

# Salesforce's Signal to Action: Agentic Pipeline, the Braze Bet, and the CMO's Reckoning

Salesforce's agent-generated pipeline and its \$5.66 billion offer for Braze point at the same ambition: turning customer signals into automated commercial action. The technology is arriving faster than the organizations meant to run it.

REPORT PROFILE

# Built to change a decision.

Salesforce's agent-generated pipeline and its \$5.66 billion offer for Braze point at the same ambition: turning customer signals into automated commercial action. The technology is arriving faster than the organizations meant to run it.

|           |                              |
|-----------|------------------------------|
| REPORT    | FR-SIGNAL-TO-ACTION-20261027 |
| AUTHOR    | Steve Mann                   |
| PUBLISHED | October 28, 2026             |
| PROFILE   | frontier-ro-0.1              |

This PDF is a publication rendering of the canonical Frontier Research Object. The web report remains the live source of truth; citations, figures, methodology, and rights metadata travel with this edition.

## Executive perspective: the marketing bottleneck has moved from capacity to human-machine relationships

Salesforce has shown that AI agents can follow up on leads its sales team never had enough people to contact. The company's own lead nurturing agents produced a reported \$88 million in pipeline from leads that previously went unanswered, work it says would have taken roughly 400 additional people [1]. An earlier engineering account, from April 2026, put the figure above \$100 million, with more than 10,000 opportunities and 1,500 closed deals in the small and mid-sized business segment [2]. On October 20, 2026, a week before this report, Salesforce agreed to pay \$5.66 billion for Braze, a customer engagement platform whose core asset is the ability to react to a customer signal within seconds [3]. On the Dreamforce stage in September, Benioff went further and claimed the Hunter agent alone generated \$500 million in pipeline in a single quarter, a figure no filing supports and which this report does not rely on (source 52). Read together, the pipeline result and the acquisition describe one strategy: ingest a signal, decide, act, and hand a human only what needs a human.

The technology is further along than the organizations expected to run it. Most marketing departments use AI every week and almost none have rebuilt a workflow around it. Seven in ten CMOs want to lead on AI and seven in ten say their own processes cannot scale it [4] [5] [6]. The gap between what vendors can ship and what marketing departments can absorb is now the central management problem.

This report takes four positions.

1. The \$88 million is credible as a directional signal and incomplete as a benchmark. Salesforce has not disclosed the period, the attribution method, or the win rate behind it [1]. Treat it as a reason to test agents against your own neglected demand, not as a number to put in a board deck.
2. Neglected demand is the right first target. Leads scored too low to work, accounts that went quiet, renewals nobody chased. These carry low opportunity cost, and what cost there is comes in deliverability and consent exposure rather than money, so an agent's errors are cheap and its wins are close to pure upside.
3. Agents are not a plug-in. The value arrives only when the workflow is rebuilt around the agent, and Salesforce's own engineers learned this the hard way [7] [8] [2]. Layering agents onto a campaign calendar designed for humans produces faster spam.
4. The marketing organization changes shape. Execution roles thin out, orchestration and judgment roles grow, market research contracts, and product marketing becomes a player and coach to a roster of agents. Labor's share of the marketing budget rose from 21.9% to 24.5% in one year [9]. Some of that is arithmetic, since martech spend fell to a five year low and paid media rose in the same budgets. My reading of the rest is that agents did not make people cheaper. They made the people who can direct agents more expensive.

The CMO's job over the next 12 months is to convert this from a technology program into an operating model, with decision rights, cost governance, and a talent plan that the Chief Financial Officer (CFO) and Chief Executive Officer (CEO) will fund.

Figure 1 translates the report's argument into a CMO operating model: customer signals pass through trusted data and bounded agents into B2B handoffs or Braze-enabled consumer activation, with human ownership, evidence and cost controls across the loop. The dashed connection depicts the proposed combined architecture, not a completed integration. [2] [10] [11] [7]




Figure 1. The CMO signal-to-action operating model: Salesforce agents, Braze consumer activation, human controls and five evidence gates.

## 1. What is the Salesforce product strategy?

Salesforce is assembling a loop: signal in, decision made, action taken, human pulled in only at the exception. The agent pipeline result suggests the loop works on business to business (B2B) leads inside Salesforce's own sales organization. The Braze acquisition buys the missing piece for business to consumer (B2C) brands, where the signal arrives as an app event or a purchase and the window to act is measured in seconds.

