



A CBC PR GUIDE · AI VISIBILITY SERIES

The PR Playbook *for AI Search Visibility*

*The Three PR lanes that determine whether ChatGPT,
Perplexity, and Google AI recommend your brand or your competitor's.*



AI doesn't Trust a Mention. It trusts consistency.

When someone asks ChatGPT for a recommendation, the model pulls from multiple sources, synthesizes an answer, and cites only a few. To be part of that answer, your brand needs consistent claims across multiple credible, independent sources.

Most PR programs are built to land the big hit. AI rewards something different: orchestrated consistency across a network of trusted sources.

5-10

sources are synthesized
per each AI answer

*"One great placement is a data point.
Five consistent placements become consensus."*

The brands building this presence now are becoming the default answers to the questions your customers ask most.

What the research tells us

89%

of links cited by AI tools come from
earned media

Muck Rack, 2025

59%

of health AI citations are from media
interpreters, not peer-reviewed research

Outcomes Rocket, 2025

30%

of all health citations come from
consumer-facing health media outlets

Outcomes Rocket, 2025

The outlets that translate research into accessible content carry more weight in AI answers than the original research papers.




AI Rewards What's Seen — and Reinforces What's Missing

Once a competitor becomes the AI's default answer to a question your customers ask, that position reinforces itself. The model sees more signal for them, which generates more citations, which generates more signal.

Your brand's absence works the same way in reverse. The model has no reason to surface you, so you generate no signal, so future versions of the model have even less reason to find you.

THE FLYWHEEL EFFECT



The brands building orchestrated presence now are locking in default positions. The window is open. It won't stay open.

Publications positioned for AI citation

HEALTH MEDIA	FOOD & NUTRITION	WELLNESS
<i>Healthline</i>	<i>EatingWell</i>	<i>Real Simple</i>
<i>Verywell Health / Fit</i>	<i>Bon Appétit</i>	<i>MindBodyGreen</i>
<i>WebMD</i>	<i>Epicurious</i>	<i>Prevention</i>
<i>Harvard Health Publishing</i>	<i>Serious Eats</i>	<i>Shape</i>
<i>Medical News Today</i>	<i>Food & Wine</i>	<i>Health.com</i>

Strategic targets based on the characteristics AI models prefer — not a verified citation ranking for every outlet.



Beyond traditional media: four platforms AI actually cites.

YouTube

39%

of social platform citations in AI answers. In-depth creator content — where a nutritionist explains your category and mentions your brand — generates real AI signal. Promotional posts don't.

Reddit

Real-world signal

Perplexity and ChatGTP draws heavily from Reddit for "has anyone tried" queries. Community-vetted discussions feed AI answers. Marketing speak gets rejected instantly.

Wikipedia

Training anchor

Significant AI training data source. If your brand or ingredient has a Wikipedia entry, its accuracy matters more than most brands realize. Outdated entries work against you.

LinkedIn

89K URLs studied

Expert citations correlate with credentials and consistency — not follower count. An RD with 400 followers gets cited the same as one with 40,000, if the content is substantive.

No single platform is the answer. The pattern across all of them is.

Three lanes. One strategy.

AI visibility isn't a new initiative. It's a lens applied to the work your team already does. These three lanes are what most communications programs already manage. The difference is running them in concert — with consistent messaging and coordinated timing — so each lane reinforces the others.

01 Media Relations

02 Thought Leadership

03 Measurement

Done in silos, each lane underperforms. Run together, they create the multi-source consensus AI models treat as authority.

Media Relations

Earned media placements in trusted publications remain the single strongest input to AI answers.

THE INSIGHT

Not all coverage is equal for AI. The publications AI engines trust most are specific — and most PR teams aren't targeting them deliberately.

What changes isn't just where you pitch. It's how you architect the message so it appears consistently across sources, in language the model can repeat. That's the discipline most programs skip — and the one that separates brands AI recommends from brands it ignores.

WHAT THIS MEANS FOR YOUR PROGRAM

1 AI citation weight isn't the same as media prestige.

A freely accessible Healthline article in Q&A format may outperform a paywalled glossy. Target accordingly.

2 Consistent core claims matter more than volume.

If Healthline frames it one way and EatingWell frames it another, the model may confidently cite neither.

3 The full playbook shows you how to build message architecture before you pitch

— and which specific publications carry the most weight for food, health, and wellness brands.

"Think like a journalist, not a marketer. Content that informs gets cited. Content that sells gets ignored."

Thought Leadership & SMEs

AI models track entities — including people. A credentialed expert appearing consistently across independent sources builds authority that transfers to your brand.

THE INSIGHT

AI models track entities — including people. A credentialed expert appearing consistently across earned media, LinkedIn, and owned content becomes a citation magnet for your brand.

Most brands have this asset and aren't deploying it strategically. The distinction between a quote in Healthline and a quote that shapes AI answers is credentials, consistency, and cross-platform presence.

WHAT THIS MEANS FOR YOUR PROGRAM

1 Credentials carry more weight than title.

An RD or board-certified dermatologist outperforms a VP of Marketing — regardless of how smart the VP is.

2 Cross-platform consistency is the signal.

The same expert in Healthline, on LinkedIn, and in a YouTube video creates multi-source reinforcement AI can't ignore.

3 The full playbook maps how to deploy your SME across every lane

— and what to do if you don't have a credentialed expert on staff.

When earned media and expert commentary point in the same direction, the compound effect on AI visibility is substantial.

Measurement & Iteration

Traditional PR metrics don't capture AI visibility. Here's what to add to your program.

THE INSIGHT

Traditional PR metrics don't tell you if you're winning in AI search. Impressions, reach, and clip counts were never designed to — and they won't show you the gaps until a competitor is already the default answer.

AI visibility requires a different measurement framework, built around share of voice across AI platforms, message consistency audits, and triangulation coverage across source types.

WHAT THIS MEANS FOR YOUR PROGRAM

1 Start with an AI audit.

Ask the top 10-15 consumer questions in your category across ChatGPT, Perplexity, and Google AI Overviews. Note who gets cited. Note who doesn't.

2 Track AI share of voice, not just traditional metrics.

How often does a competitor appear vs. you for your top category questions? This is the number that matters.

3 The full playbook includes the quarterly AI audit framework CBC uses with clients

— including the specific questions to ask and how to act on what you find.

"Brands that start measuring AI visibility now will have a significant advantage over those that start measuring it later."

The full playbook covers all five lanes.

Media Relations. Thought Leadership. Creators & Influencers. Company Content. Measurement. With the specific tactics, checklists, and quarterly audit framework your team can use this quarter.



Download the full PR Playbook for AI Search Visibility:
<https://cerconebrown.com/ai-visibility/>