

FROM PROMPT TO PURCHASE

*How B2B
decision-
makers buy in
the LLM era*

By Omniscient Digital



Omniscient

Wynter

From Prompt to Purchase: How B2B decision-makers buy in the LLM era

LLMs are now part of the buying journey, but not in the way vendors feared. AI is speeding up discovery and helping buyers frame their options. But when professional reputation, executive scrutiny, and real money are at stake, buyers still turn to humans.

The buying journey hasn't gotten simpler. It's gotten more layered.

For marketers, those layers widen attribution gaps. LLM-influenced demand often gets misclassified as "direct", creating blind spots. Many companies struggle to connect growth to revenue because of messy data, fragmented tool stacks, or weak measurement practices. Now, with ChatGPT and AI baked into search, those blind spots grow at a rate that outpaces the adoption rate of AI visibility tools.

This leaves vendors in the dark: unsure how they appear in AI answers, unclear on sentiment around LLM citations, and disconnected from the signals that turn discovery into leads.

That shift raises new questions. If AI is shaping discovery but not showing up in commonly used reporting tools, how can vendors understand where they appear, how they're represented, or what factors move buyers from exploration to action?

This report unpacks how that layered process actually works. We'll show you where buyers really start their journey, how they move between channels, what proof they rely on, and what tips the final purchase decision.

Executive Summary

- The buying journey is now a loop, not a line. Buyers move fluidly between Google, LLMs like ChatGPT, review sites, peer communities, and vendor websites. Each phase serves a specific purpose: AI for speed, humans for trust.
- AI drives discovery – but not decisions. 37% of buyers stop using LLMs after early research. Most verify AI-sourced vendors through websites, reviews, and peer input before taking the next step.
- Trust has become the real funnel. Every purchase now passes through layers of validation. LLMs surface new vendors, but peers, documentation, and transparent communication decide who makes the shortlist.
- Peer influence outweighs brand control. Buyers trust peer recommendations (85%) and third-party reviews (78%) far more than any vendor-owned content.
- Transparency beats perfection. Vendors win when they clearly articulate who they serve, how they fit, and where their limits are. Buyers reward clarity and honesty over polished, generic marketing.

Methodology

This research was conducted via Wynter's platform in September 2025. We surveyed 100 B2B SaaS leaders of marketing, technical, product, operations, or sales teams who influenced or owned a vendor purchase decision in the past 6–12 months.

The analysis was led by Omniscient Digital. We reviewed all open-ended responses, quantified recurring themes into percentages, and created visualizations to highlight key trends. From these insights, we outlined the narrative that connects the data to strategic implications for decision-makers.

Sample composition:

Company roles: 40% marketing leaders (Directors of Marketing, Demand Gen, Growth, Digital, Brand, or Marketing Ops), 25% technical leaders (Directors of IT, Engineering, Cloud, Software Development, Data Science, or DevOps), 15% product leaders (Directors of Product, Product Marketing, Product Design, or Product Platform), 15% operations leaders (Directors of Operations, Revenue Ops, Strategic Ops, People, HR, or Finance), and 5% customer and sales leaders (Directors of Customer Success, Engagement, or Sales).

Company sizes: 22% at companies with 11–50 employees, 22% at 1,001–5,000 employees, with representation across small businesses, mid-market, and enterprise organizations

Industries: 56% SaaS/software, 20% SaaS + IT Services & IT Consulting, 8% SaaS + Financial Services/Fintech, with additional representation across computer security, marketing, manufacturing, insurance, and professional services

Purchase types: Diverse range including bookkeeping software, security tools, partner attribution platforms, marketing technology, collaboration software, and infrastructure services

Question framework: 10 questions covering the complete buyer journey from initial discovery through final decision, combining quantitative rankings with open-ended qualitative responses to capture both behavior patterns and decision-making rationale.

Sample composition:

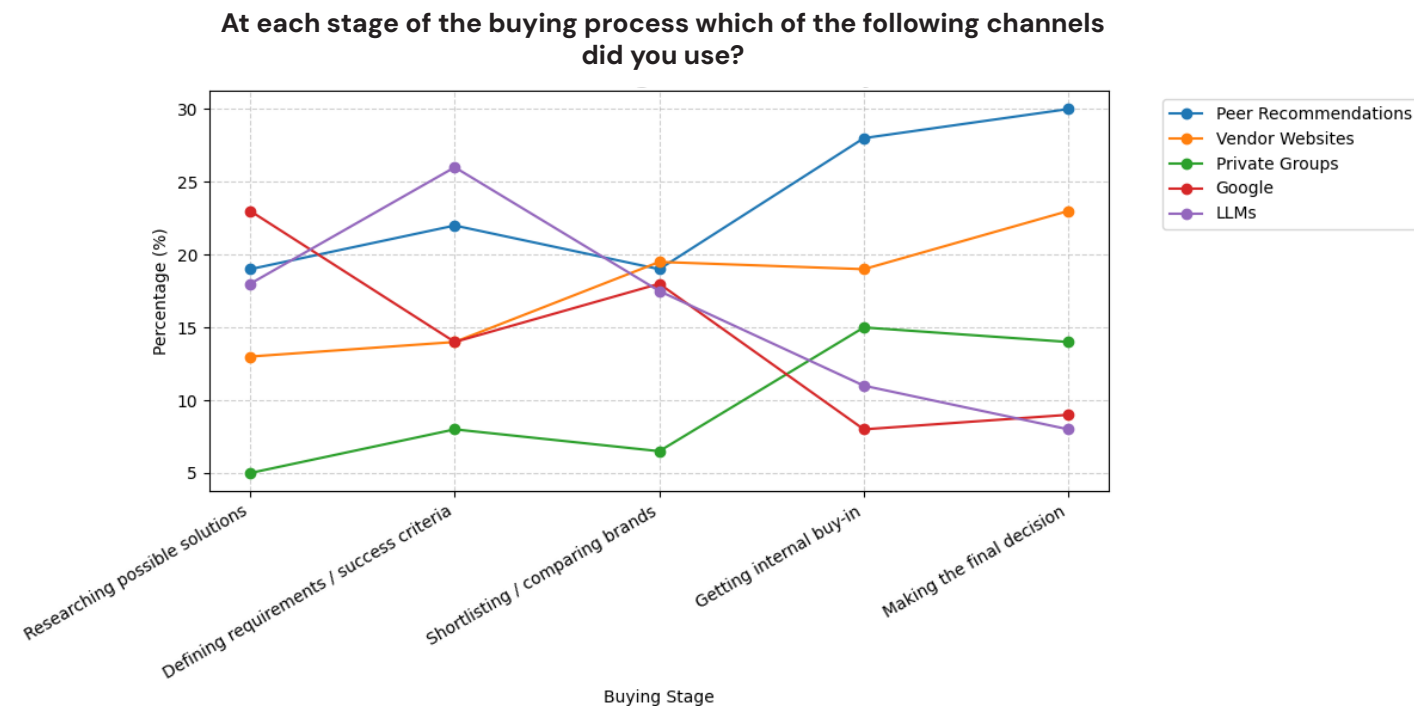
10 questions covering the complete buyer journey from initial discovery through final decision, combining quantitative rankings with open-ended qualitative responses to capture both behavior patterns and decision-making rationale.



