the AI-augmented Google surface simultaneously.

**6. Customer support absorption.** Customers are increasingly using AI as a first-line support resource. When AI cannot answer a question about a product, the cost reverts to internal support teams. Press releases are not the primary remedy here. They are the input that builds the company-level knowledge base AI engines draw from.

*“Companies paying attention to AI visibility now will own the buyer-research gap. Companies that do not are betting that nothing changes. It already has.”*

### **The data is too new for direct revenue attribution**

This report does not yet contain direct correlation of AI citation to revenue. No public research yet contains it either. The data is too new. No one has the longitudinal pipeline data to say, “AI-cited releases convert at X percent higher rate than uncited ones.” That gap is no reason to wait. SEO in 2003 had the same shape: No clean revenue attribution yet, but companies that ignored it lost the next decade.

The hidden cost most companies do not see yet is structural. Most B2B PR teams have not started measuring AI visibility. They won’t know they are losing the channel until pipeline

volume from online research traffic begins to soften. By the time that signal arrives, the competitive position has already shifted to firms that started earlier.

## Three places where this corpus already shines

Three areas the data already gets right, before the rest:

1. **Originality scored 90+ across every era and both statuses.** Across all 758 releases and 14 years, cliché counts averaged one to two per release. Originality scored 92 across all releases, including draft and final versions. The agency's discipline on cliché avoidance starts at the first draft. The fear that AI tools will catch your jargon is mostly unfounded for experienced PR writers.
2. **Top scorers came from every era.** The single highest-scoring release in the corpus came from 2017. The next four highest came from 2019, 2022, 2018, and 2024. The biggest (but not the only) constraint is the typical multi-party approval workflow surrounding each release.
3. **Releases with two named speakers scored 80 to 100 on Quotability.** The agency knows how to produce AI-citable material. Where clients and partners approved second voices in time for distribution, releases scored at the top of the AI-visibility band.

*"By the time you can measure the loss, the winners are already picked."*

## A single editorial pass moves the needle

This audit quantifies the value of agency editorial work. Two complementary comparisons informed this section. Releases scored at both draft and final stages, 156 in total, anchored the headline finding.

**The headline finding:** Releases gained a median of seven overall points between draft and final stages. Three-quarters of releases improved measurably. About one in five gained 10 points or more.

### Editorial input substantially improves releases in three categories.

- **Timeliness:** +22 median. Median draft score moved from 22 to 36. Half of all releases gained 22 or more points on this signal alone. The editorial pass replaces vague-timing language ("soon," "in the coming months") with specific dates and present-tense anchors.
- **Authority:** +22 median. Median draft score moved from 64 to 77. Most releases cross from amber into green-band territory on this signal during editing. Dateline cleanup, boilerplate verification, formal executive attribution, and contact-block additions account for most of this lift.
- **Specificity:** +14 median. Median draft score moved from 41 to 50. Editing adds numbers, percentages, dollar figures, and named entities.

