VEGA STUDIO

VEGA STUDIO
RESEARCH
PAPER 03
NOVEMBER 2026

865 Companies, One Read

What a cold buyer can't read on your first screen, and the five-question test that fixes it

51%

of 865 observations: a cold buyer can't tell who the page is for, or which problem to start with

67%

of B2B buyers prefer a rep-free experience. Your page is the rep now

Vega Studio research, October 2026 · Gartner, March 2026

1 The cost of an unclear first screen
2 What we found in 865 homepages
3 The One-Read Test
4 Four rewrites from real observations
5 Objections
6 The self-check

EXECUTIVE SUMMARY

Your homepage is now your first sales call. Gartner surveyed 646 B2B buyers in August and September 2025 and found that **67% prefer a rep-free experience** (Gartner, March 2026). 6sense found that **94% of buying groups have already ranked a preferred vendor before first contact, and buy from that favorite 77% of the time** (6sense, 2025).

Put those together. Most buyers decide whether you are worth a conversation before they talk to anyone, and the thing they read to decide is your public message.

So we read it. In October 2026, Vega Studio's research pipeline, an AI system checked by people, looked at the public message of **865 B2B companies**, mostly in Europe and on the US East Coast, all with a recent funding round or open go-to-market roles. One pass per company. One written observation each: the most visible thing a cold buyer would struggle with.

Half of the observations (51%) were about the same thing: a cold buyer can't tell who the page is for, or which problem to start with. A quarter (24%) were about a broken first step. Almost one in five (18%) were about a new hire about to inherit a message nobody has fixed.

None of this is a design problem. It is a decision problem. Somebody has to choose one buyer, one problem and one first step, and in most of these companies nobody has.

This paper gives you the tool we use to find those gaps in five minutes: **the One-Read Test.** Five questions a cold reader should answer after one read of your first screen, each with a pass or fail test and a fix pattern.

FIGURE 1. WHAT 865 OBSERVATIONS FLAGGED

Speaks to several buyers at once	35%
Catalogue instead of one door	19%
Either of the two	51%
First step broken (demo-only or competing steps)	24%
GTM hire will inherit an unfixed message	18%

Vega Studio research, 865 B2B companies, October 2026. One observation per company. Multiple categories allowed.

SECTION 01

The cost of an unclear first screen

Your page is the rep now

For years the homepage was a brochure. The real explaining happened later, on a call, with a person who could read the room and adjust.

That person is leaving the early part of the journey. In Gartner's March 2026 survey of 646 B2B buyers, **67% said they prefer a rep-free experience**, and 45% said they had used AI during a recent purchase (Gartner, March 2026). The same release names the goal of the whole exercise as "value clarity": the buyer understanding how a solution improves outcomes in *their* role and context. Confident buyers, Gartner reports, are **twice as likely to report a high-quality deal** as buyers with low decision confidence.

Read that as a description of your first screen. If the buyer can't find their role and their problem on it, nobody is there to translate.

The shortlist forms before you know the buyer exists

6sense's *Buyer Experience Report* surveyed more than 4,000 buyers across North America, EMEA and APAC. **94% of buying groups had ranked preferred vendors before first contact. They bought from that preliminary favorite 77% of the time** (6sense, 2025).

If your page does not make you the obvious fit for one buyer, you are not losing the deal on the call. You are not on the shortlist that decides who gets a call.

Then the seller contradicts the site

There is a second cost, and it lands after the hire. In Gartner's June 2025 release, **69% of B2B buyers reported inconsistencies between the information on a supplier's website and what its sellers told them** (Gartner, June 2025; 632 buyers surveyed August to September 2024).

That number is the mechanism behind one of our categories. When a page speaks to five buyers, each new seller picks one and tells their own version. The site says one thing, the rep says another, and the buyer notices.

The first screen now carries the weight the first meeting used to carry. The rest of this paper is about how often it fails, and how to tell if yours does.

SECTION 02

What we found in 865 homepages

Method. Read this before the numbers.

Sample. 865 B2B companies with a recent funding round or open go-to-market roles, researched October 12 to 28, 2026. 430 in Europe, 182 on the US East Coast, 253 unassigned.