Table of Contents

The New Validation Loop	06
The Three Entry Points	
Goals of a Buyer's Initial Search	
The Channel Choreography	
<hr/>	
The Trust Problem	10
When AI Suggests an Unknown Vendor	
What LLMs Do Well (and Where They Fall Apart)	
The Handoff Moment	
<hr/>	
What Actually Closes Deals	16
The Proof Hierarchy	
The Final Decision	
Building Content LLMs Can Actually Use	
Owning Your Own Narrative in LLM Citations	
<hr/>	
What Happens Next	22

The New Validation Loop



When a B2B buyer starts looking for a solution, the first move is obvious: Google it. See what's out there. Compare features, costs, integrations. That part hasn't changed.

What's different is what happens next. Today's buyer bounces from Google to ChatGPT for quick comparisons, then to review sites, into peer communities for real-world input, and back to vendor websites to verify the details.

The journey is no longer a straight line. It's a loop, a constant cycle of discovery, cross-checking, and validation.

That means the marketer's playbook for "meeting buyers where they're at" gets a lot more complex. Showing up in Google is no longer enough, and even showing up in ChatGPT citations may not be enough. What matters is to show up with credibility in every place the buyers look: from AI answers to peer communities to your owned site.

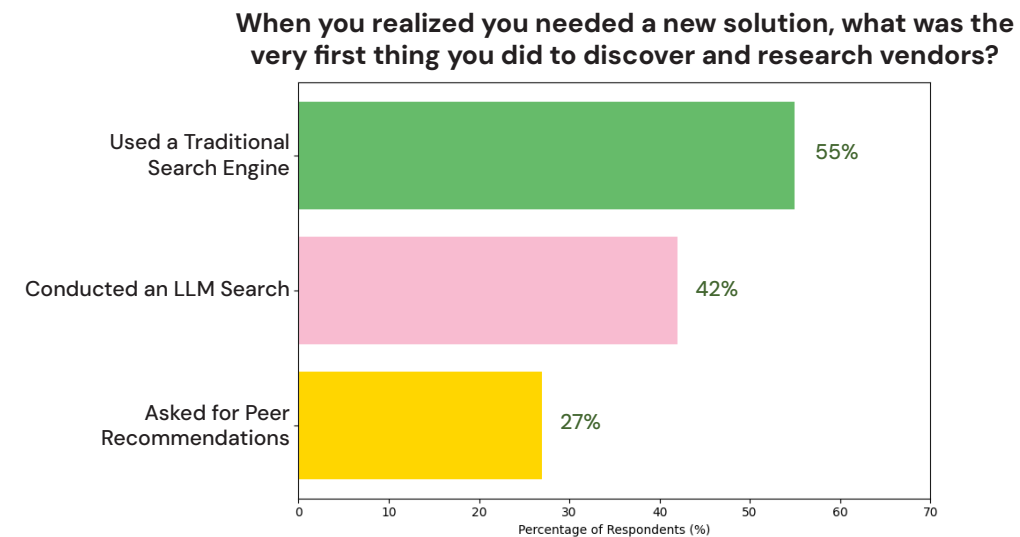
Some shifts that stand out:

- Real proof. Buyers want examples they can relate to, to see the product in action, and to trust your brand communicates the truth, not a sugar-coated version.
- Off-site influence. More of the content shaping perceptions is surfacing in spaces brands don't own, from peer threads to public communities.
- Trust as the closer. Traffic may start the loop, but trust is what closes the deal – in the product, reputation, and value delivered.

Each channel carries weight individually, but the bigger impact comes from how they reinforce one another as part of the whole.

The three entry points

When we asked 100 B2B decision-makers to describe the very first thing they did when they realized they needed a new solution, three patterns emerged:



Traditional Search (55%): The majority started with traditional search, and 46% specifically mentioned Google. It's still the default first step, which is no surprise after 20 years of ingrained internet habits.

"I started with a basic web search on the technology I was researching – I then pivoted from that search to looking at IT vendors that I recognize, either since they are current vendors, or vendors I've evaluated or have met with (at conferences, for example). From that search, I usually have a short list of potential solution providers." (Director of IT, 51-200 employees, SaaS/Software)

LLM Search (42%): Out of this group, 25% specifically mentioned ChatGPT. But only 11 started with an LLM as their absolute first step. Most followed the pattern a Director of Operations did: Google to orient, then LLM to structure and compare.

"I searched on ChatGPT and asked the model to identify the best in class tools for that category that are cutting edge and growing quickly." (Sr. Director of Demand Generation, 1001-5000 employees, SaaS/Software)

Peer Recommendations (27%): They speak to people who've solved similar problems and consult their network who've used vendors in the category. As one buyer put it: "I trust these connections more for unbiased feedback".