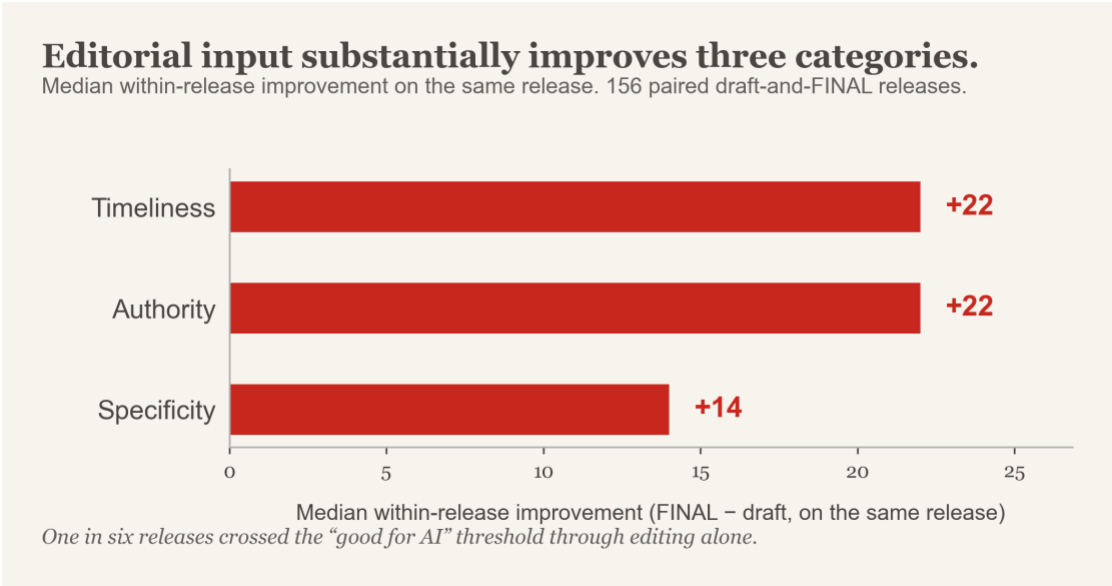


Chart 2 illustrates the median within-release improvement on three signals between draft and FINAL stages, drawn from 156 paired releases.

Temporal anchoring, attribution and authority signals, and concrete specificity are the structural elements of a release that editorial intervention can substantively shape.

*“Editorial input moves releases into anchored, attributable, and citable territory.”*

Two implications for clients and prospects:

1. **The editorial pass produces measurable, downstream-visible outcomes.** A median 22-point Timeliness gain, 22-point Authority gain, and 14-point Specificity gain on the same release moves a release from how AI engines treat undated, uncertain content to how they parse content that reads as anchored, attributable, and citable.
2. **Pitch’d catches the same gaps programmatically.** The pre-finalize audit in the Pitch’d builder flags Timeliness, Authority, and Specificity issues during the writing process. Teams without an in-house senior editor get the same structural pass before the release goes out, without the wire-window pressure that often forces compromises.

**The agency controls only some of the dials**

Several findings in this audit reflect long-standing industry conventions or client and compliance requirements rather than editorial choices. Reading the audit usefully requires distinguishing them.

| Finding from the audit   | Why it shows up                                                                                    | Who can change it                                  |
|--------------------------|----------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Citations stuck at floor | Content in complex industries writes from internal authority by long-standing convention. External | Client (must approve referencing third-party data, |

|                             |                                                                                                                                                  |                                                                                          |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
|                             | research organizations are rarely cited.                                                                                                         | competitors, or partner content)                                                         |
| Sentence length 25–30 words | Compliance review, regulatory disclosure language, and qualifying clauses push sentences long across regulated B2B segments.                     | Industry + client (compliance/legal review drives this; agency edits help at the margin) |
| Single-speaker quotes       | Customer or partner quotes require sign-off on tight wire timelines. Many releases ship before a second voice is approved.                       | Client + partner (must approve and return language in the wire window)                   |
| Boilerplate variation       | Boilerplate is usually client-supplied and rarely revised between releases.                                                                      | Client (controls boilerplate); agency (edits at intake)                                  |
| Formal AP-style attribution | A drift toward less formal attribution (“Jane said” vs. “Jane Smith, chief executive officer of..., said”) appeared across the corpus over time. | Agency (style choice within the editorial process)                                       |
| Clichés and AI-tell phrases | Originality scored 90+ across all 758 releases, both FINAL and DRAFT. The agency-controllable elements are working.                              | Agency (already strong)                                                                  |

The agency-controllable elements of editorial craft score consistently across the corpus: Originality, formal attribution, clean drafting, and structural completeness. Where scores fall short, the cause traces to signals that depend on client approvals, industry conventions, or compliance review.

This separation between agency-controllable and client-controllable signals characterizes any B2B segment with established editorial conventions, specialized vocabulary, and a multi-party approval workflow. The same patterns show up across complex, highly regulated industries: Healthcare, financial services, pharma, legal, energy, government contractors, and beyond.

## A decade of moving standards

Among final releases (the apples-to-apples cohort that matches the original 2026 audit and represents wire-distributed work across all 14 years), AI-visibility scores held remarkably stable across all three eras. The decline as AI emerged is real but small.

| Era                    | Releases (n) | Median | Mean | % scoring 65+ |
|------------------------|--------------|--------|------|---------------|
| Pre-AI (2012–2019)     | 130          | 60     | 60.3 | 21.5%         |
| Transition (2020–2022) | 193          | 59     | 59.0 | 19.7%         |
| AI era (2023–2026)     | 140          | 57     | 57.8 | 10.8%         |

## The audit measures performance against criteria that did not exist before late 2022.

The releases here served the audiences they actually had: Editors, journalists, and traditional search. By those criteria, they worked. The data describes a new audience that emerged faster than industry writing practice has adapted and quantifies for the first time what the agency's editorial pass contributes against the new criteria.

