



State of Martech 2026

by Scott Brinker and Frans Riemersma

Sponsored by:



Authors



Scott Brinker, dubbed the “godfather of martech” by *Ad Age*, has been analyzing martech and its impact on marketing organizations for more than 17 years as the editor of the [chiefmartec](#) blog. Previously, he served as VP platform ecosystem at HubSpot, and before that he was co-founder and CTO of ion interactive, a pioneering no-code platform for creating interactive content. He authored the best-selling book *Hacking Marketing* and co-authored the recent best-seller *The New Automation Mindset*. He holds graduate degrees in computer science and management from Harvard and MIT. He jointly produces the [MartechMap](#) marketing technology landscape with Frans.

Frans Riemersma is a marketing technologist, researcher, and advisor. With over 30 years of consulting experience, he blends qualitative expertise with quantitative martech data on stacks, vendors, and requirements, all consolidated in a proprietary Martech Data Warehouse. These insights power industry reports and tailored benchmarks for global enterprises through his company, [MartechTribe](#), which specializes in martech research and benchmarking. Together with Scott Brinker, he produces the [MartechMap](#), the definitive overview of the marketing technology landscape. Frans is the author of *A Small Book on Customer Technology*, co-author of *Marketing Tech Monitor* and *Hello Firstname*, and a regular contributor to Martech.org and CMSwire. He frequently lectures at Business Universities, Business Schools, and Chartered Institutes.

Sponsors

This report is sponsored by:

[GrowthLoop](#), [Hightouch](#), [Knak](#), [MoEngage](#), [Pega](#), [Progress](#), [SAS](#)

Insights from Our Sponsors



"[In contrast to coding with AI], marketing is not a verifiable domain. You can't send 100,000 emails to someone until they purchase a product you're trying to sell."

– Anthony Rotio



"Marketing is fundamentally a role of taste, judgment, and creativity. I actually think marketing teams will be among the least affected by AI headcount reductions."

– Tejas Manohar



"Underneath all the AI noise, I think we're going to see a renaissance in how we really talk to our audience — the best way, the most personalized way, the most human way."

– Brendan Farnand



"Forward-looking marketers are increasingly becoming managers of AI agents and agentic workflows rather than campaign creators."

– Raviteja Dodda



"Data silos are the enemy of AI. Your AI is only as good as your data. It doesn't matter how much money you spend on a wonderful AI tool if your data is not connected."

– Tara DeZao



"Humans are still ultimately the ones consuming the message. You need to build for agents, but understand that humans are still the end consumer."

– Sara Faatz



"The shift we're seeing is natural and what the market is demanding — that move from traditional CDP capabilities to CDPs becoming context-ready decision layers."

– Jonathan Moran

Content

Introduction	1
The Shape of the Transformation	2
The Fits and Starts of AI's Fast Lane	4
The AI-Martech Ecosystem Is Converging on MCP	5
Scarcity Mindset vs. Abundance Mindset, Redux	7
The Inner Game and the Outer Game	8
From Silent Movies to Talkies to <i>Armageddon</i>	10
 The Martech Landscape: Market in Metamorphosis	 12
The Flow Beneath the "Flat"	14
The Content Marketing Boom-Bust	16
The Landscape as a Martech Market Thermometer	18
Where Growth Is Actually Happening	20
The Reinvention of SEO as AEO (or GEO)	23
Hypothesizing the Future Shape of the Landscape	25
 Where Marketers Are Flying with AI	 29
The Through-Lines	29
Build vs Buy Is the Wrong Question	30
The Stack Is Stratifying, Not Consolidating	31
B2B Leads on Breadth; B2C Builds for Depth	32
The Trust Ladder	33
The Governance Gap	34
RAG Everywhere — The Hidden Connective Tissue	35
Advertising & Promotions	37
Content & Experience	40
Social & Relationships	43
Commerce & Sales	46

Data	50
Management	53

Context: The Creature in the Chrysalis **57**

The Spaces Between Our Contexts	59
Context Has Many Layers at Many Speeds	63
Context is a Decision-Time Game	66
The Real Butterfly Effect in Martech	69

Sponsor Interviews **71**

Causal AI and the Future of Marketing	71
How AI Agents Will Reshape the Marketing Workflow	79
The Campaign Creation Renaissance	86
Open and Explainable AI in Marketing	93
The Self-Orchestrated Customer Journey in Real Time	100
The Future of CMS in an Agentic World	106
The Middle Layer as Martech's Center of Gravity	112

Introduction

If you're wondering what that image on the cover is, it's a chrysalis — the structure a caterpillar builds around itself before it becomes something else entirely.

We chose it deliberately.

The tech industry loves its evolution metaphors. Products "evolve." Platforms "mature." Categories go through "natural selection." These are comforting framings because they imply continuity. The thing you knew yesterday is still recognizable today, just slightly improved. A caterpillar with better legs.

But that's not what happens inside a chrysalis. What happens is far more radical. The caterpillar essentially dissolves. Its cells break down into an undifferentiated soup, and from that soup, an entirely different organism assembles itself. Same DNA, unrecognizable form. The butterfly doesn't remember being a caterpillar. It doesn't need to.

That's closer to what's happening in marketing right now.

Not a graceful upgrade. A structural transformation — of the technology, the practice, the roles, and the relationship between brands and their customers. The old forms are dissolving. The new ones are assembling. And we're somewhere in the messy middle, which is exactly what makes this moment so disorienting and so full of possibility.

Welcome to the *State of Martech 2026*.

This moment in marketing is almost aggressively meta.

We're wrangling metadata¹, the contextual layer that makes AI useful. We're tagging and structuring content, the meta-markup that impacts how your brand shows up when an AI answers a question. We're all navigating metamorphosis — the structural transformation of our industry. And here we

¹ We never metadata we didn't like. (Sorry.)

are addressing this through metaphor, which is itself a meta-cognitive act of mapping one domain onto another.²

The root “meta” means “about itself” or “beyond itself.” And that’s the defining characteristic of this moment in marketing technology. The stack is becoming more aware of itself — through metadata, observability, and agent-to-agent communication. Marketing is being forced to think beyond itself. Beyond channels, beyond campaigns, beyond the org chart. And AI is pushing every tool, role, and workflow past the boundaries of what it was originally designed to do.

The Shape of the Transformation

So what, specifically, is metamorphosing? We see five dimensions of transformation playing out simultaneously — each at its own pace, each with its own messy middle:

5 Dimensions	Caterpillar (where we were)	Chrysalis (where most of us are)	Butterfly (where this is heading)
1. Who controls the conversation	The marketer — owned channels, managed funnels, controlled touchpoints, SEO as a game you could instrument and optimize	Increasingly, the customer — AI search intermediates discovery, conversations happen inside ChatGPT and Claude and Gemini where brand visibility is harder to earn and harder to measure	The customer’s agent — AI acting on the buyer’s behalf to research, compare, negotiate, and even purchase, interacting with the brand’s AI and systems directly
2. AI in marketing	Isolated task execution — generate this email, summarize this meeting, create this image	AI everywhere, integrated nowhere — every SaaS product has AI features, teams juggle multiple tools, early agentic workflows, but mostly still one task at a time	Orchestrated intelligence — agents that work across systems, maintain context across interactions, and act with increasing autonomy within guardrails

² You thought we might stuff in Mark Zuckerberg’s company too? A rare moment of writer’s restraint.

5 Dimensions	Caterpillar (where we were)	Chrysalis (where most of us are)	Butterfly (where this is heading)
3. Martech software	Deterministic SaaS — subscription platforms, rule-based workflows, the stack as a fixed machine	The great rebundling struggle — incumbents bolting AI onto existing products, AI-natives competing on model quality, and buyers hedging with all of the above simultaneously	Context-as-a-Service — platforms that deliver the right data, content, and capabilities to the right agent at the right moment, blending deterministic reliability with probabilistic intelligence
4. Marketing roles	Campaign manager — plan, execute, measure, repeat	Multimodal operator — juggling AI tools, legacy workflows, and an expanding surface area	Value engineer — designing systems that deliver measurable value across every customer interaction
5. Marketing ops roles	System administrator — configure platforms, maintain integrations, manage permissions	Stack wrangler — stitching together SaaS + AI + custom-built, often held together with duct tape and good intentions	Context engineer — orchestrating what data, content, tools, and instructions reach each AI agent at the right moment

If you're reading this as a CMO, the first row is the one to sit with. The industry has been promising a shift from seller-centric to buyer-centric marketing for at least 15 years. What's different now is that the shift may not be optional. When your customer's AI agent is the one evaluating your product, comparing your pricing, and deciding whether to surface your brand — the buyer's context isn't a strategic aspiration. It's the playing field.

The rest of the table cascades from that reality. AI is how the organization responds. The software is how the market responds. And the roles — the evolution from campaign managers and system administrators toward value engineers and context engineers — are how the people inside the building adapt.

Note the chrysalis column carefully. That's where most organizations are right now, and there's no shame in it. "AI everywhere, integrated nowhere" is not a failure state. It's a necessary stage. The caterpillar doesn't skip the chrysalis. But the organizations that mistake the chrysalis for the destination will be the ones still dissolving when their competitors take flight.³

³ The overextended metaphor police have now issued a warrant for our arrest.

The Fits and Starts of AI's Fast Lane

In our introduction to the *Martech for 2026* report back in December, we used one week of announcements from OpenAI as an example of the mad maelstrom of AI:

- Instant Checkout, which threatened to capture the full ecommerce buyer's journey, causing anxiety among many a merchant.
- Sora, a social network app for AI-generated video that rocketed to the #1 spot in Apple's App Store.
- ChatGPT Apps that could surface interactive UI from sites such as Booking.com, Canva, Spotify, Target, and Zillow directly inside ChatGPT's conversational interface.

Well, here we are six months later:

- Instant Checkout has been quietly sidelined. Only about 30 Shopify merchants ever onboarded. Walmart's EVP of AI publicly called it "a very temporary moment in time."⁴ OpenAI is pivoting to app-based commerce, where merchants own their own checkout experience — which, if you think about it, is exactly what the buyer-control row in our table above predicts. Both buyers and merchants prefer to own the customer experience, even as they acknowledge AI's expanding role in discovery.
- Sora was shut down on March 24. Downloads collapsed from 3.3 million in November to 1.1 million by February. It was burning roughly \$1 million a day in compute. And the \$1 billion Disney licensing deal — which would have made Marvel, Pixar, and Star Wars characters available for user-generated AI video — dissolved before it ever went live.⁵
- ChatGPT Apps have progressed meaningfully — an App Directory launched in December, third-party submissions are now open, and the ecosystem is

⁴ Ouch. [Walmart Inc. interview at Morgan Stanley Technology, Media & Telecom Conference 2026](#)

⁵ Given that users were already generating videos of Mario smoking weed and Pikachu doing ASMR without licensed characters, one can only imagine what they'd have done with the full Disney vault. Perhaps it's for the best.

growing. But it's still early innings for a platform that reaches 800 million weekly active users.

Whiplash, anyone?

But this too is an instructive example of life in the AI fast lane.⁶ Major AI announcements grab enormous attention, and then reality reasserts itself surprisingly fast. Sora went from world-changing to shuttered in six months. Instant Checkout went from existential threat to "let the merchants handle it" in about the same timeframe. The signal-to-noise ratio on any given AI launch is low.

Key Highlights & Insights
from Our Sponsor



GrowthLoop

"We'll move to zero-to-zero personalization faster than anyone anticipates — agents talking to agents. This will happen much faster than e-commerce did."

– Anthony Rotio


Yet zoom out even slightly, and the trajectory is remarkably consistent. The surface area where AI intersects with marketing keeps expanding. It's the individual plot points that are unpredictable, not the arc.

The AI-Martech Ecosystem Is Converging on MCP

Although the ChatGPT Apps story to date is basically a speed-run of every platform ecosystem challenge rolled into one semester, we're cautiously optimistic it will continue to evolve.

It's worth noting that one of the largest clusters of ChatGPT Apps are vendors in or adjacent to martech: Adobe, Airtable, Amplitude, Asana, Canva, Clay, ClickUp, Common Room, Conductor, Figma, Gamma, HeyGen, HubSpot, Intuit Mailchimp, Jotform, Klaviyo, Lovable, Notion, Replit, Semrush, Slack, and ZoomInfo.

⁶ [Life in the Fast Lane](#) by the Eagles, whose lyrics, upon inspection, are slightly NSFW.



Source: OpenAI ChatGPT featured Apps and Anthropic Claude sales and marketing Connectors (partial view), March 29, 2026

Similarly, Anthropic's ecosystem of Claude Connectors features many martech and adjacent integrations: ActiveCampaign, Ahrefs, AirOps, Airtable, Amplitude, Apollo.io, Asana, Atlassian, Attio, Braze, Box, Calendly, Canva, Clay, ClickUp, Common Room, Crossbeam, Customer.io, Databricks, Day AI, Docusign, Figma, Fireflies, G2, Gainsight, Gamma, HubSpot, Intercom, Make, Monday.com, n8n, Notion, Outreach, Pendo, PitchBook, SimilarWeb, Snowflake, Supermetrics, Tableau, Webflow, Wix, WordPress, Workato, Zapier, Zoho, and ZoomInfo.

Key Highlights & Insights from Our Sponsor

moengage

"Your agents need to be available everywhere work happens. Not just in your UI, but in Slack, Teams, and other tools. The future isn't a single interface; it's access from multiple places."

– Raviteja Dodda

Odds are pretty good you have one or more of those in your stack, ready to be connected if they aren't already. Martech is one of the first industries to show up in force on these AI platforms — which tells you something about where the integration gravity is heading.

ChatGPT Apps and Claude Connectors are the curated front — integrations vetted and surfaced within each AI platform. But MCP, the Model Context Protocol, is the open standard underneath both. Anthropic open-sourced it in late 2024, and then OpenAI, Google, and Microsoft all adopted it.⁷

Because anyone can build an MCP server for any tool or data source, thousands have. Independent registries — PulseMCP, Glama, mcp.so, and others — now index north of 29,000 unique servers, and counting. (It took the commercial martech landscape 15 years to reach 15,000 products. With MCP, agent integrations hit more than twice that in just 18 months.)

For martech, which has spent the better part of two decades wrestling with integration complexity, a shared protocol layer connecting AI agents to the tools they need to orchestrate is a big deal.

With that kind of fluid connectivity, the question is shifting from “can we connect these systems?” to “what do we do now that we can?”

Scarcity Mindset vs. Abundance Mindset, Redux

In our December report, we made the case that AI's real competitive potential isn't efficiency — doing the same things cheaper — but abundance: making entirely new things feasible. But six months deeper into this transformation, we want to sharpen the point.

Abundance doesn't eliminate constraints. It moves them.

When content production was expensive, the constraint was volume. How much can we create? Now that AI has made content nearly free at the margin, the constraint has shifted to relevance: which content, for which customer, in what context, at what moment?

When integrations were hard, the constraint was connectivity. Can we even get these systems to talk to each other? With 29,000 MCP servers and counting,

⁷ How often are those four companies going to agree on anything ever again?

the constraint is becoming orchestration: which connections matter, and who governs what flows through them?

When software required engineering teams, the constraint was capacity. What can we afford to build? Now that marketers can vibe-code apps and automations, the constraint is coherence: does this fit into anything larger, or have we just added another ungoverned node to the stack?

See the pattern? In every case, AI dissolves a production constraint and reveals a context constraint hiding behind it.

The organizations stuck in a scarcity mindset are still optimizing against the old bottlenecks. Producing more content. Connecting more tools. Building more automations. The organizations pulling ahead have recognized that the bottleneck has moved. They're investing in the layer that makes abundance useful rather than overwhelming: context engineering, governance, and strategic coherence.

Key Highlights & Insights
from Our Sponsor

knak.

"Efficiency unlocks optionality. It's going to unlock focus on the ability to really talk to your audience again."

– Brendan Farnand

The Inner Game and the Outer Game

Every organization is playing two games simultaneously.

The inner game is how the business works: the processes, the tools, the org chart, the data flows, the governance structures, the culture. The outer game is how it engages customers: the experiences, the channels, the brand, the speed and quality of every interaction.

These two games have always been connected, of course. But the connection used to be buffered by people. Internal silos, misaligned teams, disconnected systems — customers didn't always feel them because humans in the middle smoothed things over. But when the friction did leak through — when a customer got bounced between departments or had to re-explain their

problem for the third time — there's an old line in CX circles that captured the moment: "Oops, your org chart is showing."⁸

AI is stripping away that buffer.

When an AI agent is orchestrating a customer interaction, it can only work with the data, content, and tools it can actually access. There's no human in the middle to quietly fix the handoff between two systems that don't talk to each other, to intuit what the customer actually meant, or to compensate for the fact that the product catalog hasn't been updated since last quarter. The inner game will increasingly be the outer game.

The updated warning: "Oops, your stack is showing."

Do you have a GenAI policy?

February 2026



December 2024



■ No ■ Yes

Source: *AI Use Cases in Marketing 2026 Survey*, n=208, and *2024 GenAI Survey*, n=283
chiefmartec & MartechTribe

This has implications in both directions. For customer-facing AI — chatbots, shopping concierges, AI SDRs, agentic support — the quality of the experience is directly downstream of data quality, integration coherence, and governance. The organizations delivering genuinely impressive AI-powered customer experiences aren't the ones with the fanciest models. They're the ones who've done the unglamorous internal work of cleaning data, connecting systems, and defining guardrails.

⁸ Scott credits this to Dharmesh Shah, co-founder and CTO of HubSpot.

For the inner game, there's encouraging movement. In our AI use cases survey, 73% of respondents now report having a formal generative AI policy, up from 52% in 2024. That's progress.

But as SAS found in their year-over-year research, only 8% of organizations report full confidence in their broader AI governance readiness. A policy is a starting point, not a finish line — and the gap between “we have a policy” and “we have the infrastructure to enforce it” is where most organizations are still working things out.

From Silent Movies to Talkies to *Armageddon*

In 1927, *The Jazz Singer* added synchronized dialogue and singing to film. The technology was transformative. But consider what it demanded of everyone else in the business. Screenwriters had to learn dialogue. Actors had to learn to speak on camera. Directors had to rethink blocking, pacing, and set design. Studios had to rewire their theaters. Some careers ended. Many more were made. The industry didn't just add sound. It rebuilt itself around a new capability.

Sound didn't add a feature to silent movies. It created an entirely different industry — one that would grow from nickelodeons⁹ to a \$77 billion global business.

AI is the sound in marketing's silent movie.

The technology is here. It works. What's still being figured out is the craft. The storytelling, the strategy, the organizational design, the governance, the new roles and skills that turn a raw capability into something that actually moves customers and creates value. The models are the easy part. The hard part is everything human that wraps around them.

That metamorphosis is the thread running through everything that follows.

⁹ “Hey, Claude. What's a nickelodeon?” “It's a small, early 1900s storefront theater where you paid a nickel to watch short silent films.” So basically, the original minimum viable product.

Here's what we dug into this year:

- Our martech landscape analysis examines what's happening across 15,000+ products — and why the headline number doesn't tell the story you'd expect.
- Our AI marketing use cases chapter is the deepest analysis we've done — 70 use cases across six categories, designed to be composable so you can zoom in on what matters most to your world.
- Our final chapter explains how *value engineering* identifies where the real value lives in AI-powered marketing, while *context engineering* makes it actionable at the moment of decision.

Key Highlights & Insights
from Our Sponsor

 **hightouch**

"Think of how coding evolved. Copilot started as autocomplete in the editor, but real coding agents are now building entire projects end-to-end."

– Tejas Manohar

And throughout the next 107 pages, we keep pulling on the threads we've introduced here: context as the new constraint, the inner game shaping the outer game, and a stack that's stratifying rather than consolidating.

Thank you to our sponsors GrowthLoop, Hightouch, Knak, MoEngage, Pega, Progress, and SAS. Their support makes it possible for us to do this research and share it with you for free. The back half of this report includes in-depth interviews with each of them — not product pitches, but their thinking about where martech and AI are headed. We think you'll find their perspectives as thought-provoking as we did.

Let's go see what's taking shape in the soup.



Scott Brinker

Analyst & Advisor
chiefmartec



Frans Riemersma

Founder & CEO
MartechTribe

The Martech Landscape: Market in Metamorphosis

Pop quiz: How much did the martech landscape grow this year?


- A) 70%
- B) 7%
- C) 0.7%

If you've been following this space for a while, your instinct is probably B. In recent years, that would have been a reasonable guess. Modest, but meaningful.

