ENGINE BEHAVIOR × AI VISIBILITY

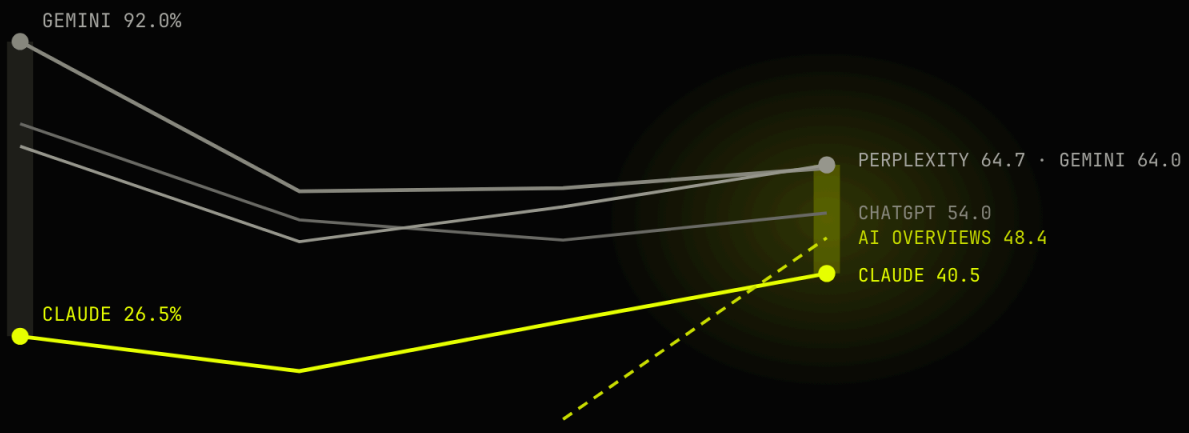
17 MIN READ

THE *CONVERGENCE.*

Six AI engines used to disagree about one contractor by 65 percentage points. Across 2,871 tracked checks, the spread closed to 24 — and Google's mainstream answer surface switched on.

MENTION RATE BY SURFACE

JUL → OCT 2026



65.5 PTS
JULY SPREAD

24.2 PTS
OCTOBER SPREAD

SIX SURFACES, ONE CONTRACTOR, CLOSING

2,871 TRACKED CHECKS

32.78°N · THE SPREAD, CLOSING

97.33°W · DFW

AUTHOR

The 6 Signal Research Team

AI Visibility · Engine Behavior

PUBLISHED

November 2026

READ

6signal.co/research/the-convergence →

WHITE PAPER · 17 MIN READ · NOVEMBER 1, 2026

The Convergence: Six AI Engines Used to Disagree About You. They Are Starting to Agree.

In July 2026 the four working AI engines we track disagreed about one contractor by 65 percentage points — 92% mention rate on Gemini, 26.5% on Claude. By October, across all six surfaces, the spread had closed to 24 points. Google AI Overviews went from citing the business in 0% of checks to 48.4%. A white paper on 2,871 tracked checks and what convergence means for a contractor's visibility strategy.

BY MATT VINCENT WALKER · 6 SIGNAL

The awkward thing about advising a contractor on AI visibility in the middle of 2026 was that there was no single answer to give. A business could be named by Gemini nine times out of ten and by Claude one time in four. Both numbers described the same company, the same week, answering the same question.

That is changing, and the change is large enough to alter strategy.

Between July 7 and October 31, 2026, our tracking system ran **2,871 automated checks** of 42 real buyer prompts against one Dallas–Fort Worth plumbing market, across six surfaces: ChatGPT, Claude, Perplexity, Gemini, Google AI Overviews, and Google Maps. This paper is about what happened to the *disagreement* between them.

