The B2B Tech shortlist is tightening where buyers start.

Brands named per ChatGPT answer fell 28% in six months, and the broadest questions lost the most.

BrightEdge referral data shows AI search volume still climbing. Against that, a shorter list of brands per answer means each named slot is worth more than it was six months ago.

BRIGHTEDGE RESEARCH | SEPTEMBER 2026

Where the list got shorter

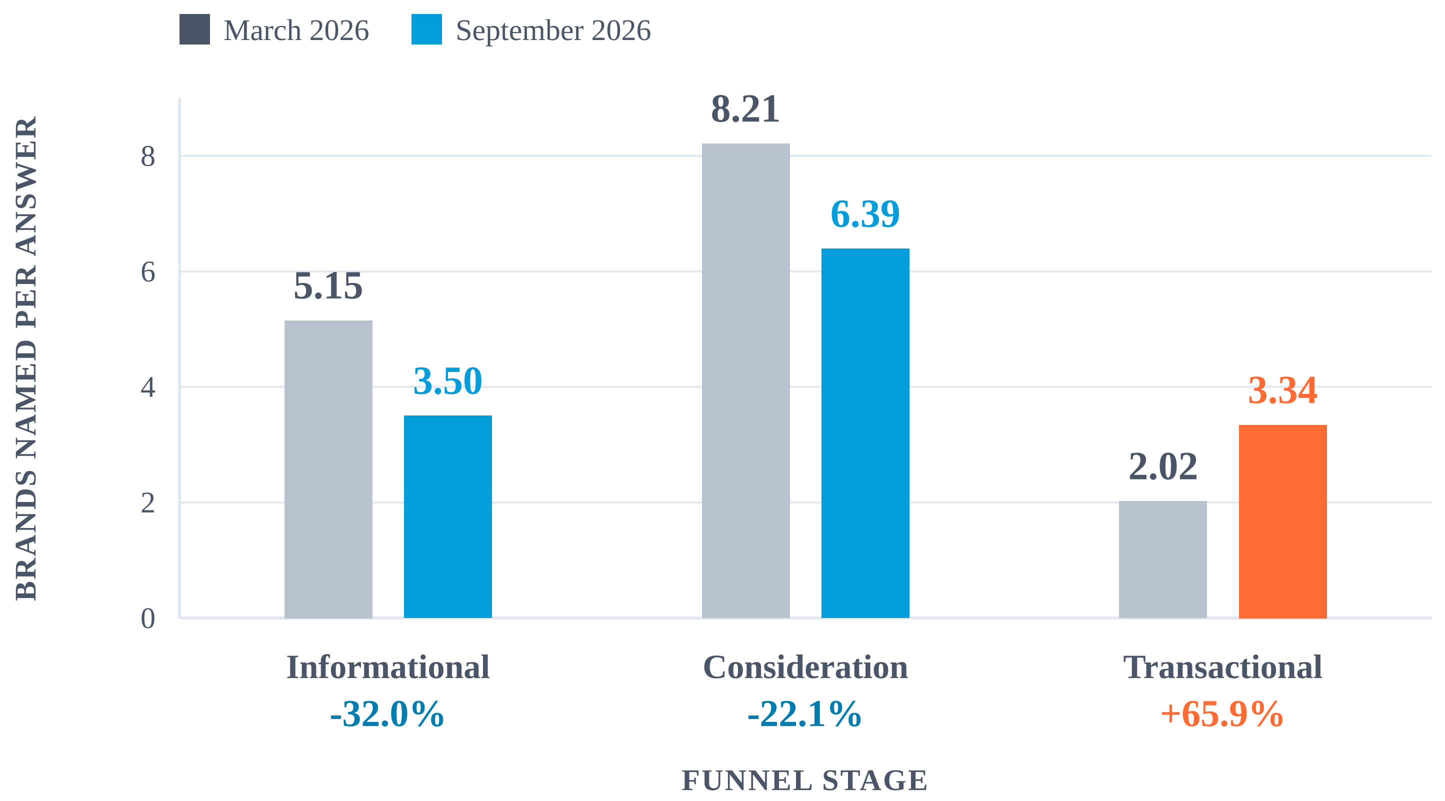

Figure 1. Brands named per ChatGPT answer by funnel stage, March against September. Vertical axis shows the average number of distinct brands named in one answer. Horizontal axis shows the funnel stage. Each stage is measured independently, so read each pair on its own. The Transactional bars are shown for completeness and are not read as a trend in this edition.
Source: BrightEdge AI Catalyst. ChatGPT only, a fixed panel of tracked non-branded B2B Tech prompts, United States. Figures cover prompts that held the same intent classification on both dates.