Then again, with the AI explosion of agentic coding and vibe coding¹⁰, answer A might just be possible. Surprising, even shocking, but not inconceivable.

The right answer is C.

Number of Martech Software Apps Since 2011



Source: MartechMap Research, chiefmartec & MartechTribe


¹⁰ We're going to adopt the standard of using "agentic coding" for professional software development with AI and "vibe coding" for non-engineers building agents, apps, and automations with AI.

The martech landscape contains **15,505 products**, up a grand total of 121 from the 15,384 we mapped last year. That's 0.79% growth, rounding to effectively zero. After 15 years of relentless expansion — from 150 products in 2011 to over 15,000 in 2025, a 100x run — the martech landscape has plateaued.

At long last, we may have hit peak martech. Or at least a plateau for now.

Either way, “flat” is perhaps the most misleading word you could use to describe what's actually happening here. Underneath that near-zero net number is a market in motion: **1,488 new products were added, while 1,367 were removed.**

Martech Landscape 2011-2026



Source: MartechMap Research, chiefmartec & MartechTribe


Like individual stacks, the martech landscape is a river, not a lake.¹¹

¹¹ *Take Me to the River* by Talking Heads (1978), originally Al Green (1974). One wrote it as a gospel meditation on renewal. The other made it art-school weird. Both versions apply to martech.

The Flow Beneath the “Flat”

The static headline number buries the real story. Let's look at the flow:

Martech Landscape: Products Added and Removed by Year



Source: MartechMap Research, chiefmartec & MartechTribe


The above chart shows a four-year view of the martech market in the post-pandemic era. From 2023 to 2025, the landscape experienced a renewed creative explosion: thousands of new products added every year, fueled by the generative AI gold rush. The market was minting new products faster than it could consolidate or kill old ones.

2026 may be the inflection point. For the first time, the removal rate has nearly caught up to the addition rate. The volume of new entrants dropped 40% from last year (from 2,489 to 1,488). The volume of exits climbed by 13% (from 1,211 to 1,367). The market is starting to balance its books.¹²


Who's leaving, exactly?

¹² "Balancing the books" is a generous framing. For those exiting, the more precise term is "closing them."


Martech Products Removed by Year Founded: 2025 vs 2026



Martech Products Removed in 2026 by Revenue



Martech Products Removed in 2026 by Headcount



Source: MartechMap Research, chiefmartec & MartechTribe

By founding year, the cohort getting hit hardest is the 2010-2019 wave of SaaS startups. That generation accounts for 51.7% of 2026's removals, nearly identical to the 56.1% of exits from that same cohort in 2025.

The exits are concentrated in small companies: 41.2% had 1-10 employees, 38.7% had 11-50. By revenue, the \$1M-\$10M band accounts for a striking 45.5% of what's been removed. These are companies that found enough early traction to survive past the zero-revenue stage, but couldn't grow their way to durability in an increasingly competitive market.

The picture that emerges is of a healthy — if sometimes painful — metabolic process. The market is not randomly eliminating products. It is systematically culling the undercapitalized, the underdifferentiated, and the unlucky from an era when martech still seemed like an easy place to make a business. Every signal in this data points to the same underlying shift: a market reorganizing itself around AI, not merely adding it.


The Content Marketing Boom-Bust

Perhaps the single most illuminating subcategory story in this year's data is the reversal in **Content Marketing**.

To understand 2026, you need to remember 2024. When generative AI went mainstream in 2023, Content Marketing was one of the first categories to feel the full force of the wave. Founders saw obvious applications — writing tools, creative production platforms, content optimization engines — and investors were just as eager to ride the new AI wave. The category nearly doubled over two years, from 575 tools in 2023 to 1,102 in 2025, making it the second-fastest-growing subcategory on the entire landscape (Sales Automation, Enablement & Intelligence beat it in that same time period, from 708 to 1,546 products).

In 2026, Content Marketing leads all other subcategories in a different, less coveted ranking: it has the **highest net product removal of any subcategory, at -37** (176 removed, 139 added). The largest wave of AI-powered martech to ever wash over a single category is now receding.

Martech Category Growth Rates, 2024-2026



Source: MartechMap Research, chiefmartec & MartechTribe

What happened? Three forces converged:

1. **The major AI labs absorbed the functionality.** The capabilities that made first-wave content martech exciting in 2023 — generate a blog post, product ad copy variations, create social content at scale — are now table stakes in ChatGPT, Claude, Gemini, and every major productivity suite. When the foundation models commoditize your core feature, the argument for a dedicated point solution becomes difficult to make.¹³
2. **Incumbent SaaS moved fast.** The Adobes, HubSpots, and Salesforces of the world were not asleep. They swiftly embedded generative AI into their existing content workflows. Many of the customers who might have paid for a standalone AI content tool already had one — attached to a platform they were already paying for.

¹³ *Video Killed the Radio Star* by The Buggles.

3. **Product-market fit proved elusive.** Generating content fast and generating content that actually works are different problems. Many first-wave tools solved the former but struggled with the latter. As the novelty wore off and organizations asked harder questions about quality, brand consistency, and downstream impact on conversion, the undifferentiated generalist tools found themselves without a compelling answer.

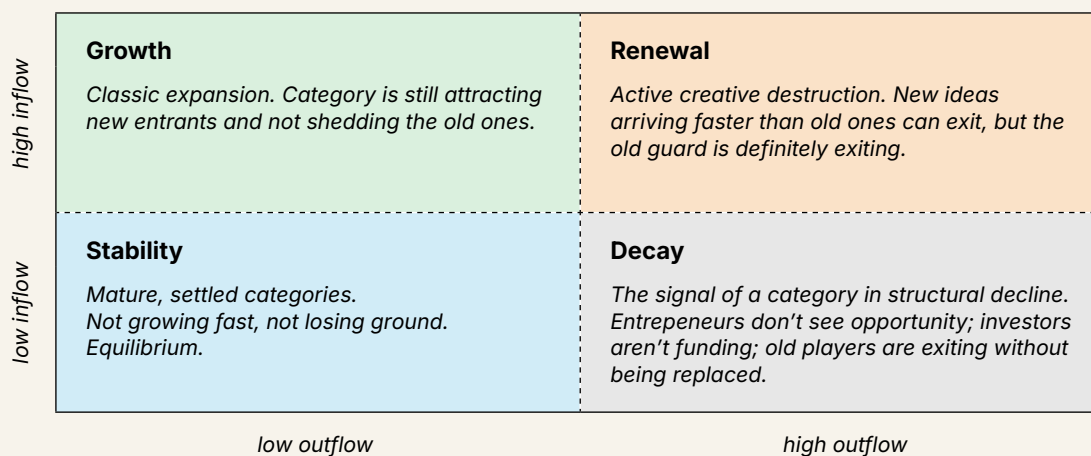
The result is a natural selection event playing out in real time. This is not the end of AI-powered content technology. But the first cohort of generative AI-powered pioneers is being consolidated, acqui-hired, or quietly wound down, making room for a more mature second generation with clearer differentiation and strong integration into the workflows that actually govern how content gets made and approved.

The Landscape as a Martech Market Thermometer

We've long thought of the martech landscape as more than a catalog. It's a living market signal — a thermometer that reflects the collective beliefs of thousands of entrepreneurs and investors about where the opportunities in marketing are emerging, stabilizing, or fading.

Reading that signal requires looking at both inflow and outflow together. Inflow — new products entering a category — reflects some combination of perceived opportunity for innovation and expressed customer demand. Outflow reflects the opposite: waning opportunity, competitive pressure, consolidation, or simply the natural end of a product's useful life.

A category with high inflow and low outflow is in classic **Growth** — still attracting new entrants without shedding the old ones. High inflow with high outflow is **Renewal** — active creative destruction, with new ideas arriving even as the previous generation exits. Low inflow and low outflow is **Stability** — mature, settled, neither expanding nor collapsing. And low inflow with high outflow is **Decay** — the signal of a category in structural decline, where entrepreneurs no longer see much opportunity, investors aren't funding new bets, and old players are exiting without being replaced.



In 2026, CMS & Web Experience Management and Ecommerce Platforms come closest to classic Growth. Content Marketing is a clear case of Renewal: 139 new entrants alongside 176 exits — heavy traffic in both directions, as a first wave of AI-era products collides with a harsher test of differentiation. CRM appears to be approaching Stability: modest movement in both directions, a mature category finding equilibrium. And a handful of legacy advertising categories now sit firmly in Decay.

The 2026 data suggests that most of martech sits in the Renewal quadrant — genuinely active churn, with fresh entrants arriving even as the previous generation clears out. That's a healthy sign for the ecosystem overall, even if it's uncomfortable for the individual companies being cycled out. But true Growth categories — high inflow, low outflow — are still rare in 2026, which is itself a signal.

Read the landscape as a whole, then, and what you see isn't a market dying. It's a market metabolizing.

The clearest examples of Decay are Native/Content Advertising, Video Advertising, and Print — categories where the entrepreneur and investor communities have largely stopped showing up. The DMP category, meanwhile, isn't merely decaying. It's six feet under.^{14,15}

¹⁴ For younger readers: a DMP, or Data Management Platform, was a system for managing third-party audience data. It had a good run, which is more than its match rates ever did.

¹⁵ In lieu of flowers, please send first-party data.

Where Growth Is Actually Happening

If Content Marketing is the cautionary tale, **CMS & Web Experience Management** and **Ecommerce Platforms** are the 2026 growth story — and they're not unrelated.

CMS & Web Experience Management grew 21.4% this year, its highest growth rate in at least three years, jumping from 504 to 612 products. Ecommerce Platforms grew 19.9%, from 547 to 656. Both are, on the surface, surprising results. These aren't new categories. They've been on the landscape since the beginning. Why are they accelerating now?

We believe CMS & Web Experience Management is benefiting from a shift in the role of the website itself.

For two decades, marketing teams built web experiences primarily for two audiences: humans and search crawlers. That list is now expanding rapidly to include AI search assistants, agentic browsers, shopping assistants, and procurement agents that arrive not to browse but to extract, evaluate, and act on behalf of their human principals. As machines become a first-class audience for digital experience, the value of CMS platforms appears to be shifting from page publishing alone to providing structured, machine-readable context that AI systems can use.

That helps explain the wave of new infrastructure forming around the category: APIs, structured content layers, and agent-facing interfaces that let AI systems query content programmatically; conversational interfaces becoming a native delivery channel rather than a bolt-on widget; and generative content assembly, where responses are dynamically composed from structured components instead of served only as pre-authored pages. If that interpretation is right, CMS growth is not about websites becoming fashionable again. It's about the fundamental job of the website being renegotiated.

Key Highlights & Insights
from Our Sponsor




"'Generative CMS' as a category label really encompasses all that this next iteration of the technology is. It is a CMS that can deliver experiences that are conversational, contextual and deeply personal."

– Sara Faatz

The ecommerce growth story is related but distinct. Our interpretation here is that ecommerce platforms are being pressured from two directions at once.

Top Martech Subcategories in 2026 by Growth Rate



Source: MartechMap Research, chiefmartec & MartechTribe

First: AI-mediated product discovery is changing where and how customers evaluate products, which increases the importance of clean, structured, machine-readable catalog data that many legacy platforms were not originally designed to expose well.

Second: customer expectations for real-time, individually tailored experiences continue to rise, pushing platforms beyond the segmentation-based personalization of the previous generation. In that sense, the category appears to be expanding because ecommerce infrastructure is being rebuilt for both machine-facing discovery and more adaptive customer-facing experiences.

The rest of the growth list follows a similar logic — categories expanding precisely because AI is disrupting adjacent parts of the stack:

The analytics renaissance. Mobile & Web Analytics grew 11.3% after years of stagnation. Call Analytics grew 8.9%. This is arguably a direct response to what's happening with AI-mediated customer journeys. As prospects move more of their research inside AI assistants and off of trackable web sessions, marketers are investing more in the instrumentation that *can* capture signal. You can't measure what you can't see, so you measure harder where you still can see.¹⁶

Infrastructure is interesting again. The iPaaS/Data Integration subcategory, now as much an AI orchestration layer as an integration tool, grew 8.0%, reflecting the reality that everyone's stack is getting more complex, not less. The connective tissue holding it together matters more when agents need to traverse it autonomously. Governance, Compliance & Privacy grew 7.1%, partly due to evolving regulation (the EU keeps busy) and partly as the natural consequence of AI adoption generating new questions about data use that didn't exist three years ago.

Marketing Automation's quiet comeback. After years of being The Category That Surely Must Be Done By Now, Marketing Automation & Campaign/Lead Management grew 5.9%, with net positive movement in both inflow and outflow. We'd attribute this partly to AI capabilities reinventing what campaign orchestration can look like. The category is mutating rather than dying, attracting a new wave of builders who see agentic marketing automation as a meaningful step change from the rule-based systems of the past decade.

¹⁶ This is, technically, the streetlight effect — searching where the light is rather than where the keys are.

The Reinvention of SEO as AEO (or GEO)

SEO has had no shortage of eulogies these past couple of years.¹⁷ AI search assistants swallowed the top of the search funnel. Zero-click searches strangled organic traffic. The content mills flooded. Pick your cause of death — someone has written the obituary.

Meanwhile, the SEO/AEO subcategory posted a **net +6** this year — 38 removed, 44 added — remaining one of the more persistently positive subcategories on the landscape despite the macro narrative of its demise. Last year it grew 24%. The year before, same story.

SEO isn't dying. It's metamorphosing. And the market is reflecting that metamorphosis in real time.

The category we've historically called SEO is becoming something broader and more complex: the discipline of making your brand findable, credible, and actionable across all the surfaces where customers and prospects are now doing their research.

Key Highlights & Insights
from Our Sponsor



"With offsite brand or third-party mentions becoming increasingly more important — as they are key signals of authority to LLMs — consistency of message and narrative is key."

— Sara Faatz

Some are calling this AEO (Answer Engine Optimization or AI Engine Optimization) or GEO (Generative Engine Optimization). Whatever you call it, there's a bevy of new tools — AirOps, Bluefish, Daydream, Evertune, Profound, Scrunch, and others — building specifically for the AI discovery layer, while the established SEO platforms like Semrush (now part of Adobe) and Ahrefs are rapidly extending their capabilities to cover it.

But here's the harder truth underneath the product innovation: AEO is a discipline where the market is growing and the marketer's visibility is shrinking.

¹⁷ SEO has been declared dead so many times it now qualifies as a prestige TV character. At this point we're just waiting for the spinoff.

For 25 years, SEO was built on a foundation of instrumentation. You could see every click, every keyword, every conversion path. Google's opacity was frustrating, but marketers developed sophisticated proxies and tools to work around it. The feedback loop was tight enough to run experiments, learn, and optimize.

AI-assisted research breaks that loop. When a customer has a conversation with Claude or ChatGPT or Perplexity about which CRM to buy or which running shoe to choose, that conversation is invisible to you. You don't know what question they asked. You don't know what answer they received. You don't know whether your brand showed up, or in what light, or in comparison to whom. You can't A/B test your way into that conversation.

What you *can* do:

Invest in brand. Brand signals — citation patterns, reputation, third-party mentions, community presence — are increasingly how AI systems decide what to surface in response to queries. The discipline of brand-building, which has always been harder to instrument than performance marketing, is now more directly upstream of AI-mediated discovery than it has ever been.

Attend to your digital body language — what you can still see. Your owned properties still generate signal. On-site behavior, search console data, analytics from users who do click through, social listening. These aren't going away, and they're worth investing in more seriously precisely because the part of the funnel you can't see is growing.

Engage where conversations happen. Reddit, LinkedIn, YouTube, review platforms — AI systems are actively indexing and learning from community discussions. The brands that have been playing the long game of genuine community participation are, somewhat unexpectedly, finding themselves well-positioned for the AI discovery era. It turns out authenticity was always good strategy. It's just now more directly measurable in its effects.

Key Highlights & Insights
from Our Sponsor



"Customers see over 6,000 messages a day now. In order to break through that noise, you have to reach them in the moment they're making a decision or need your help."

– Tara DeZao

Make your digital presence machine-readable. Schema markup. Structured FAQs. lms.txt files.¹⁸ MCP servers for customers. The standards are still being written, but the direction is clear: your digital presence needs to be legible to machines, not just humans.

Our survey last fall found that **63.1% of respondents were already publishing AI-optimized content** (structured Q&As, schema markup), by far the most adopted AEO practice. But **only 13.6% were measuring AI inclusion rate and agent-referred conversion**. That gap — between doing the work and knowing whether it's working — is one of the defining challenges of this transitional moment.

Talk to customers. This may be the most important and most underrated response to the opacity problem. When the digital journey is happening in conversations you can't monitor, the answer is sometimes as simple as having more conversations of your own. Customer interviews, sales call listening, social media conversations — human signal that doesn't require an AI assistant to mediate it.¹⁹

The new martech products in this space are getting quite clever about providing directional visibility into AI inclusion and brand presence. Think of it as a category of "AI audit" tooling. They can't replicate the granular click-by-click instrumentation marketers are accustomed to, but they can tell you whether you're showing up at all, in what contexts, and relative to competitors. That's a start.

Hypothesizing the Future Shape of the Landscape

If the 2026 data is a snapshot of a market at an inflection point, what does the trend line suggest about where the landscape goes from here?

¹⁸ An lms.txt file is a machine-readable summary of your site's content, purpose, and structure — essentially a cover letter for AI crawlers. Think robots.txt, but friendlier and more verbose, which is on-brand for the LLM era.

¹⁹ We're aware that "talk to customers" as cutting-edge strategic guidance in a 2026 AI report has a certain energy. And yet.

Our hypothesis: the martech landscape will consolidate at the head, get squeezed in the torso, and eventually re-expand at the tail. The 2026 data doesn't confirm that shape yet — net growth is essentially flat across the board — but it's consistent with a market in the early stages of that transition.

The head is getting stronger, not weaker. The AI capabilities that major platforms can bundle into existing enterprise relationships are genuinely valuable, and enterprise buyers overwhelmingly prefer to consolidate vendors rather than add new ones. Platforms that have invested seriously in AI and agentic capabilities are expanding the scope of value they deliver per contract, which makes them harder to displace and harder to work around.

Key Highlights & Insights
from Our Sponsor



"The activation layer will be more heavily agentic in the future. So the middle decisioning and orchestration becomes critical. I do see it as the new center of gravity."

— Jonathan Moran

The torso — SaaS companies at \$50M+ ARR competing in defined categories — faces pressure from both directions simultaneously. Platforms are expanding downward with bundled AI capabilities. AI-native point solutions are commoditizing specific features from below. The torso companies most at risk are those whose differentiation was always thinner than their revenue suggested: built on integrations, workflow lock-in, or category timing rather than something genuinely hard to replicate.