TRACKED CHECKS

2,871

42 prompts · 6 surfaces · Jul 7 – Oct 31, 2026

JULY SPREAD

65 pts

Gemini 92.0% vs Claude 26.5% — same business

OCTOBER SPREAD

24 pts

64.7% down to 40.5%, across all six surfaces

AI OVERVIEWS, OCTOBER

48.4%

from 8.1% in September — Google's mainstream surface switched on

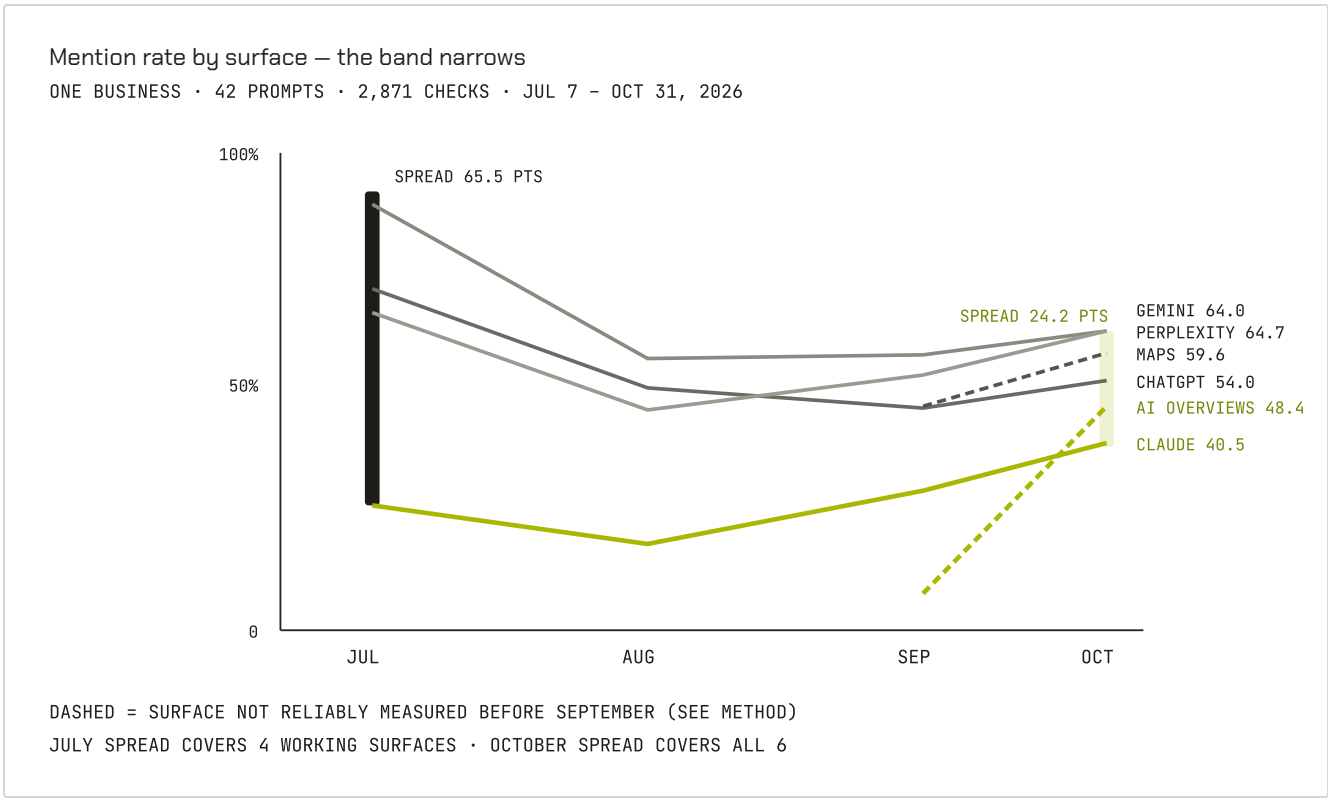
Part I — The spread, and how it closed

The cleanest way to see convergence is to stop looking at any one engine and look at the distance between them.

In July, four of the six surfaces were producing trustworthy captures. Among those four, the tracked business was mentioned in 92.0% of Gemini checks, 73.8% of ChatGPT checks, 68.8% of Perplexity checks, and 26.5% of Claude checks. Highest to lowest: **65.5 percentage points.**

In October, all six surfaces were producing trustworthy captures, and the range ran 64.7% (Perplexity), 64.0% (Gemini), 59.6% (Maps), 54.0% (ChatGPT), 48.4% (AI Overviews), 40.5% (Claude). Highest to lowest: **24.2 percentage points** — across two more surfaces

than July's figure covered.



The October band is narrower than the July band despite containing two more surfaces. That is the finding.

SURFACE	JUL	AUG	SEP	OCT	DIRECTION
Gemini	92.0%	58.8%	59.4%	64.0%	Fell hard, then stabilized
ChatGPT	73.8%	52.4%	48.0%	54.0%	Fell, then recovered
Perplexity	68.8%	47.6%	55.4%	64.7%	Dipped, then climbed
Claude	26.5%	18.8%	29.8%	40.5%	**Rose every month after August**
Google AI Overviews	—	—	8.1%	48.4%	**Switched on**
Google Maps	—	—	48.6%	59.6%	Came online

Part II — Convergence is mostly the leader falling

It would be a pleasant story if convergence meant every engine getting better at finding good contractors. That is not what the data shows.

