

WHITE PAPER

AI visibility in healthcare

Why it matters and how organizations
can lead in an AI-driven world



 **WG Content**
Strategy and Creation



Key takeaways

- AI does not rank healthcare websites the way traditional search does. It favors a small set of organizations based on key trust signals, looking for content to use as supporting evidence to explain its choice, and then builds user answers around that content.
- In healthcare, AI behaves conservatively because naming an organization functions as an implicit endorsement.
- Most AI invisibility is not caused by poor SEO, but by unclear entity identity, weak proof signals or content that is difficult for AI to interpret.
- Organizations with a site-wide SEO and AIO strategy show up more consistently than those optimizing page by page.
- AI visibility maturity can be understood across four levels: identifiable, explainable, justifiable and repeated.
- Several often-overlooked factors play a meaningful role in shaping how AI interprets a brand — including patient stories, employer reputation, organizational structure and unresolved negative publicity.
- Improving AI visibility is ultimately about making an institution easier to understand and trust (for AI and human users alike). Effective content must move from “marketing pages” to “decision support” that provides explicit proof signals.

// **We hope you find this white paper helpful as you optimize your website for AI.**

WG Content contributors



Diane Hammons
Director of Digital Engagement



Stella Hart
Content Strategist



Heather Stanley
Director of Marketing



Insights from healthcare AI testing and what it means for your brand


Healthcare marketers are paying more attention to AI visibility, and for good reason.

- **AI-driven search is growing fast.** [Research shows](#) rapid adoption of AI-powered search experiences, changing how people discover healthcare information.
- **Consumers are [getting answers](#), not links.** Some users start healthcare searches with large language models (LLMs), while others receive health information directly from Google's AI Overviews, using AI for medical questions and guidance at multiple stages of decision-making.
- **ChatGPT is part of health decision behavior.** Tools like ChatGPT are [increasingly used for seeking medical advice](#). Tens of millions of users turn to ChatGPT for healthcare questions, like symptoms, treatments and medical information — even when they understand these tools are not a replacement for clinicians.

What's surprising many healthcare organizations is who shows up in these AI-generated answers and who doesn't. While some well-known brands struggle to appear at all, others consistently break through. WG Content wanted to understand why.

Traditional [SEO best practices](#) still matter. But AI systems don't prioritize content the same way search engines do. In many cases, the criteria AI uses to select and cite sources goes beyond conventional ranking signals.

After reviewing hundreds of AI-generated healthcare answers, a clear pattern emerged: **AI is highly selective about which organizations it cites and builds responses around a small set of sources it considers safe. And in healthcare, more than in many other industries, that bar is especially high.**



“AI doesn't browse like a human. It assembles answers. And it can only assemble what it can clearly understand, verify and justify.”

- Diane Hammons, director of digital engagement



What AI visibility means in healthcare

[AI visibility](#) isn't about where you rank on a search results page. It's about whether an AI system **names, cites or references** your organization when it answers a healthcare question – especially when the user hasn't mentioned you by name.

Put simply: Are you part of the answer ... or not?

As search shifts from links to synthesized responses, that distinction matters more than most healthcare teams realize.

AI and healthcare: the “burden of proof”

Healthcare is a high-trust, high-risk domain, and AI systems know it. Because naming a hospital, health system or provider can read as an endorsement, AI takes a conservative approach, drawing from sources it considers credible before identifying entities it deems safe to cite.

How we tested AI visibility

To understand how that defensibility shows up in practice, we ran structured visibility tests across multiple AI health research assistants, chatbots and LLMs, including ChatGPT, Perplexity and Claude. All prompts run anonymously, without personalization or logged-in bias. These tests included:

- 15+ healthcare organizations, including national health systems, regional hospitals, children's hospitals, specialty clinics and senior living providers
- 20+ prompts per organization, spanning:
 - Branded vs. unbranded questions
 - Local, regional and national intent
 - Clinical and access questions: cost, access, providers
 - Experience and trust questions: reviews, reputation
 - Comparison prompts: best, top, should I choose



Our visibility tests focused on discovery behavior – not whether AI could recall a brand, but whether it would introduce one.