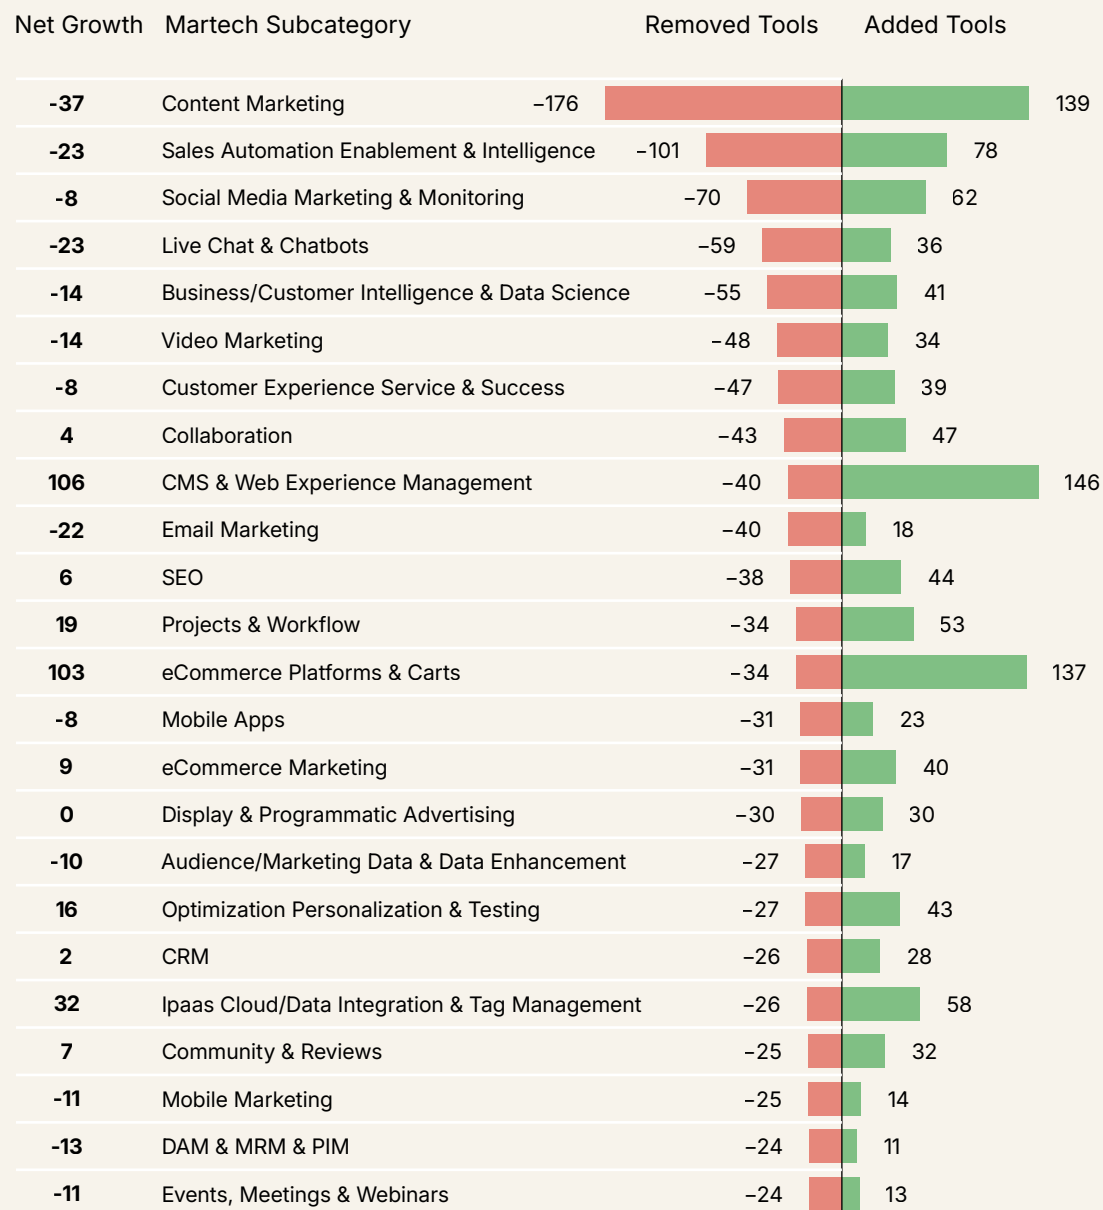
The tail tells a different story. The 2026 churn data shows that 45.5% of this year's removals came from the \$1M–\$10M revenue band. These were companies that found early traction but couldn't grow into a defensible position before running out of road. Historically, they would have been the ones graduating into the torso. The fact that they're exiting instead is arguably the most telling signal in the data. Torso compression isn't just about pressure on the companies already there. It's about fewer companies making it there at all.

But the conditions for tail re-expansion are intact. The cost to build keeps declining, AI keeps opening new niches, and the minimum viable scale for a sustainable martech business keeps shrinking. Whether that translates into net

growth depends on how fast new niches open relative to how fast old ones close.

Our bet: 2026 is the pause, not the peak.

Martech Subcategories in 2026 by Net Change (#1-#24)



Source: MartechMap Research, chiefmartec & MartechTribe ranked by number of tools removed

Martech Subcategories in 2026 by Net Change (#25–#48)

Net Growth	Martech Subcategory	Removed Tools	Added Tools
-10	Advocacy Loyalty & Referrals	-22	12
18	Marketing Automation & Campaign/Lead Management	-21	39
10	Product Management	-21	31
-11	Retail, Proximity & IoT	-20	9
-1	Influencers	-19	18
17	Interactive Content	-17	34
6	Marketing Analytics Performance & Attribution	-16	22
-4	Talent Management	-16	12
15	Mobile & Web Analytics	-13	28
-10	Print	-13	3
-4	Search & Social Advertising	-12	8
0	Affiliate Marketing & Management	-11	11
-1	Customer Data Platform	-11	10
-6	Channel Partner & Local Marketing	-10	4
6	Governance Compliance & Privacy	-10	16
-2	Dashboards & Data Visualization	-9	7
18	Call Analytics & Management	-8	26
-6	Video Advertising	-8	2
-6	Vendor Analysis & Management	-7	1
-2	Native/Content Advertising	-5	3
-2	PR	-5	3
-1	Budgeting & Finance	-4	3
-1	Agile & Lean Management	-3	2
-3	DMP	-3	0

Source: MartechMap Research, chiefmartec & MartechTribe ranked by number of tools removed

Where Marketers Are Flying with AI

In February 2026, we surveyed 208 marketing and marketing operations leaders on how their organizations are using AI across approximately 70 specific use cases. The respondent base skews toward the technically sophisticated end of the market — technology companies (36%) and professional services firms (21%) were the top industries, and 40% of respondents hold VP-level titles or above.

Given our audience, we think this reflects the leading edge more than the median. Treat the adoption rates here as where top-quartile marketing organizations are today, not where the average company is. 61% of our respondents came from pure B2B organizations and 39% from B2C and B2B+B2C companies.

For each use case, we asked respondents to report if they were using AI embedded in an existing SaaS tool, a new AI-native tool, a homegrown solution they built themselves, or none of the above. We designed it as a multi-select, because our hypothesis going in was that many organizations are doing more than one of these simultaneously.

The Through-Lines

We grouped the use cases to align with the six top-level categories of the martech landscape — Advertising & Promotions, Content & Experience, Social & Relationships, Commerce & Sales, Data, and Management. Each has its own dynamics, covered in the sections that follow. But several patterns showed up so consistently across all six that they're worth calling out before diving in.

First, the altitude view. Across every category, AI adoption jumped meaningfully from 2024 to 2026:

Category	2024	2026	Change
Advertising & Promotions	30%	50%	+20pp
Content & Experience	79%	89%	+10pp
Social & Relationships	33%	49%	+16pp
Commerce & Sales	28%	49%	+21pp
Data	61%	75%	+14pp
Management	58%	72%	+14pp

pp = percentage points

Nothing declined. Commerce & Sales nearly doubled. Content — already the furthest along — kept climbing. Two years ago, most of this was pilot-phase experimentation. The 2026 numbers reflect AI embedded in how work actually gets done.

What's less visible in the category totals — but shows up clearly in the use-case data — is *how* that adoption is happening. Several patterns repeat across every category, and they're consistent enough to be worth naming.

Build vs. Buy Is the Wrong Question

For the top use cases in this survey, total selections routinely exceed the respondent count — because respondents are checking multiple boxes. Organizations aren't choosing between building and buying. They're doing both, often for the same use case.²⁰

This is rational, not chaotic. Different tools serve different slices of the same use case — a team might use their MAP for baseline targeting, an AI-native tool for specialized segments, and a custom model on proprietary first-party data, all under the heading of “audience targeting.” On top of that, nobody wants to go all-in on a vendor who might be leapfrogged next quarter, and

²⁰ [Both Sides Now](#) by Joni Mitchell. “I’ve looked at clouds from both sides now...”

the marginal cost of running parallel solutions has collapsed with usage-based pricing and free tiers. Hedging is cheap. Conviction is expensive.

The implication: stack strategy has shifted from “standardize on a suite” to “orchestrate a portfolio of capabilities.” The build-vs.-buy debate that structured martech procurement for twenty years is a false binary. The answer, more often than not, is yes.²¹

Key Highlights & Insights
from Our Sponsor

moengage

“On the openness side, every SaaS platform should be built with composability and integrations across different data stacks and enterprise systems. The more APIs you expose, the more people can build agents themselves.”

– Raviteja Dodda

The Stack Is Stratifying, Not Consolidating

The long-running martech debate — platform consolidation versus best-of-breed fragmentation — has a 2026 answer: neither. The stack is stratifying into layers with different competitive physics, and the fault line runs between creation and orchestration.

AI-native tools are winning the creation layer: copy ideation, content strategy, pitch decks, visual production, competitive intelligence — tasks where the primary input is a prompt and some brand context, and where model quality is the product.

Incumbent SaaS tools hold the orchestration layer: lead scoring, email deliverability, pipeline management, dashboards and BI — tasks where the data already lives

Key Highlights & Insights
from Our Sponsor

 **hightouch**

“Over time, intelligence workflows will move toward upstream platforms like composable CDPs and agentic marketing platforms, because you want intelligence and context around your business to be centralized.”

– Tejas Manohar

²¹ This is also the correct answer to “should we consolidate our stack?” and “should we experiment with new tools?” Also known as the martech buyer’s paradox.

inside Salesforce, HubSpot, or Tableau, and where the incumbent simply turns on AI features within the tool that owns the workflow.

The variable that best explains this split is data gravity. Where data lives in systems of record, incumbents hold. Where a task can run on a lighter integration footprint, AI-native competes. A third layer — agentic AI that takes autonomous action — is still being negotiated.

B2B Leads on Breadth; B2C Builds for Depth

Conventional wisdom held that B2C leads technology adoption: mobile-first, social-first, personalization-first. For AI, the data inverts this. B2B shows higher adoption across more use cases and more categories, with consistently lower non-adoption rates. In Commerce & Sales, B2B non-adoption for top sales use cases sits in the single digits; B2C runs 2–3x higher.

Three things plausibly gave B2B this head start. B2B marketing teams are chronically understaffed relative to their content and operational demands. One hypothesis is that AI wasn't a strategic innovation for them so much as a relief valve, adopted out of necessity rather than vision. They also had better-prepared infrastructure: a decade of investment in CRM, MAP, CDP, and revenue intelligence created natural docking stations for AI capabilities.

But when B2C companies do adopt AI, they build deeper. The reason we believe: for B2C, the customer-facing AI output is the brand experience. Off-the-shelf tools get you 80% of the way there. The remaining 20% — brand voice calibration, proprietary guardrails, custom data integration — is where differentiation lives. That 20% doesn't get outsourced.

The Brand Surface Paradox: You can predict B2C build intensity by asking one question: *Can a bad AI output become a screenshot?* Brand Safety (high build), Customer Service Chatbots (high build), Agentic Community Management (near-zero adoption). The higher the public consequence of failure, the more B2C builds custom — or doesn't build at all.

The Trust Ladder

Across all six categories, AI adoption follows the same sequence: analytical AI first, generative AI second, autonomous AI a distant third.

The pattern reflects a trust gradient, not a capability gradient. Each step up requires more organizational confidence: "AI that tells me what happened" is low-stakes; "AI that creates something for me to review" requires editorial trust; "AI that acts with my budget or my brand" requires institutional trust.

The models are capable of more than most organizations are ready to hand them. The bottleneck is governance infrastructure, org design, and executive air cover, not the technology.

Key Highlights & Insights
from Our Sponsor



GrowthLoop

"Decisioning is crucial for helping enterprises bridge from where they are today into a more automated, agentic future with an infinite workforce — but governed and with minimized risk."

– Anthony Rotio

This is why transcription tools matter beyond their face value. They're the on-ramp. Every quarter a team finds AI's meeting notes reliable, the trust ledger grows a little. The companies furthest along on the lower rungs today will be the first to adopt more autonomous operations. Not because they're bolder, but because they've accumulated more internal evidence that their AI does what it's supposed to do.

One calibration worth noting: some decisioning AI is likely further ahead than survey responses capture. When a marketer runs Google Performance Max or Meta Advantage+, they're using algorithmic decisioning AI — but may report it as "using the platform," not "using AI." Visible AI gets counted; invisible AI gets consumed without the label. The apparent lag in bid optimization and dynamic creative may partly be a measurement artifact, which is itself an interesting finding. Marketing teams may have a significant blind spot about the AI they're already running.

The Governance Gap

In every category where AI production has hit high adoption, the governance layer is lagging badly behind.

AI copy production sits at 91% adoption in Content & Experience; Content Authenticity & AI Detection sits at 37%, with 103 out of 163 respondents doing nothing to verify the AI content they're generating at scale.

In Data, coding and automation use cases run at 75–83% adoption while Data Compliance & Governance sits at 50%, Data Lineage at 49%, and Customer Privacy & Consent Management at 47%.

In Social, companies are experimenting with agentic engagement before governance frameworks for those agents exist in any robust form.

The incentive structure explains why governance keeps losing the budget fight: there's no immediate ROI story, nobody gets promoted for implementing AI metadata cataloging, and the value of governance is avoiding a crisis that hasn't happened yet — which is perpetually next quarter's problem.

But the calculus is shifting. The EU AI Act is in effect. FTC scrutiny on AI-generated content is increasing. Platform disclosure requirements are expanding. The external forcing functions that budget conversations can ignore are getting harder to ignore.

Key Highlights & Insights
from Our Sponsor



"Marketers think they can just vibe create an email in ChatGPT and throw it over the wall, or stick it into their marketing automation platform, and it will just work."

– Brendan Farnand

Key Highlights & Insights
from Our Sponsor



"Data is siloed, but functional areas are also siloed. Without a governance board, people make decisions in a vacuum, and those tactical decisions may not ladder up to your overall strategy."

– Tara DeZao

The more acute risk is upstream. In Data specifically, governance failures don't stay contained — weak lineage and privacy controls infect every downstream agent that depends on that data.²² The governance gap in data is the governance gap in everything built on top of it. Organizations that close this gap now are building defensible infrastructure. The ones deferring it are accumulating liability.

Everyone's learned to make the sausage with AI. Almost nobody's bought a labeling machine.²³

RAG Everywhere — The Hidden Connective Tissue

The AI use case landscape looks fragmented on the surface. Customer service chatbots, sales enablement Q&A, knowledge management, and data analytics live in different departments, run on different platforms, and get measured by different teams.

Underneath, they're solving the same problem: retrieve the right proprietary context, then generate an accurate response.

This shows up in the data as a recurring structural signature. The use cases with the highest concurrent buy-and-build activity — where organizations are simultaneously deploying commercial tools and rolling their own — are Knowledge & Documentation, Chat with Data & Insights, Customer Service Knowledge Base, and Sales Enablement Q&A. Different departments, different names, same architecture: retrieval + permissions + evaluation + human-in-the-loop + logging.²⁴

Key Highlights & Insights
from Our Sponsor



"With a RAG solution, you actually have governance and control that you didn't have before. You can bring in external content that validates your conversation or story, but you control that narrative."

– Sara Faatz

²² *Chain of Fools* by Aretha Franklin.

²³ Apologies to vegetarian readers for the sausage metaphor. The tofu equivalent would be: everyone's pressing the tofu, nobody's checking the expiration date. Actually, that works too.

The vendor landscape looks fragmented because the use cases wear different departmental costumes. The capability stack underneath is converging.

The strategic implication is direct: LLMs commoditized faster than enterprise context did. The models are increasingly capable, available, and cheap. What isn't commoditized is the proprietary context wrapped around them: customer data, internal knowledge, brand voice, compliance constraints.

Competitive advantage isn't around "which model do you use" but "how good is your context plumbing."

Most organizations are discovering this empirically, one department-specific RAG tool at a time, before they have vocabulary to describe what they're building. The vendor that cracks generalized context infrastructure — retrieval across CRM, knowledge base, data warehouse, and customer profile in a single coherent layer — captures the foundational platform position. That's still an open race.

Key Highlights & Insights
from Our Sponsor

 **hightouch**

"Creative data — images, videos, email assets in Figma, Facebook, or DAMs — rarely lives in a company's warehouse. AI actually requires even more data sources than traditional personalized marketing."

– Tejas Manohar

²⁴ *Come Together* by The Beatles. Originally written for Timothy Leary's political campaign, then repurposed into something unrecognizable. Much like RAG, which started as a research technique and became enterprise plumbing.

Advertising & Promotions

Which, if any, of these functions are you using AI for?

Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Contextual Advertising		42	25	32	23	122
Audience Targeting & Segmentation		47	30	30	13	120
Ad Copy & Asset Ideation & Production		31	43	36	9	119
Campaign & Media Management		43	26	35	13	117
Video Ad Production		27	35	25	27	114
Dynamic Creative Optimization (DCO)		28	20	23	37	108
Budget & Bid Optimization		29	17	22	39	107
Brand Safety		22	13	30	40	105
answered						100
skipped						108

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option


Of all six categories, Advertising & Promotions produces the most counterintuitive finding in the entire survey: B2B is the AI advertising powerhouse, and B2C is the conservative one.

Conventional wisdom runs the other way. B2C brands, with their massive impression volumes and constant creative permutation needs, should be the AI-hungry segment. In B2B, AI for Ad Copy & Asset Ideation and Audience Targeting have effectively universal adoption. In B2C, non-adoption for those same use cases runs meaningfully higher. The bubble chart makes this visual: B2B bubbles are larger and pushed further right into the high-adoption quadrants; the B2C panel is smaller and more compressed.

The explanation connects to the through-line on B2B urgency. B2B advertising teams are small, budgets are scrutinized, and there's no 12-person creative studio

to fall back on. AI isn't a strategic innovation for them — it's table stakes for competing in channels built for B2C scale.

Build vs. Buy: Advertising & Promotions



Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

How to read the Build vs. Buy charts

Each bubble represents one AI use case. The charts appear in pairs — B2B on the left, B2C / B2B+B2C on the right — with identical axes so you can compare directly.

Position: The further right a bubble sits, the more companies are buying their way into that capability (via SaaS or AI-native tools). The higher it sits, the more companies are

building their own solution. A bubble in the upper right means companies are doing both simultaneously.

Size: *Bigger bubbles indicate higher overall adoption — fewer respondents said “not using AI / not doing this.” Smaller bubbles mean most respondents are still on the sidelines.*

Color: *Color reflects the mix within the “buy” activity. Blue means buying skews toward established SaaS platforms that have added AI features. Orange/red means buying skews toward newer AI-native tools. Green is a roughly even mix of both.*

A bubble can be large and blue (widely adopted, mostly through incumbent SaaS), small and orange (few companies doing it, but those that are favor AI-native tools), or anywhere in between. Size and color are independent signals.

Creative up, visible decisioning lagging. Ad Copy & Asset Ideation, Audience Targeting, and Campaign & Media Management sit in the high-adoption cluster. Budget & Bid Optimization and Dynamic Creative Optimization sit in the nascent zone. This is the Trust Ladder at play. However, as we noted earlier, respondents may have been unintentionally underreporting their use of AI for those use cases embedded within the Google and Meta platforms.

Brand Safety: the Brand Surface Paradox in its purest form. Brand Safety is the quietest but most structurally interesting data point in this category. Overall, 30 respondents built their own solution — more than SaaS (22) or AI-native (13). It was the most lopsided build preference of any use case in our entire survey. In B2C, that build skew is even more pronounced. Major consumer brands desire nuanced, brand-specific classifiers that encode their particular values, risk tolerances, and contextual rules.

The “all of the above” pattern in sharp relief. The top B2B use cases show total selections well above respondent counts, with Audience Targeting and Ad Copy each pulling selections across all three approaches simultaneously. These organizations aren’t choosing a single path. They’re running a portfolio: baseline targeting through the MAP, specialized segments through an AI-native tool, proprietary models on first-party data. This is rational hedging, not indecision, and it’s most visible here because advertising has the longest history of multi-vendor complexity.

The punchline for the category: in B2B advertising, the question of whether to use AI has been settled. The active question is how many parallel AI approaches to run and how to make them coherent. In B2C, AI is abundant inside existing platforms — but the moment it touches the brand surface directly, the build impulse kicks in hard.

Content & Experience

Which, if any, of these functions are you using AI for?

Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Copy Ideation & Production		79	85	41	14	219
Visual Ideation & Production (Image/Video)		68	71	27	36	202
SEO (AEO, AIO, GEO, etc.)		64	54	33	48	199
Content & Experience Personalization		70	54	32	39	195
Content Repurposing		45	48	35	65	193
Website & Landing Page Generation		52	40	27	71	190
Content Performance Optimization & Testing		64	42	27	56	189
Content Briefing / Strategy Outlines		50	60	37	42	189
Content Translation & Localization (Copy, Visual or Audio)		58	44	27	59	188
Lead Scoring & Routing		67	23	28	68	186
Brand Voice & Style Enforcement		45	41	40	60	186
Email Analysis & Deliverability		67	35	22	61	185
Audio / Podcasts Ideation & Production		35	33	18	93	179
App & Interactive Content Creation (Vibe Coding)		34	37	18	85	174
Content Authenticity & AI Detection		26	23	19	103	171
answered						181
skipped						27

Source: AI Use Cases in Marketing 2026 Survey, chiefmartec & MartechTribe
n=208; respondents could select more than one option

Content & Experience is the category where AI has moved fastest from experimental to operational — and where the stratification pattern from the through-lines is most legible in the data.