"Word of mouth, I asked other peers in the same field/different company which solutions they were using and what was working best/not working at all and took note of that and then made a list of possible contenders." (Sr. Sales Director, 501-2000 employees, SaaS/Software)

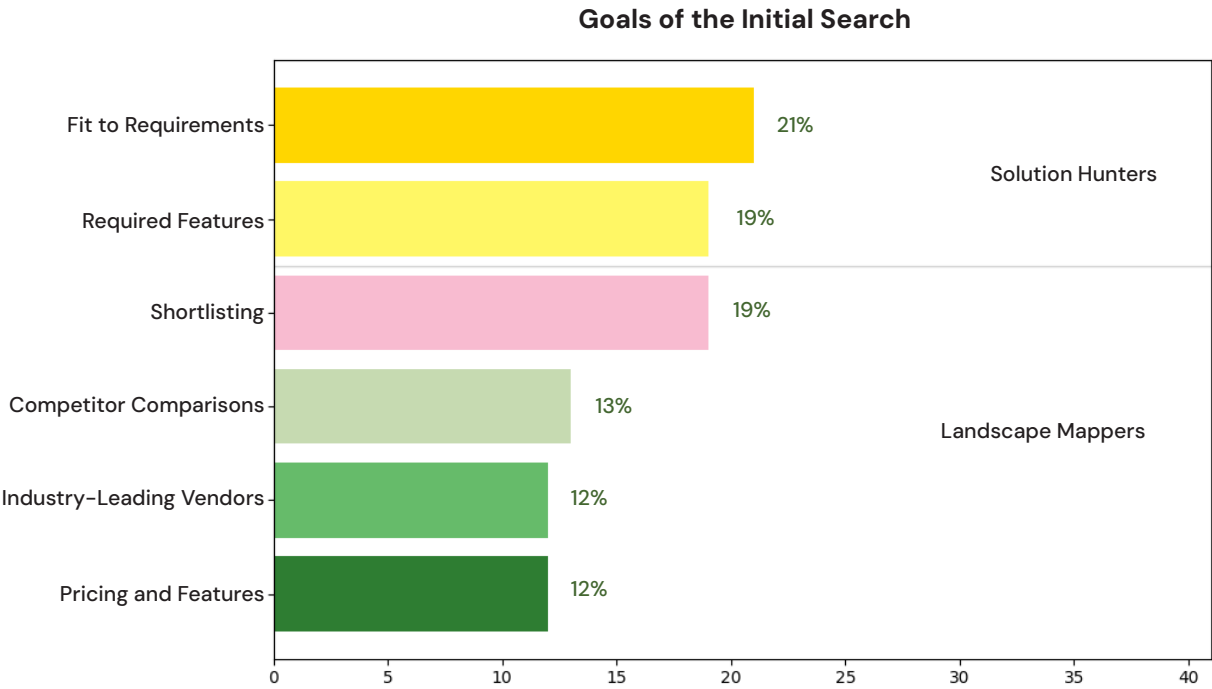
A further 21% started by looking inward, asking colleagues and leadership to define business needs, clarifying the specific problem requiring support, and identifying gaps in current tools before touching any external channels.

"When we realized we needed a new solution, the very first step we took was to conduct thorough discovery and user research. This involved engaging directly with stakeholders and end users to understand pain points, gather requirements, and identify gaps in the current process." (Director of Technical Business Operations, 501-1000 employees, SaaS/Software)

The three entry points

In those first vendor searches, whether in Google, ChatGPT, or conversations with peers and colleagues, buyers are trying to answer a deceptively simple question: What’s out there?

But what are they actually looking for in those searches? The answer splits buyers into two distinct camps, each with a fundamentally different approach to vendor discovery.



Solution hunters

They already know what they need. 21% lead with “fit to their requirements.” 19% start by defining exact features required, what’s missing in current tools, what specific solutions would solve their problem. They’re qualifying vendors against a checklist.

“Write down the list of problems we need to solve. Spend time drafting it, getting feedback and asking others to add their needs. Then come up with a set of priority principles—must have, can have, should have.” (Product Platform, 1001–5000, SaaS/Software)

Landscape mappers

This group is orienting in unfamiliar territory. 19% want shortlists, 13% search for competitor comparisons, and 12% need to understand features and pricing upfront. They are looking for “top vendors”, “best in class,” “cutting edge,” and vendors “growing quickly”.

“I searched on ChatGPT and asked the model to identify the best in class tools for that category that are cutting edge and growing quickly.” (Sr. Director of Demand Generation, 1001–5000 employees, SaaS/Software)

How their needs differ

The split between these two approaches maps directly to experience. Buyers who’ve done this before operate differently from the start.

“I started with a basic web search on the technology—I then pivoted to looking at IT vendors that I recognize, either since they are current vendors, or vendors I’ve evaluated or have met with at conferences. From that search, I usually have a short list of potential solution providers.” (Director of IT, 51–200 employees, SaaS/Software)

The key here is that potential buyers filter through existing vendor relationships and past evaluations, they validate what remains accurate. First-time buyers in a category don’t have that framework yet. They need to map the landscape before they can hunt for solutions.

This creates different needs from search tools. Landscape mappers need LLMs to compare systematically, surface options beyond top Google results, explain how vendors differ. They’re using AI to build the understanding solution hunters already have. Solution hunters need precision, exact functionality matches, integration capabilities, security

certifications, proof this vendor has solved this problem for companies like theirs. This is why 5% buyers looked for customer reviews in step one, 5% wanted demos immediately, 3% needed security information upfront.

