EVIDENCE-GRADE GTM INSIGHT FOR FOUNDER-LED B2B

The Sagentix Signal

ISSUE #3 · 2026-06-09 · MESSAGING

Your messaging sounds generic because no one built the structure underneath it

Open four competitor websites in different tabs. Read the hero sections back-to-back. If a buyer cannot tell which one is yours in five seconds, the problem is not your tagline. It is the structure underneath it.

Most founders treat messaging as a copywriting exercise. They commission a tagline, refresh the homepage, and move on. The result is a sentence that sounds clever in a deck and dissolves into noise everywhere else. Gartner's May 2025 sales survey of 632 B2B buyers found that 74% of buying teams demonstrate "unhealthy conflict" during the decision process (Gartner, 2025), and its buying research adds why that conflict is structural: 99% of B2B purchases are driven by organizational change, with 66% of buyers describing the amount of change as overwhelming (Gartner, 2023). Inside that committee, your message travels through people who do not work for you, to people who never read your homepage. The version of your story that wins is the one your champion repeats from memory, in their own words, while defending your invoice to a CFO.

That repetition only happens when the message has structure. The structure has a name: the message house. A core claim, three pillars, proof points under each pillar, vertical maps that adapt it per segment, and talk tracks that convert the pillars into objection handling. When any layer is missing, the buyer does the synthesis themselves, and they get it wrong in ways you never see.

The sharper question is not *what is our tagline?* It is *what is the one sentence the customer's champion repeats, the three pillars that sentence rests on, and the proof points that survive due diligence?* That is the unit of work. The tagline is the last paragraph, not the first.

The core claim earns the tagline

A positioning statement is not a tagline. The core claim is one declarative sentence that names the change you create: for whom, against what alternative, and why the alternative fails. The tagline is the compressed echo. Most founders write the echo first, never write the claim, and then wonder why the echo lands flat.

Forrester reports that nearly 95% of B2B buyers expect to use generative AI in their purchase decisions within twelve months, a figure Forrester states directly (Forrester, 2024). What follows from it is my reading, not Forrester's: buyers increasingly meet your homepage with a model open alongside it, asking for a comparison you did not write (Sagentix GTM Methodology, 2026). The features the model lists back to itself are the ones repeated across vendor pages. If your core claim cannot be paraphrased in two seconds and held in working memory, the LLM picks the one that can.

> **DECISION RULE:**
>
> Before approving a tagline, write the core claim it compresses. If you cannot name the buyer, the alternative, and the failure mode of the alternative in one sentence, the tagline has nothing to compress.

Pillars without proof points are decoration

Three pillars carry the core claim. Each pillar is a theme, not an adjective. "AI-powered," "comprehensive," "real-time" are adjectives. Themes are claims a champion will defend in a meeting: "we replace seven tools with one workflow," "we shorten the audit cycle from twelve weeks to four," "we are the only vendor whose data your CFO can audit on a Friday afternoon."

Each pillar carries proof points. A proof point is a specific, verifiable fact: a customer name, a measured outcome, a regulator's published threshold, a peer-reviewed study, a number from your own data. The Canadian software publishing industry generated $23.2 billion in revenue across roughly 3,300 publishers in 2026 (VerticalIQ, 2026), a market dense enough that buyers rarely take a vendor's adjective on faith. Proof points convert a pillar from marketing copy into evidence the champion defends without you in the room.

Brand voice glues this together. Brand voice is not "friendly" or "authoritative." It is the consistent way you translate your three pillars across a homepage, a sales call, a conference talk, and a one-line LinkedIn post. When the voice fragments, the message fragments, and the buyer concludes the company itself is fragmented. Storytelling is the discipline of carrying the same pillars through every channel without softening or thickening them in transit.

Talk tracks turn pillars into objection handling

A talk track is not a script. It is a 30-second bridge that lets a sales rep, a CSM, or a founder at a dinner answer a common objection — "you are too small," "we already use Vendor X," "the procurement process will not approve a new line item" — by routing the objection back to a pillar with a proof point.