Copy ideation: the first use case to reach near-universality. Only 14 out of 163 respondents aren't using AI for Copy Ideation & Production. The decision has been made. More telling is the how: AI-native tools (85) outpace existing SaaS (79), with 41 respondents also building their own. For the use case where marketers have the longest experience with AI, bolted-on copilots in existing platforms aren't winning.²⁵ Marketers have tried the built-in features and are reaching for purpose-built tools anyway.

Creation vs. orchestration — the stack stratification in plain view. The color gradient in the bubble chart tells the story directly. Creative and generative use cases skew orange/red: Copy Ideation (AI-native 85, SaaS 79), Content Briefing & Strategy (AI-native 60, SaaS 50), Visual Ideation & Production (AI-native 71, SaaS 68). Workflow and orchestration use cases skew blue: Lead Scoring & Routing (SaaS 67, AI-native 23), Email Analysis & Deliverability (SaaS 67, AI-native 35), Content Personalization (SaaS 70, AI-native 54).


The variable that explains it is the same one from the through-lines: data gravity. You can't score a lead in a standalone AI tool that doesn't know your pipeline. You can ideate a content strategy with nothing but a prompt and some brand context.

B2C builds where B2B buys. B2B adoption leads across the board in this category, with larger bubbles, further right, and more use cases in the high-adoption quadrants. But the B2C story isn't absence; it's selectivity. For brand-critical use cases, B2C build activity proportionally matches or exceeds AI-native buying. Brand Voice & Style Enforcement has the most even split of any use case in the category across all three approaches (SaaS 45, AI-native 41, Build 40) — a signal that no vendor has yet cracked this problem end-to-end. When the AI output has to sound like your brand and not a generic content generator, off-the-shelf gets you only so far.

Vibe coding for marketing is real. App & Interactive Content Creation sits at 85 non-adopters — still nascent, but less nascent than it looks. Among those doing it, AI-native tools (37) lead SaaS (34), with 18 building their own. This is the early signal that the boundary between marketing and engineering is softening. When marketers can spin up interactive calculators, assessments, or configurators without a dev queue, it changes who owns the experience roadmap.

²⁵ The sparkle emoji (✨) became the universal symbol for "we added AI to this feature." Marketers, it turns out, want more than sparkle. Playing against type.

Build vs. Buy: Content & Experience



Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe

n=208; respondents could select more than one option

The governance gap: loudest here. Content Authenticity & AI Detection has 103 non-adopters, the single highest “not using” number in the entire survey. Meanwhile, AI copy production sits at 91% adoption. The incentive structures don’t align yet. There’s no widely-enforced penalty for undetected AI content in most marketing contexts

— but platform disclosure requirements and regulatory scrutiny are building. This is the governance gap from the through-lines at its most stark.

The category summary: Content & Experience is where the *how* of AI adoption is most visible. Creation going AI-native, orchestration staying with platforms, B2C building for brand depth, and governance lagging dangerously behind production. Every other category is following some version of this same playbook.

Social & Relationships


Which, if any, of these functions are you using AI for?

Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Social Content Ideation & Production		29	28	17	25	99
Customer Service Agents (Autonomous Triage & Resolution)		25	13	18	39	95
Customer Service Chatbots (Information Only)		35	19	24	16	94
Social Publishing & Scheduling		37	13	12	31	93
Social Insights & Reporting		36	16	16	25	93
Social Listening, Analysis & Management		35	13	16	28	92
Customer Service Knowledge Base Production & Maintenance		24	13	17	34	88
Agentic Social Engagement		9	14	12	52	87
Customer Service on Social Media Channels		19	12	10	46	87
Agentic Community Management		7	12	9	58	86
Community & Review Analysis		19	13	14	39	85
Influencer Discovery & Management		18	10	10	46	84
answered						91
skipped						117

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

Social & Relationships splits cleanly into two stories depending on which side of the customer interaction you're on. For listening, publishing, and reporting — the observational layer — adoption is solid and SaaS dominates. For anything that involves the brand speaking or acting autonomously, the picture changes sharply.

Build vs. Buy: Social & Relationships



- | | |
|--|--|
| 1 Social Content Ideation & Production | 7 Customer Service Chatbots (Information Only) |
| 2 Customer Service Agents (Autonomous Triage & Resolution) | 8 Customer Service KB Production & Maintenance |
| 3 Customer Service on Social Media Channels | 9 Agentic Social Engagement |
| 4 Social Publishing & Scheduling | 10 Agentic Community Management |
| 5 Social Listening, Analysis & Management | 11 Community & Review Analysis |
| 6 Social Insights & Reporting | 12 Influencer Discovery & Management |

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

Customer Service Chatbots: the category's clearest B2C signal. With only 16 overall non-adopters, Customer Service Chatbots is one of the most widely deployed use cases in this category. But it's the B2C build pattern that stands out. B2C companies build custom chatbot solutions at roughly double the rate of B2B — and their non-adoption rate is near zero.

When a consumer interacts with your chatbot, that conversation is the brand. The stakes of generic tone, wrong escalation logic, or a hallucinated response are immediate and public. Off-the-shelf tools can handle the routine; the proprietary order systems, account context, and brand voice calibration that make the difference often require some building.

B2B social content: where AI-native has its clearest win in the category.

Social Content Ideation & Production is the one use case in Social where AI-native tools and SaaS run essentially neck and neck overall (28 vs. 29). In B2B specifically, AI-native edges ahead.

The reason is the content itself: B2B social isn't high-volume scheduling — it's thought leadership, technical narratives, and ABM-aligned messaging that requires domain fluency. The AI features embedded in publishing platforms were built for reach and scheduling, not for the specificity of B2B content strategy. B2B marketers are using AI-native tools upstream in the creative process, then pushing through their existing distribution stack. The creative workflow has quietly decoupled from the publishing workflow.

The agentic floor: high non-adoption, strong signals among early movers.

Agentic Community Management (58 non-adopters) and Agentic Social Engagement (52) have the highest abstention rates in the category — and among the highest in the entire survey.

Three specific factors are gating adoption, not abstract trust concerns. First, downside variance: a human community manager who posts something tone-deaf is a recoverable problem; an autonomous agent doing it across thousands of interactions simultaneously is a different category of incident. Second, screenshot risk: social is the highest-visibility brand surface in the stack, and a single bad autonomous response can become a PR crisis before any human sees it. Third, unclear accountability: most organizations haven't established who owns it when an agent gets it wrong, which makes sign-off hard to obtain.

Among the small cohort moving here, AI-native tools and custom builds lead, because credible agentic capabilities barely exist yet in established social platforms. The incumbents are shipping “AI-assisted” features while the earliest movers are already reaching past them for AI-autonomous capabilities. That gap is an opening, but closing it requires solving the accountability and guardrail questions, not just the model ones.

The observational layer: mature, SaaS-dominant, less interesting. Social Publishing & Scheduling (SaaS 37), Social Insights & Reporting (SaaS 36), Social Listening (SaaS 35) — all solid adoption, all firmly in incumbent SaaS territory. These are well-solved problems with entrenched vendors, and AI is being absorbed as a feature rather than driving platform change. Worth noting primarily because they form the trust base that will eventually make agentic social more palatable organizationally.

The category in one frame: B2C has gone deep on customer service AI, building custom where the brand surface is highest-risk, and is largely bypassing AI-native tools in favor of SaaS or build. B2B is buying AI-native for content creation and holding back on everything autonomous. Both segments are watching the agentic use cases from a cautious distance — but the early movers are already bypassing the platforms to get there.

Commerce & Sales

Commerce & Sales is where the B2B/B2C divide is sharpest — not just in adoption levels, but in the underlying strategic logic.

B2B has gone all-in. Sales Support for Meeting Notes (13 non-adopters), Sales Content for Pitch Decks (24), and Lead Sourcing & Outreach (31) lead the category in adoption. In B2B, the top sales use cases show near-universal uptake. The ROI math here is unusually tractable: in high-value B2B sales cycles, a single accelerated deal can pay for a year of tooling. Teams were also already instrumented — CRMs, sales engagement platforms, and revenue intelligence tools were waiting for AI to plug in.

B2C is building where it matters most. The B2B/B2C gap runs in the expected direction on overall adoption, but B2C’s build behavior is the more strategically

interesting signal. For use cases closest to the customer experience — agentic concierges, lead sourcing, pitch content — B2C build activity runs roughly equal to buying, sometimes ahead.

Which, if any, of these functions are you using AI for?


Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Sales Content Ideation & Production (Pitch Decks)		26	33	23	24	106
Sales Content Ideation & Production (Email)		30	27	21	27	105
Lead Sourcing, Prospecting & Outreach		31	23	18	31	103
Sales Content Ideation & Production (Quotes & Proposals)		20	26	19	38	103
Sales Support (Meeting Notes & Follow-Up)		42	30	17	13	102
Sales Support (Dynamic Playbook / Next-Best Action)		23	17	17	40	97
Sales Management - Pipeline & Forecast Optimization		27	13	17	38	95
Sales Content Ideation & Production (Other)		22	24	15	34	95
Sales Enablement Q&A ("Chat with Collateral")		24	17	12	42	95
Offer & Promotion Orchestration		17	9	12	55	93
Agentic Shopper Concierges / Sales Assistants		9	10	13	59	91
Webshop / Product & Merchandising Intelligence		12	10	10	59	91
Pricing & Revenue Optimization		11	16	11	53	91
Webshop / Product Catalog Management		14	9	9	58	90
Webshop / Commerce Personalization & Optimization		14	9	10	56	89
answered						89
skipped						119

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

B2B sales workflows are standardized enough that vendors can serve many companies with a common product. B2C customer-facing commerce is the

brand; a luxury fashion concierge and a home improvement assistant can't share a template.²⁶ When the AI interaction encodes your product catalog, margin rules, brand voice, and customer history, building starts to look as attractive as buying — because few off-the-shelf solutions go that deep.

Build vs. Buy: Commerce & Sales



- | | |
|---|--|
| 1 Sales Content (Pitch Decks) | 9 Sales Content (Other) |
| 2 Sales Support (Meeting Notes & Follow-Up) | 10 Sales Enablement Q&A (Chat with Collateral) |
| 3 Lead Sourcing, Prospecting & Outreach | 11 Pricing & Revenue Optimization |
| 4 Sales Content (Email) | 12 Webshop / Product & Merch Intelligence |
| 5 Sales Content (Quotes & Proposals) | 13 Agentic Shopper Concierges / Sales Assistants |
| 6 Sales Support (Dynamic Playbook) | 14 Webshop / Product Catalog Management |
| 7 Offer & Promotion Orchestration | 15 Webshop / Commerce Personalization |
| 8 Sales Mgmt - Pipeline & Forecast Opt. | |

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

²⁶ Though we'd pay good money to see an AI trained on Hermès try to recommend grout colors.

Content creation vs. workflow: the stratification pattern again. Within the category, the bubble chart color gradient tracks the creation/orchestration split exactly. Sales Content for Pitch Decks (AI-native 33 vs. SaaS 26), Quotes & Proposals (AI-native 26 vs. SaaS 20), and Pricing & Revenue Optimization (AI-native 16 vs. SaaS 11) all skew toward AI-native — blank-page tasks where the legacy tools were never AI-first.

Meeting Notes & Follow-Up (SaaS 42 vs. AI-native 30), Pipeline & Forecast Optimization (SaaS 27 vs. AI-native 13), and Sales Enablement Q&A (SaaS 24 vs. AI-native 17) all skew SaaS — workflow tasks where the data lives in Salesforce, HubSpot, or Gong, and where incumbents add AI features within the system of record.

Offer & Promotion Orchestration: a structural white space. With 55 non-adopters and zero AI-native adoption overall, this is the most underserved use case in the category. The reason is requirements depth, not lack of interest. Promotion orchestration has to be constraints-aware: inventory availability, margin thresholds, customer eligibility rules, legal compliance, and pricing logic all have to be respected in real time.

An AI-native point solution entering this space can't just generate good offers — it has to ingest and honor all of those constraints, which requires deep integration with commerce platforms and order management systems that took years to build. The 12 respondents building their own are doing exactly that: assembling custom integrations because no standalone vendor has solved the full stack. The opening here belongs to commerce platforms that can embed promotion intelligence natively, rather than asking buyers to bridge the gap themselves.

Agentic Shopper Concierges: the long game. With 59 non-adopters, Agentic Shopper Concierges is the least-adopted use case in the category, but among the small cohort that is investing here, build leads over buying. The companies investing here aren't buying off the shelf; they're building a proprietary AI embodiment of their brand's customer relationship. The B2B equivalent — an agent that navigates multi-stakeholder procurement, surfaces relevant collateral, and handles pricing queries — is nascent in AI-agentic form, even if digital deal rooms have been working toward this problem for years. The agentic version remains largely undeveloped in the data, which may prove to be a significant opportunity.

Data

Which, if any, of these functions are you using AI for?

Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Coding & Development		50	55	27	28	160
Chat with Data & Insights		51	44	36	25	156
Data Integration		49	37	34	35	155
Data Automation (Data Integration & ETL)		53	34	29	39	155
Data Automation (Data Cleaning)		53	33	32	35	153
Dashboards & BI		61	26	34	31	152
Audience & Identity Resolution		48	21	25	57	151
External Data Sourcing (Customer Intelligence)		41	36	30	40	147
Data Quality & Observability		39	24	26	57	146
External Data Sourcing (Competitive Intelligence)		32	39	30	43	144
Data / Metadata / Lineage & Cataloging		29	21	21	73	144
Data Compliance, Governance & Risk Management		31	20	18	70	139
Customer Privacy & Consent Management		31	16	16	71	134
answered						140
skipped						68

Source: AI Use Cases in Marketing 2026 Survey, chiefmartec & MartechTribe
n=208; respondents could select more than one option

The Data category has the highest overall adoption of any category after Content — and the most internal contradictions. The operational layer is well ahead; the governance layer is badly behind. B2B leads broadly; B2C builds more deliberately where it does adopt. And one use case quietly breaks every pattern in the survey.

Coding & Development: the exception that proves the rule.²⁷ Across the entire survey, data gravity reliably predicts that incumbent SaaS wins orchestration-layer use cases... except for Coding, which is the exception.

With 55 AI-native adoptions vs. 50 SaaS and only 27 builds, it's one of the few high-adoption use cases in the dataset where AI-native leads SaaS and build stays low. The reason: coding assistance was born AI-native. Claude Code, Codex, Cursor, and their peers didn't retrofit AI into an existing workflow — they rebuilt the workflow around AI from the start. And because code is a relatively universal substrate, the commercial tools work well enough that proprietary builds add little differentiation. This is a preview of what happens when an AI-native tool gets product-market fit exactly right before legacy alternatives entrench.

The operational layer: well-adopted, SaaS-dominant. Chat with Data & Insights (25 non-adopters), Data Integration (35), Data Automation for ETL (39) and Cleaning (35), and Dashboards & BI (31) all show solid adoption with SaaS leading. These are mature use case areas with entrenched platforms. Tableau, Looker, Power BI have years of embedded workflows and governance policies behind them. AI is being absorbed as a feature, not driving platform replacement. The data gravity argument is nowhere more visible than here.


Competitive Intelligence: where AI-native wins in data. External Data Sourcing (Competitive Intelligence) is the clearest AI-native success story outside of Coding, with AI-native (39) leading SaaS (32) overall. Competitive intelligence was a fragmented, workflow-light category. AI-native tools could offer step-change capability without asking buyers to migrate anything. The contrast with Dashboards & BI makes the data gravity principle concrete: AI-native disruption runs inversely to the depth of existing workflow lock-in.

B2C builds more of its data layer. Where B2B buys data capabilities — Chat with Data strongly buy-first, Dashboards & BI overwhelmingly SaaS — B2C shows a meaningfully stronger build orientation across the same use cases. The reason isn't technical sophistication for its own sake. B2C companies at consumer scale have proprietary behavioral and transactional data that off-the-shelf tools often weren't designed to query.

²⁷ The original meaning of "proves" here is "tests" — same root as "proving ground." So the phrase actually means the exception *tests* the rule, not *confirms* it. Thank you for attending our etymology seminar.

But customization depth is only part of it. When a query can touch sensitive revenue or customer data, off-the-shelf tools often can't satisfy the permissions, governance, and latency requirements that enterprises need to control the full chain. A generic LLM wrapper on a dashboard isn't useful when your differentiation lives in 50 million customer records, real-time event streams, and access policies that no SaaS vendor designed for your specific stack.

Build vs. Buy: Data



Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
 n=208; respondents could select more than one option

The governance cluster: the most actionable finding in the category.

Data/Metadata Lineage & Cataloging (73 non-adopters), Data Compliance & Governance (70), and Customer Privacy & Consent Management (71) form a distinct laggard cluster — the three lowest-adoption use cases in the category, sitting well below the operational layer running at 75–83% adoption. As covered in the through-lines, governance failures in data aren't self-contained. They propagate downstream into every agent and workflow built on top of that data. Organizations deploying AI across their data stack while leaving lineage and compliance underdeveloped are building on a foundation they can't audit or defend. The regulatory clock on this is running.

Key Highlights & Insights
from Our Sponsor



"We're seeing a lot of marketing teams drowning in data rather than acting on insight, because they don't have the right tools for the job."

– Jonathan Moran

Audience & Identity Resolution: the polarized case. With 57 non-adopters despite solid SaaS adoption (48), this use case is oddly bimodal. Organizations either have it through existing CDP, CRM, or ABM platform integrations, or they don't have it at all. A meaningful cohort still treats identity resolution as a problem owned elsewhere in the org — sales ops, IT, data engineering — rather than a marketing function. As AI gets better at fuzzy entity matching and cross-source reconciliation, that structural hesitation may prove costly.

Management

Transcription: the first use case to run out of holdouts. Only 6 respondents out of 125 aren't using AI for Transcription, Notes & Summaries — the lowest non-adoption count in the entire survey,²⁸ with Copy Ideation & Production in Content close behind at 14 out of 163. In B2B specifically, it's effectively zero.

The use case followed a clean adoption path: well-defined task, clear quality bar, low failure risk, and seamless embedding into tools organizations already

²⁸ If you're one of those six and you're reading this, we respect your commitment. Also, there's a free tier.

used — Zoom, Teams, Gong. SaaS dominates (83 vs. 53 AI-native), which tells you this was primarily a feature-absorption story: the platforms people were already in just added it. That combination of low stakes, obvious value, and zero switching cost has been the template for rapid AI adoption.

Low-risk, feature-absorbed use cases like transcription matter not because they are glamorous, but because they build organizational trust and normalize AI as infrastructure.


Which, if any, of these functions are you using AI for?

Categories	Answer Choice	AI in an existing SaaS	AI via a new AI-native tool	Built our own AI solution	Not using AI / not doing this	Response Total
Knowledge & Documentation		65	49	29	21	164
Transcription / Notes / Summaries		83	53	19	6	161
Internal Copy Writing & Ideation		54	62	26	14	156
In-Product AI Assistants		35	19	23	66	143
Agentic Project & Task Management		33	27	20	62	142
Talent Management / Recruitment		32	18	16	67	133
Vendor Selection & Stack Management		13	19	10	88	130
answered						140
skipped						68

Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
n=208; respondents could select more than one option

Internal Copy Writing: the management category's AI-native outlier. With 62 AI-native adoptions vs. 54 SaaS, Internal Copy Writing & Ideation is the only management use case where AI-native leads. The reason is reasoning quality. Writing is open-ended. The “right answer” depends on judgment, voice, and context, not workflow completion. People are going directly to the models, and the models are winning.

Build vs. Buy: Management



Source: *AI Use Cases in Marketing 2026 Survey*, chiefmartec & MartechTribe
 n=208; respondents could select more than one option

Knowledge & Documentation: high demand, active market. With the highest buy numbers in the category (SaaS 65 + AI-native 49) and also the highest build count (29), Knowledge & Documentation sits firmly in the upper-right quadrant — simultaneously the most bought and most built use case in Management.

That combination reflects a genuine organizational priority. Tools like Glean, Notion AI, and others are making real progress on connecting AI to institutional

context. But the parallel build activity suggests that the last-mile integration — getting AI to work fluently across your specific docs, history, and tribal knowledge — still requires custom work for many organizations. This is the RAG problem from the through-lines in its most visible form.