Data. One written observation per company, written in a single pass by Vega Studio's research pipeline, an AI system (Claude) that reads public pages: the most visible issue a cold buyer would hit, not an audit of the whole site. Source page: homepage in 450 rows, careers or job board in 204, news or blog in 131, other pages in 80. All are about the public message; not all are about the first screen.

Classification. Ten categories, multi-label. A Python script applies transparent rules; the 236 rows the rules could not place were read and labelled one by one by an AI model (Claude).

Validation. A random sample of 60 rows (fixed seed), labelled row by row by the same AI model, independently of the rules, and compared to the pipeline, scored on the 50 rows never seen while writing the rules. A percentage is published only with at least 85% agreement and at least five real cases in the sample. Published categories ranged from 92% to 98% agreement (86% to 94% for the rules alone, before the row-by-row review; the reviewed figure is not fully independent, because the same model did both).

Limits. These figures show the most visible problem per company, not every problem. Each row came with a research angle that steered what the pipeline looked at. The sample leans toward funded, hiring companies. One AI model did all row-by-row labelling; there was no second, human coder. People at Vega Studio set the questions and the publication rules and checked the categories, the examples and every published number. A row-by-row cross-check suggests the pipeline undercounts, so treat every number as a floor.

The headline finding: too many buyers

35% of observations said the page speaks to several buyers at once. Not a vague page. A page with a very specific list.

The typical form is a list. "The homepage speaks to banks, fintechs, marketplaces and digital asset businesses." "The homepage lists glass packaging, beverage and filling, food, consumer goods, tableware and sawmilling." "Payables, receivables, treasury and expenses are usually

different buyers inside a finance team."

Each list makes sense from the inside. But a cold reader is not reading for coverage. They are reading for themselves, and if they can't find their own name, they assume the product is for someone else.

It is the most stable pattern in the dataset: 32% to 43% of rows under every research angle we used, at similar rates in Europe and the US. It is what B2B homepages look like when nobody has chosen.

The second finding: a catalogue instead of a door

19% of observations flagged a catalogue. Five products. Nine modules. "More than 40 agents." Eleven solutions. A product described three different ways.

The buyer's question is not "what can you do" but "what do I ask for first". A catalogue makes them answer that alone.

Taken together, 51% of the 865 observations were about one of these two failures. Not style, layout or copy quality. The page had not decided who it was for or which problem opens the door.

The third finding: the first step

24% of observations were about the first step.

Half of those (**12%**) were a single heavy door: the only way in is "Book a demo", a walkthrough, a 12-week pilot or an assessment. That is a calendar commitment before the buyer knows if you are relevant. Our row-by-row cross-check put this closer to 18%.

The other half (**12%**) were the opposite: too many doors. "Get started free, book a demo and estimate the hours." "Request a demo, get access, contact sales and deploy at your firm." Four calls to action with no hint of which one fits which buyer.

Same root: nobody decided what the first step is *for whom*.

The fourth finding: the hire will inherit it

18% of observations came from the hiring side. An account executive in a new city. A product marketer. A head of growth next to two US sales roles. "Seven sales roles posted on October 26 and 27."

The company is about to add people who will repeat a message that has not been fixed, so each hire writes their own version. That is how you get the 69% from section 1.

A smaller group (**8%**) was about a new market: a round raised to expand, a seller hired abroad, and a message still built for home. A new country usually needs its own first sentence, not a translation.

What we saw but won't put a number on

Four categories were too rare in the validation sample to publish a reliable percentage. We report them as patterns.

- **Headlines that state ambition, not an outcome.** "Move more projects forward with confidence", which could describe almost any project software. "Unlock data where it lies." A headline that ends in "and beyond".
- **Proof that belongs to someone else.** Customer logos mostly from one country, on a page trying to open another. A single case study for three kinds of buyer. A savings calculator that labels its own results "indicative, pending validation".
- **Pricing that won't let you place yourself.** Three plans, all "Contact sales", no figure anywhere. A page that promises transparent pricing and shows none.
- **Founders still doing all the selling.** No commercial roles open, a demo button that books straight into the founder's calendar.

