

The AI Trust Gap: What Marketers Know That the C-Suite Doesn't

A survey of 135 marketing and brand leaders at technology-forward enterprises.



A MARKUP AI REPORT

Executive Summary

The AI Trust Gap: What Marketers Know That the C-Suite Doesn't

Marketers are caught in a bind of their own organizations' making.

AI mandates are coming from the top down. The tools are proliferating faster than control can keep pace. And marketing teams — the people responsible for every word that represents the brand — are left holding the risk.

This report is built on responses from 135 marketing and brand leaders at technology-forward organizations with 500 or more employees. The findings are clear: AI content is scaling rapidly, guardrails are not keeping pace, and the teams closest to the problem see it most clearly.

The trust gap in numbers:

92%

using more AI for content than last year

45%

believe AI can reliably check its own output

96%

say Content Guardian Agents would be valuable

The gap between those first two numbers and the third is where brands are being put at risk every day.

The Mandate Is Real — And It's Creating Pressure Marketers Feel Directly

The AI content surge isn't theoretical. 87% of marketing and brand leaders say their organization has an active AI mandate. On average, AI is now involved in generating 46% of all content their teams produce. And the volume is accelerating: 92% report using AI more for content creation than a year ago, with 42% saying usage has grown significantly.

For marketing teams, this isn't an abstract trend. It's a daily operational reality.

The primary reason marketers say their organizations turned to AI: **improved content quality and consistency (64%)** — not speed, not cost reduction. This is a critical distinction. Marketers adopted AI with quality expectations baked in. What they're discovering is that those expectations aren't being met without intentional control.

Meanwhile, 69% of marketing teams are using multiple or unapproved LLMs to generate content, with 50% using multiple approved LLMs across different use cases and an additional 19% where shadow AI usage is occurring alongside sanctioned tools. The result is fragmented output that no brand standard can survive without enforcement infrastructure.

“My peers were using Anthropic for some things, and Gemini for other things, or not at all.”

— VP of Marketing, B2B SaaS

“People [are] kind of going rogue and spinning up content.”

— Head of Content Marketing, Healthcare Tech



The Blind Spot: Marketers See It, Even If the C-Suite Doesn't

The AI content mandate usually comes from the top. The accountability for what is actually published stays with marketing.

This is the fundamental tension hiding inside every AI content adoption story. Executives see the productivity promise — faster content, more volume, lower cost — and set expectations accordingly. Marketers live with the consequences: AI output that misses the brand, terminology that's quietly wrong, claims that haven't been verified, and a publish button that still requires a human to feel confident enough to press it.

This isn't a new dynamic. Marketers have long navigated the gap between how leadership perceives their function and what it actually takes to execute it well. What's new is that AI has made this gap quantifiable, and the data makes clear that the people closest to the content have a fundamentally different view of its risks than those setting the strategy from above.

45%

of marketers believe AI models can reliably check their own content for accuracy and compliance

The majority of marketing leaders have already arrived at the correct conclusion: AI isn't a reliable self-editor.

"We drastically reduced the amount of editorial support from human editors that writers were getting. Most writers, most of the time, are not edited by a human — so we need a good tool to help check the content to keep our quality up."

— Content Operations Lead, Fortune 500 Technology Company

Yet behavior hasn't caught up with belief. **77% of marketers reported relying on a mix of manual review and spot checks to verify AI output** — an approach that creates bottlenecks and undermines the very efficiency gains AI was supposed to deliver. Without trustworthy automated oversight, human reviewers remain the final content gatekeepers: an approach that's unsustainable at enterprise scale.

There's also a meaningful gap between how marketers and the C-suite perceive their own governance readiness:

44% of marketers are “very confident” in their AI content governance | **57%** of C-suite leaders say the same

The people furthest from the content are the most confident in its quality. The people closest to it are not.

“Every time I decided to make a change, [the AI] would agree with me. That’s the part that drives me nuts — give me some resistance. You can say ‘be brutally honest,’ but it still misses the mark.”

— VP Marketing, B2B Learning & Development

This confidence gap is compounded by how organizations are rating their own guardrails:

26% of marketers rate their guardrails as strong and consistent | **40%** of C-suite leaders say the same

Marketers are more skeptical of their own governance infrastructure than their leadership is, and the data suggests they're right to be.

These aren't theoretical concerns. The risks that marketers flag as most pressing are career-level exposures for the people responsible for sign-off:

- 63% cite inaccurate or misleading information as a moderate-to-high risk
- 62% identify brand reputation damage as a moderate-to-high risk
- 58% flag non-compliance with industry regulations

Governance Is Fragmented and Accountability Is Unclear

When something goes wrong with AI-generated content, who owns it?