Surprising finding: AI selects entities, then writes the answer

Many of us think about discovery through a traditional SEO lens: Pages compete, then the best page wins. But AI doesn't work that way. In AI-driven discovery, entities compete and the safest set of entities wins. Only after it forms that shortlist does AI look for supporting evidence to explain why those organizations belong in the answer.

What that looks like in practice

Across prompts and models, the behavior was consistent:

- **“Top” or “best” prompts:** AI returns a small, conservative shortlist.
- **“Near me” or “in [city]” prompts:** There's more variance and more opportunity for local organizations to appear.
- **“What is,” symptoms” or “what to expect” prompts:** Content quality becomes the differentiator after entity selection.

In other words, you can have excellent content, but if your organization doesn't make the initial entity shortlist, it likely will never enter the conversation.



How AI builds answers: from a user prompt to a shortlist of trusted entities, supported by evidence and distilled into a final response.



If your organization doesn't make the initial entity shortlist, it likely will never enter the conversation.



Four unexpected factors shaping AI visibility

Some of the strongest influences on AI visibility were not the factors most teams obsess over. In reviewing dozens of prompts across multiple systems, a handful of less obvious signals kept resurfacing, often carrying more weight than traditional marketing efforts.

1. Patient stories: the lived experience

[Patient stories](#) matter, but not in the way many organizations assume.

When they help	They consistently show up in reassurance-seeking prompts like “what is it like?” or “what should I expect?” They also support decision-making questions about recovery, care pathways and real-world experience. In these moments, AI uses patient stories as evidence of lived experience, not as marketing.
When they often do not help	They rarely influence broad discovery prompts like “best hospital” or “top program,” where AI prioritizes proof and reputation over narrative.
What makes them AI-useful?	The most reusable stories include structure, a clear care pathway, clinician context, transcripts for video content and links back into condition or treatment hubs. Without those elements, they are emotionally compelling but difficult for AI to reference.



Explore patient stories and search visibility

Authentic [patient stories](#) build trust and strengthen visibility across SEO, AEO and AI-driven search.



2. Employer reputation: when workforce signals shape patient answers

When answering comparison or decision-oriented questions like “Should I get care here or here?”, AI frequently referenced third-party employee review platforms. Mentions of understaffing, turnover or workplace stress were signals that could affect patient experience, access to care or continuity of care. These references were not framed as employer branding. They were used as risk indicators.

Why this matters	In healthcare, AI is constantly weighing potential harm. Signals associated with staffing instability or burnout appear to influence how safe or reliable an organization feels, even when the original question is about care quality, not employment. In effect, workforce reputation becomes part of the patient experience narrative.
How this shows up in AI answers	AI doesn't quote employee reviews verbatim. Instead, it synthesizes themes like staffing challenges, morale or turnover into broader judgments about care consistency, wait times or operational strain.
What “investment” really means	This is not about managing reviews or polishing employer messaging. It's about aligning internal workforce reality with external signals and understanding that those signals now inform how AI evaluates patient-facing trust.

3. Domains, mergers and sub-brands: entity fragmentation as a silent visibility killer

Consolidation is common in healthcare, but it often introduces unintended visibility problems. As systems grow through mergers and acquisitions, domains multiply, naming conventions drift and sub-brands persist longer than expected. For humans, this can be navigable. For AI, it creates uncertainty.

In AI discovery, canonical identity matters more than it does in classic SEO. If AI can't confidently determine whether multiple sites represent one organization or several, it may become more conservative about naming any of them.



4. Negative PR and institutional memory: when AI remembers what organizations want to move past

Another unexpected finding was how long negative events continued to influence AI answers.