FIGURE 2. THE TEN CATEGORIES, AND HOW OFTEN EACH WAS FLAGGED. VEGA STUDIO RESEARCH, OCTOBER 2026. MULTIPLE CATEGORIES PER COMPANY ALLOWED, SO SHARES DO NOT SUM TO 100%.

Category	Share of 865	Status
Speaks to several buyers at once	35%	Published
Catalogue instead of one door	19%	Published
GTM hire will inherit the message	18%	Published
Several competing first steps	12%	Published
Demo or sales call is the only door	12%	Published (likely undercount)
Message not ready for a new market	8%	Published
Headline states ambition, not outcome		Qualitative only
Proof doesn't match the buyer		Qualitative only
Pricing with no self-placement		Qualitative only
Founder still sells alone		Qualitative only

SECTION 03

The One-Read Test

The categories above map onto five questions. A cold buyer should be able to answer all five after **one read** of your first screen: no scrolling, no clicking, no prior knowledge of your company.

Run it on someone who has never seen your company. Ten seconds on the first screen, close it, ask the five questions. Write down what they say, not what you meant.

FIGURE 3. THE ONE-READ TEST

#	Question	Fails when	Seen in our data
1	**Who buys?**	The reader can't name one buyer, or names three	35%
2	**Which problem?**	The reader can't name the one problem you solve first	19% (catalogue)
3	**What changes?**	The reader can't say what is different after they buy	qualitative
4	**Why believe?**	The proof on screen is for a different buyer, or absent	qualitative
5	**One first step?**	There is no small step, or there are several with no guidance	24%

Question 1 · Who buys?

The test. Ask the reader: "Who is this for?" Pass if they name one role in one kind of company, in their own words. Fail if they say "businesses", list three industries, or hesitate.

The fix pattern: one lead buyer, the rest one click down. Pick the buyer who converts best today, or the one your round was raised to win. Write the first screen for that buyer only, and name them. Give every other segment its own page, linked from one line below the fold: "Also for [segment B] and [segment C]." You are not abandoning the others. You are ordering them.

Question 2 · Which problem?

The test. Ask: "What problem does this solve first?" Pass if the answer is one problem the buyer already has a name for. Fail if they describe the platform, or list modules.

The fix pattern: one door, then the house. Pick the module that opens the most first conversations and lead with the problem it solves, in the buyer's words. Show the rest as "what you add next". The other products are your expansion revenue, not your first sentence.

Question 3 · What changes?

The test. Ask: "What is different for the buyer after they use this?" Pass if the reader names a concrete change: a number, a task that disappears, a decision that gets faster. Fail if they repeat your category or your tagline.

The fix pattern: before and after in one line. Write the headline as the buyer's situation changing: "[Buyer] goes from [current pain] to [result]." Use a real number if you have one; if not, name the task that goes away. This is Gartner's value clarity: the outcome, in the buyer's own role.

Question 4 · Why believe?

The test. Ask: "Why should this buyer believe it works?" Pass if the reader can point to proof from someone like them. Fail if the proof is generic, from another industry, another size of company, or another country.

The fix pattern: match the proof to the lead buyer. One case, one logo or one number from the buyer you chose in question 1. If you don't have it yet, say honestly what you do have: a pilot, a design partner, one measured result. Don't borrow proof from a different buyer and hope nobody looks. They will.

Question 5 · One first step?

The test. Ask: "What would you click first, and what happens next?" Pass if the reader names one step and knows what they get from it. Fail if they say "book a demo" with a shrug, or can't pick between four buttons.

The fix pattern: one primary step per buyer, sized to the moment. Choose one primary action for the lead buyer and make everything else visually secondary. Size it to someone who met you thirty seconds ago: a sample report, a short assessment, a calculator, a trial. If the step is a call, say what happens in it and what they leave with. If you run self-serve and sales side by side, split them by buyer ("For teams under 20" / "For organizations"), not by button.

Two conditions around the page

Two categories are about timing, not the screen. They make the five questions urgent.