## The agent roster

At Connections in mid 2026 Salesforce framed the pitch as giving every marketer an AI team: agents that build pipeline, create content, and run campaigns so the marketer can set strategy and own results [12]. At Dreamforce in September 2026 the company showcased seven job ready agents, announced the week before, and moved multi agent orchestration to general availability [11].

| Agent                               | Job                                                                                          | Status (late October 2026)                                                                    | Source |
|-------------------------------------|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|--------|
| Piper (Qualified)                   | Inbound SDR agent. Qualifies website and inbox leads, books meetings                         | Generally available                                                                           | [12]   |
| Hunter                              | Outbound prospecting. Researches accounts, drafts outreach, pursues multi week revenue goals | Pilot, general availability planned November 2026                                             | [11]   |
| Campaign Agent                      | Builds a full campaign from a brief: audience, content, journey                              | Generally available in Marketing Cloud Next, with fuller capabilities on the Advanced edition | [11]   |
| Lead Nurturing agents (internal)    | Engage idle inbound leads, qualify, follow up, hand off                                      | Production inside Salesforce since 2025                                                       | [2]    |
| Casey, Paige, Carter, Marshall, Fin | Service, IT and HR, shopping, back office, complex customer workflows                        | Generally available                                                                           | [11]   |

The impact has legs. In the quarter ended July 31, 2026, Agentforce annual recurring revenue (ARR) exceeded \$1.5 billion, up more than 240% year over year, and customers consumed 3.2 billion Agentic Work Units in the quarter, up 97% from the prior quarter [13]. Those figures come from Salesforce's second quarter fiscal 2027 earnings release of August 26, 2026, not from an audited 10-K, and Salesforce widened the Agentforce ARR definition that quarter to include Slackbot and Headless 360, which flatters the 240% comparison. Salesforce says 30,000 customers are now live on Agentforce [11].

## The Braze deal

On October 20, 2026 Salesforce signed a definitive agreement to acquire Braze for roughly \$5.66 billion in enterprise value, or \$47.50 per share in cash, a 58% premium to the prior close [3] [10]. Braze serves about 2,770 customers including Canva, Equinox, Instacart and Venmo [14]. Its most recent quarter showed revenue of \$227.2 million, up 26.2% year over year, with dollar based net retention of 110% across all customers and paid adoption of its AI tools at roughly one third of its large customer base [15] [16]. The deal is expected to close in Salesforce's fiscal 2028, which runs February 2027 through January 2028, and carries a \$300 million termination fee on Salesforce's side [10].

Why pay 58% over market for a company growing 26%? Because Marketing Cloud never had a real time event processing layer. Braze's architecture responds to consumer behavior in real time, a capability Salesforce had not replicated internally [17]. Bloomberg's framing was blunter: Salesforce plugged a hole in products that were increasingly seen as outdated [14].

One might ask, why hasn't this been on the Salesforce product roadmap? Is the Braze acquisition a symptom of an underlying condition, that corporate innovation and portfolio-wide strategic thinking are stymied by the machine running the business? As enterprises grow, their machines move from a scale and growth flywheel to a slowly ticking clock where incrementalism is the word of the day.

Granted it's impossible for large organizations to move with any degree of alacrity, but in this instance, the market outran the pace with which a Salesforce product roadmap could be executed.

---

## 2. Is Neglected Demand the real prize?

The best first use of agents is demand your organization already generates and then abandons or disregards: the fallow deals. The scoring here is Salesforce's, not a model of ours. Salesforce rates each inbound lead one through five on purchase likelihood, and by its own account of the funnel before agents, leads scored 3 through 5 get a human, leads scored 1 and 2 get an email blast or nothing [18]. It's a dream to be able to work such a mass of neglected leads, but typically it is not feasible because of the unit economics.

### But it shifts significantly with automation.

1. Opportunity cost is low. Nobody was going to call these leads. An agent that converts 1% of them has created something from nothing, and an agent that annoys 5% of them has cost very little, because these were contacts the company had already written off. The costs that remain are deliverability, consent exposure on stale contacts, and the sales capacity to take the meetings.
2. The counterfactual is clean. Measuring agent lift on hot inbound is contaminated by the humans who would have worked it anyway. On abandoned leads, every meeting is attributable.