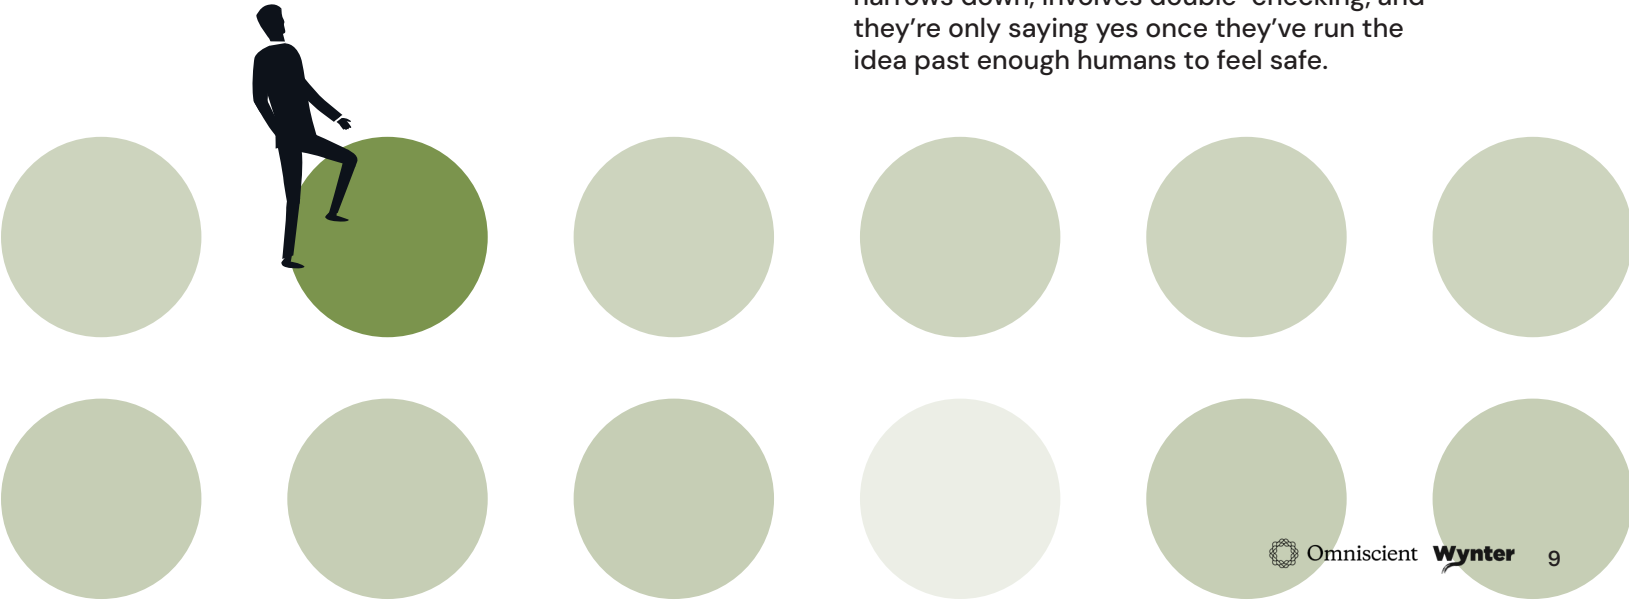
But most buyers toggle between both modes during the same search. They start in landscape mode (“show me what’s available”), shift to solution hunter mode once they understand options (“which ones integrate with Salesforce?”), then loop back to landscape mode when they discover subcategories they didn’t know existed (“wait, there’s a difference between partner attribution and partner relationship management?”).

This toggling is what LLMs solve in early stages, helping buyers move fluidly between “what exists” and “what fits” faster than clicking through twenty vendor websites.


But it’s also why LLMs fail when buyers get specific. AI struggles when a buyer needs to know whether a specific integration works with their exact tech stack, or whether “enterprise security” actually meets their compliance requirements. That’s the context LLMs don’t have, and why they can’t replace procurement experience or the vendor conversations that happen in real time.

The channel choreography

The buying journey no longer follows a linear path. Buyers move fluidly between channels based on what they need at each moment. In the initial discovery phase, Google, LLMs, and peer recommendations play the biggest role. When comparing brands, vendor websites become critical for validation. In the final decision stage, peers matter most again; the trust loop closes with humans.



Here’s the most common full sequence we observed:



The majority start broad (Google), structure it (LLMs), validate with humans (peers), verify details (brand sites), then return to humans for the final call (peers again).

The validation loop expands outwards, (external sources) then contracts inward (internal stakeholders) the closer they get to making a decision.


What’s happening here is a progressive narrowing with validation at each gate. LLMs help shortlist. Peers help eliminate obvious misfits and negative experiences. Brand websites answer specific questions about fit. Then peers, colleagues and their leadership team weigh in again on the finalists.

Every B2B purchase carries professional risk, that you’ll recommend something that doesn’t work, that costs more than expected, that creates more problems than it solves. No one wants to be the person who championed the vendor that failed.

We see here a funnel of trust. It starts wide, narrows down, involves double-checking, and they’re only saying yes once they’ve run the idea past enough humans to feel safe.

The Trust Problem

When AI suggests an unknown vendor



Vendor Website (51%): When an LLM surfaces a vendor the buyer didn't know, the majority of respondents said they go straight to the vendor website.

They're looking for:

- Does this product actually do what ChatGPT said it does?
- Is the pricing anywhere close to what the LLM suggested?
- Do they have the security certifications we need?
- Who are their customers? Anyone like us?

Review Sites (25%): Some respondents listed the site by name: G2 (8%), Capterra (2%), TrustPilot (1%). They're looking for reviews that reveal what's the actual experience with this product.

Direct Vendor Contact (14%): They reach out directly to the vendor for demos or specific pricing conversations.

Peer Consultation (13%): The group that asks peers about unknown companies is smaller here, noting that they consult peers more commonly for known vendors.

"New vendors—checked their website for high level detail as well as case studies to confirm the vendor actually aligns with the solution we are looking for. Already known vendors—focus more on cost detail (if available), actual visuals

of the product. I find for B2B that these details are generally very lacking." (Director of Product Management, 11-50 employees, SaaS/Software)

Notice what's missing for known brands: the basic verification step, as potential buyers tend to already trust they do what they say. For unknown vendors, that trust doesn't exist yet. The case studies help here prove they're a real company that has solved real problems for customers.

In regulated industries, we see how verification gets more rigorous:

"I went to the solution provider's website and reviewed the feature specifications for their offering in the space. At this phase, I am looking to see if a potential solution will meet our feature and security requirements. Cost is also an early criteria, so if a pricing model is described, I would review it... Also, I like to see the documentation (if any) for the offering—thorough, clear documentation is a sign that a solution is well-implemented for usability and supportability." (Director of IT, 51-200 employees, SaaS/Software)

The word "sign" matters here. Clear documentation is a signal of transparency, it shows how open a company is about how things work. When a company communicates thoroughly online, it signals honesty, confidence, and respect for the customer's time.



What LLMs do well (and where they fall apart)

We asked buyers to describe the quality of LLM answers they received about vendors. The results split into two camps: impressed and frustrated.



What worked:

- 63% said LLMs compared brands clearly and in detail, what they do better than any other channel
- 44% appreciated that LLMs referenced real companies or customers
- 39% found the output matched what they already knew (a validation signal, not a criticism)
- 32% got specific examples or case studies

What didn't:

- 24% found responses too generic or vague
- Multiple buyers noted LLMs couldn't grasp their specific use case
- Pricing was consistently wrong or outdated
- Newer vendors often weren't in the training data

But here's the important findings: 80% validated LLM outputs against brand websites. 60% cross-referenced with third-party reviews. 55% checked with peers. Only 22% relied directly on the output.

