SPOTLIGHT × Profound

The **AI Influence** Economy

How Earned Authority in B2B Buying Influences AI Visibility & Sentiment

OCTOBER 2026

CONTENTS

In partnership with Profound, we read more than a million AI answers across B2B software and services markets, and traced the sources behind them – what the engines cite, and who they name.

1 **Who Speaks for You**
Why AI visibility is an orchestration problem, not a technical one – and how far the answer layer moved in twelve months.

2 **What the Machines Reward**
The fifty percent cliff on owned content, which analyst firms are gaining ground, and the job each review platform actually does.

3 **The Third Voice**
Independent experts and practitioners – and how few of them it takes to carry a market.

4 **What They Say About You**
Being named is not the same as being described well – one category's biggest gain in visibility, and the two cautions now exposed at scale.

5 **No Two Markets Alike**
Four markets profiled the same way, year-over-year: vendors, domains, review platforms, analyst firms.

6 **What To Do**
The mandate by function, and the one thing to take from all of it.

WHAT CHANGED IN TWELVE MONTHS

Your buyers are asking AI systems which vendors to shortlist. The answer they get is assembled from whatever sources the engine trusts. So the first question should not be how to rank higher – it should be who the engine is listening to, and what they are saying about you. To find out, we read more than a million AI answers across B2B software and services markets, tracing the sources behind them. The year-over-year figures on this page come from four of those markets, measured twelve months apart on the same six engines.

THE DEMAND

33M

B2B buying questions put to ChatGPT alone, every single day.

Across every major AI assistant, 65 to 100 million a day on a conservative estimate – the floor, not the ceiling. The shortlist now forms before a buyer ever reaches your site.

20M → 33M
in twelve months

THE SUPPLY

10 → 8

Sources cited per answer, down a quarter in twelve months.

Four markets, four independent sets of questions, and the same fall in every one of them. Separate measurements, one result. More questions, answered from less evidence.

−23 to −24%
all four markets

THE EVIDENCE

+72%

Analyst firms and customer reviews gained share of every answer.

They gained in all four markets while the trade press fell in every one and Wikipedia dropped out of the top fifty everywhere. The pool shrank. This layer did not.

4 of 4 markets
+43% to +187%

More questions, answered from fewer sources – and the sources holding their place are the ones with a human voice of authority behind them.

WHAT YOU'LL WALK AWAY WITH

01

You'll see who is named, who is cited, and who dropped out.

The most-cited source in HR technology is now YouTube, absent from the top five a year ago. In IT services McKinsey left the top five vendors; EY took its place. We rank the vendors, sources, review sites and analyst firms in four markets, year-over-year.

02

You'll see that no two markets have the same trust layer.

Clutch is the most-cited source of any kind in IT services. In HR technology and data security it does not reach the top five at all. Below the two or three largest analyst firms, the names that carry weight change from market to market.

03

You'll see why showing up does not mean showing up well.

The vendor with the biggest visibility gain in its category is also the one the engines describe worst on price and deployment. Getting mentioned more often put a worse description in front of more people.

04

You'll see why counting citations alone can mislead you.

Most answers that name an analyst firm never link to it. Anyone measuring this channel by citations alone is seeing a fraction of what it actually does.

PART 1

Who Speaks for You

This is not an optimization problem

Most of this market has decided that AI visibility is a technical discipline. It has acronyms already – AEO, GEO – and a familiar shape: work out what the engines reward, then produce more of it. Schema, structure, freshness, more pages, better markup.

That work matters. Think of it as your résumé. You have to have one, it has to be accurate, and it has to be a fair reflection of what you actually offer. Nobody gets hired without one.

In a job search, the résumé gets you the interview and the references get you the offer – checked late, once someone has already decided they are interested in you.

The machines run it backwards. Everything said about you by someone else – analysts, review platforms, communities, creators, journalists, customers – is your reference list, and it is the first thing they read. They decide from that reading whether your résumé is worth opening at all.

Those are your market's **trust layer** – the signals the engines actually weigh when a buyer in your category asks who to buy. Every market has one. No two are alike, and yours likely moved this year.

WHAT SITS BETWEEN THE QUESTION AND THE ANSWER



When your references agree, the story is clean and the buyer moves. When they disagree – your website says one thing, your reviews say another, the analyst covering you says a third – the engine hedges, the answer gets more vague, and the buyer keeps looking.

The work is getting every channel and every voice telling a consistent story about you, so the buyer meets one version of you instead of five.

That is an orchestration problem, not an optimization problem. And the prize is bigger than a mention.

77% of the time, the business goes to the vendor the buyer contacts **first**

6SENSE · 2025 BUYER EXPERIENCE REPORT

Not because the first vendor is better, but because they frame the problem, set the criteria, and become the standard everyone evaluated after them is measured against. AI answers are increasingly where that first contact gets decided.

Counting citations is the wrong scoreboard. A citation tells you an engine linked to a page. It does not tell you whether the engine recommended you, or how it described you when it did. A vendor can gain substantial visibility while the engines caution against it in key areas. That is not a hypothetical (we'll tell you more in Part 4).

So the question this report works on is not *how do I get cited more*. It is *which voices the machines are actually listening to – and why it matters that they are all telling the same story about you.*

Ten sources became eight

We asked the same questions twelve months apart. Same prompts – identical, verified word for word. Same six engines. Roughly the same number of answers.

The engines cited a quarter fewer sources.

EX 1.1 Sources cited per answer

CATEGORY	JULY 2025	JULY 2026	
HR Technology	10.3	7.9	−23%
Data Security & Privacy	10.9	8.3	−24%
Digital Commerce	10.6	8.1	−23%
IT Services & Consulting	10.0	7.6	−24%

Four different markets. Four independent sets of questions. The same contraction in every one.

This is the fact everything else in this report sits on top of. The answer layer is not expanding to accommodate everyone who wants to be in it. It is getting smaller, and **every answer a buyer reads now rests on a thinner evidence base than it did a year ago.**

And the shortlist got shorter

The contraction is not only happening to sources. It is happening to vendors.