You are entering a new market (8%). A pass at home does not carry over. Re-run the test with a reader from the new market. Questions 1 and 4 usually fail first.

You are hiring go-to-market people (18%). A new seller answers the five questions on every call. If the page has no answer, each hire invents one. Write the five answers down *before* their first day.

SECTION 04

Four rewrites from real observations

Four real rows from our dataset, names and identifying details withheld. The "after" versions are illustrative, written from the public page alone. Every bracket is a fact only the company can fill in.

Rewrite A · An HR compliance platform

What the observation said. The homepage shows nine product blocks, from an anonymous reporting hotline to leave management and performance plans, all leading to "Book a demo". The hotline is a familiar purchase. The new capability, AI-assisted investigations, is lost among them.

One-Read Test. Q1 partial (HR is clear, which HR person is not). Q2 fail (nine blocks). Q3 fail. Q5 fail (one heavy door).

Before (structure).

> [Platform tagline]
> Nine product tiles · Book a demo

After.

> **For HR leaders who run workplace investigations: close cases in [days] instead of [weeks], with a record that holds up.**
> One line of proof from a comparable HR team.
> Primary step: *See a sample investigation report.* Secondary: Book a demo.
> Below the fold: "Also included: reporting hotline, leave, performance plans."

What changed. One buyer, one problem, the new capability in front, the rest demoted to "also included", and a first step that needs no calendar invite.

Rewrite B · A financial-infrastructure platform

What the observation said. The homepage speaks to banks, fintechs, marketplaces and digital asset businesses, with three different first steps: "Try the platform", "Book a demo" and "Talk to us". A cold buyer may not know which one is meant for them.

One-Read Test. Q1 fail (four buyers). Q5 fail (three doors, no guidance).

Before (structure).

> [Platform tagline]
> Banks · Fintechs · Marketplaces · Digital assets
> Try the platform · Book a demo · Talk to us

After.

> **For [lead segment, e.g. fintechs launching in a new country]: [the outcome they buy, in one line].**
> Proof from one company in that segment.
> One primary step for that buyer. If it is a trial: *Start building in the sandbox.*
> A segment switcher below: "Bank? Marketplace? Digital assets? See your page." Each page carries its own primary step.

What changed. The three buttons did not disappear. They were assigned, one door per buyer.

Rewrite C · An inventory platform for contractors

What the observation said. The Starter, Plus and Enterprise plans show no prices and are quoted by number of stock locations, so a contractor has to book a demo just to learn what it costs.

One-Read Test. Q1 pass. Q5 fail: a call is needed just to learn the size of the purchase.

Before (structure).

> Starter · Plus · Enterprise
> Contact us for pricing · Book a demo

After.

> **Starter: for contractors with [1 to N] stock locations.** Who it fits, what is included, and a published starting point or a range.
> **Plus:** the same, for [N to M] locations.
> **Enterprise:** "More than [M] locations? We'll quote it on a 20-minute call."
> A small tool: *How many locations do you run?* It points to the right plan.

What changed. The buyer can place themselves before talking to anyone, and the call is kept for buyers who need one. Even without published figures, the location bands tell the buyer which size of purchase this is.

Rewrite D · A Belgian B2B software company hiring in France

What the observation said. The Series A was raised for new markets, and an account executive role for France is open. A message built in Belgium may not open a meeting in Paris without changes.

One-Read Test. Re-run with a French reader. Q1 and Q4 likely fail.

Before. The new account executive starts with the home-market deck and the English homepage, and writes their own French version in week one.

After: a one-page message brief, finished before the hire's first day.

> **Lead buyer in France:** [role, company type, size].
> **First problem:** [in their words, checked with three French prospects].
> **What changes:** [one line].
> **Proof:** [the closest French customer or most comparable result].
> **First step:** [one offer sized for a cold French buyer].
> **What stays the same as at home:** [the base message].
> **What changes for France:** [the variant].

What changed. The hire inherits a decision instead of a blank page, and the site, the deck and the seller say the same thing. That closes the gap 69% of buyers in Gartner's 2025 survey said they had run into.

