



IDLAB RESEARCH · 2026

Your first audience is a machine

How AI engines represent utility programs, and what it means for enrollment. An 84-search audit of ChatGPT, Perplexity, Claude, and Gemini across three U.S. utilities.

BY JASON TURNER

IDLAB · CREATIVE, STRATEGY & MARKETING FOR ENERGY

Why I ran this test

The research, interpretation and writing of this doc is the work of humans; with minor support from AI.

I have spent the last fifteen-plus years helping utilities design, market, and measure their programs. Dozens of utilities, hundreds of programs, a great many campaigns that worked and a healthy few that taught me something instead. That kind of mileage gives you a feel for how customers behave, and lately my feel has been going off like a smoke alarm.

The behavior is shifting. Your customer used to type a question into Google, look at a list of blue links, and choose one. Increasingly they type the question into an AI engine, read the answer it writes for them, and act on it. No list. No links. Often, no visit to your website at all. The customer gets their answer and gets on with their day, and you never even know the conversation happened.

I could feel this happening. Feeling is a fine place to start and a terrible place to stop, so I decided to quantify it. This report is the result: a structured audit of what AI engines actually say when a real customer asks about a real utility program, plus a plan for what to do about it.

The numbers behind the shift are not subtle. Roughly 65% of U.S. Google searches now end without a single click, up from about half in 2019 (SparkToro clickstream data). When one of Google's AI Overviews appears, that figure climbs to around 83% (Bain & Company). Those overviews now show up on close to half of all queries (BrightEdge), and when they do, the click-through rate on the top organic result roughly halves (Pew Research measured it falling from 15% to 8%). Bain found that 80% of consumers now lean on AI-generated results for at least 40% of their searches. Gartner projects that organic search traffic to brand websites will fall by 50% or more by 2028 as customers move to AI-generated answers, a decline it expects to be well underway before then. Google told its own developer conference that AI Mode had passed a billion monthly users with query volume more than doubling every quarter.

Here is the part that should interest a marketer rather than frighten one. In the same shift, brands that get cited inside AI answers are seeing materially higher conversion, by some measures several times higher (Dataslayer). So this is not a story about hiding from the machine. It's a story about being the source it trusts. That's worth real money, and it's winnable.

THE NUMBERS BEHIND THE SHIFT ARE NOT SUBTLE

65%

of U.S. Google searches now end without a single click. SparkToro

83%

when one of Google's AI Overviews appears. Bain & Company

~50%

of all queries now show an AI Overview. BrightEdge

15% → 8%

click-through on the top organic result. Pew Research

-50%

or more organic search traffic to brand websites by 2028. Gartner

1B+

monthly users of AI Mode, query volume more than doubling every quarter. Google

SECTION 02

A note on language, and who (not what) did this work

Two quick things before we go further, because both matter.

First, throughout this report I refer to AI answer engines, ChatGPT, Perplexity, Google's Gemini, Claude, and their cousins, collectively as “the machine.” It is shorter, and it is a useful reminder that your customer’s question is now being answered by something that has never read your brand guidelines and feels no loyalty to your logo.

Second, a human did this research. Every one of the 84 searches was run, read, and scored by me and my team. We reached the conclusions ourselves, using judgment earned over years in this industry. We used AI only to help organize and synthesize what we had already gathered. Studying the machine with a little help from the machine has a certain irony to it, and I enjoy the joke, so I want to be plain about it: the findings, interpretation and write-up here are human work.

SECTION 03

What the audit found

The headline is the opposite of the one you have been sold, and I am leading with it because it is the honest one.

The machines are mostly accurate today. Utilities show up better than the panic would suggest. The real exposure is quieter, and it lands hardest on exactly the kind of mid-sized, multi-region utility that describes most of the industry.

01

Accuracy is high, uneven, and mostly self-inflicted.

A utility with clean, current, single-territory content scored a perfect 100% across every engine. A utility leaning on stale, region-split PDFs was wrong or incomplete 39% of the time, even while the machine cited that utility's own website 96% of the time. The machine performed exactly as designed. It read what the utility published and repeated it faithfully, including the parts that were out of date.

02

There are two ways to lose, and you fully control one of them.

Whether the machine states your program correctly comes down to how legible your content is, which is yours to fix. Whether the machine treats you as the authority is largely the engine's decision.

03

Getting cited depends more on the engine than on you.

The same question about the same utility earned that utility the citation 88% of the time on two engines and 43% of the time on the other two.