- 3. It does not threaten anyone's job on day one. The SDR team keeps the leads it was already working. The agent takes the pile nobody wanted. That sequencing buys time to redesign roles before anyone is asked to give one up.
- 4. It generates the data for the harder decisions. Reply rates, handoff rates, cost per meeting, and failure modes on neglected demand become the baseline for deciding whether agents should touch tier one accounts.

## Where neglected demand hides

| Pocket                           | Typical size in a mid market B2B SaaS company                                      | Signal that triggers the agent                         | Human handoff point                                                     |
|----------------------------------|------------------------------------------------------------------------------------|--------------------------------------------------------|-------------------------------------------------------------------------|
| Low scored inbound               | About 75% of inbound leads at Salesforce before agents. Your share will differ [1] | Form fill, content download, pricing page visit        | Prospect asks a question the agent cannot answer, or requests a meeting |
| Aged or dormant leads            | Months of accumulated marketing qualified leads with no sales activity             | Re-engagement: new job title, funding, site revisit    | Any reply with buying intent                                            |
| Closed lost, 12 to 24 months     | Often the largest single pool and the least worked                                 | Contract anniversary, competitor news, product release | Prospect agrees to a discovery call                                     |
| At risk pipeline                 | Deals with no activity in 30 or more days                                          | Stage stall, champion departure                        | Seller takes over the thread once the account responds                  |
| Expansion and renewal whitespace | Customers using one product with a fit for a second                                | Usage threshold, support ticket pattern                | Account manager when the customer engages                               |

Salesforce reports that a SaaS deployment saw deals close from contacts that had received no follow up for six months, with \$2.7 million in closed revenue and \$3.5 million in pipeline [19]. It is vendor reported, but it describes neglected demand, not fresh demand.

Scale cuts both ways. Salesforce had 100 million leads to work. A mid market company with 60,000 inbound leads a year has a far smaller pile but the same economics, because the agent's cost scales with conversations, not with headcount, so a pool of 40,000 unworked leads is as viable a test as one of 75 million. What changes at small scale is statistical patience. At a 1.4% positive reply rate, 40,000 leads produce about 560 replies a year, so a 90 day test yields roughly 140, enough to read direction but not to tune segments. The sizing model in Section 6 is built at that size for that reason.

## The B2C version

For consumer brands the neglected demand is the lapsed user, the abandoned cart, and the subscriber who opened three emails and bought nothing. Braze was built for this: its system watches live behavior and decides content, offer, channel and timing per person instead of running a predetermined journey [10]. Salesforce's State of Marketing found 78% of marketers need more personalized content than they can produce and only 26% are completely satisfied with their data unification [20]. Braze gives Salesforce the activation layer. Data 360 is meant to supply the unified record. Whether a given customer has the data hygiene to make that pairing work is the gating question, and it is addressed in Section 3 and again in Pitfall 4.

## How to run the first test

Start with one row of the table above and run a 90 day test.

1. Choose one pool of leads nobody is working today, for example inbound leads your sales team never calls.
2. Count how many contacts are in it and how many converted in the past, so you have a baseline.
3. Point an agent at that pool for 90 days. For the first two weeks a person reads every message before it goes out.
4. Track what each meeting costs. The comparison is against zero, because no one was working these leads before.

One test like this tells you more than a year of vendor case studies.

---

## 3. Why agents are not a plug-in: the workflow reengineering problem

There are two ways to instrument AI in a business. The first is the taxi GPS model. The GPS fitted to a taxi got it to a location more efficiently, but the fundamentals of the business did not change. Everything from the customer experience to tipping the driver stayed the same.

Contrast this with creating a ride-sharing service. The transportation business model is revolutionized. AI instrumentation must take the ride sharing model, or it will not yield enough value to justify the investment.

Choosing to overlay AI on existing processes is GPS workflow engineering. An agent dropped into a human designed workflow inherits every assumption that workflow made about who was doing the work. Campaign calendars assume a planning cycle. Lead routing assumes a rep picks up the record. Approval chains assume a person waits. None of those hold when the worker is software that acts in seconds and never sleeps. The evidence that bolting on fails is everywhere.