EX 1.2 Share of answers naming the category leader

CATEGORY	LEADER	JULY 2025	JULY 2026	
Digital Commerce	Shopify	36.9%	52.3%	**+15.4pp**
HR Technology	Workday	29.6%	42.6%	**+12.9pp**
Data Security & Privacy	Microsoft	50.1%	60.9%	**+10.8pp**
IT Services & Consulting	Accenture	30.2%	35.9%	**+5.7pp**

In all four markets we studied, the leading vendor gained presence over twelve months, and in three of them the gain was ten points or more. Shopify is now named in more than half of all digital commerce answers. Microsoft in three of every five.

Those gains came from somewhere. Adobe lost two and a half points, WooCommerce two. BambooHR led HR technology a year ago and now sits third. IBM was the third most-named vendor in data security and is no longer in the top five at all.

The top of the list is getting heavier and the bottom is thinning.

The two contractions compound. Fewer sources back each answer, and the answers themselves name a narrower set of vendors. When a buyer's shortlist starts with an AI answer, you do not lose the deal later. You were never in it.

The sources changed, not just the number

The count fell by a quarter. The mix changed far more than that.



Ex 1.3 – Share of AI answers drawing on each source type, July 2025 to July 2026. The bar spans the change and the arrowhead marks 2026.

The trade press was the backbone, and it is being dismantled. This is the single largest displacement in the data, and if your program has been built on earned media, it is the number to sit with.

Video arrived from nothing. A year ago YouTube was not in the top twenty sources in any category we studied. Today it is the most-cited domain in two of them.

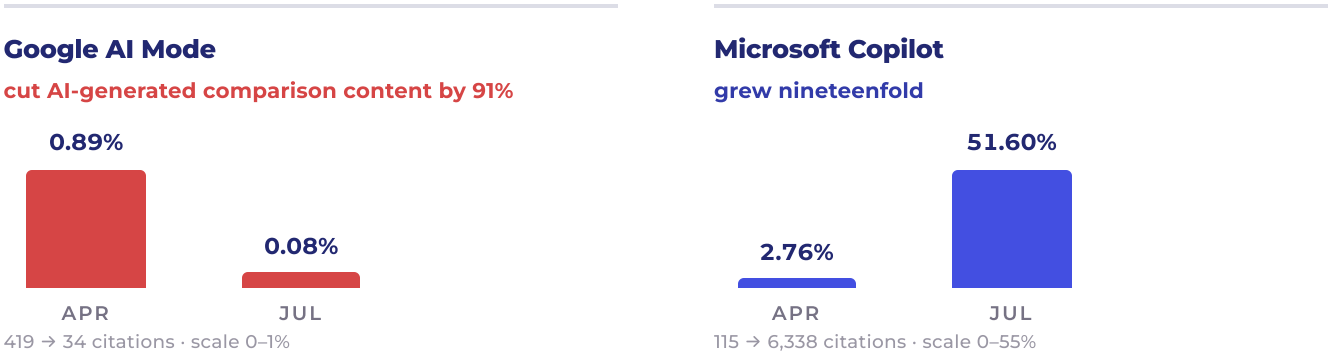
Wikipedia was deprecated everywhere. It ranked seventh in HR technology and it now ranks two hundred and tenth.

Comparison content arrived from zero. Four sites account for almost all of it – `worldmetrics.org`, `zipdo.co`, `gitnux.org` and `wifitalents.com` – each publishing close to 72,000 "best tools" lists, across ninety-plus categories apiece. All four describe an explicitly AI-assisted editorial process in their own words, and all four present as research operations with named analysts and published methodologies. They were first cited in February. By July one of them was a top-five source in every market we studied.

The engines are not being fooled here. They are being efficiently served: a buyer asks for the best ecommerce platform and these sites have a page titled exactly that, ranked, current and structured. If your answer to AI visibility is to publish more, this is what you are publishing against.

And the engines are already diverging on it

That share is an average, and the average hides a split.



Ex 1.4 – Share of each engine's citations going to the four AI-assisted comparison sites in data security, April against July 2026. Each panel is scaled to its own range, so compare the shapes and not the bar heights. The same direction held in all four markets we studied.

Every Google property moved the same direction, significantly reducing reliance on this type of comparison content at scale; Copilot, Perplexity and ChatGPT all moved the other.

A source that rises on being easy to retrieve can fall just as fast, and no amount of formatting will save it once an engine decides to weigh it differently. **What your buyer is told now depends on which engine they asked.**

One engine appears to have made that call by judgment. The industry is now building the means to make it automatically. In August 2026 Anthropic began watermarking Claude's text output with an invisible statistical mark. It is one of roughly 190 signatories to the EU transparency code. A watermark does not prove who wrote a page. It shows that a model was probably involved.

An engine does not need proof to act. For a site publishing 72,000 comparison lists, a likely mark is enough to stop citing all of them.

A prediction: an answer layer where machine-made comparison sites already reach one answer in five will not stay that way once the engines can identify them cheaply. What survives is what a machine cannot produce.

One layer gained while others dropped

The contraction did not fall evenly.

While the press fell in all four categories and Wikipedia lost between seventy-five and ninety percent of its share, **analyst firms and customer review platforms gained ground in every single one.** The third-party layer did not grow. It reshuffled – away from editorial, toward the structured and the verifiable.

EX 1.5 Analyst firms and review platforms, share of all citations

CATEGORY	2025	2026	
IT Services & Consulting	1.2	3.5	**+187%**
HR Technology	2.4	4.2	**+78%**
Digital Commerce	1.1	1.9	**+77%**
Data Security & Privacy	3.8	5.4	**+43%**

The engines cite fewer sources than they did a year ago, and the ones holding their place are the ones you do not write.