04

The single most fixable culprit is the humble PDF.

Engines reach for PDFs because a rebate chart in a PDF is often the cleanest version of the fact a utility has published, which quietly turns your least-maintained document into your most-quoted source.

None of this needs new technology to fix. It needs program information made legible to a machine: findable, extractable, specific, current, corroborated, and region-clear. I call that structured, verifiable evidence your **proof layer**, and building it is how a utility wins the customer's decision even when the machine, and not the utility, is the one making the introduction.

What SEO was, and what it is becoming

To see where this is going, it helps to remember where it came from.

For twenty years, the game was Search Engine Optimization. You structured your site so Google would rank your page near the top, because the top of a list of links is where the clicks were. An entire industry grew up around earning position one. It worked because there was a list, and a human chose from it.

The list is dissolving. The machine now reads the same web SEO was built for, then writes a single synthesized answer and hands it over directly. The customer reads the answer. Two new disciplines have grown up to influence what that answer says, and you will hear both names, so here they are in plain language.

SEO · AEO · GEO

SEO (Search Engine Optimization) is the practice of structuring your website so a search engine ranks it highly in a list of links. AEO (Answer Engine Optimization) and GEO (Generative Engine Optimization) are the two labels the industry now uses for optimizing to be the trusted source inside an AI-generated answer. AEO comes from the older world of being the answer rather than the link, the world of featured snippets and voice assistants. GEO is the newer term coined for generative systems specifically. In practice the two are used almost interchangeably, and the field has not fully settled the definitions, so treat them as two overlapping lenses on the same shift rather than two separate disciplines.

WHY IT MATTERS SEO won you a click in a list your customer got to choose from. AEO and GEO, by whichever name, win you a place inside an answer your customer never sees a list for, whether that answer is composed inside a search engine or inside a chat assistant. The skills carry forward. The target moved from “rank on the page” to “be the source the machine believes.”

If you have done real SEO, you are already halfway here. Clean structure, clear content, and authority still matter enormously. What changed is the reader. You used to optimize for a person deciding which link to click. Now you also optimize for a machine deciding which source to trust and repeat. This will only grow more central. As search volume shifts toward AI engines and the click keeps getting rarer, the answer becomes the storefront. Showing up inside it, accurately, becomes the whole ballgame.

The proof layer

Your proof layer is the structured, verifiable, machine-readable evidence about your programs, the rebate amount, who qualifies, the deadline, how to apply, that an AI engine can find, trust, and repeat when it answers a customer for you. In most cases it is a single page, or a small, consolidated collection of pages, that carries these facts cleanly. Think of it as the difference between telling a machine a story and handing it a fact it can lift cleanly.

WHY IT MATTERS The machine will answer your customer with or without your help. A strong proof layer means it answers with your facts, from your source. A weak one means it fills the gap with whoever published the fact more clearly, which is usually somebody else.

How I ran it

I chose three utilities to span the market on purpose, so the findings would generalize rather than flatter anyone:

CALL IT THE MODEL

CA muni

A single-territory California municipal utility. Well-resourced, strong digital presence, simple rate structure.

VERY LARGE, MULTI-REGION

FL IOU

A major Florida investor-owned utility. Very large, multi-region.

ITS OWN REFLECTION

Midwest IOU

A mid-sized, two-state Midwest IOU. Multi-region, with a typical program footprint. This is the one most of the industry will recognize as its own reflection.

For each, I ran real customer-voice program questions, heat-pump and HVAC rebates, water-heater rebates, “is it worth it,” “how do I get it and who installs it,” bill savings, and a federal-tax-credit control question, across four engines: ChatGPT with search on, Perplexity, Google’s Gemini, and Claude. I ran each query three to four times and scored it only where the answers held steady, so the results reflect the typical response rather than a lucky or unlucky roll.

My team and I scored every capture by hand against each utility’s actual published rebate terms, on four dimensions: accuracy, whether the utility’s own site got cited and who won the citation, freshness, and whether the answer sent the customer to a third party to act. Eighty-four scored captures in total, every one timestamped, August 2026.

A word on scale, offered up front because it is the honest thing to do. This is a small, deliberate audit, the kind you run to see clearly rather than to publish a p-value. Its value is that the measurement exists at all, and that the method repeats. Per-utility results rest on 28 captures each, and the cross-engine patterns on the full 84. Read the percentages as strong directional signal.

