

BRANDS2LIFE RESEARCH · REPORT #4
SEPTEMBER 2026

THE TECH INFLUENCE GRAPH

INFLUENCE BY WHO
GETS CITED AND WHERE,
NOT BY CROWD SIZE

BRANDS2
LIFE

A PARITEE AGENCY

B2B technology buyers have a longer, more considered path to purchase than consumers. The market has not resisted the influencer movement, but experts leading tech conversation are a very different breed from the creators working with consumer lifestyle brands. To rank the most valuable authorities across B2B tech, we must quantify influence by who gets cited and where, not by crowd size.

What follows is an AI-scale analysis of 10,138 citations extracted from 1,612 items published by 92 podcasts, newsletters, tech mainstays and trade outlets across AI and AI infrastructure, cybersecurity and developer platforms, from a full quarter: June – August, 2026.



CONTENTS

- 04. Executive Summary
- 05. Five Critical Findings
- 06. Takeaways for Communications Teams
- 07. What We Measured
- 08. Breadth Beats Volume
- 10. Three Categories, Three Separate Communities
- 13. The Beef Coefficient
- 15. Bridge Voices and Consistent Output
- 15. Putting the Analysis to Work
- 16. Methodology
- 19. About Brands2Life

| EXECUTIVE SUMMARY

Every influencer ranking measures audience size – typically by aggregate followers across active social media channels and usually throwing in visitors to a website or Substack. It's a formula that makes sense for consumer brands, and consumer influencers are wise to value brand collaborations around their potential audience.

But audience size is the wrong measure for B2B brands, and especially B2B technology brands. The buying cycle is long. The material is dense. A Chief Information Officer is hopefully going to spend considerably more time researching a tech partner than teenagers spend choosing their skincare routine.

The influence potential of the referring audience matters more in a major technology adoption than the sheer popularity of the brand. Put simply, tech buyers are not so impressed by “I’ve heard a lot about you” but instead “I’ve heard from several people I greatly respect about you.”

We set out to quantify this trusted influence, borrowing academia's favorite metric: citations. We analyzed 1,612 pieces of content published across a full quarter by 92 respected publications — newsletters, podcasts, trade press, tech stalwarts like WIRED and MIT Tech Review, and independent blogs — extracting every instance of these influential voices citing an expert by name: 10,138 citations in total. From there, we could zoom out to spot the patterns and similarities

10,138

CITATIONS EXTRACTED

1,612

PIECES OF CONTENT
ANALYZED

92

PUBLICATIONS, JUNE –
AUGUST 2026

FIVE CRITICAL FINDINGS

1

Being widely cited is rare, and valuable.

Of more than five thousand named experts in the corpus, only 277 were cited by two or more unaffiliated publications. Four were cited by ten or more. In other words, routinely putting out citation-worthy thought is tough. And getting the attention of a cross-section of influential press consistently is even tougher.

2

Breadth and volume are different.

Whether they admit it or not, every podcast host has their favorite guests. One prolific newsletter can cite the same person a dozen times in a month. But it takes independent editorial judgment, repeated time and again across publications and media types, to produce breadth.

3

Two mechanisms make trusted influence difficult to quantify:

- **Organizational masking**, where the organization gets the credit and the expert goes unnamed
- **Audience mismatch**, where a voice is cited heavily by outlets your buyers don't see

4

Nobody's wrong more often than an expert on the internet.

1,199 citations in the corpus are disagreements. For some experts in our study, a third of their citations are rebuttals. Being argued with is a form of authority that every share-of-voice tool scores as neutral coverage. Boldness is fundamental for a quotable spokesperson, as any PR pro will tell you. But it also invites disagreement, which is not a bad thing.

5

Our three categories barely overlap, and they credit differently.

AI, cybersecurity and developer platforms are three distinct communities. Only six people out of the field of 5,000+ were cited in all three.

- Cybersecurity outlets cite organizations, like a private lab or government agency, instead of named experts 62% of the time, versus just 31% in AI and 27% in developer platforms.

TAKEAWAYS FOR COMMUNICATIONS TEAMS

01

Treat cited experts like analysts, not press.

Academics and researchers are not in the business of breaking news, so you can invite them to pre-release briefings, offer them access to data, and collaborate often.

02

Maximize spokespeople as trusted authorities in the field.

Be generous with influential outlets in your space. The goal is not to appear in one article, but to establish a trusted sourcing relationship.

03

Aim for breadth, not volume.