This suggests buyers have learned that AI is great for structuring options but unreliable for final decisions.

Some don't trust LLMs to search on their behalf, instead using them to synthesize information that's already been validated.

"I only use LLMs to summarize details I already know. I know they've increased in resolution and capability, but the results are only as good as the source material. I'd prefer to select my material, add it to Google Drive and use Gemini to summarize what I've collected, for example. Or use a public summarization to identify terms or concepts that I didn't include in my original vetting criteria." (Director of Operations, 11-50 employees, SaaS/Software)

Despite LLMs ability to speed up the decision making process, some still don't trust LLM comparisons at all:

"I don't trust the technical specification comparisons and I don't trust LLMs to provide me with an actual objective and comprehensive list of vendors. LLMs are good for cross-checking, brainstorming and finding possible fits for our unique requirements." (Director Paid Digital, 1001-5000 employees, SaaS/Software)

This specific buyer assumes LLM outputs are biased, either by the training data, by the vendor's SEO strategy, or by the companies that have the resources to flood the internet with content.


Others agree:

"It's not accurate as it is based off of SEO. Some really great companies cannot afford SEO." (Sr. Director of People and Culture, 201-500 employees, SaaS/Software)

This is the paradox vendors face: you need to show up in LLM outputs to be discovered, but buyers know visibility doesn't equal credibility.

Earning citations matters, but it's only part of the picture. Visibility gets you into the conversation and it's what makes you one of the vendors buyers compare. What actually moves buyers is when the LLMs' evidence they find matches the promise you made.

Buyers want customer stories that feel real, documentation that's clear, and integrations they can trust. That comes from consistency, transparency, and proof that holds up under scrutiny.



What LLMs do well (and where they fall apart)

It's clear from our findings that LLMs help buyers get moving, but they don't close the deal. Interestingly, 37% of buyers say they stop using AI after the early research phase.

It's the classic "handoff moment": Buyers lean on AI for speed, fast comparisons, quick sense-making, shortcuts to clarity. But as soon as the stakes rise, they turn back to human-validated sources.

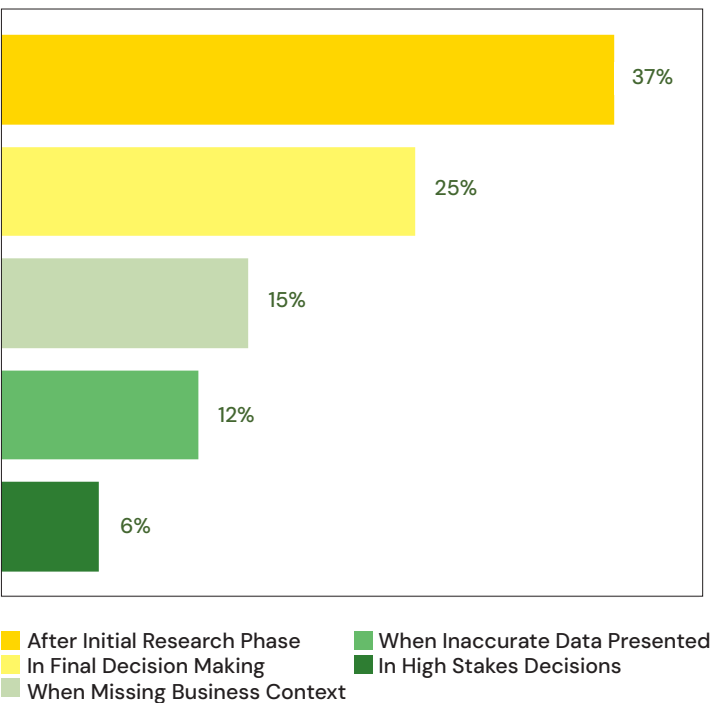
Why? When reputations, careers, and company money are on the line, no one is going to tell their executive team, "I trusted ChatGPT."

So, it's less about the tech, and more about professional risk. 25% specifically said they would never let an LLM make their final decision.

"Let's be real. I'm not going to put my professional reputation on the line for an LLM recommendation for a corporate purchase... by the time I put real money on the line I will leave nothing to chance. An LLM can help me to get informed in the broad sense about a tooling category, but it doesn't come near me when making a purchase decision." (Director of Information Security, 1001-5000 employees, SaaS/Software)

When we looked deeper into their vendor search behavior, their reasons for stopping LLM use fall into clear categories:

When do you avoid using LLMs in purchase decisions?



Missing business context (15%): Users find LLMs lose track of the discussion, not fully understand the specific use case, and can't grasp the nuances of organizations and their constraints.

"I don't feel it provides the realistic data I need, it's just the kind of research they do and combine the data to show an output, you need to keep asking questions to get the details and at some point, you lose the track of discussion." (Director of Cloud Infrastructure/Security Expert, 51-200 employees, SaaS/Software)

"Once it gets to the point of my specific company needs and overall environment or any other issue about company specifics that a LLM would not know." (Sr. Director of Strategic Finance, 1001-5000 employees, SaaS/Software)

Inaccurate data (12%): When LLMs hallucinate features, get capabilities wrong, or churn out outdated information, buyers stop using them.

"I don't typically trust LLMs to have the most current/up-to-date information, especially when it comes to pricing. If I'm looking at a newer brand, I'll trust the LLM less." (Director of Marketing, 51-200 employees, SaaS/Software)

"I generally avoid LLMs in any situation in which hallucination is a concern—since I'm looking to review specific product offerings and features, the risk of a mythical product is very real. I have no interest in reading a fictional product description." (Director of IT, 51-200 employees, SaaS/Software)

High stakes (6%): When the decision is too important, involves too many stakeholders, requires human judgment, this is when they'll steer clear of LLMs.

"I avoid using LLMs when the decision is mission critical, or the decision must rely on recent data." (Director of Marketing, 51-200 employees, SaaS/Software)

"Never really, unless I know the products so well that I do not need additional information. I find LLMs can always be useful, but depending on the "weight" of the purchase, there needs to be more double or triple checking of the information with brand websites or other 3rd parties." (Director of Strategic Operations, 51-200 employees, SaaS/Software)

But 21% of buyers said they would "not avoid using LLMs in purchase decisions". They've just learned how to double-check the answers and build that into how they work.