Agentic Project Management: more experiment than deployment. With 62 non-adopters and a mix of cautious SaaS and build activity among those who are, Agentic Project Management sits in the middle of the chart as a careful experiment rather than an operational commitment. The gap between “AI summarizes my meeting” and “AI assigns tasks, tracks deadlines, and escalates blockers” is not a technology gap. It’s a trust and accountability gap. Most organizations haven’t worked out who’s responsible when an agent makes a wrong prioritization call. Until that question has an answer, autonomous project management stays in the pilot column.

Vendor Selection & Stack Management. The use case with the most non-adopters in the entire survey — 88 out of 130 — is using AI to evaluate and manage a martech stack.²⁹ SaaS management platforms and procurement tools are beginning to incorporate AI, but adoption here lags well behind every other management use case. For organizations managing complex stacks and integration decisions, this gap will become harder to justify.

²⁹ The martech industry’s practitioners are not using AI to manage their own martech. The cobbler’s children, as always, go barefoot — though in this case the cobbler also has 15,505 shoes to choose from.

Context: The Creature in the Chrysalis

Marketing has always been about context, even when we didn't call it that.

The best ad reaches the right person, with the right message, at the right moment. The best salesperson walks into a meeting knowing what the prospect actually cares about. The best customer experience feels like the company remembers you.

Marketers have been chasing context for as long as there have been markets.

What's new is that the word "context" has exploded across the AI martech landscape, attaching itself to a dizzying range of concepts — technical, strategic, operational — that all share a name but mean different things. Context windows. Context graphs. Context engineering. All swirling in the soup of marketing's AI metamorphosis.

We believe that context — in all of its contexts — is the structure forming inside the chrysalis. It's what will give marketing's AI transformation its shape, its strength, and its ability to fly.


But first, let's disambiguate all these different contexts:

Term	Meaning
Context (colloquial)	The background, circumstances, and framing that make information actionable. Without context, a data point is just a number, a message is just words, and a relationship is just an outdated record in your CRM.
Context window	The finite amount of information an AI model can hold and reason over in a single interaction. Currently ranges from tens of thousands to millions of tokens, but always bounded. A bigger window helps, but more input doesn't automatically mean better output. ³⁰

³⁰ Anthropic's Claude Sonnet 4.6, Google's Gemini 3 Flash, and OpenAI's GPT-5.4 models all sport a 1 million token context window, equivalent to about 1,500 pages of text (only slightly more than the length of this report). But a 1,500-page tome filled with irrelevant information is just largely irrelevant.

Term	Meaning
Context engineering	The discipline of assembling the right data, content, tools, and instructions for an AI system to act effectively in a specific situation. The successor to “prompt engineering.” ³¹
Context graph	A structured, living record of decision traces across entities and time — capturing not just what happened, but why. The emerging “system of record for decisions, not just objects.” ³²
Context-as-a-Service (CaaS)	The evolution of SaaS platforms from selling features and interfaces to packaging their accumulated domain expertise, governance logic, and operational intelligence as a contextual foundation that any app or agent can tap into. ³³ Also called “systems of context.” ^{34 35}
Context stack	The layered time scales of context, from durable attributes at the base (industry, role, company size) to ephemeral signals at the top (real-time intent, session behavior, in-the-moment sentiment).

These definitions range from the colloquial to the deeply technical. But they all sit inside a larger picture: the intersection of what a company is trying to achieve, what a customer actually needs, and what systems can connect between them.



³¹ The term was popularized by Andrej Karpathy, who argued that “context engineering” more accurately describes what practitioners actually do than the earlier term “prompt engineering.”

³² *AI's Trillion-Dollar Opportunity: Context Graphs*, by Jaya Gupta and Ashu Garg, Foundation Capital, December 2025.

³³ Coined by Scott in *The SaaS moats are crumbling, but the opportunity is bigger*, February 2026.

³⁴ Coined by Scott and Frans last year in *The State of Martech 2025*, May 2025.

³⁵ **Bit coining** (verb): doing a little bit in the footnotes about coining things. Cryptic currency?

Company context is the company's own universe of knowledge: its goals, strategies, brand, processes, capabilities, governance, and priorities. This is everything the company and its people know — or should know — about who they are, what they're trying to accomplish, and how they operate.

Customer context is the customer's universe: their situation, goals, intent, preferences, history, jobs-to-be-done, and moments-that-matter. This is everything you would want to know about why a specific person is engaging with you right now — and what would help and delight them.

Systems context is the technical reality: what the company's systems can actually access, connect, and deliver. You might have deep customer insights in your data warehouse and a sharp brand strategy in a Google doc. Systems context is what determines whether either of those reaches the right AI agent, the right touchpoint, or the right decision at the moment it matters.

Each circle matters on its own. But the real action happens in the overlaps.

Key Highlights & Insights
from Our Sponsor

moengage

"LLMs are excellent. But on their own, they aren't enough. You need to feed them the right context. Business context, campaign context, historical campaign data. Once you do, the output gets genuinely useful."

– Raviteja Dodda


The Spaces Between Our Contexts³⁶

The overlap between company context and customer context is where marketing creates value. In our recent [*State of Marketing Attribution 2026*](#) report, we called this discipline **value engineering**: identifying where a company's goals and a customer's needs align in ways that drive disproportionate revenue, then focusing resources there.

Value engineering reverse-engineers from outcomes — who are your most profitable customers, what do they buy, where are the margins — rather than

³⁶ [*The Space Between*](#) by Dave Matthews Band. Missing lyric: The space between your context and mine.

spreading effort evenly across channels and hoping the math works out.




It's the *who*, *what*, and *why* of marketing — positioning, segmentation, messaging, and go-to-market strategy. When done well, it concentrates resources on the customers and journeys that matter most, rather than optimizing for activity across the board.

The difficulty is that company context and customer context live in different heads. Product knows capabilities. Sales knows conversations. Finance knows margins. The customer knows their own world — their desires, motivations, hurdles, and alternatives — but they aren't necessarily willing or able to spell it out to you.

Value engineering synthesizes these perspectives into a shared picture. That's a leadership and alignment challenge before it's a technology challenge.

Which brings us to the technology challenge. Value engineering identifies *where* to focus. **Context engineering** figures out *how* to deliver.


Context engineering covers the overlap between systems and the other two circles. It's where company knowledge becomes machine-readable and customer understanding becomes actionable. The product catalog that an AI agent can query in real time. The brand voice skills that shape its tone. The governance rules that stop it from promising something you can't deliver. The customer's journey history that informs what to say next — and what not to.



All of that is context engineering — assembling the right inputs for a specific situation so that technology can act with the judgment it doesn't inherently have.

The temptation is to throw everything at the problem. But more input doesn't automatically mean better output. In fact, it can be worse (and more expensive). Context engineering is curation, not cramming.

But the real prize is where all three circles converge.



The triple intersection — where company context, customer context, and systems context all converge — is what we call **Golden Context**.³⁷

³⁷ If "golden context" reminds you of "golden record," that's not entirely a coincidence. The key difference: a golden record is static; golden context is dynamic. It evolves as customer needs do.

It's the moment when the company's goals and capabilities, the customer's needs and situation, and the systems' ability to connect and deliver all align perfectly. These are the moments where revenue is generated.


However, in most organizations today, those three circles have limited overlap — especially as a function of misaligned systems. Improving this alignment is one of the central jobs of marketing ops and martech leadership in the next 12-18 months.

Key Highlights & Insights
from Our Sponsor



"Agents need four critical components — the four C's: context, constraints, compromise, and cognizance."

– Jonathan Moran




This requires investment in value engineering to sharpen the strategic overlap between company and customer. And investment in context engineering to make that overlap accessible to every relevant AI agent, automation, and touchpoint in the stack.

Data infrastructure, integration protocols like MCP, and CaaS platforms are the practical mechanisms. But the starting point is simpler than any of those: decide what context matters most, and connect it first.

Context Has Many Layers at Many Speeds

Context is multi-layered. And each layer operates on a different clock.

Borrowing from Stewart Brand's pace layers model³⁸ — where slow-moving layers like culture and governance stabilize fast-moving ones like commerce and fashion — context in marketing operates across multiple time scales simultaneously.




The slow layers at the bottom are relatively stable. Market structure, company strategy, brand positioning, core capabilities. These form the foundation that most marketing has historically operated on: firmographics, demographics, broad segmentation. Reliable, persistent, and available to anyone with a decent CRM.

³⁸ Stewart Brand introduced the pace layers framework in *How Buildings Learn* (1994) and expanded it in *The Clock of the Long Now* (1999). His insight: in any healthy system, the fast layers innovate and the slow layers stabilize.

However, in today's environment AI is forcing even the slow layers to move at speeds they weren't designed for. That's part of why this moment feels so disorienting.

The fast layers at the top are the hardest to capture and the fastest to decay — which makes them the most underutilized, and the ripest for competitive differentiation. An intent signal from this browsing session. A question asked in a chat five minutes ago. A surge in a prospect's hiring activity that suggests a budget just opened up. This context has a half-life measured in hours, sometimes minutes.



Context reduces friction, for customers and for the company. The more context an interaction carries, the less work both sides have to do: less repeating, less guessing, less starting from scratch.

The relationship follows a convex curve. The first increments — knowing who the customer is, what segment they fall into — often eliminate the most friction. Each additional increment still helps, whether that means connecting layers that were previously siloed, enriching the layers you already have, or harnessing the faster layers at the top of the stack for the first time.

In finance terms, context is alpha.³⁹ In a world where tools, platforms, and AI capabilities are broadly available to everyone, context generates excess returns by knowing things your competitors don't, or by connecting what you know faster than they can.

But like alpha, context decays. The more granular and timely the context, the faster it loses value. That intent signal from this morning? By this afternoon the customer may have moved on, found an alternative, or simply lost interest.

This creates a context distribution problem most organizations have yet to solve.

Key Highlights & Insights
from Our Sponsor



GrowthLoop

"The agentic context graph allows you to make marketing more verifiable by running tests in simulation against all the different customer states you've seen before."

– Anthony Rotio

Consider an example with a SaaS company. A customer visits the pricing page, spends four minutes on the enterprise tier, and then opens a support ticket about API rate limits. Ideally, that context reaches the account team within minutes. But if the website visit sits in one analytics tool, the support ticket sits in another, and the account exec works in another, the customer's context goes unrecognized and unserved.

Or consider a B2C retailer. A customer tells an AI shopping assistant they're training for a half-marathon and need running shoes under \$150. That conversation lives in the chat platform. The customer's site activity — five visits to the running shoe category in two weeks — sits in the analytics tool. All the loyalty program knows is that they haven't made a purchase in six months. Three systems, three signals, one customer. The email that goes out that evening? A generic clearance sale. Context fail.

The question isn't, "Do you have the data?" You probably do somewhere. *The context game is a matter of connecting the dots before they decay.*

The practical response to the context distribution problem isn't to boil the ocean. Start with the context that matters most for your best customers — an

³⁹ Alpha, in investing, is the return a portfolio generates above a benchmark index. A fund that returns 12% when the market returns 10% has generated 2% alpha. (Slightly oversimplified.)

exercise in value engineering. For those high-value use cases, map the three or four highest-value signals your teams are missing, identify where they originate and where they need to arrive, and build the connective tissue between those specific points. Real-time event streaming, AI orchestration layers, and protocols like MCP are making this kind of targeted context distribution increasingly feasible, without having to rearchitect your entire stack.

Key Highlights & Insights
from Our Sponsor



"Your data freshness is measured in milliseconds because the window of opportunity to get in front of your customer is so incredibly short with how fast they move through our channels."

– Tara DeZao

Marketing has spent decades getting good at the bottom pace layers of the context stack. The competitive frontier is at the top.

Context is a Decision-Time Game

Context has no value in storage. It only creates value when it arrives at the moment of decision — when a system, an agent, or a person has to choose what to do next for a specific customer in a specific moment.

A brilliant customer insight sitting untouched in a data warehouse is not an asset. It's inventory that's spoiling.

In the previous era, applying context at decision time meant hardcoding rules into each system: *if customer = enterprise tier AND last purchase > 90 days, then send re-engagement email.*

Every decision point got its own thin slice of context, manually configured and painfully maintained. Scale that across hundreds of segments, dozens of channels, and thousands of possible states, and the whole thing turned into a combinatorial hairball. Most companies eventually surrendered to broad segments, generic messages, and the quiet resignation of "good enough."

AI agents don't eliminate that complexity. They relocate it.

Instead of hardcoding every branch of decision logic into every system, you package the context — the data, the instructions, the tools, the governance guardrails — and let the agent work through the decision at runtime.⁴⁰ The complexity moves out of brittle if/then trees in the application layer and into context that is structured well enough to be interpreted dynamically.

But that shift is what makes context engineering so important.

Context engineering is the work of making context usable at decision time. It means assembling the right data, policies, tools, permissions, and instructions so that an agent or application can act intelligently in a specific use case. Not just more data. Better-prepared context.

This matters because every company's context is different. Your products are different. Your policies are different. Your margin constraints, escalation paths, compliance requirements, territory rules, service model, and brand voice are different. So are the moments that matter most: the exceptions, opportunities, judgment calls, and edge cases where real differentiation lives.

For a long time, that uniqueness was too expensive to encode. You could blunt it with generic software and live with the mismatch, or build custom systems the hard way and pay dearly for the privilege.

That tradeoff is changing. AI is dramatically reducing the cost of building software tailored to a company's specific context. The hypertail grows because the cost of tailoring software to your context has dropped faster than the cost of living with generic software.

But that only increases the importance of context engineering.

When every team can spin up specialized apps, agents, and workflows, the real bottleneck becomes preparing the context that makes them reliable. The problem isn't building software. It's making sure that software knows enough, is allowed enough, and is constrained enough to make good decisions.

⁴⁰ Robin Williams reportedly improvised so much of the Genie's dialogue in *Aladdin* that the original script was barely recognizable by the time recording finished. The best AI agents work the same way: they can riff, adapt, and deliver something genuinely tailored to the moment — if they have the right context.

We believe this is a strategic opening for commercial martech platforms looking to escape the SaaSpocalypse.⁴¹

The winning vendors won't be the platforms that try to build every app and agent themselves. They'll be the ones that make it easier for their customers to build their own by doing the heavy lifting of context engineering within their domain: structuring the data, codifying the domain knowledge, enforcing the governance rules, and exposing it all through APIs and protocols like MCP.

This is the evolution from SaaS to CaaS (context-as-a-service).

A SaaS platform says: here's our application, use our interface, work within our workflows. A CaaS platform says: here's our contextual foundation — our domain expertise, our data models, our governance logic — build whatever you need on top of it. The application may be ours, yours, or an agent composed on-the-fly. The enduring source of value is the context layer underneath that makes all of them useful.

This reframes personalization too.

For two decades, personalization mostly meant segmentation: sort customers into buckets and give each bucket a slightly different experience. But that was always a crude answer to a more important question.

The real question was never "which segment is this customer in?" It was always "what is the best action we can take for this person right now, given our goals, their needs, and our constraints?"

Personalization, done right, is decisioning. And decisioning depends on context from all three circles — company, customer, and systems — converging at the point of action. This is what context engineering ultimately delivers: not more information, but better judgment at scale.

⁴¹ **SaaSpocalypse** (*noun*): the predicted mass extinction of SaaS companies unable to articulate a value proposition beyond features, interfaces, and data lock-in. Prognosis: exaggerated in the short term, underestimated in the long term. See also: every other technology transition in history.

The Real Butterfly Effect in Martech

We opened this chapter with a claim: context is the structure forming inside the chrysalis.

Value engineering tells you *where* the value is — which customers, which journeys, which moments matter most. Context engineering makes that value *accessible* — getting the right knowledge to the right system at the right time. Together, they are the two disciplines that expand the golden center of the Venn diagram.

Neither alone is sufficient. Value engineering without context engineering is strategy without execution. You know what to do but can't make it happen. Context engineering without value engineering is plumbing without purpose. Your systems are beautifully connected, but nobody asked whether the data flowing through them actually matters.

Key Highlights & Insights
from Our Sponsor

knak.

"Marketing ops teams that should be focused on strategy and optimization instead spend their days executing technical production tasks."

– Brendan Farnand

It's often said that you don't have an AI strategy if you don't have a data strategy.

We'd expand that: *You don't have an AI strategy if you don't have a context strategy.* AI is a commodity. Context is differentiation.

This is the real butterfly effect in martech. Small differences in context — what you know, how you structure it, where you deliver it, and which constraints shape it — can produce very large differences in outcomes.

Speaking of butterflies, the Monarch butterfly navigates 3,000 miles from Canada to a specific grove of trees in central Mexico. It does this using a multi-layered context stack of its own: a sun compass for direction, a magnetic field sensor for calibration, circadian clocks for timing, and inherited genetic memory that encodes a destination the individual butterfly has never visited.

No single layer is sufficient. The navigation works because the layers integrate.

Marketing's context stack works the same way. The slow layers provide direction. The fast layers provide precision. The systems in between connect them. And the agents at the edge — your AI-powered touchpoints — are the ones that actually make the journey, carrying the right context to the right customer at the right moment.

Content was king. Context is the Monarch.



Causal AI and the Future of Marketing



Interview with *Anthony Rotio*
Co-Founder & Co-CEO, GrowthLoop

Introduce us to GrowthLoop and your role.

GrowthLoop is an agentic, composable CDP that drives compound growth by uniting your cloud data and AI into one Compound Marketing Engine. It's a system that plugs into all your marketing data in the data warehouse and drives compounding results in the outcomes you care about most.

We started as a composable CDP and have evolved by adding AI to our product in the form of decisioning and AI Studio, which helps us drive compounding outcomes for enterprise organizations.

I've been with the company for about six years. Prior to that, I was at Anheuser-Busch InBev for seven years as a marketing director. Now I get to solve a lot of the challenges I saw there — running the retail business unit — with technologies I know and love. I'm a computer scientist by training, so I get to mix that background with marketing here at GrowthLoop where I serve as co-CEO, leading our product partnerships, working on our AI product, and helping tell our story.

Let's talk about causal AI. Most AI in martech today is essentially pattern matching on historical data — correlation. Can you paint a picture of what that looks like in practice and where it breaks down for marketers?

A lot of solutions in martech have been doing predictive AI for some time. This is where you look at observational data — what happened — and try to predict what might happen next.

In practice, this usually manifests as a propensity score: how likely is someone to convert? You might have some folks with zero or low likelihood to convert, and others with nearly 100% likelihood. That tells you what will happen if the status quo stays the same. The best you can do is say, "For customers in that 70% to 80% range, maybe we should try to push them over the edge."

But it can't answer the crucial question: What should I do? What will actually cause the change I need to happen for this customer?

Can you give a concrete marketing example where confusing correlation with causation can really lead you astray?

A common example is churn prevention with offers. You see that customers who received an offer didn't churn, so you conclude the campaign worked.

But there are two big problems. First, from a business standpoint, do you want to send offers all the time to prevent churn? Is that really the best way to maximize long-term customer value? Second, the churn prevention might have had nothing to do with the offer for specific subsets of your audience. Maybe you hit someone on their birthday when they were in a good mood, or they browsed an item the night before and your offer just brought it to mind.

If instead you could say, "This customer was in this state, we tried this intervention, and that directly caused this outcome," you could identify the key features of that state where the action drove the outcome. Then, when you see customers in a similar state, you can actually determine what needs to change to drive the business outcomes you need. That's the power of causal AI.

How would you define causal AI, and how is it fundamentally different from the predictive, correlative AI that most martech has been built on?

The key is snapshotting points in time. For every experiment you run, you capture three things: the customer state at that moment, the action you took (the campaign or intervention), and the outcome it drove across your business metrics.

This is the foundational piece of causal AI — it comes from reinforcement learning, which I studied in college. It was mind-blowing to realize how similar that is to how humans actually learn and make decisions.

For marketing, the customer state is any information you have about them at that time. The action space is all the different campaigns and messaging you can send. The metrics you care about are the “reward” for the system.

Only if you’re tracking this over time can you answer counterfactual questions. “What if” questions. Like, “What if we try this campaign with this customer? What outcome will it drive?” With causal AI, you can answer that with confidence. With predictive AI using observational data, you can’t say with certainty that a specific action will produce a guaranteed outcome. It’s a fundamentally different type of machine learning.

Let’s talk about composability. You’ve made a strong claim that composability isn’t just a nice architectural preference. It’s actually a prerequisite for causal AI decisioning to work well. Why?