Thirty mentions in one outlet is not nearly as valuable in your buyers' discovery journey as five citations spread across five different, respected outlets. Invest in relationships across publications and tastemakers, rather than a few deep relationships.

04

Don't obscure your spokespeople in the name of brand visibility.

Podcast hosts book guests, not labs.

01 WHAT WE MEASURED

We work with brands transforming our world, with a high concentration of clients in AI and AI infrastructure, cybersecurity, and developer platforms. For each category, we assembled a universe of publications that shape its conversation, and read everything they published across June, July and August 2026.

From every post, article and episode, we extracted each reference to a named third party and classified it two ways. First by form: quote, paraphrase, link with commentary, framework adoption, guest appearance. Second, and more consequentially, by role:

Expert Authority

Cited for their analysis, research, framework or expertise, or for something they built that the citer passes on as instructive.

Subject

The named expert, or their work, is the topic of the piece.

That distinction is the methodological core of this research. Share-of-voice platforms count both as simple “coverage”. But conflating them gives popularity the same weight as influence. Drop the split from this corpus, and the most widely cited AI voice would be Elon Musk, on fourteen publications, ahead of Dario Amodei, Sam Altman and Jensen Huang. That is a ranking of who the press writes about, but not whose expertise they trust to frame the fundamentals of a given industry.

We also separated commentators from principals. A founder cited about their own company is doing something categorically different from an analyst, researcher or developer cited for their read on the industry. Both sit in the data; but we rank only the expert commentators. Principals and political figures cleared the raw numbers, so we moved them to a separate layer. Again, the likes of Dario Amodei, Jensen Huang, Sam Altman and Demis Hassabis are cited often and substantively — but for their own companies’ positions, which is a different thing from being cited for your read on the industry.

Breadth beats volume

Legend: ● Cited in 2+ categories ● One category

Name	Total citations (volume)	Distinct citing publications (breadth)	Category
Simon Willison	21	15	Cited in 2+ categories
Dylan Patel	16	12	Cited in 2+ categories
Andrej Karpathy	14	10	Cited in 2+ categories
Dean Ball	30	11	Cited in 2+ categories
Ethan Mollick	17	8	Cited in 2+ categories
Dwarkesh Patel	12	8	Cited in 2+ categories
Linus Torvalds	5	7	Cited in 2+ categories
Ben Thompson	7	7	Cited in 2+ categories
Richard Sutton	6	7	Cited in 2+ categories
Ryan Greenblatt	12	7	Cited in 2+ categories
Charity Majors	12	6	One category
Gergely Orosz	10	6	Cited in 2+ categories
Brian Krebs	15	5	One category
Tyler Cowen	12	4	Cited in 2+ categories
Linus Torvalds	4	4	Cited in 2+ categories
Ben Thompson	8	4	Cited in 2+ categories
Richard Sutton	5	4	Cited in 2+ categories
Ryan Greenblatt	7	4	Cited in 2+ categories
Charity Majors	8	4	One category
Gergely Orosz	6	3	Cited in 2+ categories
Brian Krebs	5	3	One category
Tyler Cowen	4	3	Cited in 2+ categories
Linus Torvalds	3	3	Cited in 2+ categories
Ben Thompson	6	2	Cited in 2+ categories
Richard Sutton	4	2	Cited in 2+ categories
Ryan Greenblatt	5	2	Cited in 2+ categories
Charity Majors	11	2	One category
Gergely Orosz	10	2	Cited in 2+ categories
Brian Krebs	11	2	One category
Tyler Cowen	12	2	Cited in 2+ categories

Volume is much easier to manufacture than breadth. A single outlet citing the same person repeatedly produces a high appearance count, but can miss most buyers. Breadth is hard-won. It means separate journalists, analysts, podcasters, and writers, working independently, each concluded the same person was worth citing.