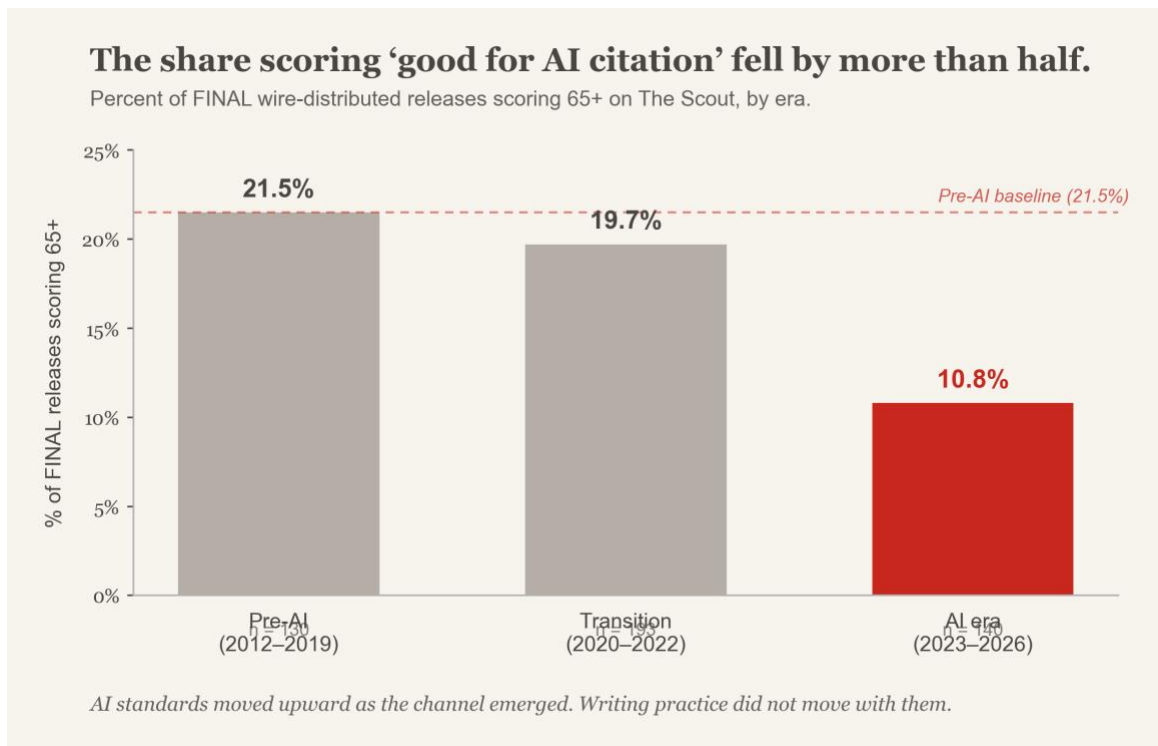


Chart 3 illustrates the share of FINAL releases scoring "good for AI" (65 or higher) across three eras of the corpus.

The decline in absolute terms runs about two and a half points of overall score across a decade. The more striking pattern shows up in the rightmost column: The share of releases scoring "good for AI visibility" fell from 21.5 percent to 10.8 percent between the Pre-AI and AI eras. Roughly half. The explanation is single and simple: AI standards moved upward as the channel emerged. Writing practice was optimized for editors and search, and it did not move with them.

Until late 2022, none of these signals were AI signals. They were editor signals and search signals. The corpus performed appropriately for those audiences. The data does not show that PR has gotten worse. It shows that an entirely new audience emerged with different criteria.

## Three factors drove the decline

The drop in "good for AI" traces to three contributors that, taken together, explain the shift:

1. **Authority eroded as formal attribution softened.** Consistent name-and-title attribution is how AI engines tie a quote to a verified, identifiable source, and its absence

pulls Authority scores by ten points or more on affected releases. The decline runs directly through the data: From a mean of 81 in the Pre-AI era to 71 in the AI era.