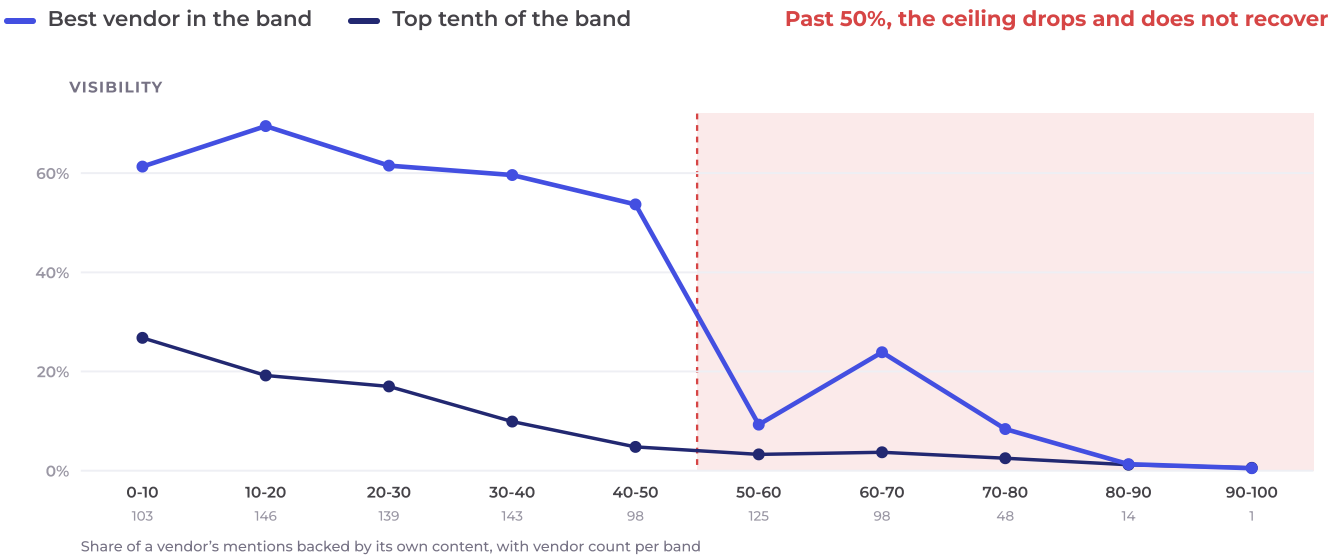
So how far can your own content take you? Part 2 starts there.

PART 2

What the Machines Reward

Nobody wins on their own website

We plotted 915 vendors across ten B2B categories against a single variable: how much of what the engines say about them is backed by their own content.



Ex 2.1 – Past roughly 50% reliance on owned content, visibility collapses. 915 vendors, ten B2B categories, six AI assistants, vendor-shopping questions only.

Below roughly half, owned content is no obstacle at all. The strongest vendor in each of those bands reaches between 54% and 70% visibility. Plenty of them lean hard on their own material and still win, which is why the obvious conclusion here – that owned content is a liability – is the wrong one.

Past half, it stops. The best performer in the 50–60% band reaches 9% visibility, and no band above it recovers past 24%. By 80–90%, the strongest vendor sits at 1.3% visibility.

Not one vendor in 915 is both heavily reliant on its own content and highly visible. That is not a gentle correlation someone can dispute. It is a cliff, and it sits at about fifty percent.

If more than half of what backs your AI mentions is your own material, you are not approaching a visibility ceiling. You are past one.

The clearest illustration is a vendor whose content the engines reward heavily.

In data security, `sentinelone.com` is the single most-cited domain in the entire category. Not the most-cited vendor site – the most-cited source, ahead of every publication, every analyst firm and every review platform. Almost all of it is one library: 94.3% of those citations point at `/cybersecurity-101/`, a set of category explainers on what ransomware is, what a SIEM does, how zero trust works.

The engines use it constantly. It answers the definitional question and the vendor's name is on it.

In twelve months, SentinelOne's visibility went nowhere.

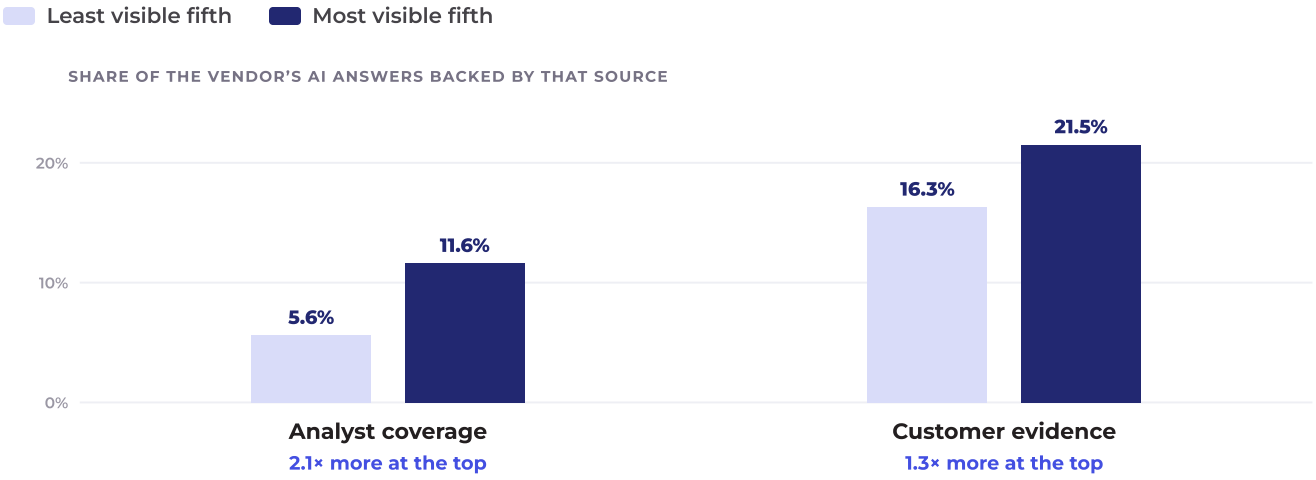
The explainers get cited when a buyer asks what something *is*. They are almost never there when a buyer asks who to *buy* – not one page in that library names a competitor.

Owned content can succeed as content and fail as influence. The distinction is worth applying to your own best-performing pages: are they turning up when a buyer asks what something is, or when a buyer asks who to buy? Those are different jobs, and only one of them moves a shortlist.

So what is carrying the story instead?

The most visible vendors have analyst coverage and customer evidence

Split every vendor into five bands by AI visibility, then look at what backs the top band against the bottom.



Ex 2.2 – Analyst backing roughly doubles from the least visible vendors to the most. Customer evidence is the larger channel at both ends.

Two things are true here and both matter.

Customer evidence is the bigger channel at both ends. More vendors have it, and it is present in far more answers than analyst coverage is, at every level of visibility. If you are choosing where to start from nothing, this is the larger surface.