The failure pattern is the same in every study that has looked: the organization bought the tool and kept the process. Roughly 95% of enterprise generative AI pilots show no profit and loss impact, and the cause is a learning gap between tools and workflows rather than model quality [21]. Fewer than three in ten marketing organizations are rewiring teams and workflows rather than layering AI on top [5]. More than 40% of agentic projects are expected to be canceled by the end of 2027 for cost, unclear value or weak controls [8]. Salesforce's own experience fits: its engineers describe the agent program as turning Sales Cloud from a system of record into a system of action (a description that has been around for twenty years), and getting there meant re-architecting lead engagement rather than adding a feature to it [2].

## What has to be rebuilt

The rebuild starts with a taxonomy of every activity in one priority workflow, then a small set of agent archetypes (content generator, knowledge, localization, analyzer, planner, operator), then a map of the specific agents and the systems each must reach. One consumer brand that did this found nearly 100 modular agents inside its content workflows alone [7] [22]. The binding constraint is usually integration, not reasoning. Salesforce's early agents failed 50% to 60% of the time on infrastructure, duplicate sends and rate limits, before anyone questioned the model [2].

| Workflow element | Designed for humans                      | Redesigned for agents                                                        | Who owns the redesign                        |
|------------------|------------------------------------------|------------------------------------------------------------------------------|----------------------------------------------|
| Trigger          | Calendar, quarterly plan, campaign brief | Customer signal: behavior, stage change, time since last touch               | Marketing operations with revenue operations |
| Routing          | Lead assigned to a rep queue             | Lead engaged by agent at once, routed to a human on a defined condition      | Revenue operations                           |
| Content          | Approved assets produced in advance      | Generated per contact from an approved claims library and brand rules        | Product marketing and brand                  |
| Approval         | Manager reviews before send              | Policy encoded as guardrails, human review on a sample and on exceptions     | Marketing, legal, compliance                 |
| Measurement      | Campaign level: opens, clicks, MQLs      | Agent level: handoff rate, positive reply rate, cost per meeting, error rate | Marketing analytics                          |
| Cost control     | Seat licenses, fixed                     | Consumption: conversations, work units, tokens                               | Marketing operations with finance            |

## Data is the gate

Every agent is only as good as the record it reads. Salesforce's State of Marketing names siloed systems and poor data quality as the top barriers to personalization, and 46% of marketers say they lack data on customer preferences [6]. The new risk is simple to state: bad records used to produce bad reports, and now they produce actions taken on bad data. Duplicate contacts, loose sharing rules and a pipeline nobody trusts become agent behaviors delivered with confidence [23]. Braze's value to Salesforce depends on the same gate. Real-time decisioning on a fragmented customer record is real-time error.

## AI cost governance is a new discipline

Agentic pricing is consumption-based, and that changes how marketing budgets behave. More than half of CMOs have already shifted spend to consumption pricing, and governing that spend, not choosing tools, is becoming the defining procurement task [24] [25]. Salesforce's Koa model, a CRM tuned reasoning model still in pilot, points at the intended fix: route routine work to a cheap specialized model and reserve frontier models for judgment, which would change the unit economics of running agents at scale [11]. A CMO who cannot forecast agent consumption will discover the bill the way cloud computing customers did a decade ago. One question without a current answer is, "Why not treat the agent as a user seat?"

## Any automation should foster workflow redesign.

The sequencing that works: one workflow, mapped end to end, rebuilt with the agent as the first worker and the human as the exception handler, instrumented at the agent level, then extended. Companies that skip the mapping step buy an agent and get a faster version of the process they already had.

---

## 4. The marketing organization with agents

The marketing department does not get smaller first. It gets more expensive per head, then it changes shape, then in some functions it gets smaller. The first year of agentic marketing already shows it. Labor's share of the marketing budget went up, not down, and the organizations furthest along with AI were the least likely to cut people [9] [25]. Meanwhile the share of marketing work that AI performs is on track to roughly double by 2028 [26], while training spend is under 4% of budget [27]. Activity is automating faster than people are being re-skilled.

The shape that results is a diamond rather than a pyramid: a thin layer of executors, a wide band of people who orchestrate, judge and integrate, and a leadership layer that has to manage anxiety as actively as it manages budget. Most marketers are worried about what agents mean for their jobs, most CMOs are worried about their own, and most CMOs name that worry as a barrier to experimentation [5] [26]. Role design and change management are the same project.