2. **Press releases drifted from news announcement to executive feature.** As newsroom-driven coverage contracted, press releases shifted toward executive-centric storytelling. The lead paragraph increasingly served to feature the C-suite by name rather than to announce news. The trade-off is structural: Concrete dates, sourced claims, and specific numbers gave way to brand voice and executive presence. The Scout penalizes the trade because answer engines parse specifics, not aspirational tone.
3. **Wire-window compression squeezed out editorial discipline.** Faster turnarounds in PR mean less time for the editorial steps AI citation rewards: Second-voice approvals, dateline verification, formal attribution, and citation of external sources. As release cycles compressed, those steps got shorter or dropped. The data shows the consequences: Quotability went bimodal as second voices missed the approval window, and Citation Potential never lifted off its floor.

A closer look at the Authority decline, since it carries the most weight of the three factors:

| Era                    | Specificity | Quotability | Authority | Originality |
|------------------------|-------------|-------------|-----------|-------------|
| Pre-AI (2012–2019)     | 50          | 89          | 81        | 93          |
| Transition (2020–2022) | 50          | 91          | 74        | 92          |
| AI era (2023–2026)     | 49          | 86          | 71        | 90          |

The Scout's Authority component scores datelines, boilerplate, formal executive titles, contact blocks, and the "FOR IMMEDIATE RELEASE" timing tag. Datelines and boilerplate held steady across the corpus. The component that declined was **formal executive attribution**. The corpus drifted toward shorter, less formal constructions ("Jane said," "Smith added") and away from the AP-correct "Jane Smith, chief executive officer of Company X, said ...." This is a broader B2B writing trend, not specific to any single agency or segment. It is also among the easiest patterns to reverse. A style choice within the editorial process fixes it.

*"For the first time, we can measure what the editorial pass contributes against criteria that did not exist before late 2022."*

## Four signals worth a closer look

### Citations are the silent killer, by industry convention

Mean Citation Potential across the corpus runs about 11 of 100. The Scout's formula floor is 6. Across all 14 years and all three eras, B2B content in this segment rarely cited external research organizations like Gartner, Forrester, or Deloitte. It rarely used sourced-claim language ("according to," "research shows"). It rarely embedded external URLs.

This is by industry convention. Complex, highly regulated B2B segments write from internal authority. That serves editors well, since well-sourced internal claims read as credible. It puts AI engines at a structural disadvantage, since they weigh citation density heavily as a way to confirm a claim is verifiable. The fix requires client approval to reference third-party data or partner content. Agency execution alone cannot solve it.

### **Quotability is bimodal, and constrained by the approval window**

Releases with two or more named speakers and AP-style attribution scored 80 to 100 on Quotability. Releases with one quote, no quote, or non-AP attribution scored 40 to 60. The constraint is rarely editorial. The agency knows how to write a strong second-voice quote. The constraint is the wire-distribution timeline: Customer or partner quotes require sign-off, and many releases ship before a second voice clears the approval window.

This finding aligns with broader industry data. Press releases that include direct quotes receive approximately 40 percent higher pickup rates from media outlets than releases without quotes. Quotes serve as a coverage signal too, and they always have. The audit's strongest single recommendation lands the same way on both criteria: Get the second voice in.

### **Readability rarely exceeds 50, a real industry tension**

Average sentence length across the corpus runs 25 to 30 words. AI engines reward sentences in the 18-to-22-word range. Compliance review, qualifying clauses, and disclosure language push sentences longer in highly regulated B2B segments. This is the most stable structural feature in the corpus across all 14 years. It is a real tension between client-side review constraints and AI parseability, and no individual editorial pass fully resolves it.

### **Timeliness and Entity Disambiguation offer meaningful improvement room**

Two newer scoring categories joined The Scout in May 2026, and they surface signals legacy frameworks miss. Timeliness measures dateline freshness, present-tense markers, future-date specificity, and vague-timing penalties. Entity Disambiguation measures consistency of canonical names, formal name-and-title attribution, and pronoun discipline.