FINDING 01

The machine is mostly accurate, which is exactly why you should pay attention


Across the audit, the engines were accurate the large majority of the time, and utilities’ own sites appeared in 92% of answers. If you went hunting for utilities being savaged by AI, you would mostly come home empty. I did.

The averages, though, hide the interesting bit. Accuracy tracked one variable with almost eerie precision, legibility, and it did so consistently on every engine:

Accuracy tracks legibility, not size

■ ACCURACY VS. PUBLISHED TERMS

■ STALE OR REGION-INCOMPLETE



84 SCORED CAPTURES • 28 PER UTILITY • AUGUST 2026

The California muni, single territory, clean current HTML, simple rates, was accurate on every engine, every time. The Midwest IOU, two states, split program charts, key facts entombed in aging PDFs, came back wrong or incomplete on roughly four questions in ten. The machine was doing its job impeccably. It read what that utility published and repeated it, dusty corners and all.

My favorite example, and the one I would frame on a wall, is this. When I asked whether the Midwest IOU offered a heat-pump water-heater rebate, all four engines got it wrong, and all four cited the utility's own website as the source.

The utility was the source. The source was wrong.

That single row is the whole report in miniature: the machine repeating you, mistakes included.

The reassurance, “we are mostly accurate,” is precisely what makes this dangerous. The gaps stay invisible until a customer walks into one, and when they do, it is confident, official-looking misinformation about money.

Two ways to lose, and you own one of them outright

The audit splits cleanly into two failure modes that behave like different animals.

YOURS TO FIX

Accuracy is yours.


Whether the machine states your program correctly is a function of how legible your content is. The California muni proves the ceiling and the Midwest IOU proves the floor, and the fix in both cases is content, well within any utility's power.

TREAT IT AS WEATHER

Citation belongs largely to the engine.

Whether the machine treats you as the authority, versus a contractor, an aggregator, or a news story, is driven far more by which engine answered than by anything you did. A utility can be perfectly accurate and still lose the citation.

Plotting the three utilities on those two axes tells the strategic story:



The California muni is the model.

Accurate and cited, 100% accurate and cited as the anchor 75% of the time. This is what winning looks like.

The Florida IOU is the interception case.

Highly accurate at 93%, yet it earned the lowest anchor-citation rate at 54%. It is correct, and it is often overlooked, because a large, complex, multi-region footprint pushes engines toward navigators and aggregators.

The Midwest IOU is the legibility case.

Cited constantly, its own site appeared 96% of the time, while landing the least accurate at 61%, because what the machine faithfully read was stale and region-split.

Here is the line to take home. You can win the accuracy axis every time, because it holds across every engine. So win it. Treat the citation axis as weather: real, powerful, and worth planning around, and largely outside your control.

Across all 84 captures, the utility's own site appeared in 92% of answers, yet the utility was the source the machine actually built its answer on only 65% of the time. So a third of the time, on a question about a utility's own program, somebody else was the authority.

92%

of answers included the utility's own site.

65%

of the time the utility was the source the machine actually built its answer on.

FINDING 03

Who gets cited depends on the engine

That 65% is an average of two populations that behave nothing alike. Split by engine, it gets vivid:


Who the machine cites as the authority depends on the engine — not you

88%

UTILITY CITED ON
CHATGPT + PERPLEXITY

43%

UTILITY CITED ON
CLAUDE + GEMINI



Who the machine cites depends on the engine. The utility was cited 88% of the time on ChatGPT and Perplexity, and 43% of the time on Claude and Gemini.

Same utility. Same question. Cited anywhere from about a third of the time to nine times in ten, depending purely on which engine happened to answer. ChatGPT and Perplexity leaned hard on the utility's own source. The other two reached for aggregators and third parties far more often.

The lesson lands in one place. Being the cited source is volatile, because it hangs on an engine your customer picks and you do not. Which is exactly why the durable strategy is to win the accuracy axis, the one that holds steady across all of them, and to build the kind of legibility that improves your odds everywhere at once.

The PDF paradox

Engines reach for PDFs constantly, and there is nothing quirky about it. A rebate chart in a PDF is frequently the cleanest, most data-dense version of the fact a utility has published. The web page wraps the number in warm marketing prose. The PDF just states it in a tidy table. So the machine, being sensible, takes the PDF.

Sit with what that implies for a moment: the machine is quietly telling you that your own PDF beats your own website.