"I always use them for research, but prompt specifically asking the LLM to browse Reddit, G2, social media, and pretend that the LLM is 'a buyer for a B2B company of X' (put the LLM in my shoes) evaluating for X,Y,Z criteria. I never avoid using an LLM." (Sr. Director of Demand Generation, 1001-5000 employees, SaaS/Software)

"I don't avoid LLMs at all. They are a very helpful information source. They don't hallucinate much anymore." (Sr. Director of Product Management, 11-50 employees, SaaS/Software)



The Takeaway

LLMs help buyers get moving, but they don't close the deal. A third of buyers stop using AI after the first research phase. They turn back to websites, reviews, peers, and direct conversations once the stakes rise.


The higher the risk, the less buyers are willing to trust an algorithm. Professional credibility, budgets, and careers are on the line. No one wants to explain to their executive team that a major purchase was made on ChatGPT's word.

This creates a clear division of labor. LLMs are accelerators: they surface options, compare features, and help buyers frame the problem. But the closer buyers get to a decision, the more they demand human validation.

What Actually Closes Deals

The proof hierarchy

When we asked buyers which types of proof or evidence make them trust vendors most, the rankings revealed something uncomfortable for marketing teams:



Buyers trust peer recommendations (85%) and third-party reviews (78%) above all else, yet those are the areas you can influence the least. Interestingly, most of the lower-impact trust signals are the ones directly within your control.

Buyers know that everything on your website, every testimonial, case study, and stats are curated. But still, these website features are influential for building buyer trust:

- 65% say case studies build brand trust
- 43% like to see customer logos or testimonials
- 39% turn to analyst reports
- 32% browse through through leadership and company content

- 16% check vendor ROI calculators
- 15% seek out the founder's or executive leadership content

They're already aware your website shows your best customers, your biggest wins, your cleanest implementations.

"Customer references and 3rd party reviews. All other material is biased or influenced by the brand. Even though the customers we'll have access to aren't unhappy customers, they provide a frame of reference to see if they had similar challenges and needs with the software." (Director of Product, 1001-5000 employees, SaaS/Software)

Some layer validation sources and aren't checking marketing materials at all:

"I utilized my professional network and IANS subscription to learn more about the vendor, its offerings, and reviews. My professional network was most helpful in providing real life feedback on the procurement, configuration, and use of the tool as their feedback lacked a marketing tone. After reviewing this content I usually dug into a vendor's support documentation... The level of detail in customer documentation helps me gauge the level of support I may expect." (Director Information Security, 1001-5000 employees, SaaS/Software)

This creates a challenge for vendors: the content that matters most for trust is the content over which you have the least direct control. You can't write your own G2 reviews. You can't control what peers say in Slack channels. You can't dictate what analysts publish.



Your Move:


Make it easy for buyers to find independent validation and give them and their team proof points.

Buyers put the most trust in peers and third parties. Your marketing isn't their first stop. It sets the stage, but real proof comes when they hear from people they see as credible.

That's why peer voices need to be easy to find. Reviews, communities, reference calls, they're the foundation of trust in a market where skepticism comes standard.

The final decision

When it comes to the final deciding factor, buyers aren't looking for the cheapest option or the flashiest pitch. They're weighing what they get against what it costs. Sometimes that means price. Sometimes it's how well the product actually solves the problem. And often, it comes down to how honest and clear the vendor is in the process. Here's how the core themes netted out:



Price (37%): The majority of respondents said price, but when we read their explanations, 23% immediately clarified they meant "value for money." It's not always the lowest price – it's the best tradeoff between cost and capability.

"The single factor that decides my final vendor pick is overall value for money because it balances cost with features, reliability, and support." (Company Owner, 201-500 employees, SaaS/Software)

"I could not narrow it down to a single factor. Price is always a factor. Ability to fulfill requirements of my needs was the major factor, I would rather pay a higher price knowing it has everything I would need." (Director of Operations, 11-50 employees, SaaS/Software)

"Unfortunately, for a small business, it's mostly price (but only if the product actually fulfills the requirements), the government calls this 'LPTA'—lowest price technically acceptable." (Director of Operations, 51-100 employees, SaaS/Software)

Fit to Requirements (25%): The biggest deciding factor is how well the product solves their actual problem.

"There is no single factor, it depends on the need. If it is an urgent need, then the product that best addresses it usually wins (regardless of price). Otherwise, it is a compromise of price and offering." (Director of Strategic Operations, 51-200 employees, SaaS/Software)

Vendor Engagement (18%): They referenced the quality of conversations during the buying process, transparency, how the relationship feels, customer support quality.

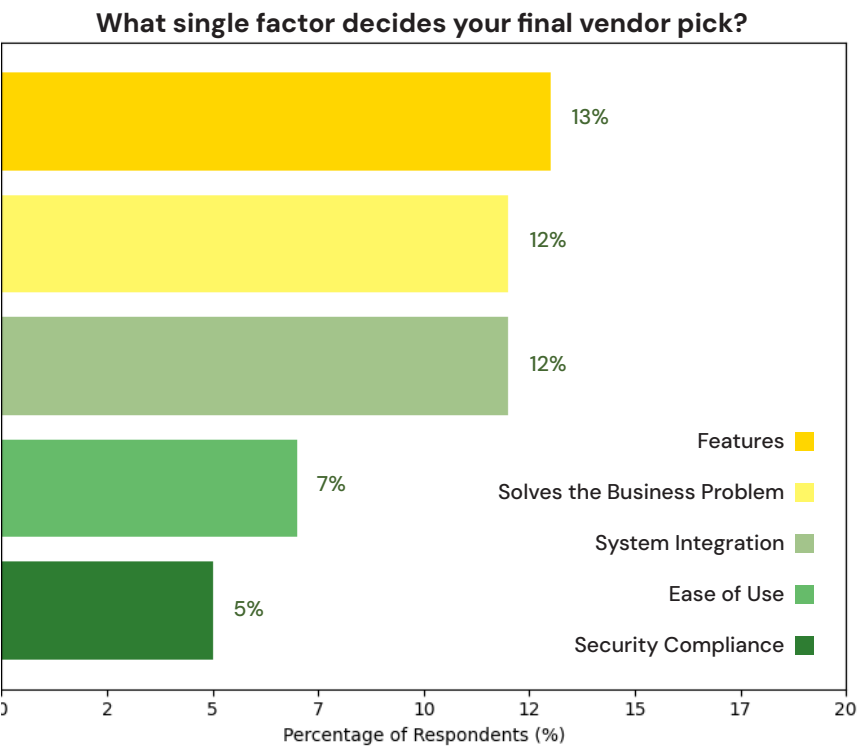
"Self awareness. I deal with so many vendors who don't know their sweet spot or 'lane' and it is refreshing when I deal with one who can confidently say 'that isn't our problem space.' The best vendors I deal with are equally adept at saying what they excel at as they are adept at saying what isn't their strong suit. If a vendor can't tell me where they don't belong or in which vertical they don't fit, this is a vendor that hasn't really understood their core value in my view." (Director of Information Security, 1001-5000 employees, SaaS/Software)