What we saw	In patient decision and comparison prompts, AI often referenced older news about hospital closures, care deserts, data breaches or operational failures. Even when the events were years old, they resurfaced as context for trust, access or safety. In many cases, the original incident mattered less than the absence of a visible response.
Why this matters	AI systems are designed to minimize risk. When negative events are widely documented but not visibly addressed, they remain unresolved signals. From an AI perspective, silence looks like uncertainty. Without clear evidence of change, past issues continue to shape present answers.
How this shows up in AI responses	AI does not accuse or editorialize. Instead, it weaves past events into broader assessments of reliability, access or community impact. This can quietly influence whether it frames an organization as stable or risky.
What “doing something about it” looks like	Organizations that appeared more resilient did not erase the past. They documented their response. Public-facing narratives described what changed, where investment occurred and how care access or safety improved over time. In AI terms, remediation becomes proof.



Press releases and visibility

See why [press releases](#) still matter and how they contribute to discoverability and trust online.



AI invisibility: what's going wrong (with examples)

AI invisibility is usually not about poor performance. It is about gaps in clarity, proof and structure that prevent systems from confidently including an organization in the answer.

Mistake 1

Treating content like marketing pages instead of decision support

Service-line pages that describe offerings, but don't help answer real questions.

Example: An orthopedics program page focused on provider bios and accolades, useful for booking, but not for AI answering "how do I choose the right treatment?"

Mistake 2

Thin or outsourced "education" that donates authority elsewhere

Health libraries point users to national publishers instead of building owned expertise.

Example: A children's hospital routed most educational queries to an external dictionary. Helpful for users, but it prevented the organization from becoming the source AI learns from.

Mistake 3

Missing a public "proof layer" AI can repeat

AI needs defensible reasons to name you, and it can't infer what isn't visible. Common gaps include buried or missing accreditations, outcomes, program recognitions or research participation.

Example: In the research, the organizations that surfaced most often made proof explicit and specific. They clearly named programs, certifications and safety or quality distinctions. They explained accreditations in [plain language](#) and linked directly to accrediting bodies. In specialty care, research recognition and clinical trial scale were stated outright, not implied. These details gave AI something concrete to cite, elevating the organization to the shortlist of "safe" examples.



Mistake 4

Entity confusion from domains, sub-brands and mergers

Multiple domains and inconsistent naming split credibility signals.

Example: A regional system with several post-merger web identities diluted what should have been a single canonical entity.

Mistake 5

Schema used as a checkbox, not identity infrastructure

Structured data exists, but only on the homepage, or doesn't connect organization → locations → providers/services.

Example: One health system had Organization schema on its homepage and basic Local-Business markup on facility pages, but no structured relationship between the system, its hospitals and its physicians. From an AI perspective, the system appeared as a collection of loosely related entities rather than a single, coherent organization. As a result, AI struggled to reliably associate services and providers with the parent brand.

Mistake 6

Patient stories not treated as usable experience signals

Emotion-heavy patient narratives lack the elements AI can reference: pathway steps, "what to expect," clinician context, transcripts or links back to condition hubs.

Example: A health system published strong patient stories, but they were isolated from condition and treatment pages and did not document the full care journey. The stories were compelling for humans, but difficult for AI to reuse in decision-support answers.



The common thread: AI can't reuse what it can't clearly understand.



What AI-visible organizations do differently

After reviewing where organizations fall short, the next question is obvious. What are the organizations that do show up doing differently?

Across AI systems and prompt types, the organizations that surfaced consistently were not chasing AI tactics. They were practicing a small set of habits that made them easier to understand, justify and reference.

Below are four habits shared by AI-visible healthcare organizations.

Habit 1

They publish answer-shaped content (and organize it into hubs).

They structure their content around how people actually ask questions.

Pages use clear headings and plain-language sections that define the condition, explain options, outline risks and answer “what happens next.” Instead of isolated articles, content is clustered around conditions, treatments and patient decisions.

This gives AI something it can reuse, not just something it can read.