SECTION 05

Objections

"We serve all those segments. Picking one means losing the others."

You are not deleting segments. You are choosing which one gets the first screen. The others keep their pages, proof and first steps, one click away. You lose the comfort of listing everyone. You get a page one buyer reads as written for them, instead of one every buyer reads as written for someone else.

"Our buyers are enterprise. They expect a demo."

Many do, eventually. The question is whether the demo has to be the *first* step. Gartner's 2026 survey found 67% of buyers prefer a rep-free experience. Even an enterprise buyer wants to know if you fit before giving up an hour. Keep the demo. Put a smaller step in front of it for the buyer who isn't ready yet, and say exactly what happens in the demo for the buyer who is.

"Our homepage converts fine. This is a cold-outbound problem."

Possibly. Warm visitors already know you. But if most of your pipeline comes from referrals and founder networks, your page has never been tested on a stranger. The moment you open a new market, hire sellers or start outbound, strangers are who reads it. Run the test on one. If they pass, you're done.

SECTION 06

The self-check

Show your first screen to someone who has never seen your company. Ten seconds. Close it. Ask:

FIGURE 4. THE ONE-READ TEST, ONE PAGE

#	Ask them	Pass	Fail	Your result
1	Who is this for?	One role, one kind of company	"Businesses", a list, or a pause	
2	What problem does it solve first?	One problem the buyer already names	The platform, or a list of modules	
3	What is different after they buy?	A concrete change: a number, a task gone	Your category or tagline repeated back	
4	Why would that buyer believe it?	Proof from someone like them	Generic, or from another buyer or country	
5	What would you click first, and what happens next?	One step, and they know what they get	A shrug at "book a demo", or four buttons	

Then answer two questions yourself:

- Are we entering a new market in the next six months? If yes, re-run the test with a reader from that market.
- Are we hiring anyone who will repeat this message? If yes, write the five answers down before their first day.

Scoring. Five passes: your first screen is doing its job. Three or four: you know exactly which line to rewrite. Two or fewer: the problem is not the copy. It is a decision nobody has made yet.

Next step

If you ran the test and got two or fewer passes, the fix is not a redesign. It is a decision about one buyer, one problem and one first step, written down so the site, the deck and every new hire say the same thing.

That is the work we do. If you want a second pair of eyes, we run a **free 20-minute diagnostic**. No pitch, no commitment. We run the One-Read Test on your first screen with you, and you leave with your five answers scored and the one line we would change first, whether or not we ever work together.

Write to fran@vega-studio.com.

ABOUT THIS PAPER

Vega Studio works on go-to-market for B2B companies: the positioning, the messaging and the execution that make marketing and sales compound. External figures come from primary sources only, dated next to the number. Company identities are withheld. AI did the reading and the labelling at scale; people chose the questions, set the rules for what gets published and checked the result. The classification script, rules, labelled samples and agreement table are available on request.

Sources

- Gartner, *"Gartner Sales Survey Finds 67% of B2B Buyers Prefer a Rep-Free Experience"*, press release, March 9, 2026. Survey of 646 B2B buyers, August to September 2025. Source of the 67%, the 45% AI use and the "twice as likely to report a high-quality deal" finding.
- Gartner, *"Gartner Sales Survey Finds 61% of B2B Buyers Prefer a Rep-Free Buying Experience"*, press release, June 25, 2025. Survey of 632 B2B buyers, August to September 2024. Source of the 69% website-versus-seller inconsistency figure.
- 6sense, *2025 Buyer Experience Report*, newsroom release, November 2025. More than 4,000 buyers in North America, EMEA and APAC. Source of the 94% and 77%.
- Vega Studio research, 865 B2B companies, October 2026. Source of all category percentages. Method, validation and limits in section 2.

This is paper 03 of the Vega Studio Research series. Paper 02, The Messaging Capital Leak, is about the targeting decision behind the page: who your acquisition money is pointed at.

VEGA
STUDIO

RESEARCH · PAPER 03 · NOVEMBER 2026

VEGA-STUDIO.COM