Buyers want vendors who understand their limitations. In this instance, self-awareness signals that a vendor understands their customers deeply enough to know who they serve well and who they don't.

This is where transparency comes in to play:

"I only knew about our existing vendor, as this is an area that I am not very familiar with. My first step was to check if their solution met our needs... My second step was to review pricing, as we were not looking for an enterprise-level solution and the price was an important factor. In fact, I did engage a salesperson from a leading vendor, but their solution ultimately ended up being way out of our budget, which was a waste of both of our time." (Director of Ecommerce Product Management, 501-1000 employees, SaaS/Software)

It's clear withholding information about fit, pricing, or limitations doesn't protect your position. It disqualifies you. Buyers would rather hear "we're not right for you" in the first conversation than discover it three calls later.



When we dug into the data about functional factors, the decision-makers pointed to:

- 13% want to see the product's features
- 12% say it must solve their current problem
- 12% need the solution to integrate with their current toolstack
- 7% prioritize ease of use
- 5% need the vendor to meet security and compliance requirements

Within these functional needs, a clear pattern emerges. Buyers are comparing products and weighing tradeoffs. They want the option that feels worth it for their use case, the tool with the right balance of cost and capability.

"I typically need software that is specifically scoped to my use case. I avoid software that does too much ('all-in-one'). Price is typically NOT my deciding factor." (Director of Marketing, 51-200 employees, SaaS/Software)

3% of buyers mentioned this explicitly: they don't want vendors who claim to do everything. They want vendors who do specific things exceptionally well.

"The vendors who claim to do everything and or stick AI into everything that needs it or doesn't are the ones we either reject or do additional due diligence on. How are the vendors with the good initial reputation and lack of high pressure sales tactics are the ones who are seriously looking at." (Director of Software Engineering/ Platform Engineering, 1001-5000 employees, SaaS/Software)

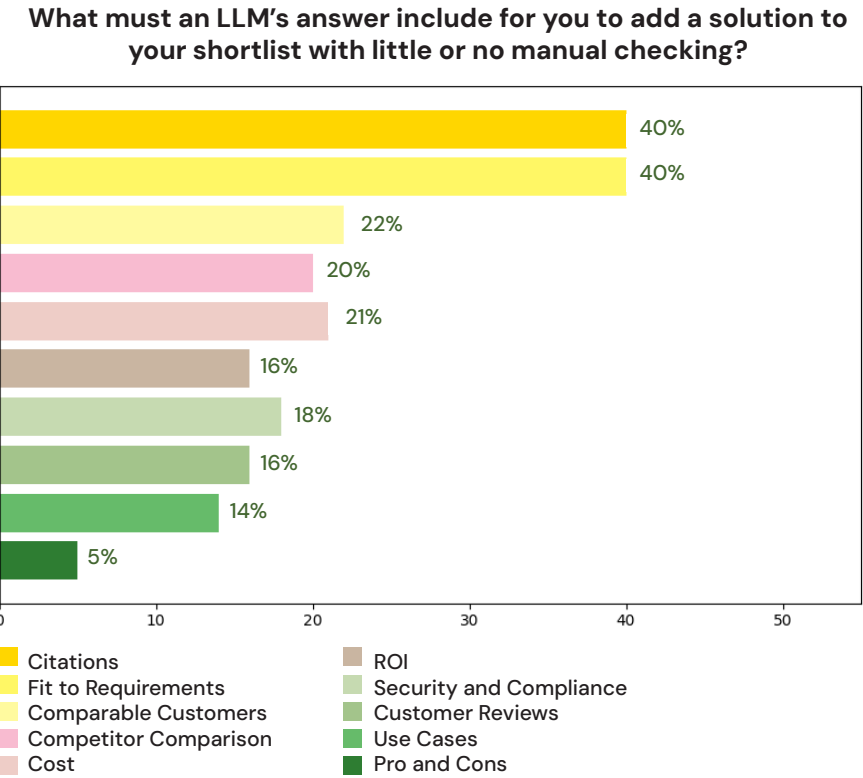
The Takeaway

Potential buyers are seeking out vendors who strike the right balance: a product that clearly fits their needs, priced fairly for the value it delivers, and backed by a vendor who is honest about their strengths and limits.

What tips the scale isn't a single factor but the confidence buyers gain that they're making the right tradeoff. That confidence comes less from polished marketing and more that the vendor understands both the problem and where they fit in solving it.

Building content LLMs can actually use

Buyers don't trust LLMs at face value. To take an answer seriously, they need certain details spelled out. We tested what those details are, the elements that would actually earn a vendor a spot on their shortlist.



Here's what came out on top:

Citations (40%): This is the foundation. Buyers want clickable links to reputable sources to dig deeper on their own.

"I generally don't trust the LLM without fact-checking, but if I needed to, I'd require citations as the #1. I don't think an LLM would be truly capable of assuring fit to requirements, because clients often have a hard time articulating what the requirements are—it's often an 'I'll know it when I see it.'" (Director of Operations, 11-50 employees, SaaS/Software)

"I have not had a situation where an LLM has given me a short-list and I have not done any additional manual research to validate it. I'd imagine at the very least it would need to cite a highly trusted analyst report (like a Gartner Magic Quadrant or Forrester Wave) that is written by experts." (Director of Product Marketing, 5001-10000 employees, SaaS/Software)

Fit to requirements (40%): The LLM needs to explicitly spell out how the vendor meets their specific needs. Generic capabilities aren't enough, potential buyers are seeking detail.