Across final releases, Timeliness averaged 34 to 39 across eras and Entity Disambiguation averaged 60 to 64. Across draft releases the picture is starker: Timeliness averaged just 16 to 20. The paired sub-analysis confirms this directly. Releases scored at both stages showed the editorial pass adds a median of 22 Timeliness points to the same release. Both categories sit firmly in mid-range for final work, and Timeliness, in particular, represents the largest pool of agency-controllable score uplift in the audit. A handful of fixes lift most releases in the corpus by five to ten overall points: Replace "soon" with a date, repeat the canonical company name four or more times, and tighten exec attribution.

## **The findings track broader AEO research**

Two points of external context anchor this audit's findings:

**Press release pickup is already a small-numbers game.** Industry data places original media coverage from a press release at roughly two to three percent of distributed releases, with about 18 percent receiving any pickup from at least one outlet within 24 hours. Releases that include quotes see meaningfully higher pickup rates than those without.

### **AI engines do not lean on press releases the way they lean on Reddit and**

**Wikipedia.** 5W Public Relations' 2026 AI Platform Citation Source Index analyzed more than 680 million citations across ChatGPT, Claude, Perplexity, Gemini, and Google AI Overviews. It found Reddit cited at roughly 40 percent frequency across all engines and Wikipedia at 26 to 48 percent on ChatGPT. Press releases compete for the remainder. The implication for PR teams is structural: AI visibility hinges on being citable from a brand surface AI engines will lean on, such as a newsroom, blog, or partner page. Press releases feed that surface.

The audit's findings track what other observers see across the PR and AEO research community. The originality of this dataset is its scale and longitudinal range. The conclusions are the same.

## **Five fixes for PR teams, ranked by leverage**

Five practical adjustments, ordered by leverage and ranked from data:

1. **Push for at least one external link.** A research source, partner page, or product page. This single change moves Citation Potential off the floor and lifts overall scores noticeably across every era and genre tested, but it requires client buy-in.
2. **Build the second-voice quote into the draft.** When the second speaker is identified early, approval timing usually works. The audit's widest single score gap traces to whether the second voice made it through the approval window.
3. **Restore AP-correct executive attribution.** Use "Jane Smith, chief executive officer of Company X, said ..." rather than "Jane said," "noted," or "stated." This is the only category that meaningfully declined as AI emerged, and it sits fully within agency control.
4. **Fight for one short, declarative lead sentence.** The compliance battle for the rest of the release may be unwinnable. The lead sentence usually is winnable. AI engines weight the first 100 words heavily.
5. **Stop assuming professional polish equals AI visibility.** These are different competencies. Professional polish protects you from looking unprofessional. AI visibility determines whether ChatGPT cites you when a prospect asks an industry question. Both still matter. Only one is on most teams' checklists today.

## **The audit's known limits**

- Whether AI-visible releases convert to coverage at higher rates. No one has clean data on this yet. Anyone who claims they do, does not.
- Whether one wire service indexes faster into AI engines than another. Anecdotally, the major wires surface in answer engines within 48 to 72 hours; smaller wires take two to three times longer. This needs systematic study.
- Whether the correlation between Scout scores and actual AI citation rates is causal or merely descriptive.

- How releases prior to 2012 would score. The dataset begins in 2012 because complete file archives prior to that year are not available.

## **Pitch'd is the solution**

This audit catalyzed everything. We ran it on our own portfolio first because we wanted to know how far the goalposts had moved before we built tooling to address it. The five fixes above ship as defaults inside Pitch'd's release builder. The Scout makes them a checking step before distribution. The Ump catches releases that lack newsworthy substance before they enter the writing process at all.

We made the toolkit free because the gap between traditional PR craft and AI visibility is structural, not segment specific. Structural problems get easier to solve when the tools land in every team that needs them. Most PR teams write the 99 percent of releases that never see a wire-service AI optimizer, since they do not carry an enterprise wire contract. Those teams matter.



*The agency that produced this dataset is also the firm  
that built the response.*