And that is a trap with a lovely bow on it, because the PDF is usually the asset you update least and would least like a customer to wander into. So the document the machine reads most eagerly is simultaneously your stalest and your most dead-ended. Two harms follow. Stale PDFs become confident misinformation, which is precisely where the Midwest IOU's accuracy failures came from. And a customer who lands inside a PDF has nowhere to go, no enrollment path, no next step, which quietly raises your acquisition cost and lowers satisfaction.

The remedy is cheerful and boring. Make the web page as clean and specific as your best PDF chart, and lift the operative fact up onto it in plain, current text.

Regionality multiplies every problem

The clearest predictor of trouble was serving more than one territory. The Midwest IOU's two-state structure, with different programs and amounts in one state versus its neighbor, produced error after error as engines blended and mismatched the regional facts. Where a fact was clearly tagged to its territory, the engine handled it gracefully. Where region was ambiguous, the engine guessed, and it guessed like a machine, confidently and often wrongly.

This matters because most of the industry serves multiple regions. The single-territory clarity that made the California muni easy is the exception. For everybody else, region-clarity earns its keep: it is the difference between the machine getting your program right for a given customer and inventing a blended answer that fits no customer at all.

SECTION 07

Why the machine does this

I also asked the engines directly why they represent some utilities better than others, and why they sometimes prefer a contractor's page to the utility's own. Their answers were remarkably consistent, and they amount to a to-do list written by the very systems you are trying to please:

01	Thin, vague content.	Utility program pages often narrate a program in soft marketing copy while the concrete fact, the amount, the eligibility, the deadline, hides somewhere in paragraph four.
02	Aggregators read more cleanly.	Third-party databases present rebates as tidy, structured, machine-legible data, so the machine leans on them.
03	Crawler barriers.	Some utility sites block or hinder the crawlers, which all but guarantees the machine reaches for a source it can actually read.
04	PDFs preferred, freshness unknown.	The machine happily takes the dense PDF and has no way to tell whether it is current unless the utility gives it a signal.

The aggregator trap, and why fragmentation quietly costs you

Here is a pattern in the data that deserves its own flag, because it runs against a comforting assumption.

Implementers, trade allies, 3rd party incentive aggregators (not part of your site) can either kill or enhance your authority and citations.

When the machine answered from a third-party source that lives off your domain, a rebate aggregator, an advocacy site, a trade ally or contractor page, it felt harmless in the moment. The customer still got a roughly correct answer. It is not harmless. Every one of those answers is a customer who formed their impression of your program somewhere you do not own, and it costs you four things at once.

The visit

You lose the visit, so the customer never lands on your site and never enters your funnel.

The content

You lose control of the content, because you cannot edit an aggregator's page and cannot fix it when your program changes and theirs goes stale.

The brand experience

You lose the brand experience, the reassurance, the design, the cross-sell, the human tone that makes a customer trust the program enough to enroll.

The relationship

And you lose the relationship, because the next time that customer has a question, the machine sends them back to the third party, not to you.

The remedy is to make the authoritative source of your program facts something you actually own. Your own domain is ideal. A branded subdomain works perfectly well. The point is control and continuity: a place where the fact is correct, current, and yours, so the machine has an owned source to trust and the customer has a path that leads home.

There is a second, subtler tax, and it is one utilities pay without noticing: **fragmentation**. Program information at a large utility tends to get scattered across a landing page, a separate rebate portal, a PDF, a microsite built by one implementer, and a schema implementation added by an agency. Each piece was reasonable on its own. Together they hand the machine a puzzle. When the same rebate appears three times with three slightly different numbers, or when your identity is described inconsistently across pages and vendors, the machine cannot resolve which version is canonical, so its confidence in all of them drops. Duplicate and conflicting content dilutes authority. Inconsistent structure muddies extraction. A splintered identity makes it harder for the machine to connect “your program” to “your utility” at all.

Consolidation is not a tidiness preference. It is how you give the machine one clear, consistent, authoritative source to believe. Platforms built for this challenge can do that consolidation for you, structuring the verifiable, machine-readable evidence about your programs and pulling your fragmented information into one owned, on-brand experience the machine can trust and the customer can follow, Incenva and Maven among them.

SECTION 09 • WHAT TO DO

Build the proof layer

Every failure in this audit traces back to a legibility gap, and legibility is buildable. For a machine to treat a program fact as the source of truth, that fact should be:

01	Findable.	On a crawlable HTML page with a clean, stable URL, reachable by the AI crawlers.
02	Extractable.	Stated as real HTML text and tables the machine can lift cleanly.
03	Specific.	Exact amount, exact eligibility, exact deadline, stated plainly.
04	Current.	Accurate and dated, with a machine-readable “last updated” signal and a human who owns keeping it right.
05	Corroborated.	The same fact stated consistently on the trusted third-party sources the machine already reads, pointing back to your owned source, so the wider web reinforces you.
06	Region-clear.	Every amount tagged to the territory it applies to, so the machine keeps your regions straight.