Analyst coverage is the sharper differentiator. It roughly doubles between the bottom band and the top, while customer evidence rises by a third. Analyst backing is the thing the most visible vendors have that the least visible ones do not.

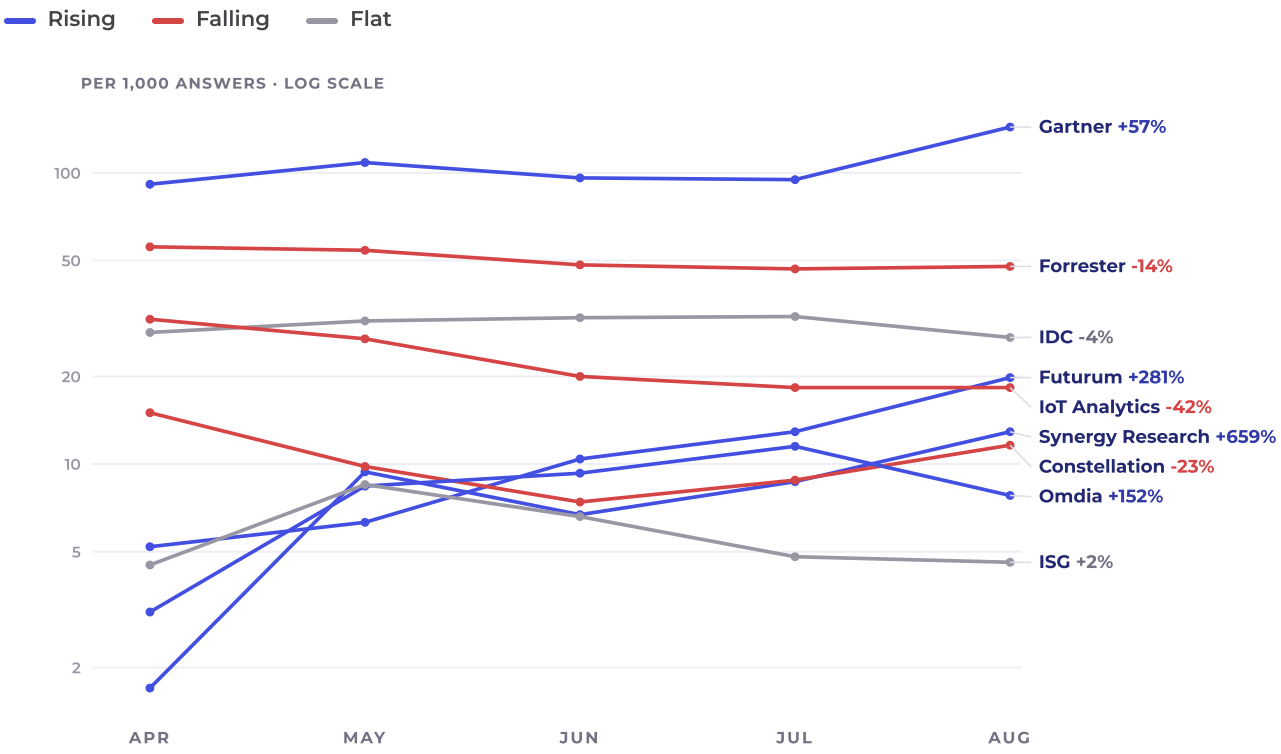
The pattern is not one relationship or the other. Vendors in the top quintile have both, and the programs that treat these as competing line items are arguing about a question the data has already answered.

So which analyst firms, and which review platforms?

Analyst influence is not a fixed ranking

To see how the analyst channel actually behaves, we read 68,299 AI answers in a single category – digital workplace – across six platforms and ten topics, and counted how often each research and advisory firm was named in the text or cited as a source. Thirty-one firms turned up. Nine carry almost all of it, and they are not moving in the same direction.

Answers naming or citing each firm



Ex 2.3 – Nine firms, three trajectories. Digital workplace only: 68,299 answers, six platforms, ten topics, April to August 2026. Log scale, so a steeper line is faster proportional change, not a larger absolute one. Gartner Peer Insights excluded throughout. A firm counts once per answer whether it is named in the text, cited as a source, or both.

One firm is almost as big as the other eight combined. Gartner finishes at 143.6 mentions per thousand answers, up 57% in five months, and three times Forrester. Three firms hold 74% of all analyst presence, and the other quarter is where the movement is: the three fastest-growing firms in the set all sit outside that group. The useful question is not which relationship to fund, but which firms are moving in your category, and which of them carries the narrative you need to shift.

Below the top two, the order is being rewritten inside a single quarter. Gartner and Forrester hold first and second in all five months. Every other firm changes rank at least once: Futurum moves five places, Synergy three, and four more move two. Synergy Research Group is up 659% over the period, Futurum 281%, Omdia 152%, all three from under 6 per thousand. Futurum passed IoT Analytics in August; Synergy passed Constellation. A tier list built on last year's rankings does not have them on it.

A firm's standing is not the same in any two markets. Forrester fell 52% in data security while rising 173% in IT services – same firm, opposite directions, same twelve months.

There is no such thing as a firm's overall AI influence. There is only its influence in your category, this quarter. That is now something you can watch month by month, which means the return on an analyst relationship can be evidenced while there is still time to act on it.

Why citations undercount analysts

> *Most answers that name an analyst firm never link to it – two-thirds to four-fifths of them, depending on the market.*

The engine says *Gartner positions them as a leader*, and the citation list points somewhere else entirely. Measured by citations alone this channel looks like a rounding error. Measured by what the engines actually say – which is what EX 2.3 counts – it is substantial, and it is rising.

How analyst wins reach the engines

The research sits behind a paywall, so the engines cannot cite it. But the conclusions travel anyway – written up in the trade press, posted on LinkedIn, quoted in analyst blogs, announced by every vendor named in the report. The engine knows perfectly well who the analysts rated as leaders. What it does not have is the firm's own page to point at.

So amplification is a mechanism, not a victory lap – and there is a number for it.

In the same category, Google is one of the vendors. **Of every citation pointing at a Google URL, 17.5% goes to a page about what an analyst said.** Analyst-validation pages are the best-performing content on Google's own domain. One press release about a Magic Quadrant placement pulled 641 citations on its own, outperforming nearly everything else the company publishes.