## Sources and references

- Muck Rack, 2024 State of Journalism (1,890 journalists surveyed across the U.S., U.K., Canada, and India).
- 5W Public Relations, AI Platform Citation Source Index 2026 (analysis of 680+ million AI citations across ChatGPT, Claude, Perplexity, Gemini, and Google AI Overviews. Published April 2026.)
- Edelman Trust Barometer 2024 (annual global trust survey, public data).
- Forrester / SiriusDecisions B2B buying journey research (60 to 70 percent of buying journey complete before sales contact).
- USC Annenberg Center for Public Relations, Global Communication Report (annual practitioner-side data).
- Institute for Public Relations (IPR) research on earned-media trust and PR ROI.
- HubSpot State of Marketing 2024 (AI adoption among marketers, public data).
- OpenAI public statements and industry trackers on ChatGPT weekly active users (early 2026).
- Industry-reported press release pickup rate benchmarks: About 18 percent pickup within 24 hours; two to three percent original media coverage; about 40 percent pickup uplift on releases with direct quotes (multiple PR-industry sources).
- The Scout scoring methodology, available at [pitchdapp.com/tools/the-scout.html](https://pitchdapp.com/tools/the-scout.html). Source code for the scoring functions is available on request.

## Methodology

**Dataset and sources.** The 758 releases analyzed in this audit represent the editorial output of St. Nick Media Services across 2012 to 2026. They were sourced from the agency's Dropbox and iCloud archive with zero overlap. 463 were marked FINAL and distributed via paid wire service. 295 were substantive drafts for which no FINAL version exists in the archive. Versioning duplicates were merged before scoring, and each unique release appears once. Within the FINAL set, 156 releases had at least one earlier substantive draft preserved alongside the FINAL, and those formed the paired sub-analysis. Industry concentration: The entire archive comes from one complex, highly regulated B2B segment (insurance, insurtech, and insurance-adjacent technology) with specialized vocabulary and multi-party approval workflows. This single-segment concentration was chosen because complete archive files are available across the full 14-year span. Findings are expected to generalize to other regulated B2B segments with similar conventions, and the Q3 2026 cross-industry follow-up audit will test that directly.

**Scoring tool.** All releases were scored against The Scout, an eight-signal AI-visibility scorer (Specificity, Quotability, Authority, Originality, Citation Potential, Readability, Timeliness, Entity Disambiguation) with genre-aware weighting. The Scout was built in early 2026 and updated in May 2026 with the Timeliness and Entity Disambiguation signals to align with how answer engines parse and cite content. Era binning follows the AI-search timeline: Pre-AI 2012 to 2019 (n=215), Transition 2020 to 2022 (n=340), AI era 2023 to 2026 (n=201). For era comparisons in the longitudinal tables, the FINAL-only cohort preserves apples-to-apples comparison with the original 2026 audit. The paired sub-analysis uses the within-release delta.

**Disclosure.** The Scout is a Pitch'd product. The dataset is sourced from St. Nick Media Services, the parent firm of Pitch'd. The audit and the tool are both Pitch'd work. Methodology, raw aggregated scores, and source code for the scoring logic are available on request.

## About Pitch'd

Pitch'd is a B2B SaaS press release builder and AI-visibility scoring tool, live at [pitchdapp.com](https://pitchdapp.com). St. Nick Media Services operates it. The Pitch'd toolkit centers on a paywalled release builder that integrates a pre-finalize AI-visibility audit. The same eight-signal Scout scoring used in this audit runs automatically before any release exports. Surrounding the builder are five free standalone tools: Home Plate (boilerplate scoring), The Ump (a seven-question newsworthiness check), the Headline Analyzer, the AP-Style Coach, and the Cliché Detector. A free preview of The Scout is also available. All five free tools work without a Pitch'd account. Pitch'd serves PR pros, agencies, and in-house communications teams that do not carry enterprise wire contracts but still need to ship AI-optimized releases. Pricing starts at \$149 per release with paid plans at \$199 per month and \$349 per month for agency teams.

## About St. Nick Media Services

St. Nick Media Services has operated as a B2B public relations firm in continuous operation since 1998. The firm specializes in communications for companies in complex, highly regulated industries: Sectors where specialized vocabulary, multi-party approval workflows, and compliance review shape the editorial process.

*Pitch'd is part of the firm's 2026 product line, built from 14 years of editorial output and the methodology behind this audit. Both the dataset and the response are agency work.*

## Contact

**Founder:** Jennifer Overhulse

**Email:** [jen@pitchdapp.com](mailto:jen@pitchdapp.com)

**Web:** [pitchdapp.com](https://pitchdapp.com) · [stnickmedia.com](https://stnickmedia.com)

For interviews, editorial requests, or questions about publishing this report's findings with attribution, please reach out.

---

© 2026 St. Nick Media Services / Pitch'd. May be quoted with attribution.