## Role by role

| Role                                      | Direction                                                       | What changes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------|-----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Go to market (GTM) engineer               | New and evolving as the organization grows in AI sophistication | Builds and runs the agent and data plumbing behind outbound and lifecycle programs: enrichment, signal triggers, routing, sequencing, and the instructions agents follow. Postings for the title grew 205% from 2024 to 2025 and passed 3,000 open roles in early 2026, with 95% of them mentioning AI, large language models or agents. Prompt engineering lives inside this role rather than as a separate job: standalone prompt engineer postings fell roughly 79% from their 2023 peak while postings that list it as a required skill grew by multiples. |
| Agent operations or AI orchestration lead | New, growing                                                    | Owns the agent roster, guardrails, evaluation, consumption budget and the handoff rules. Reports into marketing operations or revenue operations. The hire that most often does not exist yet. Integrates into MarTech stack management.                                                                                                                                                                                                                                                                                                                       |
| Product marketing manager (PMM)           | Radical Evolution                                               | Becomes a player/coach: still writes the narrative, enablement, programs and runs interfaces to product and sales. But the PMM now extends to the claims library and brand rules they draw from. The PMM reviews agent output, and coaches agents the way a content marketing lead coaches sales enablement teams. 41% of PMMs hired in 2026 said the scope was not realistic for one person.                                                                                                                                                                  |

|                                        |                                        |                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Demand generation and campaign manager | Shrinks in count, rises in seniority   | Campaign assembly moves to agents such as Campaign Agent. The human performs the metacognition: directing the brief, setting goals, approving the execution and reading the results. This is where the cuts land first. One 2027 outlook has marketing leaders eliminating as many as a quarter of roles, and execution roles are the ones agents reach soonest.                      |
| SDR                                    | Shares the role with agents            | Agents take first touch and qualification on neglected demand. Humans move to discovery, objection handling and relationships.                                                                                                                                                                                                                                                        |
| Market research and insights           | Shrinks                                | Agents now run the research itself: they pull signals from CRM, call transcripts, web and survey data, synthesize them and generate the insights a research team used to deliver in weeks. The headcount that goes is the study production layer. The remaining team sets the questions, validates what agents find, and governs where synthetic respondents may and may not be used. |
| Content and creative                   | Volume moves to agents, judgment stays | What agents produce has to be carefully metered. Half of U.S. consumers say generative AI has made content quality worse, and the human role exists to protect quality, voice, brand and personality.                                                                                                                                                                                 |
| Marketing analytics                    | Changes focus                          | Shifts from campaign reporting to two jobs. The first is real time analytics that feed campaign adjustment and orchestration while the campaign is running, so spend, audience and message change in flight rather than in the post mortem. The second is agent evaluation: positive reply rate, handoff rate, error rate, cost per outcome, drift over time                          |

|                       |                |                                                                                                                                                                                                                                          |
|-----------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Early career marketer | Under pressure | Stanford's payroll analysis found employment of young workers in AI exposed occupations 19% below trend, with no comparable gap for experienced workers. The entry level rung that trained the next generation of marketers is thinning. |
|-----------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## The GTM engineer

The closest analogy is what is happening to software engineers, who are moving from writing code to directing agents that write it. Marketing splits the same way, along one line: who the agent faces. Agents that face systems, the enrichment, routing, sequencing and data plumbing inside the martech stack, belong to the GTM engineer, whose job is now managing those agents more than building the pipes by hand. Agents that face the market, the ones that write to a prospect, answer a customer or adjust a live campaign, belong to marketers, because the judgment those agents need is about positioning, tone, claims and customer context, not about infrastructure. A CMO who lets the market facing agents drift into the MarTech operations organization has outsourced the brand voice to people who were never hired to hold it.

## The product marketer as player and coach

Product marketing is the role most changed, because the agents need what PMMs produce. A Campaign Agent drafting outreach needs approved positioning, a claims library with evidence behind each claim, competitive boundaries, and tone rules. Those are PMM artifacts, and they now have to be machine readable and kept current rather than published once in a deck. The PMM is increasingly asked to design context, challenge evidence, set boundaries and orchestrate decisions across product, marketing, sales, legal and privacy [28]. In practice the PMM writes the playbook, reviews a sample of what the agents produce, corrects the instruction set when output drifts, and still carries the launch. That is a player and coach job, and it is one reason PMM scope is already judged unrealistic by 41% of recent hires [29].