The announcement, the blog post, the LinkedIn write-up, the customer-facing summary: that is the route by which a paywalled finding reaches a machine at all. An analyst win you do not activate is a win the models may never see.

That is the most useful thing in this section, because it points at work you have already paid for. **Every placement you have earned is an asset you can still activate** – the mention in a report from eighteen months ago, the rating nobody wrote up, the quote that went into one deck and nowhere else.

The opportunity is continuous activation rather than a launch moment. A placement announced once gives the engines a single surface to reach it through. The same placement written up, summarized, quoted and returned to across owned and third-party channels gives them many – and those surfaces keep working long after the report itself has gone back behind the paywall.

A prediction: the firms that open more of their research will convert reputation into citation, and they will do it while citations are getting scarcer. Analyst authority today arrives as something the engine remembers. It could arrive as something the engine can show – and the first firms to make that trade will be the ones that gain.

Analyst influence moves, and much of it never shows up as a citation. Review platforms are the reverse: impossible to miss in the data, and easy to treat as interchangeable.

Not all reviews do the same job

Review platforms need no argument made for them. In July 2026 they appeared in **27% of all data security answers** and 21% of HR technology answers – nobody is under the impression that the engines ignore them.

The useful question is a different one: **which platform, doing which job, in your market.**

EX 2.4 How AI uses each review platform

PLATFORM	HOW AI USES IT	WHAT TO OPTIMIZE FOR
G2	Volume validation. Precise star ratings and ease-of-use pros and cons, quoted at high review counts – "4.8/5 across 14,000 reviews." Cited most often, and most precisely.	Review volume and recency; ease-of-use and satisfaction Grid categories. The review count itself is leverage, because the engines cite the scale.
Gartner Peer Insights	Credentialing and shortlisting. Willingness-to-recommend percentages, "Customers' Choice," and structured sub-scores for integration, deployment and support.	Clearing the review-count threshold – thin counts leave niche vendors invisible. Customers' Choice, and the enterprise sub-scores.
TrustRadius	The justification engine. Balanced pros-and-cons synthesis plus verbatim reviewer quotes the engines lift directly into their reasoning.	Detailed, quotable, use-case-rich reviews. This is where an answer gets the language to explain why you fit.
PeerSpot	Numbers and mindshare. Supplies mindshare percentages and willingness-to-recommend data no other source has. Strong in enterprise, security and IT.	Presence and mindshare for enterprise and technical products; richer reviews, since it is aggregate-heavy and depth helps.
Capterra	SMB corroborator. Almost always co-cited to reinforce G2 or Gartner, and rarely drives the narrative alone. Skews small business and ease of use.	Presence and rating for corroboration, especially with SMB buyers. Not a standalone driver.

The platforms are not interchangeable, and the difference is not quality – it is function.

The horizontals get you onto the list. G2 and Gartner Peer Insights carry the breadth and the volume, and they are what the engines reach for when assembling a shortlist. If you are starting from nothing, start here.

The category specialist can outrank them in its own market. In IT services, Clutch is not merely the leading review platform – it is the **single most-cited domain in the entire category**, ahead of every publication, every analyst firm and every vendor site, at twelve times the volume of the next platform. A strategy built only on the horizontals will miss the source that matters most in the market where that happens to be true.

The depth platforms shape how the answer reads. TrustRadius and PeerSpot sit near the bottom on raw volume – between 0.04% and 0.21% of citations – and their content is long-form and specific about deployment and support. When an answer explains *why* a vendor suits a particular buyer, this is often where the language comes from. Low frequency, high influence on wording. In the category panels at the back they look negligible; that measures how often they are cited, not how much they shape the answer when they are.

Capterra corroborates. It rarely leads and frequently confirms.

And the leading platform is different in every market

G2 leads in HR technology and digital commerce. Gartner Peer Insights leads in data security. Clutch leads IT services at twelve times the volume of the next platform.

The leading platform changed in two of these four markets in a single year.

> *That turns review generation from a volume exercise into a targeting one.*

Between the role table and the market rankings you can say which platform does which job, and which one leads in your category – and that is the basis for a precision advocacy program. The right customer, on the platform that does the job you need done: breadth to reach the shortlist, depth to justify the choice. Rather than asking every willing reference for a review everywhere.

That is a trust layer moving. It moved measurably, in a single year, in all four markets we tracked.

But there is a third voice in that layer, and it is neither an institution nor a platform.

PART 3

The Third Voice

The independent experts your buyers already read

Analysts and review platforms are two of the three voices this report keeps naming. The third is a person rather than an institution – the independent expert, the practitioner, the creator with outsized standing in a market. The people your buyers were already reading, and who the machines are now reading too.

They do not sit in one place. The same individual is the byline in a trade publication that turns out to be a one-person shop, the practitioner answering a question in a subreddit, the outside voice in a video a vendor co-produced.

In ITSM – 100 buyer prompts across eight AI platforms – each channel does a different job. Some shape what a buyer thinks the category is, others decide who makes the list, others supply the reason one vendor beats another.

No channel carries a buyer the whole way, and two of them are near mirror images. Reddit is at its strongest on comparison and its weakest on implementation. LinkedIn is the reverse: almost absent when a buyer is comparing, and the strongest voice in the set when they ask what it costs.

EX 3.1 Channel presence by buying stage, ITSM

BUYING STAGE	TRADE PRESS	REDDIT	YOUTUBE	LINKEDIN
Education – what is this category	1.12	1.23	0.68	0.56
Shortlisting – which tool should we buy	1.20	1.24	0.97	0.86
Comparison – this one or that one	1.27	**1.38**	**1.28**	0.36
Cost – what will it run us	**1.31**	0.96	0.68	**1.38**
Lived experience – what was it actually like	–	**1.28**	0.96	1.03
Implementation – how do we run it	1.05	0.62	0.49	0.55

Presence index: 1.00 is each channel's own average, so Reddit at 1.38 is 38% more present in comparison questions than it is overall. A voice that carries at one stage is often absent from the next. The trade-press figure for lived experience rests on too few citations to report.

So an independent voice is something to activate on purpose, channel by channel, against the narrative you are trying to land or counter. Comparison and cost are not the same problem, and they are not answered on the same surfaces. A roster that treats every expert and every channel as interchangeable is spending against an average that describes none of these stages.