“To be the answer, you must be the most interpretable authority. Building that defensible identity can’t be piecemeal; it requires a thoughtful, comprehensive content strategy. User-centricity has always been the true driver of success, and the AI tools of today and tomorrow exist and iterate to make finding and acting on the right information as easy and intuitive as possible for humans. The most competitive brands, in and outside of healthcare, are consumer-obsessed. AI visibility is interwoven with an obsessive approach to ensuring your content anticipates the needs of your target market at every step of the patient journey.”

— Stella Hart, content strategist



Answer-shaped content delivers real results

See how [Franciscan Health](#) built organized content hubs that support SEO and AI-driven discovery.



Habit 2

They make credibility easy to cite.

AI needs reasons it can repeat. Consistently surfaced organizations maintain public-facing proof pages that clearly explain:

- Designations and recognitions
- Accreditations
- Volumes or outcomes where appropriate
- Specialized program differentiators

This information isn't just implied through marketing language. It is explicit, accessible and easy to justify in an answer.

Habit 3

They build navigable systems.

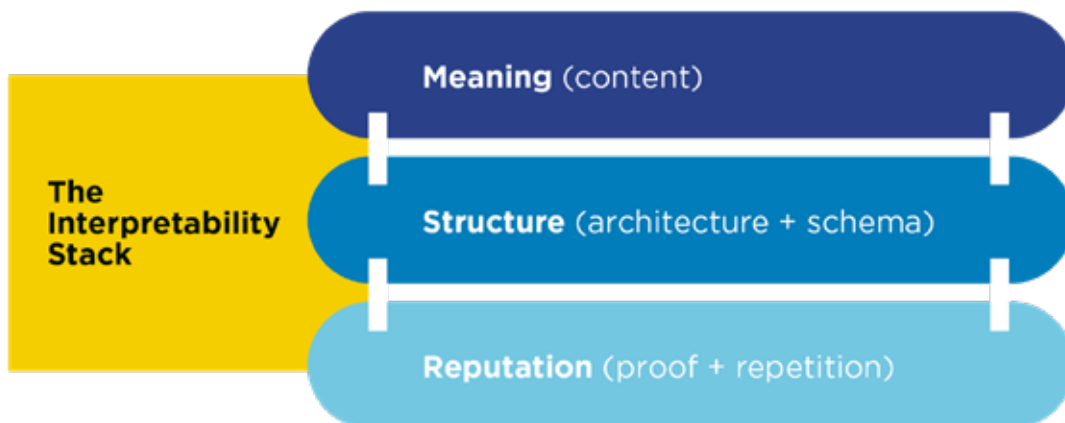
These organizations treat content as a system, not a collection. Strong internal linking connects condition pages to treatments, providers, locations and next steps. Both humans and machines can move through the experience quickly and understand the full scope of care without guesswork.

Habit 4

Their entity identity is coherent.

Their organizational identity is clear and consistent. They use standardized naming. A canonical domain is respected. Relationships between the system, facilities, clinics and physicians are easy to follow. Structured data supports those relationships.

When a user asks, "Who is this organization?", it's easy for AI to answer.



Interpretability consists of meaning, structure and reputation.




The 4 levels of AI visibility maturity

One of the biggest challenges with AI visibility is knowing where to start. Most healthcare organizations are not failing outright. They are simply at different stages of maturity, often without realizing it. This framework will help you diagnose what's missing and prioritize investment.

Level 1: Identifiable

At this level, AI can tell who you are, where you're located and what you do. Your organization has a recognizable name, a clear service footprint and a digital presence that machines can interpret.

Symptoms if you aren't here	<ul style="list-style-type: none">• Inconsistent naming across domains or platforms• Unclear service geography• Scattered domains, sub-brands or legacy sites
Use this prompt to test	<p>“What is [Organization Name]?”</p> <p><i>Follow-up prompt if not answered:</i></p> <p>“Where do they operate and what services do they offer?”</p>
Not getting the answer you want?	<ul style="list-style-type: none">• Clarify your entity foundation.• Standardize naming.• Define service geography .• Consolidate domains where possible.• Make relationships between organizations, locations and services explicit.





Level 2: Explainable

At this level, AI can use your content to help answer common healthcare questions. Your organization publishes content that explains conditions, treatments, risks and next steps in plain language.