Gemini named this business in 92% of July checks. It has not been near that since — 58.8%, 59.4%, 64.0%. The business did not get worse. Its site, its license record, and its reviews were stable throughout; our October paper shows its own website holding as the most-cited domain across the whole period. Something changed inside Gemini's answer behavior, and the effect was to stop naming a specific local business as reliably as it once did.

Meanwhile Claude climbed from 18.8% in August to 40.5% in October. ChatGPT and Perplexity both dipped mid-window and recovered. The net is a band forming in the 40% to 65% range, approached from both directions.

For a contractor, the practical reading of that asymmetry matters more than the convergence itself:

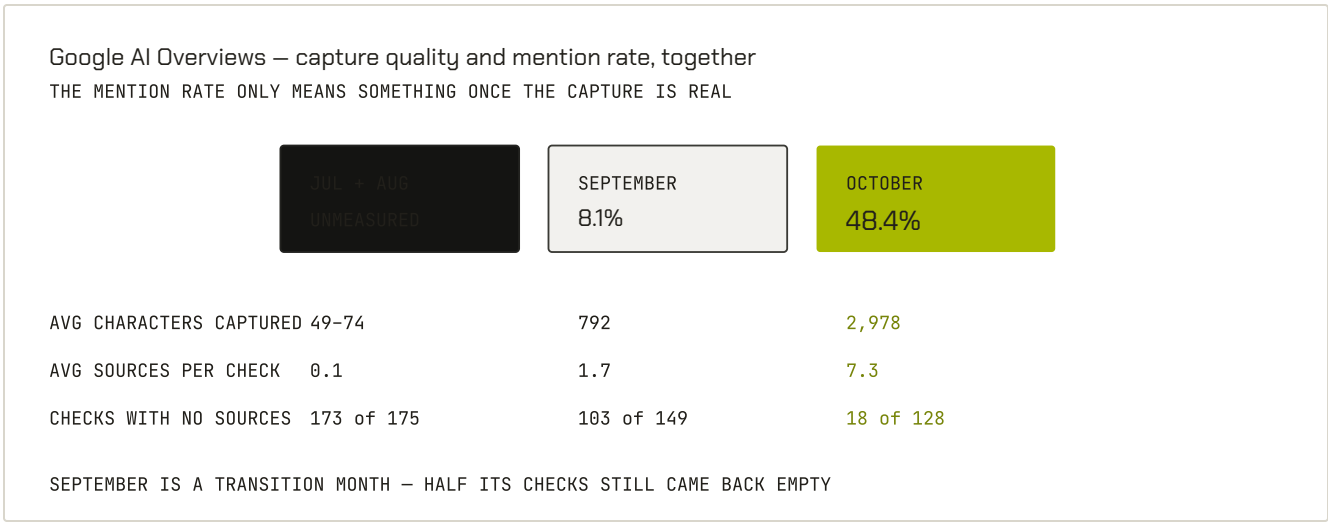
- **A strong early position on one engine was never durable.** The 92% was real when we measured it and it was gone six weeks later, without the business doing anything wrong. Any strategy whose success condition is a single engine's behavior is renting its results.
- **The floor is rising.** The weakest surface we track went from 18.8% to 40.5% in three months. Being invisible on a given engine is becoming less common — which also means it is becoming less excusable.
- **The ceiling is falling.** Nothing in this dataset now sits above 65%. The reflex to chase a perfect score on a leaderboard is chasing a number that the engines themselves appear to have stopped producing.

Part III — The Overviews event

The largest single change in the window is Google AI Overviews going from functionally absent to naming this business in roughly half of all checks.

We have to be careful here, and the care is not optional. As documented in The Directory Collapse, our Overviews probe was broken in July and August: it captured an average of 74 and 49 characters per check, with no sources attached in nearly every check. Those months measure our instrument, not Google. The first trustworthy numbers are **September at 8.1%** and **October at 48.4%**, by which point the probe was capturing nearly 3,000 characters and seven sources per check — the profile of a real Overview.

So the honest claim is narrow and still significant: *between September and October 2026, in this market, Google AI Overviews went from naming this contractor rarely to naming it in about half of all checks.* Whether the true July baseline was zero, or something we simply failed to read, we cannot say and will not guess.



We are publishing September as a transition month rather than a data point. Two-thirds of its checks still returned nothing.

Why this matters more than the other five surfaces combined: **AI Overviews is the only one of them that appears in front of someone who did not choose to use an AI product.** A homeowner who types "water heater replacement near me" into Google gets an Overview whether or not they have ever opened ChatGPT. Every other surface in this study requires the buyer to have adopted something. This one arrives by default, at the top of the page that local search has run through for twenty years.

Which makes the September-to-October move the single most consequential thing in this dataset. Whatever share of local buyers an AI answer reaches, that share just went up sharply, on the surface with the broadest reach.