Third-party presence beats publishing

Part 2 showed the ceiling on your own content. This is the other side of it, and it is where these voices come in: **third-party content does not write itself.** It is produced by the people this section has been describing, and being present in their work is measurable at the vendor level.

EX 3.2 Own publishing against third-party presence, two ITSM vendors

	INVGATE	MANAGEENGINE
Content on its own site	24,560	18,259
Presence in third-party content	15	**1,239**
Named in AI answers	7.7%	**23.2%**

InvGate publishes more on its own site than ManageEngine does, and is named in a third as many AI answers. The difference is not output. **ManageEngine appears in other people's content roughly eighty times more often.**

Across all 23 vendors we measured, presence in third-party content predicts AI visibility better than a vendor's own publishing does.

So an influencer program is how a vendor gets into that content. These are the people writing it, and the way you work with them decides where a vendor gets named.

What to do, and why

EX 3.3 Working with independent voices

RECOMMENDATION	WHY
Do not select influencers on follower count alone	On one vendor's own influencer list, the person with the ninth-largest following is picked up by AI nearly eight times more than the other eleven combined.
Point influencers toward comparisons, not think-pieces	Thought leadership gets read, but lands in answers that name no vendors. "Which tool" and "how to choose" pieces are what put a vendor in front of a buyer.
Get influencers bylined in trade press	Trade press is the single largest thing AI reads in this market, bigger than Reddit, YouTube and LinkedIn put together.
Ask for LinkedIn Pulse articles, not just feed posts	Pulse articles are picked up roughly eight times more often than timeline posts.
Publish podcast transcripts, or put the episode on video	Podcast audio is almost never picked up. The same conversation on video, where transcripts are generated automatically, appears hundreds of times more often.
Support the publications influencers own	One independent expert's own articles are picked up by AI six times more often than Forrester, IDC and Everest Group's research combined.
Place your own experts in independent publications, including categories you do not yet lead	A competitor's employees published on IT operations integration, a category they are not known for. On that exact buyer question their mention rate more than doubled, from 7% to 16%.

None of that requires a bigger program, only a more precise one: the same budget, aimed at specific people, specific formats and the surfaces where buyers are actually asking.

Getting experts and customers to write about you, somewhere an engine can retrieve it, is one half of the problem. What they say when they do is the other. Part 4 starts there.

PART 4

What They Say About You

Everything so far has been about being named. This part is about being described – and they are not the same achievement.

Winning the mention, losing the description

Workday posted the largest visibility gain in HR technology: 29.6% to 42.6% of all answers in twelve months. Same questions, same six engines, July to July. By the measure most of this market uses, that is the year's best result in the category.

It did not come from the question everyone competes on.

EX 4.1 Workday visibility by topic

TOPIC	2025	2026	
Employee engagement	16.3%	44.4%	**+28.1pp**
Workforce analytics	50.4%	77.2%	**+26.8pp**
Remote workforce management	3.7%	15.9%	+12.2pp
Time and attendance	1.2%	11.7%	+10.4pp
HR software solutions	45.2%	47.2%	+1.9pp

The gains are all in specific use cases – and on the generic category question, "HR software solutions," it barely moved at all. Workday is winning the buyer who already knows what they need to do, not the one still asking what to buy.

The company it keeps changed with it. In answers that name Workday, the employee engagement specialist Culture Amp now appears alongside it 13.0 points more often than a year ago, while BambooHR – its long-standing rival in core HR – appears 10.8 points less. Workday is turning up in a different conversation rather than winning more of the old one.

Now look at what those same answers say about Workday.



Ex 4.2 – Share of sentiment that is positive, by theme. Point-in-time from 15 May 2026.

> *Workday has the biggest visibility gain in its category, while it loses on two narratives.*

A buyer asking which HR platform to shortlist now meets Workday more often than anyone. A buyer asking what it is like to actually deploy one meets a description that is four percent positive. Both of those are the same engine, in the same market, in the same week.

This is why visibility on its own is not influence. **You can win the mention and still lose the description** – and more visibility does not fix it. It puts a worse description in front of more people.

The fix is narrower and more useful than a bigger program: two attributes are named. That is the brief. Find the sources producing those two descriptions and work on those, rather than on the number that already looks like a win.

Which sources those are depends on the market you sell into.

PART 5

No Two Markets Alike

We profiled four B2B markets as samples – HR technology, data security, digital commerce, and IT services & consulting – to see how much the trust layer varies from one to the next, and what that does to the shortlists the engines build.

For each one: the vendors the engines name, the domains they draw on, the review platforms they cite, and the analyst firms they mention – every list a top five, both years, with what moved.

The names will be different in your market. The pattern will not.

No two trust layers are shaped alike. Clutch is the single most-cited source in all of IT services, at twelve times the next review platform – and it does not appear in the top five of HR technology or data security at all. The analyst firms in third, fourth and fifth place barely overlap between any two markets.

A program designed for one of these markets would be misallocated in the other three.

A year moves more than most planning assumes. The leading review platform changed in two of the four. Video entered the top five in every one. Wikipedia left the top five in three of the four and fell out of the top fifty in all of them. Even the sources that held rank moved share by multiples. Whatever your category's trust layer looked like when you last audited it, it does not look like that now.

The vendors move with it. These are not separate stories. In three of the four markets the leading vendor added ten points or more of presence while the sources behind those answers were being reshuffled underneath them – and vendors that had held top-five positions for years dropped out of them entirely.

One thing to know before you read the panels.

Units differ by column and do not compare across them. Vendors are a share of *answers naming the vendor*. Domains and review platforms are a share of *all citations*. Analyst firms are a share of *answers naming the firm in response text*. A vendor at 42% and an analyst firm at 1.5% are not forty times apart – they are measured on different instruments.