Symptoms if you aren't here	<ul style="list-style-type: none">• Service-line pages exist, but decision-support content is thin• Educational content is shallow or points outward• AI explains the topic well, but does not reference your organizations
Use this prompt to test	“What are treatment options for [condition]?” <i>Follow-up prompt:</i> “How do patients decide on a treatment option?”
Not getting the answer you want?	Create answer-shaped content and organize it into hubs around real patient and caregiver decisions, not just services.

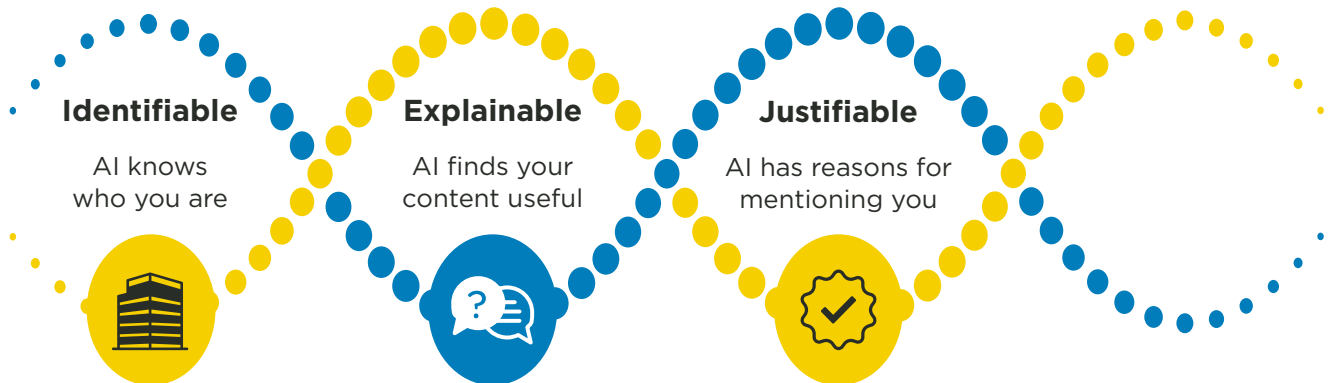




Level 3: Justifiable

At this level, AI can defend mentioning you. Your organization has a visible proof layer that explains why you're a credible example, not just a possible option.

Symptoms if you aren't here	<ul style="list-style-type: none">• Strong clinical or research work, but proof is not public or easy to cite• Accreditations and recognitions exist, but are buried• AI uses vague quality language instead of naming you
Use this prompt to test	“Why should I choose [Organization Name] for [service or specialty]?”
Not getting the answer you want?	Publish clear, plain-language proof pages.