## Market research contracts

Research is the function where the shrinkage is most visible and the most contested, and researchers themselves are the group most hesitant about the synthetic respondent pitch [30]. They are right to be. Synthetic panels are useful for stress testing a discussion guide or pre-testing a concept, and useless for predicting what real buyers will do, because a language model has no preferences, only an average of what people have written [31]. The headcount that disappears is the data wrangling layer: open end coding, transcript cleaning, cross tab production. What remains is smaller, more senior, and increasingly a governance role over what agents and synthetic panels are allowed to claim.

## What someone does with this

Run the table above against your own org chart. For each row, name the person who owns the redesign, the date by which the role description changes, and the skill test that gates promotion into it. Do this with HR, on spans of control, layers and reporting lines, rather than leaving the redesign to attrition [5].

---

## 5. Pitfalls

Seven failure modes recur across the deployments examined for this report. Each has a control.

### Pitfall 1: Volume becomes the strategy

The cheapest thing an agent can do is send more, and that is what most teams let it do. Outbound volume per rep has gone up roughly sixfold since AI arrived while reply rates have fallen by about a third, and a growing share of what is sent never reaches an inbox at all [32] [33]. The buyer has learned the texture of generated text and deletes it on sight. My read is that the inbox is now a commons being overgrazed, and the first vendor whose agents send less, on a real signal, will stand out precisely because everyone else's send more. Even Salesforce's own marketing leadership concedes the pattern: the most powerful technology in history is being used to send more one way spam, faster [6].

Control: cap agent volume per domain, require a trigger signal for every outbound touch, and measure positive reply rate and meetings rather than sends.

### Pitfall 2: Reliability is assumed rather than engineered

Left to itself, a language model agent completes about a third of multi step business tasks [34]. Reliable production agents are therefore an engineering achievement, not a model property. Salesforce's own went from failing more than half the time to producing 10,000 opportunities only after a control plane stopped duplicate sends and overloads [2]. Buyers should assume that any vendor's demo reflects the model and that their own results will reflect the plumbing.

Control: deterministic checks before any action with consequences, a persistent queue so each contact is touched once, and a human review sample that never drops to zero.

## Pitfall 3: Headcount is cut before the agent has earned it

Klarna is the cautionary case. In 2024 it announced that its AI assistant did the work of 700 customer service agents and froze hiring. By 2025 it was recruiting humans again after service quality declined, and it settled on a hybrid model with AI on routine cases and people on the rest [35]. The error was not using agents. It was taking the headcount saving before the agent had demonstrated, on live customers, that it could hold quality at the volume the saving assumed. Marketing has the same temptation: a Campaign Agent that drafts in minutes makes the campaign team look expensive long before anyone has measured what the drafts do in market.

Control: redesign the role before you reduce the role. Run the agent on neglected demand first, where no human loses work, and only then move it onto work people currently do.

## Pitfall 4: Bad data becomes bad action

Agents act on the record as it stands. Duplicate contacts, stale ownership, and loose sharing rules that used to produce bad reports now produce outreach to the wrong person, two agents working the same lead, and confident answers from a pipeline nobody trusts [23]. Only 26% of marketers are completely satisfied with how unified their customer data is [20]. For Salesforce, Braze's real-time decisioning raises the stakes: a wrong signal becomes a wrong message and lost opportunity within sixty seconds.

Control: a data readiness gate per workflow. No agent goes live on a record set that has not passed deduplication, ownership and consent checks.

## Pitfall 5: Consumption cost outruns the budget

Agentforce is priced on usage, in conversations and work units, and most CMOs have already moved some spend onto consumption terms without yet having a way to forecast it [25]. Cost is one of the three reasons agentic projects get canceled, alongside unclear value and weak controls [8]. The pattern will be familiar to anyone who ran a cloud migration: the pilot is cheap, the production bill arrives in month four, and the CFO asks who approved it.

Control: a monthly consumption forecast owned jointly by marketing operations and finance, with per workflow budgets and automatic throttles.