HR Technology

20,910 answers / 214,823 citations **('25)** → 21,018 / 165,549 **('26)**

MOST-SHORTLISTED VENDORS

share of answers naming the vendor

1	Workday	▲2 #3 · 29.63%	42.56%
2	Rippling	– #2 · 32.71%	40.60%
3	BambooHR	▼2 #1 · 35.53%	35.08%
4	Gusto	– #4 · 25.21%	26.02%
5	ADP	– #5 · 22.57%	24.66%

TOP CITED DOMAINS

share of all citations

1	youtube.com	NEW new	4.23%
2	peoplemanagingpeople.com	▼1 #1 · 7.42%	3.03%
3	worldmetrics.org	NEW new	2.11%
4	g2.com	NEW new	1.85%
5	reddit.com	▼2 #3 · 1.86%	1.79%

DROPPED OUT

~~forbes.com~~ #2 · ~~selectsoftwarereviews.com~~ #4

REVIEW PLATFORMS

share of all citations

1	G2	▲1 #2 · 0.66%	1.85%
2	Gartner Peer Insights	▼1 #1 · 0.75%	0.92%
3	Capterra	– #3 · 0.31%	0.49%
4	GetApp	▲1 #5 · 0.16%	0.48%
5	Software Advice	▼1 #4 · 0.28%	0.25%

ANALYST FIRMS NAMED

share of answers naming the firm

1	Gartner	– #1 · 0.36%	1.35%
2	Forrester	– #2 · 0.13%	0.09%
3	ISG	NEW new	0.06%
4	Info-Tech	▼1 #3 · 0.06%	0.05%
5	Brandon Hall	– #5 · 0.01%	0.05%

DROPPED OUT

~~Nucleus Research~~ #4

Data Security & Privacy

18,837 answers / 205,306 citations **('25)** → 18,596 / 153,748 **('26)**

MOST-SHORTLISTED VENDORS

share of answers naming the vendor

1	Microsoft – #1 · 50.11%	**60.86%**
2	CrowdStrike – #2 · 28.84%	**34.54%**
3	Palo Alto Networks ▲1 #4 · 24.05%	**31.04%**
4	SentinelOne ▲1 #5 · 20.63%	**19.96%**
5	Okta NEW new	**17.68%**

DROPPED OUT

~~IBM~~ #3

TOP CITED DOMAINS

share of all citations

1	sentinelone.com – #1 · 3.61%	**2.75%**
2	worldmetrics.org NEW new	**2.53%**
3	gartner.com ▲1 #4 · 2.03%	**2.51%**
4	youtube.com NEW new	**1.81%**
5	g2.com NEW new	**1.76%**

DROPPED OUT

~~microsoft.com~~ #2 · ~~expertinsights.com~~ #3 · ~~wikipedia.org~~ #5

REVIEW PLATFORMS

share of all citations

1	Gartner Peer Insights – #1 · 1.98%	**2.35%**
2	G2 – #2 · 1.01%	**1.76%**
3	GetApp NEW new	**0.36%**
4	TrustRadius – #4 · 0.10%	**0.21%**
5	PeerSpot ▼2 #3 · 0.25%	**0.20%**

DROPPED OUT

~~Capterra~~ #5

ANALYST FIRMS NAMED

share of answers naming the firm

1	Gartner – #1 · 4.59%	**5.58%**
2	Forrester – #2 · 2.08%	**1.00%**
3	IDC – #3 · 0.20%	**0.13%**
4	Info-Tech – #4 · 0.04%	**0.04%**
5	Frost & Sullivan – #5 · 0.04%	**0.02%**

Digital Commerce

18,852 answers / 199,132 citations **('25)** → 18,600 / 150,988 **('26)**

MOST-SHORTLISTED VENDORS

share of answers naming the vendor

1	Shopify – #1 · 36.90%	**52.34%**
2	BigCommerce – #2 · 27.57%	**26.63%**
3	Amazon – #3 · 21.55%	**23.29%**
4	Adobe – #4 · 16.18%	**13.67%**
5	WooCommerce – #5 · 14.84%	**12.80%**

TOP CITED DOMAINS

share of all citations

1	youtube.com NEW new	**2.11%**
2	shopify.com ▼1 #1 · 4.13%	**1.71%**
3	bigcommerce.com ▲1 #4 · 1.05%	**1.70%**
4	worldmetrics.org NEW new	**1.31%**
5	linkedin.com NEW new	**1.24%**

DROPPED OUT

~~wikipedia.org~~ #2 · ~~reddit.com~~ #3 · ~~forbes.com~~ #5

REVIEW PLATFORMS

share of all citations

1	G2 ▲1 #2 · 0.32%	**0.65%**
2	Clutch ▲1 #3 · 0.20%	**0.55%**
3	Gartner Peer Insights ▼2 #1 · 0.33%	**0.42%**
4	TrustRadius ▲1 #5 · 0.04%	**0.04%**
5	GetApp ▼1 #4 · 0.05%	**0.04%**

ANALYST FIRMS NAMED

share of answers naming the firm

1	Gartner – #1 · 0.52%	**1.50%**
2	Forrester ▲1 #3 · 0.14%	**0.32%**
3	IDC ▲1 #4 · 0.03%	**0.04%**
4	IoT Analytics ▼2 #2 · 0.22%	**0.04%**
5	Info-Tech NEW new	**0.02%**

DROPPED OUT

~~Frost & Sullivan~~ #5

IT Services & Consulting

13,197 answers / 132,367 citations **('25)** → 13,019 / 98,888 **('26)**

MOST-SHORTLISTED VENDORS

share of answers naming the vendor

1	Accenture – #1 · 30.20%	**35.88%**
2	Deloitte ▲1 #3 · 26.67%	**31.63%**
3	IBM Consulting ▼1 #2 · 28.45%	**27.51%**
4	Capgemini – #4 · 17.29%	**18.54%**
5	EY NEW new	**16.54%**

DROPPED OUT

~~McKinsey~~ #5

TOP CITED DOMAINS

share of all citations

1	clutch.co NEW new	**2.65%**
2	linkedin.com ▼1 #1 · 2.01%	**2.17%**
3	youtube.com NEW new	**2.01%**
4	corsicatech.com NEW new	**0.91%**
5	worldmetrics.org NEW new	**0.79%**

DROPPED OUT

~~managementconsulted.com~~ #2 · ~~wikipedia.org~~ #3 · ~~reddit.com~~ #4

REVIEW PLATFORMS

share of all citations

1	Clutch – #1 · 0.73%	**2.65%**
2	G2 – #2 · 0.23%	**0.21%**
3	Gartner Peer Insights – #3 · 0.14%	**0.19%**
4	TrustRadius – #4 · 0.03%	**0.07%**
5	Capterra NEW new	**0.01%**

DROPPED OUT

~~PeerSpot~~ #5

ANALYST FIRMS NAMED

share of answers naming the firm

1	Gartner – #1 · 1.01%	**2.91%**
2	Forrester – #2 · 0.40%	**1.08%**
3	IDC – #3 · 0.28%	**0.63%**
4	Everest Group ▲1 #5 · 0.02%	**0.22%**
5	ISG NEW new	**0.08%**

DROPPED OUT

~~IoT Analytics~~ #4

PART 6

What To Do

Start with the shape of your own market

Every category has a trust layer, and no two are shaped alike. Some run on peer review, some on analyst authority, some on trade press and community with almost no analyst signal at all. Find out which one you sell into before you spend another quarter against it.