The data suggests no one is sure. Only 44% of marketing leaders report having AI content policies that are fully documented and widely communicated. The remaining 56% are operating with policies that are either inconsistently enforced, informal only, or nonexistent. Teams are scaling AI output without a shared definition of what “good” looks like.

Accountability is equally diffuse. When asked who is primarily responsible for ensuring AI-generated content meets standards:

- 37% of marketers say IT owns it
- 34% say marketing and brand management owns it
- 13% point to an AI committee
- Only 6% cite risk or compliance

Marketing and IT are essentially in a tie for ownership of a problem neither function is fully equipped to solve alone. With responsibility split across functions and policies inconsistently communicated, effective enforcement is structurally impossible.

“We’re very clear that we talk about getting people to practice in a risk-free environment, and [the AI] would come back with ‘a safe environment.’ We don’t like to use ‘safe.’ If you don’t know to look for it, you miss it — and then you’ve got to take the blog down and fix it and put it back up.”

— VP Marketing, B2B SaaS

“In highly regulated industries, using the word ‘safe’ is usually not allowed. If you’re on a content marketing team not connected to product marketing, you may not get that message. You may have to really, really dig for it.”

— Content Strategy Consultant

34% of marketers say their department owns AI content standards

25% of C-suite leaders agree marketing owns it

This ambiguity is compounded by the fact that the majority of organizations haven’t communicated or enforced formal policies for AI use in content creation. Teams are scaling AI output without a shared definition of what “good” looks like.

This gap matters. When marketing teams believe they own the governance problem but their leadership doesn’t fully agree, resources, tools, and budget won’t follow. Marketers are left accountable without authority.

The Opportunity: Consistency at the Speed AI Demands

The case for Content Guardian AgentsSM isn't a risk argument. It's an efficiency and quality argument.

Marketers adopt AI to produce more content, faster, with greater consistency. Manual review loops negate that efficiency. The only path to realizing the full productivity promise of AI content — without sacrificing brand integrity — is automated control that operates at the same speed as the content generation itself.

“At the end, where I am right now, is basically vibes — do I feel good enough that if the CEO sees this and asks what happened, I would be able to stand up for myself.”

— VP Demand Generation, B2B Marketing Technology

96%

of marketing and brand leaders say a dedicated Content Guardian Agent would be valuable to their organization

That near-unanimous signal is worth pausing on. Marketers aren't just open to the idea — they're asking for it. 50% say Content Guardian Agents would be “somewhat valuable,” reflecting a pragmatic, show-me orientation that's consistent with how marketers evaluate any new tool. The remaining 46% say they would be “very valuable” — a figure that will grow as proof points emerge.



The three options are clear:

<p>1 Fly Blind</p> <p>Scale AI content and accept the brand and compliance risk.</p>	<p>2 Manual Review</p> <p>Scale AI content and rely on manual checks — surrendering the efficiency gains.</p>	<p>3 Content Guardian Agents</p> <p>Automated content control that scans, scores, and remediates AI output at the speed of production.</p>
--	---	--

What they do:

- **Scan, score, and rewrite:** Automatically evaluate content against brand, compliance, and accuracy standards, assigning a readiness score.
- **Trigger oversight intelligently:** When a threshold is crossed, the agent either rewrites content or flags it for human review.
- **Integrate everywhere:** Built on API-first and MCP architectures, Content Guardian Agents work within enterprise ecosystems ensuring oversight happens wherever content creation is happening.

Analysts agree: Gartner predicts that 40% of CIOs will demand Guardian Agents within the next two years. The time to prepare is now.

Guarding the era of safe AI-generated content

Generative AI has made content creation effortless. But without the trust and scalability needed to provide enterprise-grade outputs, that content quickly becomes a liability. Manual review can't keep up, policies can't enforce themselves, and AI can't check its own work.

For organizations looking to fully and confidently unleash efficiency gains made possible by LLMs across content generation, Content Guardian Agents are the new standard for AI content control.

AI writes your content. Content Guardian Agents make sure it's right.

To learn more about Content Guardian Agents, visit Markup AI's website or sign up for early access.

LEARN MORE



markup ai

Content Guardian AgentsSM

Methodology

About This Research

Markup AI surveyed 266 respondents across two audiences: 135 marketing and brand leaders, and 131 C-suite executives. All respondents work at organizations with 500 or more employees, are involved in AI-related purchasing decisions, and identify their organizations as technology-forward. The survey was conducted online between September 11–19, 2025.

Unless otherwise noted, all statistics cited in this report reflect responses from the marketing and brand leader cohort specifically (n=135). Percentages may not sum to 100% due to rounding or multi-select responses. An asterisk (*) denotes responses below 0.5%.