State, action, and reward are the foundation of the entire system. The more data you have access to for each component, the higher resolution you’ll have and the better predictions you can make.

If you try to do decisioning in a non-composable system — like an email platform or marketing cloud with limited customer data because you have to sync it there — you lose resolution. You might think two customers are in a similar state when you only see five to thirty fields, but if you have access to hundreds of attributes in your warehouse, you get a much more nuanced view of that customer state.

This works similarly with actions. In a paid ads channel, you can only try certain interventions when you’re working solely in that channel. But a composable architecture that connects all your data, channels, and campaigns, you can say, “Maybe it’s not an ad — maybe it’s an email and SMS combination that will work.”

On the metrics (or reward) side, if you're only tracking opens and clicks in an email platform or clicks in ads, you're optimizing for vanity metrics. But if you can see customer lifetime value and the entire transaction history in your warehouse, your system makes decisions that build valuable long-term relationships, not just short-term click maximization.

This is why composability is the best architecture for this work. And in regulated industries like banking or healthcare, you definitely don't want to load all that sensitive data into third-party systems. You want the processing to happen in your warehouse, where you control it.

What is an agentic context graph, and what makes it different from the customer data models and identity graphs marketers already work with?

An agentic context graph is exactly the snapshot of data we've been describing. All these different points in time capturing your customer state, the interventions you've tried, the metrics that you care about as a business, and how those metrics moved as a result of a campaign or intervention. And it all lives in your cloud data warehouse.

It's remarkably different from just adding an identity spine to your warehouse data or connecting tables in a CDP or CRM schema. This is truly about building a foundational layer — a competitive moat — that differentiates you from competitors who might be doing predictive analytics or connecting systems with agents but aren't able to answer the counterfactual questions: "What should we do? What will work? What is guaranteed to work based on what we've seen before?"

We power this entirely through experiments you run in GrowthLoop. To build an agentic context graph, you need a composable platform that can run cross-channel experiments, snapshot that data, and build this foundational layer that powers both decisioning and the agents working on your behalf.

Is decisioning then the real heart of martech? Should marketing leaders reframe their stacks around decisioning?

To me, it's overwhelmingly obvious that agents are going to eat most of the jobs-to-be-done in martech. We're already seeing this in our own

go-to-market stack. We have agents that look across our data, HubSpot, Salesforce, Slack, and work management tools. They can assemble information about an account or prospect and decide what action to take next.

Everyone's experimenting with this — enterprises and smaller companies alike — and it's becoming easier to connect these systems and have agents do work. But the key missing piece is having this type of context graph to reference. Do they actually know that what they're deciding is grounded in truth?

Decisioning is crucial for helping enterprises bridge from where they are today into a more automated, agentic future with an infinite workforce — but governed and with minimized risk. These systems work well with guidance today, but they don't scale to the enterprise yet without oversight and significant cost.

But everything is accelerating. Enterprises are rightly cautious about letting things rip immediately, which is a very savvy business decision. GrowthLoop helps enterprises bridge this gap smoothly and quickly, but also in a governed way that minimizes risk. Decisioning is a key piece of that bridge. It gives them a way to view where agents are making decisions in their journey graph, controlling which specific decisions they want automation on while keeping guardrails around actions they want to test.

And here's the good news: decisioning isn't a side quest on the path to full automation. The experiments that power decisioning build the agentic context graph, and that same history becomes a first-party data source for agents when you have dozens or hundreds of them working on your behalf.

You mentioned that a lot of martech problems are being “solved” by AI in what you call verifiable domains. Why is marketing different, and how does the context graph help?

In 2026, we're seeing problems fall to AI at or above human level in verifiable domains — places where you can test outcomes at scale. This includes math, coding, and science. Coding is a great example: you can write unit tests and tell an AI to build a program that passes those tests, and it will iterate until it succeeds.

Marketing is not a verifiable domain. You can't send 100,000 emails to someone until they purchase a product you're trying to sell. You'll burn the relationship, and if you get engagement, it might be an unsubscribe, which is the opposite of what you want to maximize.

The agentic context graph allows you to make marketing more verifiable by running tests in simulation against all the different customer states you've seen before. You can answer "What will work?" before you ever contact a real customer. That's how decisioning becomes key to pulling forward the timeline — helping enterprises move toward automation in a responsible way that keeps their brands alive and maintains customer relationships firmly.

Let's talk about agentic commerce. When AI agents start making purchasing decisions on behalf of consumers or negotiating with other agents, how does that change the marketing equation?

It's a change in control. Today, if you're Walmart selling on your website, you control what the customer sees. You control the messaging. You can even buy ads to place yourself in their journey. But when agents are buying on behalf of customers or when other agents are selling on behalf of companies, you lose that control.

One version is human-to-agent: you ask Claude for a t-shirt recommendation, Claude shows you three options and recommends one with a buy-now button. But if you're Walmart, how do you ensure your shirt is the one recommended? If you have the context graph and can interface with systems like Claude in a governed way with all your data about customers you know, you can say, "If someone looks like Scott, offer them what worked for Scott before." That maintains some control.

The more interesting shift is agent-to-agent buying. I was talking with Yasmeen Ahmad at Google Cloud about personalization at scale, and I said we'll move to zero-to-zero personalization faster than anyone anticipates — agents talking to agents.

All it would take is email spam filters evolving into agentic filtering. You could have an agent in your inbox deciding what gets through to humans. There's so much AI-generated content slop landing in inboxes that none of

it will get through once an agent is screening for relevance and quality. If you're trying to get marketing through, you'll absolutely need all this context and history about the customer. And when agents are buying on behalf of humans, those agents will have human preferences built in, so knowing more about that human and activating that knowledge will be key.

Companies that aren't thinking about this now will be behind their peers when this shift happens. It feels far off and small right now, but this will happen much faster than e-commerce did. When I was at AB InBev, we talked about e-commerce when Amazon owned only 1% or 2% of the retail market. That happened fast. This will happen much faster.

The best outcome for every customer, every time.

Book a demo to see how Composable AI Decisioning helps marketing teams drive better outcomes and more revenue.

GrowthLoop

AI Decisioning

Metric to maximize:
Customer lifetime value

Hey Angela!

15% OFF

Last question: What should marketing leaders be thinking about today to prepare for this future?

Start building your agentic context graph now through controlled experimentation. Get your data structured in your warehouse with all the high-resolution detail you can access. Run cross-channel experiments that capture state, action, and outcome.

This gives you the foundation for everything else — decisioning, agents, and ultimately understanding what actually drives customer behavior and business results. The competitive advantage will go to companies that have this ground truth built in, because agents will need something reliable to

reference. And move thoughtfully toward automation, with governance and trust as central principles, not afterthoughts.

These are fascinating times we're living in.



How AI Agents Will Reshape the Marketing Workflow



Interview with *Tejas Manohar*
Co-CEO, Hightouch

Can you give us a brief introduction to yourself and Hightouch?

I'm Tejas, one of the founders of Hightouch, and we're building the future of agentic marketing. We really think the marketing workflow is about to drastically change with agents, where marketers will be able to go end-to-end on campaigns — from data analysis and briefing to creative development, launching, and executing them repeatedly at scale.

I started my career as a software engineer at Segment, one of the first CDP companies that Twilio acquired for about \$3 billion. We actually invented the Composable CDP model and started Hightouch about five years ago, and that still represents a very significant part of our business today.

Why has AI transformed software engineering so dramatically while marketing has lagged behind?

There are a couple of key differences. Coding is inherently a verifiable task. You can write code, test it, and see if it works. Coding agents like Claude Code or Cursor can run tests and validate their work against actual system outputs.

But with marketing, you can't put it in a black box and say, "Does this work or not?" The outcomes are far more subjective and dependent on context, audience, and market conditions.

What marketing really needs is context — a lot of it. Foundation models like OpenAI, Anthropic, and Google don't have the context they need to succeed at marketing tasks out of the box. They hallucinate products that don't exist, generated images don't look right, and they don't know your brand, colors, or how you message things. That gap is precisely what we're solving.

What exactly is a marketing context layer and how does it unlock agentic marketing?

Context solves the AI gap. The marketing context layer includes several essential components.

First, your brand guidelines, tones, colors, and visual styles — the foundational brand context. But brand guidelines alone don't make AI output truly feel like it came from your company.

Second, we connect to all your past campaigns — everything you've ever created across Facebook, Google, and TikTok. AI can see how they performed and what they looked like.

Third, we access all your creative assets in Figma, DAMs, and photo libraries so AI can learn what's on brand through pattern recognition.

And finally, your data — customer data, business data, performance data, inventory, website traffic.

A marketing context layer makes all of that accessible to AI simultaneously. It's brand guidelines, creative context, and data working together in a unified way. That combination is what enables agents to generate marketing that's genuinely on brand and likely to perform.

How does agentic marketing differ from just bolting AI features onto existing marketing platforms?

When you bolt AI features onto existing CDPs, ESPs, and DXPs, you get tools like GitHub Copilot helping with autocomplete — faster, but fundamentally the same workflow. An agentic marketing platform rethinks how marketing is done from the ground up.

Think of how coding evolved: Copilot started as autocomplete in the editor, but real coding agents are now building entire projects end-to-end. Similarly, an agentic marketing platform should enable a marketer to say, "I just saw that the federal government increased interest rates, and I work at a banking app. I want to send everyone a push notification, launch ads, and email the customer base about higher savings rates."

Agents will help execute all those tasks in a human-in-the-loop way. It's a completely new workflow powered by AI that uses your existing marketing tools to send emails, push notifications, and ads, reimaged from the ground up.

What does the day-to-day workflow look like for a marketer managing an agentic marketing platform?

Marketers are held responsible for campaign success — customer acquisition costs, ROAS, lifetime value — but they rarely have direct control over all the steps. They often need to go to a data team for audience building and to the creative or design teams for content. The marketer becomes a coordinator, managing all these dependencies rather than actually doing marketing.

That's what's fundamentally changing. A performance marketer can log in to Hightouch and request 20 new marketing campaign ideas for a product launch. Within 10 minutes, AI comes back with dozens of drafted campaigns — something that would take days and weeks before. Marketers can then edit them with AI or a visual editor and launch them across platforms, seeing how they perform, all within a single agentic workflow.

It puts autonomy and speed into a marketer's hands so they can actually start doing things at scale and spend time on what brought them to marketing in the first place.

How do agents handle the full lifecycle from identifying an opportunity through deployment?

The human stays in the loop for every single part. That's key. There are two ways to operate.

One is that you come to the AI with an idea you're ready to execute. Maybe you're talking with a coworker and notice an interest rate change, a social trend, or you have a product launch planned. You punch that into the agentic platform, request twenty campaign ideas, and within ten minutes, you have drafts.

The other way is that agents are always running inside the platform, proactively analyzing every ad, email, and lifecycle campaign you're running and figuring out what's working. We have proactive agents that are always on, analyzing performance and surfacing insights around competitive activity and emerging social trends. You can click an insight on performance or social trends to generate new campaign ideas. There's also a huge library of every single ad every company is running on Facebook and Meta Ads, which we can access to inspire new ideas.

So it's two operating modes: either you drive the agent with what to do, or you look through a list of proactive ideas and riff on those.

How does Hightouch's approach to agentic advertising differ from Google and Meta's AI-driven tools?

We're focused on enterprise and mid-market companies doing hundreds of millions to billions in revenue. At that scale, there's enormous detail that goes into making marketing feel on brand and perform better than what you've already tried.

Tools like Google Ads, Meta, and TikTok add their own AI features, but we believe those don't have enough context on the business to meet the enterprise bar. We're focused on bringing together all that enterprise context: the data, customer data, performance data, creative context from every campaign you've created, every asset in Figma or your DAM, your taste, style guides, and past feedback on assets.

We're learning from all of that to help enterprises move faster and better. We think they'll want one central tool across all channels and programs that has that comprehensive context, rather than relying on the generic AI features built into each individual platform.

How does an agentic layer change the composability story?

Composability is a core principle you need to be an innovator rather than a copier. You can't do the same thing as everyone else, so you need to benefit from the broader ecosystem of marketing technology tools.

With our CDP, we don't store all the data — it's composable because we connect directly to data systems like Snowflake, Databricks, BigQuery, and Redshift. We also have over 300 out-of-the-box connectors and integrate deeply with Salesforce, Adobe, Braze, Iterable, and any platform out there.

Agentic marketing expands this story significantly. We can connect to all the context in your data warehouse and all your marketing channels that we can target customers with and send marketing through. We make campaigns and content accessible to agents so when brainstorming, they can see what you've already run and how it performed. We help you launch campaigns across any number of channels or martech tools you have. We can code the HTML of an email and launch it directly on any ESP or ad platform you use. It's completely composable and platform agnostic, with no bias toward any tool.

Over time, intelligence workflows will move toward upstream platforms like Composable CDPs and agentic marketing platforms, because you want intelligence and context around your business to be centralized and not siloed in your different last-mile channel tools.

Is the data warehouse's role expanding?

The data warehouse remains super important as a source of data for all your proprietary business knowledge — sales data, website traffic, inventory, customer behavior. For the CDP, it's still the number one data source.

But when agents are actually doing marketing, it wouldn't be right to say the warehouse is the only source they need. It's far from it. Creative data — images, videos, email assets in Figma, Facebook, or DAMs — rarely lives in a company's warehouse. At Hightouch, we don't require it to be there. Our agentic marketing platform integrates with those creative tools and channels to read that data directly.

There's also so much learning that comes from the actual interactions of a marketer with the platform. When you say you like an ad or don't like it, or make revisions, the agents are learning from that feedback. That's not really structured data that fits well in a warehouse.

So the warehouse remains critical for business and customer data — you shouldn't do it any other way. But AI actually requires even more data sources than traditional personalized marketing.

What are the use cases you're seeing for MCP servers in marketing workflows?

There are two key areas.

First is user experience. Marketers spend a lot of their day brainstorming, coming up with ideas, writing emails, and writing docs — and a lot of that work is moving to AI. Marketers use tools like Gemini, ChatGPT, or Claude all day. I see this in my own workflows: when I need data analysis, I open Claude; we have an enterprise plan at Hightouch, and I've even had it make presentations for me that are way better than what I could make myself. Day-to-day work is moving to AI tools, and you want that work to be as connected as possible. The number one use case for MCPs is being able to interact with the Hightouch platform — both our CDP and agentic marketing platform — in the primary place you're working. That means switching from bouncing between browser tabs to talking to AI all day.

The second important thing is that enterprise companies are building custom internal workflows in ChatGPT or Claude to brief campaigns or generate reports. Our vision is that Hightouch provides valuable infrastructure that enterprises can use within their internal workflows via our MCP, so they don't have to switch contexts or tools.

What will marketing organizations look like in three years if agentic marketing delivers on its promise?

My belief might be a little contrarian. Because marketing isn't verifiable like support or coding, where you can easily check if something's right, marketing is fundamentally a role of taste, judgment, and creativity. I

actually think marketing teams will be among the least affected by AI headcount reductions.

Every marketing team I talk to is overwhelmed — not able to get to all their ideas and campaigns they want to execute. AI will unlock more time to actually do those things. But the way marketing teams are organized will undoubtedly change with AI, just as every other function will.

The big shift will be greater autonomy, power, and control for full-stack marketing roles. Today, you see marketers as Google Ads channel managers or marketing ops teams in the critical path of every campaign. In the future, you'll see business-minded, high-taste, creative marketers who can go end-to-end — from an idea in their head to live campaigns impacting customers, learning from the results.

If I were a marketer right now, I'd be focused on learning AI tools really well and adopting AI into every aspect of my workflow while practicing creativity, taste, and judgment. I think that's what's really going to matter. I'm very excited about the future of marketing.

The logo for hightouch, featuring a small square icon to the left of the word "hightouch" in a sans-serif font.

Data + AI for Marketers

Brilliant marketing*

*With a little help from your agents.



The Campaign Creation Renaissance



Interview with *Brendan Farnand*
Co-Founder & CCO, Knak

Could you give folks a background on yourself and Knak?

I'm Brendan, one of the co-founders of Knak. I've spent my entire career as an enterprise marketer, working in product marketing, demand gen, field marketing, and marketing operations. Throughout all these roles, I kept running into the same fundamental problem: creating content and campaigns was always a lengthy, difficult process.

One of my core passions is building great campaigns that actually convert. And honestly, creating campaigns that drive real results? That's even harder than the creation process itself.

Knak exists because of this frustration. We're a campaign creation platform that helps marketers go from ideation and campaign briefs all the way to execution. Rather than forcing you to swap out your entire martech stack, we integrate with everything you already have — your marketing automation platforms, your data systems, your design tools. That way we can supercharge your marketing and empower marketers to build incredible campaigns that convert, without costly rip-and-replace implementations.

You've argued there's a critical creation gap in the industry that's been hiding in plain sight. What is this creation gap?

There are really two dimensions of creation to think about. The broader one is the entire process — going from a campaign idea or brief all the

way to actually getting that campaign into the hands of your audience. That includes execution through your marketing automation platform and delivery to your customers.

The problem is that a lot of marketers can't do the full end-to-end creation process themselves. There's a production gap between that initial idea and actually getting it executed.

For enterprises, this gap is particularly painful because the process involves so many technical steps: designing emails and landing pages, coding them in HTML, using clunky templates in your marketing automation platform, localizing content across dozens of languages, getting QA approvals, and handling all those technical intricacies. It's a messy workflow that compounds at scale.

Can you walk us through what this process actually looks like at enterprise scale — the teams involved, handoffs, tools?

At an enterprise level, you have lots of different functions requesting campaigns. Field marketers need regional campaigns for their specific markets. Demand gen teams have their own lead generation initiatives. Product marketing has theirs. All these groups generate campaign ideas, and they all funnel down to one team: the marketing operations team that runs the execution platform.

The challenge is the sheer volume and complexity. When you're localizing a campaign across a region or the globe, or personalizing it for multiple audiences, the amount of content that needs to be created compounds rapidly. You end up with incredible operational challenges — getting all those different groups' campaigns to market efficiently and effectively.

Marketing ops people are particularly squeezed. Their sweet spot is in platform operations, data analysis, and segmentation — they're the systems experts. But they end up getting pulled upstream into the creation process because they're the only ones who know how to navigate all the technical complexity. They become bottlenecks. As the volume of content requests proliferates from different teams, that bottleneck becomes real very quickly. Marketing ops teams that should be focused on strategy and optimization instead spend their days executing technical production tasks.

How common is the creation bottleneck? Is it a constraint on the entire stack?

This is a huge problem industry-wide, and it's been hiding in plain sight.

I'll give you a concrete example: a pharma customer did a major marketing automation platform implementation on their B2B side. They invested heavily in Marketo, completely revamped their underlying tech stack, and looked at all the powerful capabilities they'd just unlocked. Then they looked upstream and realized something startling: "There's no possible way we can take full advantage of all these capabilities because we can't possibly create enough content to feed our nurture campaigns, target all our different segments, and deliver the personalized content we need."

They called it "feeding the beast." They couldn't feed the beast with enough content.

And here's what shocked them. That company has two major platforms. They have Marketo on B2B and Salesforce Marketing Cloud on B2C. When they examined both platforms, they didn't identify a sending problem, a segmentation problem, or even a data problem. Their biggest constraint was the creation problem: they couldn't produce the amount of content needed to take full advantage of these platforms.

This is common everywhere. Marketing automation platforms are incredibly powerful. They can do highly targeted campaigns, nurture journeys, audience segmentation at scale. But companies are constrained by their ability to create the content to fill them. Every company invests in these very powerful platforms, but they're constrained by the ability to create enough content to use them effectively.

Does generative AI solve this creation gap, or is it making things worse?

That's a great question. I was just in Silicon Valley talking to some of our most AI-forward customers — Meta, Google, OpenAI, Nvidia — and they're all dealing with this exact issue too.

The concern is what I call "vibe creation." Marketers think they can just vibe create an email in ChatGPT and throw it over the wall, or stick it into their marketing automation platform, and it will just work.

But OpenAI told us directly: we could never let our marketers use ChatGPT to create emails. There aren't enough guardrails around what they can do to create. The technical issues downstream are very real.

What our most sophisticated customers are actually doing with AI is using it to automate their workflow, not to create unguarded content. They're using AI to drive efficiency into their existing workflow, with proper guardrails and controls in place.