The average enterprise buying group runs to five to 11 stakeholders representing roughly five distinct business functions (Gartner, 2023). Many of them will hear the pillars second-hand rather than from you. A talk track ensures the second-hand version does not collapse into *they said they're AI-powered*. Talk tracks make the message portable. They are the storytelling that survives transfer between people — and the mechanism is better evidenced than the word "storytelling" usually suggests. A Warwick doctoral study of B2B opportunity generation finds that buyer intuitive thinking corrupts established sales methodologies through selective attention, confirmation bias and risk aversion, and that the durable countermeasure is not forcing the buyer to reason harder but **redirecting that fast thinking** with story and social proof (Al-Sened, 2025). Its framing of the purpose is the useful part: storytelling there is a discovery technique for encouraging disclosure and extracting needs, not a pitching technique.

Do this this week. Open a blank document. Write the core claim. Write the three pillars. Under each pillar, write three proof points and a talk track for the objection you hear most often. If any pillar has fewer than three proof points, that pillar is not yet defensible. The first version will be uncomfortably specific. That discomfort is the work, and it is the test the message has to pass before a buyer ever sees the homepage.

The Second Signal

One core claim, many segments, never many claims

The most expensive failure in messaging architecture is not a weak tagline. It is the moment a founder decides each vertical needs its own message and ends up with three brands, three sales decks, and three sets of proof points that do not reinforce one another.

Vertical maps adapt the message; they do not replace it. The core claim stays the same regardless of who reads it. The three pillars stay the same. What shifts per vertical is the proof point density, the named reference customers, the regulatory vocabulary inside the talk tracks, and the objection that earns the most airtime. The claim and the pillars are identical. The evidence around them changes.

Buyers detect the seam. Gartner finds that B2B buyers value interactions with third parties (peers, communities, analysts) more than they value direct supplier interactions (Gartner, 2023b). When your champion forwards your one-pager to a peer in a different vertical, that peer should recognize the company. If they read it and ask, "wait, are these the same people who pitched us?", the message has been rebuilt per segment instead of adapted.

The fix is structural, not creative. Build it once. Document the pillars and proof points centrally. Write a vertical narrative document per segment that reuses the pillars verbatim and swaps only the proof points and the talk-track objections. The discipline is dull and load-bearing. It is also what lets a five-person company sound like one company across six channels, a state most twenty-person companies never reach (Sagentix Phase 03 Messaging Architecture, 2026).

The Third Signal

Paid demand generation went flat for three separate reasons

If paid demand generation that carried your pipeline for two years has gone flat inside about six months, with no change to the offer, the budget, or the team, the instinct is to audit the team. Check the environment first. Three things in it changed, all outside your account, and only the third is yours to fix.

The auction repriced. Meta reports average price per ad up 12% year over year in each of the first two quarters of 2026, against 10% and 9% in the same quarters a year earlier, while impressions grew 19% then 14% (Meta Platforms, 2026). More inventory, still costing more.

The open web that display budgets run on is contracting. Google Network impressions, the inventory on publisher sites rather than Google's own results pages, fell 12% year over year in Q2 2026 while cost-per-impression rose 13%. Google Search moved the opposite way in the same filing: paid clicks up 13%, cost-per-click up 3% (Alphabet, 2026). The squeeze is in display, not search.

Pew's browsing panel shows why. Across 68,879 Google searches by 900 U.S. adults in March 2025, users meeting an AI summary clicked a traditional result in 8% of visits against 15% for those who did not, and clicked a source inside the summary in 1%. Users ended the session entirely on 26% of pages carrying an AI summary, against 16% of pages without (Pew Research Center, 2025). Fewer sessions on the open web, fewer display impressions to buy.

Neither of those is yours to solve. The third one is.

The third reason is the one you control

The problem. Everyone's messaging converged at the same time, and the research on why is unusually clean.

Doshi and Hauser gave some writers story ideas from a large language model. Those stories were rated more creative, better written and more enjoyable, especially for less creative writers, and were measurably more similar to one another than stories written without. The authors call it a social dilemma: with generative AI, writers are individually better off while collectively a narrower scope of novel content is produced (Doshi & Hauser, 2024).

Moon, Green and Kushlev quantified it across three preregistered studies of 2,200 admissions essays. In the first, each doubling of human-written essays produced a 0.40 increase in collective diversity against 0.13 for GPT-4, about 31% of the human rate. The third study put GPT-4 at 11%. They tested three mitigations, and a gap survived every one (Moon et al., 2025).

