

They Say GEO Isn't Real. Watch a Website Answer Questions About Itself.

Tygart Media — October 22, 2026. Companion to the screen-capture video at tygartmedia.com.

People keep telling me AEO and GEO are the new snake oil. Acronyms invented to sell audits. "SEO is SEO," they say, "and the rest is conference talk."

Fine. Let me show you the thing itself, on a site I run, with nobody's hands on the scale.

[Video: five minutes of unedited screen capture — healthcarefacilityhub.org answering from its own content, and refusing what it does not have.]

What you're watching

That's Healthcare Facility Hub — a site for people who run hospital facilities. In the corner of every page sits an AI that knows one thing: the site itself. It's trained on the site's own content and nothing else. No open internet. No general model knowledge doing the heavy lifting. If the answer isn't in the site's pages, the AI doesn't have it.

I spent five minutes asking it questions on camera. Three things happened, and they're the whole argument.

First: it answers like the content is a database

Ask it what an infection control risk assessment requires and it doesn't write an essay. It returns the site's own material as structure — the risk classes, the types, and then, when I push on negative air pressure, the actual working numbers: the target pressure, the alarm setpoints, the HEPA specs a facilities team would use on a job.

That's not a chatbot being clever. That's what happens when every page is built as a clean, standalone record — one question, one complete answer, the facts in the open — instead of a pile of marketing prose with keywords folded in. The AI can pull from it because the page was built to be pulled from.

Second: ask it something the site doesn't have, and it says so

I asked it for stories about Chicago. It said: we don't have any content or stories about Chicago in our current site resources.

Sit with that for a second. The big chatbots, asked the same question, would give you Chicago stories — real, invented, blended, who knows. This one refused, because its world ends at the edge of the site's content, and it knows where the edge is.

That's not a limitation of the demo. That IS the demo. An answer engine is only ever as good as the corpus it can draw on. When the corpus is disciplined, the answers are disciplined. When there's nothing there, an honest system says nothing's there — and a visitor, or an engine deciding what to cite, can trust everything it did say.

Third: the answer ends at an asset

Follow the thread to the end of the video and I'm not reading a blog post. I'm opening a one-page field sheet — the article distilled into a printable PDF with a QR code, made to be carried onto a hospital job site.

That's the shape of the whole thing. A question gets asked. The structured content answers it. The answer hands you something usable. Search used to end at a blue link and a hope. This ends at a tool in a facility manager's hand.

Why this is what the engines are doing to you

ChatGPT, Perplexity, Copilot, Google's AI — different companies, same basic move. They read corpora, they extract answers, they cite the sources that could answer. Your website is either a corpus they can extract from, or it's a brochure they skim past on the way to somebody else's.

We measure this instead of arguing about it. We run a fixed battery of questions a buyer would actually ask, score what our sites' AI answers from their own content, and log where it has nothing — those gaps become the next pages we write. This morning, on a sibling site in a different industry, that battery came back 24 passes out of 36, with the misses itemized as a content list. The scoreboard is not a vibe. It's a work queue.

And the citation side gets measured too — which engines cite which sites, on which questions, over a trailing thirty days. When the numbers move, we know why, because we know what we published and what it answered.

The part nobody wants to hear

None of this is a trick you bolt onto an existing site. The AI on that page answers well because the site underneath it was rebuilt to be a knowledge base — treated like a headless database that happens to have pages attached. Every article answers one real question, in the words the asker would use, with the numbers in the text where an extractor can find them.

If an AI trained only on your own website can't answer your customers' questions, the big engines can't either. They'll answer from a competitor whose site can. That's not a threat I'm selling. It's arithmetic you can run on your own site this afternoon — most sites have a chat box or a search bar you can interrogate. Ask it your customers' five best questions. Count what it actually knows.

We already know how that count comes out for most sites. It's why the phone rings here.

The recording is unedited screen capture from healthcarefacilityhub.org, October 22, 2026 — the site's own AI, its own content, my own questions. Narration on the video is Glint, my AI. This is the work, on tape.