8 All content © Brands2Life | Brands2Life.com

The qualifying voices, ranked by breadth

#	Voice	Pubs	Cites	Categories	Affiliation
1	Simon Willison	15	21	AI · Cyber · Dev	Independent developer
2	Dylan Patel	12	16	AI · Dev	SemiAnalysis
3	Dean Ball	11	30	AI · Cyber	AI policy writer / OpenAI
4	Andrej Karpathy	10	14	AI · Cyber · Dev	Eureka Labs
5	Ethan Mollick	8	17	AI · Dev	Wharton / One Useful Thing
6	Dwarkesh Patel	8	12	AI · Dev	Dwarkesh Podcast
7	Ryan Greenblatt	7	12	AI · Dev	Redwood Research
8	Ben Thompson	7	10	AI · Dev	Stratechery
9	Linus Torvalds	7	8	Cyber · Dev	Linux
10	Richard Sutton	7	7	AI · Dev	Openmind/University of Alberta
11	Charity Majors	6	12	Dev	Honeycomb
12	Gergely Orosz	6	10	AI · Dev	The Pragmatic Engineer
13	Kent Beck	6	8	Dev	Tidy First?
14	Kevin Beaumont	6	8	Cyber	DoublePulsar
15	Mitchell Hashimoto	6	8	AI · Dev	Ghostty
16	Matthew Green	6	7	Cyber · Dev	Johns Hopkins
17	Nathan Lambert	6	7	AI · Cyber	Ai2 / Interconnects
18	Brian Krebs	5	14	Cyber	KrebsOnSecurity
19	Dex Horthy	5	7	AI · Dev	Humanlayer
20	Jasmine Sun	5	7	AI	The Atlantic/Substack

Ranked by distinct citing publications, then volume. Affiliations are editorial context added by hand, not extracted data. Full table of all 277 qualifying commentators available as data.

A note on robustness. Single-source concentration posed the main risk to this table: one AI newsletter contributed a disproportionate share of raw records, and one security newsletter did the same inside its category. Because the ranking is breadth-first rather than volume-based, no single publication can overpromote any one expert. Multiple citations by the same source only count once.

03 THREE CATEGORIES, THREE SEPARATE COMMUNITIES

The headline table mixes all three categories. Split them apart and they barely overlap.

AI and AI infrastructure

#	Voice	Pubs	Cites	Note
1	Dylan Patel	11	14	The most widely cited AI voice; disputed in 31% of his citations study-wide
2	Dean Ball	9	28	Highest citation volume in the study, 30 across all categories; 29% disputed here
3	Ryan Greenblatt	6	11	Safety research, cited across the AI and developer divide
4	Dwarkesh Patel	6	10	Interviewer whose framings get argued with — 33% rebuttal rate study-wide
5	Andrej Karpathy	6	9	Cited in all three categories; almost never disputed
6	Ethan Mollick	5	14	Present in all three months, in this category and overall
7	Jasmine Sun	5	7	Cultural and policy writing on AI adoption
8	Nathan Lambert	5	5	Open models and post-training

Cybersecurity

#	Voice	Pubs	Cites	Note
1	Kevin Beaumont	6	8	Vulnerability analysis; present in all three months
2	Brian Krebs	5	14	The category reference point, as a domain expert would predict
3	Matthew Green	5	6	Cryptography; also cited in developer platforms
4	Marcus Hutchins	5	5	Malware analysis
5	Katie Moussouris	4	8	Disclosure policy — the quoted expert on bug-bounty economics
6	Vangelis Stykas	4	4	Device and automotive security research
7	Andy Greenberg	3	5	Investigative reporting, cited by peers
8	HD Moore	3	5	Network discovery; Metasploit's author, still the reference


Developer platforms

#	Voice	Pubs	Cites	Note
1	Simon Willison	9	13	Leads this category and appears in the other two
2	Charity Majors	6	12	Observability; cited for practice as much as for tooling
3	Kent Beck	6	8	Method and practice, forty years on
4	Linus Torvalds	6	7	Cited in both developer platforms and cyber
5	Gergely Orosz	5	9	Cited as The Pragmatic Engineer three times here, but always named too
6	Dex Horthy	5	7	Credited coiner of "context engineering"
7	Steve Yegge	5	6	Essayist; present in all three months
8	Mitchell Hashimoto	4	6	Cited for what he builds rather than what he argues

READ THE VOLUMES CAREFULLY

The tables above count only citations to the 277 ranked commentators: 559 in AI, 207 in cybersecurity, 274 in developer platforms. What separates them is who receives the citation. In cybersecurity, 62% of analyzed authority citations name an organization rather than a person, against 31% in AI and 27% in developer platforms.

Share of analyzed authority citations that name an organization rather than a person



Security communications teams are inadvertently hiding talent.

A cyber spokesperson competes with their own employer's brand for the citation, because the field's convention is to credit the research body. An AI or developer spokesperson does not have that problem, because those fields tend to credit people by name. Note also which organizations cybersecurity publications cite most: Anthropic 48 times, OpenAI 31, Microsoft 30. The AI labs now sit inside the security canon.