The diagnostic. Better prompting improved the output without restoring the variety. That is what you would expect if convergence is happening across the population of firms using the same few models rather than inside any single company's prompt (Sagentix GTM Methodology, 2026). It follows that no amount of rewriting gets you out, because the rewriting is the thing that converged. If your differentiation lives in your *language*, it is now a commodity input. You need differentiation that sits somewhere a model cannot reach.

The solution. Christensen's Jobs-to-be-Done is that somewhere, and he states the mechanism directly: "When you see a product or service that no one has successfully copied, the product itself is rarely the source of the long-term competitive advantage." What is defensible is the insight into the job, which "allowed the company to differentiate its offering in ways competitors weren't likely to copy—or even comprehend" (Christensen et al., 2016).

The reason that holds against AI is structural. A job is the progress a customer is trying to make *in a given circumstance*, and Christensen is explicit that the circumstance matters more than any buyer characteristic. A circumstance is a fact about one real buyer at one real moment. It is not in the training corpus. A model can generate your adjectives, your three pillars and your tagline. It cannot generate the job your customer is hiring you for, because it has never met them.

That is the moat. Not the product, and not the words: the job, plus proof you can serve it that somebody other than you validated.

How we run it. Phase 02 is Value Proposition Design and its unit of work is a JTBD statement, not a positioning line. Phase 03 builds the message house on top. Phase 10 then grades every claim by who validates it, on a five-level scale from Verified to Unsubstantiated (Sagentix Phase 10 Evidence Discipline, 2026).

That last step is where the sameness becomes visible and measurable. In one recent engagement, five of the twelve claims the company could make were things every competitor could also say — and a sixth was a credential whose own issuer disclaims it as evidence of service quality. That firm did not have a copy problem. It was leading with the half of its story that anyone could generate.

Then it compounds in your favour, because a Digitas UK pilot on B2B fintech payment solutions found more than 80% of the sources large language models drew on came directly from the brands themselves (Joshi et al., 2026). Buyers increasingly arrive through a model rather than a results page, and the model reads your owned content. Converged content gives it nothing to tell you apart by. A named job and an external validator give it something no competitor's prompt can produce.

> **DECISION RULE:**
>
> Take your three pillars and ask two questions of each. Could a competitor produce this with the same prompt and no access to my data? And does it name a job a real buyer is trying to get done, in a circumstance I can describe? A pillar that fails both is not a differentiator, and paying to amplify it is paying to sound like the category.

From the Field

The pattern shows up in every artifact a founder ships: the website, the deck, the board update, the cold email.

- **We Have the Best Product. We Just Can't Explain Why.** A CEO with a measurably better product was losing five of seven competitive deals because the sales team described the platform with the same adjectives as four competitors. The pattern is structural: when the pillars are adjectives the whole market shares, the champion has nothing distinct to repeat.

- **The Series A Board Deck Has an Evidence Gap — and It's Killing Founder Credibility** Board members who used to nod through narrative slides now interrogate the proof points behind every claim. Slides survive that shift when the proof points exist before the slide is built, when the messaging carries the evidence and the deck just renders it.
- **Canadian B2B SaaS: The $23B+ Opportunity Window** A reminder that 94% of Canadian software firms have fewer than 100 employees and 54% sit in Ontario (VerticalIQ, 2026). When the field is that crowded and concentrated, undifferentiated messaging is not a soft problem; it is the reason the market does not see you.
- **From Market Intelligence to Evidence Discipline: The 10 Phases of a Complete GTM Strategy** Phase 03 (Messaging Architecture) is sequenced after Phase 01 (Market Intelligence) and Phase 02 (Value Proposition Design) for a reason: a message house built without market evidence collapses under buyer scrutiny. This piece walks the full sequence and shows where each phase feeds the next. Read this if your team is debating which phase to start with.

The same pattern reappears in capital markets, the part of the buying environment most founders do not see directly.

What We're Watching

Two things over the next 14 days. First, whether the Google Network contraction continues into Q3 filings, because a second consecutive double-digit impression decline would make display repricing structural rather than cyclical, and would change where a seed-stage budget belongs. Second, whether Canadian founders pitching at the Toronto-Waterloo cluster can name a proof point that a competitor could not generate from a public prompt. That second one is not measurable from a filing. It is measurable in a room, and it is the faster diagnostic.