The principle underneath all six is **answer-first**: lead the page with the concrete fact, in plain text, at the top. Give the machine, and the customer, the answer before they have to go looking for it.

What a machine-readable program page looks like

The six criteria on a single page.

yourutility.com/rebates/heat-pump

Heat pump rebate

YOUR UTILITY · RESIDENTIAL

Up to \$1,000 back on a qualifying heat pump.

Updated January 2026

NORTH

SOUTH

Rebate amount	\$1,000
Eligibility	SEER2 ≥ 15.2, full system
Deadline	December 31, 2026
How to apply	Via an authorized contractor

How much is the heat-pump rebate?

Also verified on:

DSIRE

ENERGY STAR

Why the machine trusts it

- ★

Answer-first
Lead the page with the concrete fact, in plain text, at the top.
- 01

Findable
On a crawlable HTML page with a clean, stable URL, reachable by the AI crawlers.
- 02

Extractable
Stated as real HTML text and tables the machine can lift cleanly.
- 03

Specific
Exact amount, exact eligibility, exact deadline, stated plainly.
- 04

Current
Accurate and dated, with a machine-readable “last updated” signal and a human who owns keeping it right.
- 05

Corroborated
The same fact stated consistently on the trusted third-party sources the machine already reads, pointing back to your owned source.
- 06

Region-clear
Every amount tagged to the territory it applies to, so the machine keeps your regions straight.

A closing word on the goal, because it is easy to overreach here. You are not trying to banish every other voice from the answer. The machine will include contractors and third parties by design, and no amount of cleverness changes that. You are trying to be the most authoritative, most accurate, most consistently cited voice in an answer that will always have company, so the customer's decision finds its way back to you.

SECTION 10

Where this is going

Two forward-looking notes, flagged clearly as informed extrapolation, so you can weigh them as such.

From reading to acting.

Today, a lost citation costs you a click, and my audit found customers were almost never routed to a third party to actually transact, just 1 of 84. Soon, AI agents will begin acting on customers' behalf, enrolling them, choosing the installer, selecting the rate. When they do, the entity the machine cites as the authority becomes the entity the agent transacts with. Today's citation gap is an early-warning light for tomorrow's enrollment gap, and proof layers take time to build. The utilities that start now buy themselves a lead.

Emerging signposts.

Conventions like `11ms.txt`, a plain-text file that points AI crawlers toward a site's authoritative content, are worth piloting as a cheap experiment. They help the machine find your good content. Keeping that content accurate and current is still the work that wins.

How much to trust this study

I would rather tell you the limits than have you find them, because the credibility of this work depends on it.

It is a small, deliberate sample.

Eighty-four captures, 28 per utility. Directional signal, offered as such.

It is a snapshot.

AI engines change weekly. These results reflect August 2026 and will drift, and that drift is itself a finding: visibility here is something you monitor, like a garden, forever.

It is three utilities.

Chosen to span model, large-multi-region, and mid-market, and selected to be useful rather than nationally representative.

Scoring involved judgment.

We scored by hand against published terms, and where a program was genuinely region-dependent we counted a blended answer as a failure, which is a defensible call and an arguable one.

The right way to read this report is as the first systematic look at a real and growing shift, with patterns clear and consistent enough to act on, and an open invitation for every utility to run the same test on itself.

METHOD NOTE AND ABOUT THE RESEARCH

This audit was designed, conducted, scored, and interpreted by IDLab, by human beings with long experience in this industry. The method, a structured, repeatable, hand-scored capture protocol across engines and utilities, is proprietary to IDLab and forms the basis of an ongoing benchmark of AI findability for utility programs.

RUN IT ON YOUR OWN UTILITY

Curious what the machine is telling your customers about your programs?

The fastest way to find out is to run the audit on your own utility: the same protocol, applied to your programs, scored against your published terms, delivered as a ranked list of the specific gaps and what to fix first. IDLab can run your proof-layer audit and hand you your Monday to-do list.



IDLAB. CREATIVE, STRATEGY, AND MARKETING
FOR THE ENERGY AND UTILITY SECTOR.