Part IV — What convergence changes about strategy

When engines disagreed by 65 points, the rational contractor response was uncertainty: nobody could tell you which surface mattered, so nobody could tell you what to fix first. Convergence resolves that, in a way that is good news for operators and bad news for tactics.

One substrate, six outputs. If six independently built systems increasingly agree about a business, they are increasingly reading the same underlying evidence. Our source-trail data says what that evidence is: the company's own website first, then institutional records, then the commercial directory layer — the re-weighting documented last month. Improving that substrate is what moves all six numbers. There is no longer a plausible engine-specific play, if there ever was.

The distribution of outcomes is tightening, not the ranking. Convergence does not mean every business scores the same. It means the *same* business scores similarly across engines. A contractor with a thin site, a stale license record, and inconsistent

listings now gets a consistently poor result everywhere, where six months ago they might have caught a lucky 92% on one surface and called it a strategy.

Measurement gets cheaper and more honest. If the engines agree, you need fewer checks to know where you stand. The corollary is that a single flattering screenshot from one engine is now worth even less as evidence than it was — and it was never worth much. Read your results as a spread, not a score.

Timing has changed for Google specifically. Any contractor who concluded in August that AI Overviews was not naming local businesses and deprioritized accordingly is working from a conclusion that is now two months stale — and, as it happens, was never reliably measured in the first place.

Part V — The work order

01 **Stop optimizing per engine.** If you have been maintaining a mental model where ChatGPT needs one thing and Gemini another, retire it. Spend on the shared substrate.

02 **Audit the substrate in the order the engines weight it.** Your own site, then your license and municipal records, then the directory layer. The Directory Collapse has the per-check evidence for that ordering.

03 **Assume Google's Overview now sees you — or does not.** It is the default surface for buyers who never adopted an AI tool. Test your own market directly rather than relying on anything published before September.

04 **Measure a spread, not a number.** Run the same prompts across at least four surfaces and record the range. A 24-point band is now normal; a 60-point band means something is broken on one surface and worth diagnosing.

05 **Re-measure quarterly, and expect the band to keep moving.** Everything in this paper describes seventeen weeks. Use the measurement playbook to run it yourself, or let the $27 AI Visibility Audit score all six signals for your trade and city in one pass.

Method, and what this paper does not claim

The data comes from our internal tracking system: 2,871 automated checks of 42 buyer prompts for one Dallas–Fort Worth plumbing market, July 7 – October 31, 2026, run against ChatGPT, Claude, Perplexity, Gemini, Google AI Overviews, and Google Maps, with each engine's answer and source list stored at check time. Mention rate is the share of checks in which the tracked business was named in the answer.

Google AI Overviews and Google Maps captures before September 2026 are excluded from every figure: both show the signature of a capture failure on our side — average captured answer lengths of 49 to 74 characters for Overviews and roughly 330 for Maps, with sources absent in nearly every check — rather than genuine silence from the engines. This is documented at length in The Directory Collapse, and it is why July's spread figure covers four surfaces while October's covers six. Comparing a four-surface spread to a six-surface spread is conservative in the direction of our own argument: adding surfaces can only widen a range, and the range narrowed anyway.

This is **one business, one trade, one market, over seventeen weeks.** Convergence in this dataset is not evidence of convergence everywhere, and a four-month window cannot separate a structural shift from a long swing. We can also see only outputs, never the systems producing them, so every explanation offered in Part II is a hypothesis. What we can state precisely is the measurement: the gap between the most and least generous

surface narrowed from 65.5 points to 24.2 while the surface count grew from four to six, and Google's mainstream answer surface went from 8.1% to 48.4% in a single month. We will publish the next window whether or not it agrees with this one.

The Convergence is part of 6 Signal's visibility research. For the source-trail analysis underneath it, read The Directory Collapse; for the trust-stack model both papers build on, read The Supplier Signal; for the original baseline, read We Ran 700 AI Visibility Checks in One Texas Market. To see where your business stands today, start with the free AI visibility check or the $27 AI Visibility Audit.

Sources

- 6 Signal internal tracking system — 2,871 probe records with stored answers and source trails, July 7 – October 31, 2026 (method above)
- Probe capture-quality diagnostics (answer length, source counts, empty-capture rates by engine and month) — same system, queried 2026-10-31

NEXT STEP

See where you *actually stand.*

The AI Visibility Audit runs your company through all six layers — GEO, AEO, LEO, VEO, PEO, IEO — and delivers instant results. $27. Specific to your business, your trade, and your market.

GET THE AUDIT

6signal.co/visibility-check