Simon Willison tops one category and appears in both others.


He leads developer platforms outright and ranks ninth in cybersecurity, on the strength of work that began as neither. His coinage of "prompt injection" moved an AI concern into the security lexicon.

04 THE BEEF COEFFICIENT

For each voice, we measured the share of their citations that dispute them by name. We expected a curiosity and found one of the clearest signals in the data. 1,199 citations in the corpus are disagreements: 604 dispute a named person, 261 dispute an organization, 328 dispute a commonly held wisdom or trend.

The Beef Coefficient

Rebuttal share (rebuttals / cites)



Share of a voice's substantive authority citations that dispute them by name. Minimum five citations to qualify.

Tyler Cowen is disputed in two-thirds of his citations. Dylan Patel, the second most widely cited voice in the study, is argued with in three out of ten. Neither is a marginal figure — Patel is second on breadth, Cowen twenty-sixth of 277 — and both are the subject of frequent disagreements. The industry argues with the people it takes most seriously.

A traditional share-of-voice model would count all of these as either neutral or negative coverage. But nobody spends column inches or airtime rebutting a person who does not matter. A higher Beef Coefficient indicates that a given expert provokes thought and elicits strong feelings among their peers.

“

Nobody spends column inches or airtime rebutting a person who does not matter.

”

The distribution supports that reading. Of the 865 rebuttals we tracked:



“OLD MAN YELLS AT CLOUD”

328 rebuttals in the corpus target nobody in particular. These include phrases like “the uninformed”, “the critics,” or “voters.” We capture them as a separate measure of argumentative climate and keep them out of any individual’s score.

05 BRIDGE VOICES AND CONSISTENT OUTPUT

98

OF 277 CITED IN TWO OR MORE CATEGORIES

6

CITED IN ALL THREE CATEGORIES

52

CITED IN ALL THREE MONTHS (19%)

98 of the 277 qualifying commentators were cited in two or more categories. Only six — Simon Willison, Andrej Karpathy, Tyler Cowen, Alex Stamos, Daniel Miessler and Dawn Song — were cited in all three. These bridge voices are disproportionately valuable targets: one relationship reaches two or three ecosystems at once, and there are so few to begin with.

Only 52 of 277 qualifying experts (19%) were cited in all three months of the study. A further 123 appear in two, and 102 in only one. A shorter analysis would have downplayed the importance of consistent, thoughtful commentary over one-hit wonders, and comms teams should include at least three months of visible expert commentary when choosing outside partners.

06 PUTTING THE ANALYSIS TO WORK

Of course, our three chosen industry categories are highly volatile, and new expert voices emerge every day. Our rankings clearly display current thought leaders, and surfaced rising stars, yielding a constellation of experts whose work we love, and with whom we'd love to work. The tables will eventually become time capsules of the summer of 2026, but the underlying tools we've trained to spot these voices are available to our in-house team. Please reach out if you'd like to unpack this data further, or apply this methodology to other sectors.

METHODOLOGY

- 10,138 citation records extracted from 1,612 items published by 92 mastheads between 1 June and 31 August 2026, against a signed-off universe of 98 named mastheads.
- Each record carries 20 fields, including the authority/subject split, a topic domain, rebuttal target, direction and severity, and coinage role, plus two processing flags for gated pages and alias resolution.
- Every ranking in this report is computed on 7,405 analyzed rows, not on all 10,138. A row is analyzed when the citation is an authority citation, engages with the substance, and resolves to an identifiable entity. The other 2,733 are subject citations, passing name-drops, and references we could not resolve. They stay in the corpus and out of the scores.
- Where a citation names a publication instead of a person, a reviewed alias map resolves it to the person behind the masthead. That resolution is what makes Mechanism 1 measurable.
- Organizations carry one of five kinds. Across the analyzed set: other organization (1,523), publication (677), company (467), news publication (273) and research body (140) — 3,080 in all, and none of them is ranked. The kind is recorded so the exclusion can be argued with rather than assumed.
- Ranking requires citation by two or more distinct publications. 277 commentators qualify.
- Organizations stay out of the rankings, regardless of their research output. 3,080 of the 7,405 analyzed citations name an organization rather than a person as the authority — 42% of everything we analyzed. 1,835 of those go to an organization cited by two or more publications, which is the population that would have entered a ranking. The largest are AI labs and news organizations (Anthropic 150, OpenAI 102, Reuters 46, Microsoft 46, Bloomberg 41), with security vendors and consultancies behind them. Their findings are real, but a company is not a voice, and ranking them would put communications budget on the same axis as editorial judgment. Where an organization's write-up credits a named researcher, that researcher is recorded and counts.
- One masthead, one name. A publication's newsletter, articles and podcast merge into a single value, as do sibling titles that syndicate each other.