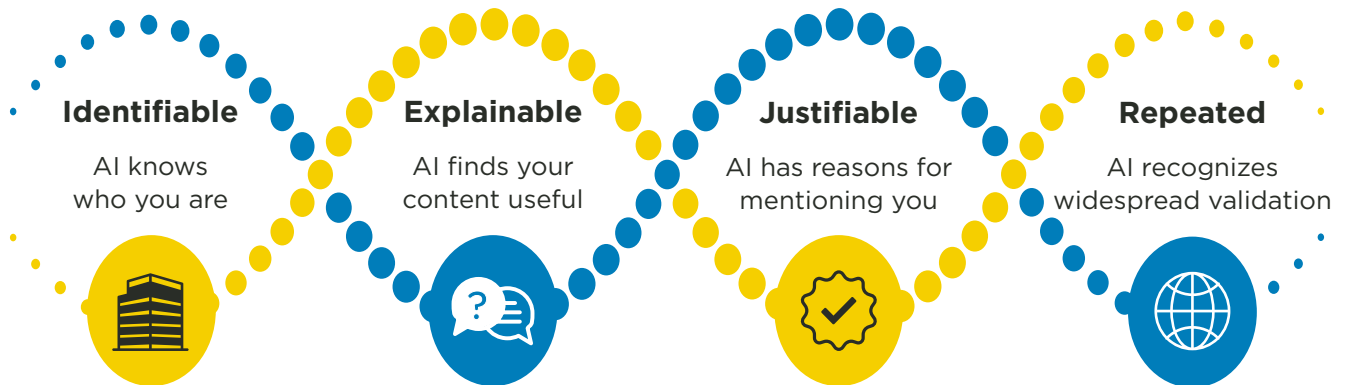




Level 4: Repeated

At this level, the broader web consistently corroborates you. AI encounters your organization across rankings, partnerships, citations, press coverage and reputation signals, reinforcing confidence through repetition.

Symptoms if you aren't here	<ul style="list-style-type: none">• Strong local reputation, but limited visibility in broad “top” or “best” prompts• Mentions appear inconsistently across AI systems
Use this prompt to test	“What are the best hospitals for [service] in [region]?”
Not getting the answer you want?	Invest in sustained reputation signals through partnerships, thought leadership, earned media and third-party validation.



The four levels of AI visibility maturity



How to measure and monitor AI visibility for your brand

For healthcare marketers, AI visibility measurement is becoming an essential part of modern AI healthcare marketing analytics — complementing, not replacing, traditional SEO and performance metrics.

AI visibility isn't something you optimize once and move on from. Because AI systems continuously retrain, reinterpret sources and adjust how they cite organizations, visibility needs to be evaluated over time, not treated as a one-off fix.

This raises a common question: *Why should I track AI visibility at all?* Because if AI systems influence discovery, trust and choice, visibility becomes a leading indicator, not a vanity metric.

Key indicators to track include:

- **Frequency of inclusion:** how often your organization is cited or referenced across different AI tools and question types
- **Prompt coverage:** whether you show up only in basic definitions or also in comparison, “best of” and decision-stage prompts
- **Context and confidence:** how AI describes your organization and whether mentions feel authoritative, neutral, qualified or cautious
- **Repeat visibility:** whether your brand appears consistently across related questions or only surfaces occasionally

To track these signals, marketers should run structured prompt testing regularly, using the same questions, geographic modifiers and comparison sets. This creates a baseline that makes changes in visibility easier to spot.

From there, optimization becomes an iterative process. Limited visibility often points to specific gaps: an unclear entity definition, insufficient proof or validation, weak third-party reinforcement or content that doesn't establish the organization as safe to cite.

Closing those gaps strengthens AI visibility while also improving fundamentals, like user trust, engagement and overall search performance.



“AI exposes what’s missing. Prompt ChatGPT from the patient’s point of view, identify where you fail to appear, ask why and build a plan to close the gaps.”

- Diane Hammons, director of digital engagement



Improve content strategy to improve AI visibility

AI visibility is not about chasing a new algorithm or reacting to another platform shift. It reflects something more fundamental. A move away from optimizing individual pages and toward building a digital footprint that machines can clearly understand, explain and defend.

The goal is not to “win” AI. It is to make your organization easier to understand, easier to trust and easier to include. And that work pays off everywhere, not just in AI answers.

If your organization is navigating this shift, WG Content helps healthcare teams build content strategies that are clear, credible and ready for AI discovery.

“AI visibility is a multiplier, not a tactic. It exposes gaps in clarity, credibility and connection. Fixing those gaps improves SEO, accessibility and AI-driven discovery alike. The goal is not to optimize for AI over search. It is to optimize for interpretability.”

- Heather Stanley, director of marketing

About WG Content

More than 20 years ago, WG Content became the first company to focus solely on content for the healthcare industry. Year after year, our ultimate mission is to put relationships at the center of all we do. As savvy communicators with combined decades of content strategy and creation, we partner with clients to drive traffic to websites, optimize search, capture leads, provide bold marketing solutions and much more. All while cultivating unique soft skills that create strong connections. Every day, our team works to shape meaningful content experiences and build lasting relationships for caring clients — one word at a time.

// Visit wgcontent.com or email hello@wgcontent.com