Here's a real example of how this works. OpenAI is ingesting all the marketing requests coming at their operations team — and for the fastest-growing company on earth, that volume is insane. They use AI to structure these unstructured requests into proper tickets in their project management system. Then they call our MCP to spin up a brand-approved email template based on content they've already approved. The marketer gets dropped into that campaign to tweak and edit it. Then it goes through QA for brand voice and email best practices, through a collaboration workflow, through approvals, gets translated into however many languages they need, and then it's ready for execution.

What they're seeing is AI embedded within their workflow, with marketers still at the center — not replaced. The AI accelerates the process, but everything stays on-brand and follows their guidelines.

What about “shadow creation” — marketers bypassing processes and approvals when AI lets them create content quickly?

If shadow creation takes hold in enterprises, a couple of real problems emerge. First, there's design and copy creep — or as I prefer to call it, brand voice creep. If marketers aren't guided by specific brand guidelines and email best practices, things fall apart downstream.

Look, I wasn't an email expert as a marketer. I didn't know that an email over 102K gets clipped on Gmail, or that preview text over 15 characters gets truncated, or the critical importance of alt text on images for accessibility and screen readers. These technical details weren't at the forefront of my mind. But they matter tremendously for deliverability and performance.

When marketers vibe create HTML and throw it at a marketing operations person or jam it into the marketing automation platform, it creates technical problems downstream — not better outcomes. That's the real risk of shadow creation: brand voice creep, design degradation, and technical failures that hurt performance. You lose consistency across campaigns and regions, and customers start seeing wildly different brand expressions depending on where they're receiving messages.

The answer isn't less process control — it's smarter integration of AI into existing processes. Yes, let AI generate an email, but within the guardrails of your brand and existing designs that have performed well in the past. Those are the things marketers need to think about as we move into this AI age. The goal is to embed AI into your workflow, not around it.


What does this kind of campaign creation capability actually unlock for CMOs? What are the business outcomes?

Two big things.

First is speed to market. I work with AT&T, and they realized their speed to market was costing them hundreds of millions of dollars. When there's a 6-to-8-week lag responding to a competitor's campaign, you lose inbound leads and customers during that entire window in a competitive market that moves quickly. You're essentially ceding market share while waiting for internal processes to complete. We're helping customers collapse that timeframe by 95%, especially where AI is involved.

Second is performance. When you can create an email in weeks or months, you can't possibly test it. You get one shot and you hope it works. But some of our customers are now creating between 3 and 15 versions of every email they send. That allows them to test every send, pick the winner, learn what works, and iterate continuously. This is how you improve campaign performance over time.

The other piece is personalization, which is really about content volume and speed. When you can't create content fast enough, you're forced to choose: either you keep things generic, or creating personalized content takes weeks. Many enterprise companies literally can't create a single personalized email in 4 weeks. That's not an exaggeration. I talk to these teams regularly. With faster creation, they can finally personalize at scale.

A promotional graphic for Knak, a marketing production platform. The background is a dark blue gradient. On the left, the Knak logo is in white. Below it, the text 'The Marketing Production Platform for Enterprise Teams' is written in a bold, sans-serif font, with 'The Marketing' in blue and 'Production Platform for Enterprise Teams' in green. Underneath this, a paragraph in white text states: 'The system that takes your emails and landing pages from first draft to launch-ready in minutes.' Below that, another paragraph in white text says: 'With AI built into the production workflow, teams produce more, faster, and better without adding headcount or losing brand standards.' At the bottom left, there is a green button with the text 'See how Knak works'. On the right side of the graphic, there are several overlapping screenshots of the Knak software interface. These screenshots show various features: a 'Post Editor' with a 'Publish' button, a 'A complete data platform, built for growth' section, a 'STAGE ASSIGNED TO' table with names like 'Sam Young' and 'Alan Agosta', and a 'REQUEST CHANGES' button.

If you had to bet on one thing that will change enterprise campaign creation over the next 2 years, what is it?

Here's what I think: enterprise marketing has been obsessed with the mechanics of the process — how do we make the workflow better, how do we go faster, how do we be more efficient? And those questions are important. They have to be.

But when you have AI-enabled systems giving you essentially unlimited efficiency, that's going to unlock something different. Efficiency unlocks optionality. It's going to unlock focus on the ability to really talk to your audience again.

I think we've kind of lost that over the last 15 to 20 years as we went digital and optimized for all these new systems. At large-scale enterprises especially, we've lost focus on what audiences actually want to hear about, what real engagement looks like, what actual performance looks like, rather than just driving more volume or faster processes.

When you can create a campaign in minutes and personalize it to the individual level in minutes, it unlocks a different way of thinking. I think the quality of campaigns is going to go through the roof. People are going to think much more deliberately about engagement.

Some of that will be forced by inbox changes happening right now. Gmail, Yahoo, and others are putting real pressure on senders to improve quality. So underneath all the AI noise, I think we're going to see a renaissance in how we really talk to our audience — the best way, the most personalized way, the most human way. That's what really excites me about this shift.



Open and Explainable AI in Marketing



Interview with *Raviteja Dodda*
Founder & CEO, MoEngage

Can you give our readers a brief introduction to MoEngage and what you do?

MoEngage is an AI customer engagement platform. We help consumer brands understand their customers and engage with them across every channel they use. Today we work with over 1,350 brands globally, from Fortune 500 enterprises to fast-growing internet companies, and we focus on helping them improve retention and customer lifetime value. We're about 800 people across multiple regions, so the work is genuinely global.

Marketers have heard for years that AI will transform everything, yet many teams are still using it as a glorified autocomplete. When MoEngage talks about AI decisioning, what does that actually mean?

We're after two things with our AI strategy. The first is efficiency, so marketers can run more campaigns in less time. The second is effectiveness, and that's where AI decisioning really comes in.

AI decisioning helps marketers make smarter decisions at scale. Better retention, better conversions, better engagement. The shift from traditional AI in marketing is pretty fundamental. Older approaches sat inside a single campaign or journey. We rethought the workflow from the ground up.

Instead of the old model where marketers schedule campaigns, define segments, and manage every piece manually, you now create a decisioning

agent or policy. You tell it your business goals and your guardrails. What outcome you want it to improve, whether that's driving purchases or retention. You set limits on how often it can message people. From there, the agent uses customer context, past interaction history, and ongoing learning about which customers respond to which value propositions. It matches the right offering to the right customer, and it figures out channel and timing too.

This isn't only about outbound messaging. It applies to personalizing web and app experiences in the same way. Take T-Mobile, one of our customers. Their app has several content slots, and they could show any number of things in them. Adding a line. A new product launch like T-Satellite. An incentive. The possibilities are wide open. Our AI decisioning figures out which customer should see which value proposition, and it keeps learning and improving over time.

There's a philosophical divide right now between teams that bolt on general-purpose LLMs onto their stack versus platforms that build purpose-built AI into specific workflows. What's your take? Why does that distinction matter?

You simply can't put LLMs to all use cases. Take decisioning, where you're figuring out which customers should receive which value propositions. That work needs more traditional machine learning combined with LLM-based context extraction. You can't just ask an LLM to make that call at scale. Consumer brands have millions of customers. Running an LLM for every single decision isn't realistic.

For content creation, campaign writing, and campaign analysis, though, LLMs are excellent. But on their own, they aren't enough. You need to feed them the right context. Business context, campaign context, historical campaign data. Once you do, the output gets genuinely useful.

One thing we've learned building agents is that getting to 80% quality is fast and easy. Getting to 99%, where the thing quietly works and people actually enjoy using it, is the hard part. That's where fine-tuning, the right guidance, and the right context matter most. And that hard work gets baked into your workflow, which is where real adoption happens.

When we rolled out an in-app agent inside the campaign creation flow, adoption moved quickly. Marketers were already in the product building campaigns, so the context the AI needed was right there. Purpose-built AI sitting inside the workflow is what drives adoption. We pair that with the intelligence of LLMs to give users what they actually need.

How are you thinking about embedding AI so it shows up naturally in campaign creation, audience segmentation, creative suggestions, channel selection, and timing, without feeling like a separate thing marketers have to learn?

From the beginning, we've focused on two factors: quality and adoption. If you get both right, monetization follows naturally.

Getting quality right takes real work. You need agents that do the job well without constant back-and-forth or a big time tax on the user. Adoption means embedding those agents inside workflows and making them easy to find. People also need to see social proof. How these agents have worked for other customers and what kind of value they've created. Once a marketer tries one and sees what it can do, going back to manual workflows feels painful.

The other piece is that agents need to show up everywhere work is happening. Not just inside our UI, but inside Slack, Teams, Cloud, and the other tools people use daily. The future isn't one interface. It's access from many places. That's how you meet users where they actually are.

Most AI in martech today is a bit of a black box. You feed it data, it gives you output, and if it doesn't work, good luck understanding why. You've been passionate about open and explainable AI. What does that mean in practice? How do you make an AI explain its choices in terms a marketer can understand and act on?

This is absolutely critical, especially for decisioning and recommendation agents. Marketers and enterprises need to trust your AI, which means they need to understand why a given customer is seeing a given offering.

We've built a few things to help explain that. One is a simulator where you can input a customer profile and see exactly what offerings would

be shown and why. The reason might be tied to customer context, past interactions, or the patterns the system is currently learning from. One thing to remember is that these decisions are moment-based. If you check again a day later, both the decision and the reasoning can shift.

We also give full visibility into every communication and decision that went to each customer. For enterprises, that's important. If something looks off, they can go back and see that a month ago this customer got a specific message with a specific value proposition, and understand the thinking behind it.

Transparency extends to outcomes too. Did the model help or hurt? Was the lift positive or negative? That kind of visibility matters for any AI product, and it matters most for anything driven by decisioning.

On the openness side, every SaaS platform should be built with composability and integrations across different data stacks and enterprise systems. The more APIs you expose, the more people can build agents themselves. We recently had a CIO at one of the largest Fortune 500 companies in the US reach out because he was already using our APIs to build his own custom solutions. That's the kind of future we want to enable.

There's real tension between autonomy and control in AI. Marketers want AI powerful enough not to surprise them, but predictable enough not to embarrass them. How do you think about guardrails and letting marketers refine AI outputs while maintaining hard rules about what the AI can and cannot do?

Forward-looking marketers are increasingly becoming managers of AI agents and agentic workflows rather than campaign creators. When they do, they're working with AI based on the inputs they provide: goals to optimize for, eligibility criteria, guardrails about communication frequency, and so on.

But there is a real tension we're seeing. While those guardrails exist, marketers also want control around business goals. They might have specific targets for certain product categories and want the AI to favor those. So the customer interaction history might suggest one offering, but business goals require preference for another.

We're in a transition phase where optimization must balance both customer outcomes and business outcomes. To address this, we give marketers significant levers — inputs they can use to weight certain value propositions or categories. This lets them guide the AI while still benefiting from its decisioning power.

Of course, seasonal campaigns and tactical moves will always exist. But we believe the crawl-walk-run approach we take with customers helps them navigate this transition successfully as the world changes week by week.

You recently launched an MCP server that lets marketers connect their MoEngage data directly to AI systems like Claude and ChatGPT. For readers new to MCP, what's the thinking, and how are people actually using it?

The MoEngage MCP is a connector that exposes tools and interfaces to read data from our platform or create entities within it. All usable by any AI tool or agent people are building.


For example, someone using Claude with the MoEngage connector could ask: "Give me a summary of all campaigns that ran in the last week and insights on how to improve them." The agent goes through everything and provides both analysis and recommendations.

We started with read-only tools so people wouldn't accidentally create something incorrect. We expose comprehensive read capabilities: analyze campaigns, review segments, examine content. We're now expanding to write capabilities next.

One thing we've noticed is that open access like this actually drives more adoption of your product. People go deeper into campaigns. And senior leaders, VPs and directors who don't log into a dashboard every day, can now pull insights whenever they need them. Without the MCP, logging in and doing that analysis would take too much effort to bother with.

Enterprises are also plugging MCPs into tools where work already happens. They write deployment briefs in Confluence, and now they can ask an agent to turn those briefs into MoEngage campaign structures and create the campaigns directly through our connector. Collaboration across teams gets easier, and everyone gets simpler access to all the data sitting inside MoEngage.

We don't worry as much anymore about people not logging into our dashboard. We think this will drive more usage overall and surface real value from the context in our platform. Your campaigns, segments, customers, all accessible wherever the work is happening.



moengage

The Customer Engagement Book
Adapt OR Die

Real world stories from
Meta, Hulu, Grammarly and WealthSimple.

www.moengage.com/book



If AI is handling more decisioning, optimization, channel selection, and timing, what's left for the marketer? What should they do with all their free time?

Actually, AI isn't creating free time. It's making us all do a lot more work than before. All of us are automating parts of our jobs so we can do a lot more overall.

What we observe is that lifecycle and CRM teams are usually small, even at large companies. As those teams adopt agentic workflows and decisioning, they can do significantly more with the same headcount. They can run more experiments. The optimization and feedback loops tighten considerably.

The job shifts from creating campaigns to managing AI agents and agentic workflows. Reviewing them. Making sure they're pointed at the business strategy. Marketers will spend more of their time on strategy and on the creative side of the work. More time understanding customers and tuning the goals they set for these agents and workflows.

The future is collaborative: human plus AI. You can't leave AI agents fully autonomous. You need human-in-the-loop approaches. But when done right, these agents become tools that make marketers ten times more effective.



The Self-Orchestrated Customer Journey in Real Time



Interview with *Tara DeZao*
Sr. Director of Product Marketing, Pega

Can you introduce us to Pega and your role there?

Pega is the AI-powered enterprise transformation company. We help organizations transform their operations, customer service, customer engagement, and modernize their legacy applications and technology.

I lead product marketing for Pega's Customer Decision Hub, an AI-powered brain that sits at the center of your martech stack. It ensures that across all your channels, your interactions are relevant, perfectly timed, and value-added to customers. Essentially, it's a way to scale personalization across all your channels.

What does "AI decisioning" actually mean, and why should leaders consider it a foundational martech capability?

AI decisioning is essentially a way for your martech stack to get smarter. It combines adaptive and predictive analytics to sense the signals coming from your customer and then decide the next best action, offer, or interaction to have with them. The goal is to move away from batch-and-blast campaigns toward true one-to-one interactions that happen in real time.

Consider this scenario: you're watching TV while scrolling on your phone and you change direction in that scroll. If you're in a brand's channels, that

brand needs to be able to give you the right offer in the right channel at the right moment. If it's connected TV and then you switch to email, they need to do that in real time.

That speed is critical because the window of opportunity is remarkably short. Customers see over 6,000 messages a day now. In order to break through that noise, you have to reach them in the moment they're making a decision or need your help. And you have to change direction almost instantly when they do, because now even an hour is too late.

In the early days of adtech, fresh data meant your cookie was less than 30 days old. Today it has to be in milliseconds. Your data freshness is measured in milliseconds because the window of opportunity to get in front of your customer is so incredibly short with how fast they move through our channels. Marketers recognize this. The real imperative is that your tech has to be able to keep up with the customer on their journey.

How has the customer journey fundamentally changed, and what does that mean for engagement models?

The way we think about customer journeys has undergone a major shift. In the past, we pushed customers on linear journeys. They were scripted, like a yes-or-no-then situation. Now I describe it more like an octopus. Your customer is on simultaneous journeys. You might be sitting on the couch watching TV while scrolling on your phone, doing two or three screens at once. That's not a linear customer journey anymore.

One of the key capabilities we need is adaptive AI. At Pega, we call this "re-decisioning." You need to be able to sense the data that's coming in real time. Predictive AI tells us where a customer was going to go, but if they actually go somewhere else, we need to adapt immediately. For example, if a customer was heading toward one product but suddenly shifts their behavior toward something completely different, we need to detect that shift and change our engagement approach on-the-fly.

The fundamental relationship has shifted. Customers are at the center of their journey; they control it. We are no longer in control. Our job is to show up with help and meet their needs in the moment, wherever they go. But it has to be seamless. If I'm shopping for internet service and I

see one price in the mobile app and a different price on Instagram, I don't trust you anymore. You've added friction to my journey instead of making it seamless. And sometimes the best action is saying nothing at all — recognizing when silence is better than an intrusive offer.

Can you walk through a concrete example of how “next best action” works in real time?

Next best action takes in a pre-built conversation library — think of it like dynamic creative optimization on steroids. You load it with any number of messages you want to say to customers: onboarding messages, retention, nurture, or empathy-based messages for customers struggling financially. In a real scenario, let's say you call a contact center because your internet is really slow. The agent's console should tell them in real time: this customer has this product, maybe needs an upgrade, and they're only using half the data they're allowed. They could upgrade to a faster plan with less data for the same price.

All of that analysis should happen in milliseconds. The next best action either pops up for the agent to review or goes directly to the customer in their channel of choice. If you have the right technology at the center of your martech stack, it changes the whole dynamic. Most people will do everything humanly possible not to talk to a human — I know I will. So the fact that you've solved all of a customer's problems in a five-minute interaction, even before you've had a full conversation, is the essence of next best action. That's the next best experience. You're not just being efficient. You're actually making your customers' lives better.

What does a proper data foundation look like for AI-powered marketing, and what goes wrong when teams skip this step?

Data silos are the enemy of AI. Your AI is only as good as your data. It doesn't matter how much money you spend on a wonderful AI tool if your data is not connected and it's stale — those are the outcomes you're going to get. We've always said in martech and adtech: garbage in, garbage out. It's the same with AI. And when you layer AI on top of data that's not connected, well-organized, or clean, you're asking for trouble.

You want your data to be connected across channels, as fresh and up to date as possible, and you need to be able to articulate to customers and

regulators how you're using their data and how your AI is decisioning off of it. When data is stale, it's not good. For marketers, if you're sending someone an impression in channels they're not even in, that's wasted budget. When you're doing that across a million or hundred million customers, a hundred thousand times a day, you're wasting enormous amounts of money.

Where do AI agents add real value, and where is human oversight essential?

A human in the loop is essential, particularly because we work across various industries with different regulations and we're trying to protect customer trust. Healthcare and banking are sensitive — these decisions affect people's lives. You need to be able to control the AI when it doesn't do what you want. Human oversight isn't about slowing things down. It's about ensuring safety and building the guardrails that make broader adoption possible.

We recommend AI agents work within a defined workflow, which in itself is a guardrail. We want agents to act autonomously within reason. For example, with our marketing agent, you upload a campaign brief. If it's the fourth or fifth time you've uploaded a brief for the same product and you're missing email when you previously included it, the agent should flag that: "Are you sure that's what you want to do?" That's helpful and safe. The agent is adding value by catching potential mistakes without micromanaging the process.

Agents should also intelligently monitor your campaigns. If a message isn't resonating and a customer has ignored it 3 times, but they interact well with different types of offers, agents should recommend changing the message now. Or if you thought someone was interested but they're really not, the agent should tell you this was a mistake and they're not the right person for this product. This kind of agent intelligence also helps with measurement and reporting — something that has been a nightmare for marketers. Agents can see what happened and why, and recommend what would be better.

What warning signs indicate tactical versus strategic AI adoption, and what does a holistic AI strategy look like?

The simplest answer is having a governance board. Data is siloed, but functional areas are also siloed. We have little fiefdoms all over the corporate enterprise. Without a governance board, people make decisions in a vacuum, and those tactical decisions may not ladder up to your overall strategy and can cause trouble later.

Here's a concrete example: if you build several agents to help find documents without starting from the bottom of your data layer or architecture, and they're pulling from both internal and external sources, you'll inevitably leak internal information that shouldn't be external. You said "we're having trouble surfacing documents, let's build agents and launch them," but that doesn't ladder up to a holistic strategy where you've decided on the 3 AI tools you're using, which sources they connect to, what they should be used for, and your overall goal — like becoming more nimble in the sales process. Real strategy is implementing for the wider good and outcomes-based results, not just fixing one thing.

What does "responsible AI" actually mean in marketing, and why does it matter for customer trust?

This is my favorite topic because we haven't been paying enough attention to it unless we were in finance or healthcare. But this is at the core of consumer trust. If you break your consumer's trust even once, they will churn. One bad interaction, one mistake, and it's over.

AI needs to be responsible and ethical. Just because you can do something doesn't mean you should. If you're a gaming company and you realize you're pushing customers into financial ruin, you should not keep doing that if you want them to remain customers.

Responsible AI has several dimensions. You want it to be fair and unbiased, and this has a lot to do with how your data is labeled. You want it to be robust, tested against real-world exposure. Microsoft's Tay bot is a cautionary tale. They released it on Twitter without proper testing, and Twitter users taught it to be racist within 4 hours. That's what happens when you don't harden AI for real-world conditions.