The board-room standard for a defensible claim is rising faster than founder messaging is catching up. PitchBook reports B2B SaaS captured more than billion in venture capital globally in 2025, and the boards behind that capital now interrogate positioning the way they used to interrogate the financial model (PitchBook, 2025). Messaging architecture has become a diligence artifact, not a marketing one.

Work With Us

If you are heading into Q2 planning and your homepage, deck, and one-pager would not survive being placed on a wall next to four competitors, our Phase 1 Market Intelligence engagement runs in 5–7 days at CA$4,500 and produces the evidence base a defensible message house rests on. Money-back if the evidence does not hold. Reply with "Phase 1" or book at sagentix.ca/contact.

— Stephane Raby CMC · CISSP · P.Eng. · 25+ Years · Executive MBA

References

- Al-Sened, D. (2025). *Thinking fast; selling slow: How buyer intuitive thinking stalls B2B opportunities, whilst also providing the antidote in storytelling* [Doctoral thesis, University of Warwick]. Warwick Research Archive Portal (repository intermittently unreachable 2026-08; record indexed at OpenAlex). https://api.openalex.org/works/W7165297879
- Alphabet. (2026, July 23). *Quarterly report (Form 10-Q) for the quarterly period ended June 30, 2026*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1652044/000165204426000071/goog-20260630.htm
- Christensen, C. M., Hall, T., Dillon, K., & Duncan, D. S. (2016, September). Know your customers' “jobs to be done”. *Harvard Business Review*. https://hbr.org/2016/09/know-your-customers-jobs-to-be-done
- Doshi, A. R., & Hauser, O. P. (2024). *Generative AI enhances individual creativity but reduces the collective diversity of novel content*. *Science Advances*, 10(28). https://doi.org/10.1126/sciadv.adn5290
- Forrester. (2024, December 4). *Forrester: To master B2B buying mayhem, providers must prioritize buyers' needs* [Press release based on *The State of Business Buying, 2024* report]. Forrester Research. https://www.forrester.com/press-newsroom/forrester-the-state-of-business-buying-2024/
- Gartner. (2023). *The B2B buying journey*. Gartner, Inc. https://www.gartner.com/en/sales/insights/b2b-buying-journey
- Gartner. (2023b, June 8). *Gartner marketing survey finds B2B buyers value third-party interactions more than digital supplier interactions* [Press release]. Gartner, Inc. https://www.gartner.com/en/newsroom/press-releases/2023-06-08-gartner-marketing-survey-finds-b2b-buyers-value-third-party-interactions-more-than-digital-supplier-interactions

- Gartner. (2025, May 7). *Gartner sales survey finds 74% of B2B buyer teams demonstrate "unhealthy conflict" during the decision process* [Press release]. Gartner, Inc. https://www.gartner.com/en/newsroom/press-releases/2025-05-07-gartner-sales-survey-finds-74-percent-of-b2b-buyer-teams-demonstrate-unhealthy-conflict-during-the-decision-process
- Joshi, A., Buche, I., & Schwaer, C. (2026, June 12). *How gen AI is disrupting B2B buying decisions*. Harvard Business Review. https://hbr.org/2026/06/how-gen-ai-is-disrupting-b2b-buying-decisions
- Meta Platforms, Inc. (2026, July 30). *Quarterly report (Form 10-Q) for the quarterly period ended June 30, 2026*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1326801/000162828026050705/meta-20260630.htm
- Moon, K., Green, A. E., & Kushlev, K. (2025). *Homogenizing effect of large language models (LLMs) on creative diversity: An empirical comparison of human and ChatGPT writing*. *Computers in Human Behavior: Artificial Humans*, 6, 100207. https://doi.org/10.1016/j.chbah.2025.100207
- Pew Research Center. (2025, July 22). *Google users are less likely to click on links when an AI summary appears in the results*. Pew Research Center. https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/
- PitchBook. (2025). *SaaS fundraising trends 2025: Venture capital deployment in enterprise software* [Industry analysis]. PitchBook Data. https://pitchbook.com/
- Sagentix Phase 03 Messaging Architecture. (2026). *Message house construction methodology* [Internal methodology]. Sagentix Advisors Inc. https://sagentix.ca/how-it-works
- VerticalIQ. (2026). *Canadian software publishing industry profile*. VerticalIQ. https://app.verticaliq.com/