"Fit to requirements is #1. I also need to see drawbacks/cons for each option: What are its limitations? LLMs are uniquely suited to find this because brand websites won't always acknowledge it." (Director of Marketing, 51-200 employees, SaaS/Software)

"I need clarity on why out of 4 email platforms this one is the best for contact management with tacit proof points of features or capabilities that the others lack and what the tradeoff is. Good LLM outputs query back for clarity on what I need to help narrow the decisions down vs. making an assumptive guess to what the best is for my need." (Director of Marketing, 501-1000 employees, SaaS/Software)

Beyond citations and fit, buyers also flagged a set of details they look for:

Comparable customers (22%): Who else uses this? Companies like mine?

- Competitor comparison (20%): How does this stack up against alternatives?
- Cost and ROI (21% for cost, 16% for ROI): Does the cost and ROI work for us?
- Security and compliance (18%): Do the certifications, risk assessments, and compliance requirements match out needs?
- Customer reviews (16%): Is the feedback real? Are the sources controlled by the vendor?
- Use cases (14%): How has the product been used successfully?
- Pros and cons (5%): What are the tradeoffs?

Owning your own narrative in LLM citations

When tailoring your content for LLM search discovery, make sure you have the technical elements to earn citations, that your content is readable crawlers to pull out information like comparable customers, ROI, customer reviews, and real use cases. The more content your brand can give to the LLM, the more information you give the buyer to put you on their shortlist. If you are not creating this story, LLMs will do it for you, and could get it wrong (or you risk competitors controlling the narrative).


And yet, 16% of buyers said they'll never fully trust an LLM.

"There isn't anything that an LLM's answer could include that would have me add it to the shortlist with little or no manual checking. I simply don't trust LLMs to have accurate information all the time and don't want to risk my reputation on the answer given by an LLM." (Director of Marketing, 201-500 employees, SaaS/Software)

"I don't know that I would blindly trust an LLM alone, but I suppose if it concludes that a solution checks all of my boxes, including cost, security, purpose, ease of use and adjustability, I would add the solution to my list. However, if a peer used it and gave a negative review/experience—it would be off the list instantly." (Sr. Sales Director, 501-1000 employees, SaaS/Software)

What matters is whether the output gives buyers the confidence to move forward, and that confidence hinges on transparency. Citations, clear fit to requirements, and honest tradeoffs make the difference between content that earns a shortlist spot and content that gets ignored.

An LLM can put you on the list, yet a single peer review can kick you off it instantly. That means marketing must continue to go beyond giving LLMs the right signals, and focus equally on backing those signals with proof in the real world.



What Happens Next

LLMs have made buyer discovery faster, but also broader and more complex.

It's not as simple as some in the industry say, i.e. that LLMs are replacing Google or other sources. It is, however, very clear that LLMs are shaking up and rearranging the customer journey.

LLMs are helping buyers find vendors they never would have encountered through traditional search or word-of-mouth. They've made it possible to compare options more systematically and identify potential fits more quickly.

But they haven't necessarily made the decision to purchase easier.

If anything, they've made it more complex. Every new vendor surfaced by AI creates another validation gate. Every LLM comparison needs to be verified against brand websites, third-party reviews, and peer experiences. Every claim needs to be checked because in its current state, buyers know AI can hallucinate, misunderstand context, or present outdated information.

In this study, we see buyers using AI for speed in the early stages, for structure in the middle, but never for the final call.

*"I'm not going to put my professional reputation on the line for an LLM recommendation."
(Director of Information Security, 1001-5000 employees, SaaS/Software)*

When professional reputation, executive scrutiny, and real money are on the line, AI can't erase the friction.

What does this mean for vendors?

You need to be discoverable in three key places: Google, LLM outputs, and peer networks. Miss any one of these and you'll be out of the shortlist for a third of buyers.

Your content needs to work for both AI and humans: structured enough for LLMs to crawl and cite, detailed enough for buyers to verify and trust.


Independent validation is key: reviews, case studies from customers willing to go on record, analyst reports, peer recommendations. The proof buyers trust most is the proof you control least.

And most importantly: transparency beats perfection. Buyers don't expect you to do everything. They expect you to be honest about what you do well, who you're built for, and where your limitations are.

AI has added another layer of validation to navigate. The ones still optimizing only for Google, or only for brand awareness, or only for sales conversations are missing the new opportunities where deals are won and lost.

The question is: when a buyer asks an LLM about vendors in your category, what shows up? And when they verify that information with peers, reviews, and your website, does it hold up?

That's where deals are made now. In the gap between what AI suggests and what humans trust.



Omniscient Wynter

Why Omniscient: Omniscient helps B2B companies drive attributable revenue through SEO, AI search, content, and digital PR. Through data-driven insights and a strategic approach to organic growth, we help you capture demand, increase brand visibility and sentiment in AI search, and accelerate pipeline growth. By aligning your organic strategy with how buyers actually research and make decisions, we help marketing teams build a durable, compounding growth engine.

Why Wynter: Wynter's verified B2B network gave us real decision-makers with recent buying experience, people who've actually influenced or owned a purchase in the past 6-12 months. That means the data reflects what's happening right now, not theory. Wynter helps B2B companies see how buyers perceive their messaging, positioning, and brand. With a 70,000+ verified panel, it delivers fast, reliable insights across messaging tests, ICP research, and brand tracking.



AUTHOR Cate Dombrowski

Cate Dombrowski is a Content Data Analyst at Omniscient Digital, where she blends storytelling with statistics. With a background in marketing and data analytics, she's driven by a curiosity to uncover hidden insights or validate ideas with data. Outside of work, she enjoys cycling, trying new restaurants, and reading in Central Park.



AUTHOR Martha Whiting

Martha Whiting is a Content Marketing Manager at Wynter, helping B2B companies uncover how buyers perceive their messaging and brand. Combining her background in brand strategy, research, and product development, she brings both analytical rigor and creative insight to her work. Based in Portugal, she enjoys traveling and spending time with her dog.



Wynter