ChatGPT named 5.36 brands per answer in March and 3.87 in September, a 27.9% decline. Informational fell hardest at 32.0%, Consideration 22.1%.

Transactional moved differently over the same window and is not read as a trend here. A third measurement point reverses that movement.

The broadest questions got more selective

LOST THE MOST GROUND

Presence change, March to September

Microsoft	-9.5 pts
Google	-6.7 pts
Amazon	-6.4 pts
Apache	-4.5 pts
dbt	-2.8 pts
Snowflake	-2.5 pts

GAINED THE MOST GROUND

Presence change, March to September

IBM	+9.9 pts
NIST.gov	+8.1 pts
TechTarget	+6.9 pts
FBI.gov	+1.7 pts
Cisco	+1.6 pts
FTC.gov	+1.3 pts

Figure 2. Presence change in Informational prompts, March to September, for the six domains that lost the most ground and the six that gained the most. Presence is the share of tracked prompts naming that brand at least once.
Source: BrightEdge AI Catalyst. ChatGPT only, a fixed panel of tracked non-branded B2B Tech prompts, United States. Brand-level movements are measured across all prompts at each date, a wider cut than the fixed cohort used for the stage averages.

The largest platforms lost the most, and a cluster of developer and data-infrastructure names went with them. Institutional and government sources moved up, and so did TechTarget, a publisher. The expected publisher squeeze does not appear in this data.

The short list is made of specialists

BRAND GAINING IN CONSIDERATION	MARCH	SEPTEMBER
Zoho	8.5%	**12.8%**
Wave Apps	2.8%	**5.7%**
SAP	3.5%	**5.7%**
Squareup	2.8%	**5.0%**
Malwarebytes	1.4%	**4.3%**
ProtonVPN	2.1%	**4.3%**
NordVPN	2.1%	**4.3%**
Kubernetes.io	2.1%	**4.3%**
ServiceNow	1.4%	**3.5%**
Terraform.io	Not present	**2.1%**

Figure 3. Brands whose presence grew in Consideration prompts, the stage where someone is comparing named options. Presence figures sum past 100% across brands because one answer names several, so this is not a market-share table.
Source: BrightEdge AI Catalyst. ChatGPT only, a fixed panel of tracked non-branded B2B Tech prompts, United States. Brand-level movements are measured across all prompts at each date, a wider cut than the fixed cohort used for the stage averages.

Consideration lost 22.1% of its brands per answer, so the list got shorter. Every name that grew inside it is a single-purpose tool.

Where visibility is worth fighting for

✓ **Fight for the winnable Consideration prompts**

Specialists are filling this stage as it narrows. Find the one use case your product owns, and build content that answers that exact question.

✓ **Audit where you appear stage by stage**

A single visibility score hides which stage you are missing from. Split your tracked prompts by intent and find the stage where you are absent, because that is a different fix from being ranked low.

✓ **Hold your top-of-funnel investment**

Awareness answers hold their value. They just got harder to win, and the goal there is now a place on a shorter list, so pick one to three prompts where you have a right to win and go after those.

✓ **Google is still the channel to plan around**

The canonical, specific content that earns organic visibility is also what gets a brand named in an AI answer. Build it once and it works on both.

KEY TAKEAWAY

The moment AI is least likely to name your brand is the moment most buyers start looking. The moment it is most likely to name you is later, and more crowded, than most plans assume.

CHECK THIS IN YOUR OWN CATEGORY THIS WEEK

Three checks that show where your funnel moved

1 **Split your prompts by funnel stage**

In AI Catalyst, open Mentions, filter to your tracked prompt groups and split by intent. Compare two dates six months apart, which is long enough for a shift to show.

2 **Check the buying-intent answer first**

Ask ChatGPT a buying-intent question for your category and see who it names. Compare that list against the one a broad category question returns.

3 **Compare across engines before you act**

In AI Hyper Cube, pull the mentioned-brand list across funnel stages for the cross-engine view. This measurement is ChatGPT only, and Google may not behave the same way.

This check needs enough tracked prompts at each stage to compare. A category with very few Consideration or Transactional prompts will show noise in place of a trend.

Request a demo