Extraction precision: audited twice

An independent human reviewer audited two stratified random samples, once on the pilot corpus and once on the final.

Population	Correct	Rate	What it affects
Analyzed set (round 2)	12 / 16	75%	Every ranking in this report
— entity identification	15 / 16	94%	Whether the right person was named
— topic classification	13 / 16	81%	Category assignment only
Analyzed set (round 1, pilot)	11 / 14	79%	Superseded corpus

Both audits used stratified random samples reviewed against source text. Round 2 covered 16 rows in the analyzed set, so read every rate here as indicative rather than precise: 12 of 16 carries a 95% interval of roughly 48–93%. Round 2 also tested the fields introduced for this report — organization classification, rebuttal severity and coinage role.

Topic classification is the weakest field, at 81%. Topic assignment affects only the category breakdowns in Section 3, not the rankings themselves. Entity identification, which the rankings depend on, held at 94%.

Known limitations

- Thirteen high-volume trade titles were sampled, not read in full. The daily security and technology trade press publishes hundreds of items a quarter each, so the spec capped them at roughly twenty items apiece: 301 items read of 2,125 published, 14%. Everything else — 79 mastheads, 1,318 items — was read depth-first. The effect is concentrated in cybersecurity, where eight of the thirteen sit and where sampled titles supply 44% of the category's items. A cyber voice's true breadth is therefore understated more than an AI or developer voice's, and the category's totals are not comparable with the other two as raw volumes.

Known limitations (continued)

- We read 93 of the 98 named mastheads in the signed-off universe. Read every breadth figure as a floor: each voice's true breadth is at least what we show, and probably higher. The sign-off sheet listed 99 rows. Four of those name a masthead another row also names (−4), one names no publication at all (−1), and two bundle three newsletters each into a single line (+4) — so 98 named mastheads is the honest denominator, and the 93 we read resolve to 92 once DevClass and The Register merge. Five we did not reach: Malicious Life, which stopped publishing in November 2024; Fabricated Knowledge, which published twice in the quarter, both behind a hard paywall; and Ben's Bites, TLDR AI and The Rundown, three daily aggregator newsletters that would have taken a disproportionate share of the budget for citations that are mostly bare links.
- 681 records, from 163 items, come from gated pages where we could read only the free portion. We flag every one. They skew toward people introduced early in a piece and away from those engaged with later: link-only citations run 12% in this subset against 5% in full-text items. We never inferred anything past a gate.
- 190 rebuttals sit on subject rather than authority rows, 120 of them against a named target. All are kept out of the Beef Coefficient. We checked: the targets are overwhelmingly chief executives, who fall outside the table regardless. Five touch a ranked commentator — Chamath Palihapitiya, Gavin Baker, Jeff Dean, Leopold Aschenbrenner and Miles Brundage.
- Single-source concentration is real and disclosed. Two newsletters contribute disproportionately to raw record counts. Breadth-first ranking neutralizes this for the published table; read any volume figure with that in mind.
- We record pseudonymous voices exactly as printed and let the two-publication threshold do the filtering. Genuinely influential pseudonymous figures therefore appear under their handles.

Read more of our AI-scale media research reports

THE RISE OF INDIE JOURNALISM



THE TECH PODCAST OPPORTUNITY



About Brands2Life

Brands2Life's **BETTER STORIES | BIGGER IMPACT** approach enables us to deliver exceptional creativity and outstanding results, whether it's a PR programme or a multi-discipline, multi-channel campaign.

We run these campaigns all around the world, led out of our offices in the UK and US.

In 2023 Brands2Life was acquired by the Scandinavian comms holding company, Paritee.

Learn more about Brands2Life at
www.brands2life.com

Get in touch



Rich Gallagher

HEAD OF COMMUNICATIONS INNOVATION

Rich.Gallagher@brands2life.com

T: +1 415 610 7500

M: +1 347 229 2482

THE AGENCY FOR THE BRANDS TRANSFORMING OUR WORLD



BRANDS2 **LIFE**

A PARITEE AGENCY