The mandate, by function

FUNCTION	WHAT TO DO	WHAT IT GETS YOU
CMO	Stop funding output. Fund the evidence you cannot manufacture – and read each engine separately, because a composite hides the variation.	A defensible answer to *"why does the AI recommend them and not us?"*, and a reason to move budget that survives the CFO.
Analyst Relations	Measure what the engines *say*, not only what they link. Then activate every analyst win on a surface the models can reach.	The budget argument you have never had evidence for, and a use for placements you already own.
Customer Marketing	Find which platform leads *your* category before funding review volume. Breadth to reach the shortlist, depth to justify the choice.	A specific platform to fund and a specific reason – and a move from reference desk to shaping the answer.
Influencer Relations	Build the roster from what the engines cite, not from follower counts. Reach describes an audience; citations describe influence.	The first quantified evidence that your channels move AI answers, and a roster you can defend.
Product Marketing	Audit the narrative, not the mention. Find the two attributes where you are described worst and treat those as the brief.	You stop guessing which objections are killing deals, and learn where each one came from.
Content & SEO	Stop competing on output. Publish to be cited as a source – and know that being the source is not being the choice.	An exit from a volume race you cannot win, and budget redirected to work that compounds.
Communications / PR	Change what goes into trade press, not how much. Put independent experts and analyst findings into the placements you already earn.	Trade press is still the largest single source the engines read, and the largest decline in this report. The channel is not the problem; what fills it is.

The window is open now

Ten sources became eight. Fewer slots, far more questions, and the sources that kept their place were the ones with a person behind them – analysts, customers, and the practitioners who already work in your category.

The slots that opened went to content produced at a scale no marketing team can match, and there is no point racing it. But the layer that gained anyway is the one a better-funded competitor cannot simply buy more of. Most of what it takes, you already have.

The analyst briefings, the reference customers, the experts who know your product: that work is already sitting in your organization, spread across five or six budgets and measured five or six different ways. Briefing counts, review volume, follower reach – every channel keeps its own books, and none of the numbers convert.

The answer layer is the first place all of them show up together – the same answers, in front of the same buyer, where you can see what each one is actually worth. That is the measure that has been missing – not how many analysts you briefed or how many reviews you collected, but which voices your market runs on and what they say about you when the shortlist forms. The honest answer to *what did it get us* has always been a story rather than a number. Now it can be a number.

So read every layer, not just the one you own, and put one person in charge of whether all of it tells the same story.

There are fewer slots than there were a year ago, and no extra credit for arriving late. Find out what the engines already say about you – then decide what it is worth to change it.

No single voice moves a market. Perception is assembled – and decisions made – from every voice your buyer already trusts. It moves when your story holds up in all of them.

METHOD

Prompt volume. The 33 million figure is from Profound, an AI visibility platform that measures prompt activity and answer content across the major LLM assistants. It represents daily B2B buying-related prompts on ChatGPT – questions about vendor selection, product comparison, category evaluation and purchase decisions. The prior-year figure of 20 million comes from the same source on the same basis, so the year-over-year change reflects growth in query volume rather than a change in method. Prompt volume at this level is not published by the platforms themselves and exists only through third-party instrumentation; it is best read as a measure of scale and direction. It is a separate dataset from the citation panels below and shares no sample with them.

The cross-platform estimate. The 65 to 100 million range extends the measured ChatGPT volume across the rest of the market using ChatGPT's share of assistant usage. Third-party estimates of that share run from 54% (Similarweb, May 2026) to 77% (StatCounter, June 2026). We adjust downward from the low end, on the basis that web-traffic measurement structurally undercounts enterprise-embedded assistants – Microsoft Copilot's roughly 420 million monthly users sit largely inside Microsoft 365, Gemini's inside Workspace – which are precisely the surfaces where B2B buying research is most likely to occur. The range assumes ChatGPT accounts for between one third and one half of all B2B buying prompts. A higher multiple is defensible; we have not used one. We report the floor.

The year-over-year panel. July 2025 against July 2026, across four B2B markets – **143,029 AI answers and the 1.32 million citations behind them**. Restricted to the six engines present in both windows: ChatGPT, Google AI Mode, Google AI Overviews, Google Gemini, Microsoft Copilot and Perplexity. The same questions were asked in both windows, so what moved is the market and not the measurement.

Other datasets. The analyst-firm trajectories are digital workplace, April to August 2026 – 68,299 answers across six platforms and ten topics, with Gartner Peer Insights excluded throughout. The buying-stage and independent-expert analysis is ITSM, November 2025 to August 2026 – 197,944 answers across eight AI platforms, drawing on 199,423 distinct cited URLs. Each is a self-contained window; figures are not comparable across them.

Units differ by measure. Vendor figures are a share of answers naming the vendor. Domain and review-platform figures are a share of all citations. Analyst-firm figures are a share of answers naming the firm in response text – text-derived, where everything else is citation-derived.

ABOUT SPOTLIGHT

Spotlight orchestrates influence across every credible market voice – analysts, customers, independent experts and AI systems – to put clients on the shortlist and move buyers to buy. More than 190 B2B organizations across software, tech and services.

spotlight-io.com

ABOUT PROFOUND

Profound is the full-stack marketing platform for the marketer of the future – understand, analyze, build and measure how your brand appears in AI search. Coverage spans ChatGPT, Google AI Overviews, Gemini, Copilot, Perplexity, Claude, Grok and DeepSeek.

tryprofound.com