You want your AI to be explainable and transparent. You should be able to explain every decision your AI made, especially in regulated industries. Why


did one customer get a credit card offer and another didn't if they have the same credit score? It might be they already have that card, but it could also be the AI is discriminating based on where they live.

We have a responsibility to our customers to do the right thing and not harm them just because we're allowed to. Building consumer trust is the foundation of long-term success. Yes, we're here to make money and want the stock price to go up, but we have to balance that with the ethics of our actions.

What's one key piece of advice for marketing ops leaders planning their AI strategy for the next 18 months?

I can't stress enough the importance of the real-time and adaptive nature of the technology. It's critical that you look under the hood of what you're buying. A lot of folks are saying, "let's just build whatever we want on top of an LLM," but LLMs hallucinate. You're going to need to explain your decisioning and make accurate decisions. If you want to run accurate marketing at scale, you can't have hallucinations.

So you want a real-time, adaptive solution that can unify all your channels, connect all your functional areas, and make your marketing team's life easier. That's the capability that will set you apart in 2026 and beyond.

A dark blue rectangular graphic with a woman in a blue dress on the right side, holding a glowing green square with a red hand icon. A white line connects the square to a play button icon on the right. The Pega logo is in the top left corner.

PEGA

Move your martech stack from siloed to synchronized

We're moving out of the era of more tools and into the era of more control. Centralized AI helps connect data and decisions, reducing waste and enabling consistent activation across channels.

[Download the eBook](#)



The Future of CMS in an Agentic World



Interview with *Sara Faatz*
Senior Director, Community, Progress

Could you give readers a quick introduction to Progress and your role?

Progress is a software company that's been around for a little more than 40 years, and we build products that empower organizations to develop, deploy, and manage responsible AI-powered applications and experiences. I lead our community and awareness programs predominantly for our digital experience products, which includes Sitefinity, our CMS platform.

We've seen rapid evolution from web content management to digital experience platforms to now "generative CMS." Are we talking about CMS 3.0, or something entirely new?

It is safe to say we are looking at CMS 3.0, and "Generative CMS" as a category label really encompasses all that this next iteration of the technology is. It is a CMS that can deliver experiences that are conversational, contextual and deeply personal. It is a CMS that understands intent, not clicks; adapts in real time; delivers clarity instantly; and is visible where discovery happens.

There are caveats, of course. The biggest is that your data — both structured and unstructured — has to be in good shape, and you need a team that understands it at a molecular level. The data issue isn't new. However in an era of AI, it becomes hyper-critical.

Many marketers are running mission-critical websites right now. What does the journey from today's CMS to this generative future actually look like without breaking what's working?

Yes. Absolutely. I think the smartest marketers will look for opportunities to adopt generative CMS elements incrementally in a crawl-walk-run fashion, if you will.

I would (and we have) start with search. The search experience has been less than stellar for many years but we have all accepted it for what it is. A generative CMS can exponentially improve that by replicating the experiences people have come to know and love with their favorite LLMs like ChatGPT or AI-powered answer engines like Perplexity.

This is something quite honestly any marketer can do today, regardless of the CMS they are on, by embedding an Agentic RAG-based search in their digital experience. There are huge benefits to this.

It allows them to deliver a modern search experience that people are hungry for – one that is conversational and deeply contextual.

It allows them to set up guardrails and select content they want to show, while maintaining brand compliance and governance.

And, if implemented correctly, you can capture the prompts that people are using (without violating privacy). With that, you suddenly know how your customers are searching, what they are searching for, what their pain points are. This is a treasure trove of information that can be used to improve messaging and content and potentially even influence product.

Can you walk us through what that actually looks like in practice?

Absolutely. I have two examples I can share.

As a first example, we've been drinking our own champagne at Progress. Last fall, we embedded Progress Agentic RAG into our Sitefinity instance that runs progress.com to provide a conversational search experience. The RAG-based search box sits on the homepage and the results we are seeing are really impressive. First, people who use the new search functionality

are spending an average of 27 minutes on our site. Almost a half an hour. That is gold. And second, we've seen 3.6x more engagement. Plus, as I mentioned before, we are learning so much more about our customers from the prompts they are entering.

The second example is from a customer of ours that embedded Progress Agentic RAG into their Sitefinity instance not just for general search but for product discovery. They manufacture doors and windows and have many different products and three distinct personas. Before the implementation, a customer would have to go through a fairly rigid 7-step form to be served potential products. Now, their customers use natural language to search for products and they are served with relevant and contextual choices.

The beauty of this is that a native integration can be deployed in 60 days. For those who are being tasked with finding ways to leverage AI, something like this is low-risk and very high-reward - and quick to implement and start seeing success.

With RAG capabilities, you're not constrained to just your website. You can tap into any external or internal knowledge source. How should marketers think about what they feed into that system?

With a RAG solution, you actually have governance and control that you didn't have before. You can bring in external content that validates your conversation or story, third-party information, but you control that narrative. You have an expanded knowledge graph that allows you to build a more predictable path based on prompts and customer understanding. It's incredibly compelling.

As a marketer, you should be thinking about the ownable conversations your organization can have and make sure the data — both structured and unstructured — your system is ingesting is supporting that narrative. You have more control than ever before, and that is very powerful.

How do you map customer intent signals — the language they use in search — to segments and personalized responses?

Let's take a bank site as an example. Today, most financial institution sites provide every visitor the same structure, the same content and the same

experience. Regardless of who they are or why they're there. A consumer navigates to what they are looking for and pokes around to get the answers they need.

With a fully generative experience, a dynamic page could be generated from prompts and customer profile data — and personalized in real time based on the visitor's intent.

So if a homeowner arrives looking for a mortgage, they will write a specific prompt, "I'm a first time home buyer and I want to know more about the mortgage process, what the current rates are and how I go about applying." The system would automatically assemble a page with the right banner, content blocks, rate calculators and resources. All tailored to the end user's goal.

Understanding onsite behavior through search and conversation — does that change how you think about your strategy for showing up in ChatGPT, Gemini, Claude, and other external AI engines?

Absolutely. Once you understand how your customers are searching and what their pain points are, you understand the language they're actually using and you can begin to identify ownable conversations that are unique to your brand. With that, you will be able to show up consistently across all channels you own or influence. With offsite brand or third-party mentions becoming increasingly more important — as they are key signals of authority to LLMs — consistency of message and narrative is key.

You mentioned agents are researching and comparing on behalf of users. What does it mean to build a site that's agent-friendly, and does that cannibalize the human-friendly version?

Humans are still ultimately the ones consuming the message. You need to build for agents, but understand that humans are still the end consumer. There are things you can do on the backend: use the right page structures, comparison tables, FAQs, and clear, succinct language. On long-format content, add table of contents, time-to-read indicators, and solid metadata. All of those enhancements make it easier for agents to scrape and also make it more accessible and available to humans.

Change management comes up often. Where do you see the biggest resistance when working with marketing teams?

The biggest fear is still the unknown. Unpredictability is challenging and uncomfortable for a lot of people. But here's something I've heard from everyone at conferences and from customers: everybody's afraid they're behind, that somebody else is ahead of them. And they're afraid they'll ask the wrong questions.

I'd start by leveling the playing field. Everybody is still trying to figure this out. Start with what we call "boring AI" — things like search, which is low risk and high reward. You can implement it right away without breaking anything. Be vulnerable. Know that everybody's in the same boat. Experiment, but start with things that aren't flashy. If it sounds too good to be true, it probably is. Don't start with something mission-critical. Start with small stuff, see how it works, learn from it, and implement more.

As AI handles more execution, how do the next generation of marketers build critical thinking and judgment about outputs?

This question keeps me up at night, literally. I've been talking to younger generations trying to understand what they're learning and how they're being taught to use these tools. One thing I've learned is that subject matter expertise is still incredibly important.

You need checks and balances from a human perspective: Is this information valid? Is it new? Is it unique thought? And I think the only way we do that is by being uniquely human — talking to one another, brainstorming together, not just relying entirely on AI. AI only knows what we've already told it.

The muscle we have to build is the spirit of experimentation and collaboration.

This is the year of agents. But as you automate at higher velocity, validating both inputs and outputs becomes crucial. How do people build that judgment when AI answers sound so confident and polished?

You still have to understand the process and workflow, because you're the one instructing the agent on how to get there. It's a lot of checks and

balances. I'd treat agent development as I would user testing. Do a lot of agent testing.

"Trust but verify" is probably one of the best ways to think about it. We're not at a point where you can set it and forget it. Understand the workflow and what you're trying to accomplish, and then lean into the agent as a coworker. Treat them as somebody new to the organization who doesn't have all the institutional knowledge. Guide that process. That's the only way you'll have agents that are accurate and ultimately productive.

As everyone adopts similar generative CMS, conversational search, and agent-ready sites using largely the same AI engines, won't every website converge to the same experience? Where does differentiation live?


The answer lies with humans. Humans today are the only element capable of unique thought, critical thinking, and creativity. If you just let AI do everything and scrape what's out there, you'll have a lot of noise and also-rans.

The organizations that will excel and succeed are the ones who continue to provide unique thought and perspective and embed that into everything they do. Yes, that unique thinking will eventually be picked up and scraped and used elsewhere. But as humans, we can't stop thinking and creating. We have to stay on the forefront to keep doing more of that. I think that's the only way we move forward as a society.



A Generative CMS Designed for the Way Your Work.

Progress Sitefinity
An AI-powered CMS that simplifies
complexity without compromising
control or governance.



www.progress.com/sitefinity-cms



The Middle Layer as Martech's Center of Gravity



Interview with *Jonathan Moran*
Senior Marketing Manager, SAS

For readers who may not know SAS as a marketing technology company, can you give a quick introduction?

At SAS, we're a company mostly known for data and AI, but we also sell solutions into several different areas — martech, adtech, fraud, risk, IoT, and more. We have industry-specific solutions and industry-agnostic ones.

I'm fortunate to work with a team of marketers across our horizontal solutions that sell into a variety of industries, including martech, which is my first love. We're increasingly recognized as a player in the space because of our strength in that middle decisioning and orchestration layer — the layer that many marketing teams are realizing they need to invest in.

The martech stack conversation has historically focused on two layers: data platforms like CDPs and cloud warehouses, and activation channels like email and social. Is there a critical middle layer that's been underappreciated?

When we talk about martech stacks, there are three main layers. The foundational data and data application layer — where cloud data warehouses like Snowflake, BigQuery, and Redshift reside. The top panel delivery and activation layer where messages are sent through email, mobile, social, and other channels. And then there's the middle decisioning

and orchestration layer — that insight layer needed to turn data into action in the form of offers, messages, and interactions.

Historically, brands have invested very heavily in those outside layers. Why? Because that's where most vendors have traditionally excelled. But where brands haven't invested as heavily is in that middle layer — the layer where orchestration happens based on decisions about who gets what engagement at what point in time.

In this age of AI, as brands adopt technologies throughout their stack, they're realizing they need that middle layer. They need AI and analytics to make decisions about who to contact. But many brands don't have the capabilities in the middle layer that they need. Most of the biggest martech vendors don't really excel there either, because it's not what they were founded on. You've got companies very good at CRM automation, companies very good at content creation, companies very good at mobile messaging — but not as many that are deep and well-versed in analytics and decisioning.

You talk to folks out in the industry and they'll tell you the middle layer is becoming the new battleground as AI advances. Every vendor wants a piece of it. But here's the thing: they can latch onto the term "AI decisioning," but are they really truly doing it? Do they have the tools and technologies to make it succeed? We're seeing a lot of marketing teams drowning in data rather than acting on insight, because they don't have the right tools for the job.

We've moved beyond simple adoption of AI to brands figuring out how to derive and use AI with precision and utilization.

CDPs are evolving from unified customer data platforms into context-ready decision layers. How does SAS see this transition, and where does customer intelligence fit?

The shift we're seeing is natural and what the market is demanding — that move from traditional CDP capabilities to CDPs becoming context-ready decision layers. I'm seeing a bit of a squeeze in the traditional CDP model. The traditional model involved the CDP managing its own proprietary data store. Customer data was ingested, unified, and activated from within

the CDP's own database. But this created a frustrating dynamic because the golden customer profile lived inside the CDP and was siloed from the broader ecosystem.

Then we saw the shift toward composable or warehouse-native CDPs, where CDP vendors started sitting on top of cloud data warehouses rather than replacing them, using the warehouse as the system of record for customer data. The CDP provides identity resolution, segmentation, decisioning, and activation on top.

There's pressure on CDPs from the cloud data warehouse coming up and the activation layer solutions coming down. This is forcing CDPs to evolve. Organizations don't want a second copy of their data living in a CDP silo. They want their data science and analytics teams working in the warehouse. Reverse ETL vendors have accelerated this concept by proving you could activate warehouse data without a traditional CDP.

Today's spectrum includes packaged traditional CDPs with their own data layer, increasingly with good warehouse connectivity; purely warehouse-native composable CDPs with no proprietary store; and hybrid approaches where warehouse vendors are building CDP-like capabilities directly into their platforms.

SAS's embedded CDP capabilities fall fully into the composable category. Customer Intelligence 360 can connect to external data sources and cloud environments and doesn't require or maintain a proprietary store. This positions us as a context-ready decision layer with a bunch of surrounding capabilities.

What does true orchestration look like when the middle layer is doing its job well? What are all the moving parts?

Orchestration in its simplest form is the coordination of decisions along a customer journey, across channels, time, and context — so that every interaction is the right one, not just the scheduled one or the one that the brand forces upon the consumer. Many moving parts work simultaneously in harmony to get orchestration right.

Think of orchestration like the human body. First, the unified customer profile is the fuel for the body. Everything depends on that continuously

updated, identity-resolved profile — not batch refreshed, but a live view that incorporates behavioral signals, transactions, and interactions.

Second, decisioning logic is the brain. Real orchestration requires a decisioning layer that can evaluate, for any given customer at any given moment: What's this person eligible for? What's the best action, offer, or experience across all competing actions? What channel and timing maximize the likelihood of the right outcome? This is where logic, business rules, suppression rules, fatigue controls, and AI-driven propensity models all have to coexist. A true journey builder has all these things included; if it doesn't, it's just a simple flowchart in a GUI.

Third, channel connections and integrations with bi-directional feedback form the nervous system. Good orchestration doesn't just push to channels — it listens back. Did an email get opened? Did a push notification get dismissed? Did the customer call after receiving an SMS? That feedback has to be present for true orchestration to happen. If it's just one-way integration, that breaks the feedback loop and nullifies true orchestration.

Fourth, event-driven triggers are the reflexes — the ability to respond to real-time events like a browse, abandonment, failed payment, loyalty tier change, or geofence entry, firing the right response in seconds or milliseconds. This requires event streaming. You can't rely on nightly batch ETL jobs.

Fifth, suppression and arbitration is the referee. When multiple journeys, campaigns, and triggers fire simultaneously at scale, who decides what actually wins? Decisioning arbitration requires good governance ensuring a customer doesn't get three messages in one hour from three different campaigns.

Sixth, cross-channel continuity is the memory. A customer who ignored an email, then saw a paid ad, then visited the website should be treated as one continuous experience, not three separate interactions in three separate channel tools. This requires coordination across channels and departments—brands have to break the traditional model of siloed functions by channel.

Finally, optimization — AI and ML integration at decisioning time. Can your decisioning logic call live models? Can it call churn propensity models,

next-product-to-buy models, optimal send time, message variant selection at the moment of decision? Not in batch or based on a score computed last week. This is the difference between model-informed orchestration and true adaptive learning orchestration over time.

The point: orchestration is hard. True, deep, advanced orchestration requires a lot of tools, technologies, processes, and capabilities working simultaneously all at the same time.

AI agents are a major theme in *State of Martech 2026*. How does an intelligent middle layer interact with autonomous AI agents?

The middle layer is critical for autonomous agents to function. It's the brain that provides intelligence to those agents. We know autonomy without direction often results in chaos. But even more than direction, agents need grounding and continual monitoring, at least right now.

Consider what happens if you give an agent a direction like 'optimize click-through rates' and the agent starts hammering a single customer with messages, cannibalizing other marketing efforts, violating suppression rules, and making offers a business can't honor. That's trouble.

The middle layer is where agents get four critical components — the four C's. First, context: all relevant data about the customer, their history, stage, preferences, interactions. Second, constraints: the middle layer defines eligibility rules, suppression rules, fatigue limits, channel preferences, and business rules.

Third, compromise — essentially arbitration. In a multi-agentic world, agents will compete to act on the same customer simultaneously. The next-best-action agent might want to talk to a customer before the retention agent. The email agent might say it can save the day with a message instead of a push notification. The middle layer becomes the referee ensuring the customer receives a coherent experience.

And fourth, cognizance — closing the loop. After an agent acts, the middle layer records what happened and updates customer context. Agents have to be able to learn from outcomes. If that middle layer doesn't close the loop, if the agent acts but that outcome isn't recorded and the context isn't updated, the agent isn't well informed for the next interaction.

Does the middle layer become the new center of gravity of martech stacks?

I honestly think it does. It's moved from being an afterthought to being not just the center of the martech stack, but really the intelligence layer itself. In the age of agentic AI, this intelligence layer becomes more than a decisioning engine — it's a sensory nervous system that coordinates the entire customer experience.

The middle layer has long been overlooked due to the dominance of certain players born in and of the outside layers. But as agentic AI comes more into view, it's going to change the activation layer. The activation layer will be more heavily agentic in the future. So the middle decisioning and orchestration becomes critical. I do see it as the new center of gravity.

SAS has long been known for code-level adaptability. As AI makes custom code easier than ever, how does SAS position itself across the spectrum from out-of-the-box UI to fully custom solutions?

To some degree, this is a "have your cake and eat it too" moment. SAS has long positioned itself as providing capabilities for both the do-it-yourself persona and the do-it-for-me persona. If you want to code, great. If not, that's okay too.

What's awesome is you can create custom code and solutions outside of SAS — in Python, R, or whatever — and drop it right into SAS solutions for use. We position ourselves as very much open and agnostic, working with providers of all types. That's one of the reasons so many brands look to SAS to provide that middle decisioning and orchestration layer: we work well with both upstream and downstream technologies.

What's your advice for martech practitioners implementing and managing these stacks?

A lot of conversations come back to data. As with traditional martech, you've got to have good data for good campaigns. You also have to have good data for good AI. Identity resolution is critical to martech data management, and you've got to get that right before anything else. The orchestration layer is only as smart as its ability to recognize the same person across channels. Most decisioning failures we see aren't really logic failures — they're identity failures that make logic look broken.

The next big component is treating this middle layer as infrastructure, not as a campaign tool, a bolt-on, or an afterthought. The orchestration layer should hold the rules and logic, and campaigns and journeys should consume those rules. Suppression logic, frequency governance, and channel priority should be standing infrastructure, not things configured by individual campaigns. That separation is what lets you scale without having to rebuild over and over again.


And my last piece of advice: the process component of decisioning and orchestration. You've got to know who owns that middle layer. It needs an owner. Processes have to exist where handoffs are explicitly detailed, and people are held accountable. Orchestration most commonly falls apart at the boundary between marketing automation and sales and service systems. If orchestration actions just get logged in the CRM and forgotten, the loop is broken. There's got to be a closed loop. That handoff is a relationship design problem with sales operations before it's a technology problem.

Documenting and really understanding processes — how decisioning and orchestration will actually work inside your organization — is so important before you start buying tools and technologies and trying to fix any decisioning and orchestration problems you might have.

SAS CUSTOMER INTELLIGENCE 360

“The orchestration layer is no longer just decisioning – in the age of AI, it’s becoming the intelligence system that coordinates the entire customer experience.”

Jon Moran, Senior Marketing Manager, SAS



Additional Credits

A special thank you to these 19 martech experts from 14 countries. They contributed to validating thousands of new martech tools.

Beste Akkaya in Turkey
David Blackburn in the USA
Dunja Riehemann in Germany
Edwin Ng in Singapore
Gerben Busch in the Netherlands
James Schlesinger in the UAE
Jim Groo in Switzerland
Johannes Schulz in Germany
Jordie van Rijn in the Netherlands
Kristina Huddart in the United Kingdom
Małgorzata Marcinkowska in Poland
Michael Klazema in the Netherlands
Niels van Meerte Janse in the Netherlands
Odd Morten Sørensen in Norway
Reynier Matos Padilla in the USA
Sean Broderick in Ireland
Sinan Günal in Canada
Stefania Massenza in Italy
Yogita Wadhwa in India

Stay connected



www.chiefmartec.com



www.martechtribe.com



www.martechmap.com