## Pitfall 6: Trust erodes on the buyer's side

Buyers are adjusting faster than sellers. They already trust AI generated information less, they are adding human conversations to verify what agents tell them, and they will walk from vendors whose claims they cannot check [36] [37]. Half of consumers think generative AI has made content worse [38]. The conclusion for a marketing leader is uncomfortable: every message an agent sends draws on a trust balance the brand built with humans, and an unverifiable claim overdraws it.

Control: every claim an agent can make traces to an approved, evidenced claims library owned by product marketing. Agents disclose that they are agents. Adversarial testing agents should be deployed to assess content quality before being shared publicly.

## **Pitfall 7: The organization is not redesigned, so nothing scales**

This is the pitfall under all the others, and it is the taxi GPS model from Section 3 at the scale of the whole department. The AI is installed in the taxi like a GPS unit on the dashboard. The dispatcher, the fare structure, the shift pattern and the way a customer hails a ride are untouched, so the company gets a slightly faster taxi and none of what the ride sharing model made possible. Every failure statistic cited in Section 3 describes that organization: one that bought the GPS and kept the taxi company [21] [5] [4]. Worse, adding AI to an initiative where decision rights are already unclear amplifies the confusion rather than resolving it [5]. A ride sharing company cannot run on a taxi service business model, and an agentic marketing function cannot run on top of existing workflows.

Control: decision rights first. For each workflow, write down who approves the agent's instructions, who reviews its output, who owns its budget, and who can switch it off.

---

## **6. Opportunities**

The upside is concentrated in five places, and most of it is available without betting the department or your career.

### **Opportunity 1: Work the demand you already paid for**

The largest and cheapest gain is the pile from Section 2. Salesforce's figures, whatever their exact accounting, came from leads the company had already generated and abandoned [1] [18]. Any marketing organization with a lead scoring model has an equivalent pile. Every meeting from it is incremental, and the acquisition cost was already sunk.

## Opportunity 2: Collapse campaign cycle time

Speed is the first dividend, and it compounds. A campaign that took ten weeks to build takes hours when an agent assembles brief, audience, content and journey from one instruction [11]. The brands that have done it report campaign creation two to three times faster and, more important, two to three times as many tests in market [12] [5]. I discount every vendor and consultant number here by half and the conclusion still holds: the team that runs twelve tests a quarter learns faster than the team that runs four, and the learning rate, not any single campaign, is what moves the return on marketing spend [20].

## Opportunity 3: Move from campaigns to conversations

The strategic prize is the one Salesforce is paying \$5.66 billion to chase: responding to each customer's signal rather than scheduling messages at segments. Customers already expect a reply, not a broadcast, and the marketing teams that outperform are the ones that have put agents in the conversation rather than in the content queue. High performing teams are about 1.7 times more likely than the rest to have unified their data and to use it for relevant experiences, which is the precondition for agents in the conversation [20]. The upper bound of the prize is large: as much as two thirds of today's marketing activity is automatable and the modeled revenue lift from true one to one engagement runs 10% to 30% from hyper-personalized marketing [7]. The revenue figure is a model, not an observation, and should be treated as a ceiling.

## Opportunity 4: Reset the CMO's standing with the CEO and CFO

Agents hand the CMO the metric the CEO already uses. CEOs grade marketing on revenue and margin. Most CMOs still report on reach and engagement, and only about a third track revenue and margin at all [5]. Agents produce outcome data by construction: meetings booked, pipeline created, cost per outcome. A CMO who runs the first neglected demand test and reports it in pipeline dollars against a zero baseline has a story the CFO understands, and the organizations that have made that case are already getting a larger share of company revenue as budget [4].

## Opportunity 5: Hire and keep the people who can direct agents

Labor's rising share of the marketing budget, after allowing for the fall in martech spend, is my signal that the scarce input is now judgment, not hands [9]. Organizations that define the orchestration roles early, write the skill standards, and promote into them will hold the people everyone else is about to try to hire. That means a baseline skill standard for every role, with the investment going into strategic thinking and business acumen rather than tool training [26], delivered as continuous coaching rather than a one time course, with a change champion in each team [5].

## Sizing the opportunity: an illustrative model

The table below is a worked example for a hypothetical B2B software company with \$200 million in revenue, 60,000 inbound leads a year, and a sales team that works the top 30% of them. Every input is an assumption to be replaced with your own.

| Input                                              | Assumption                                                                               | Basis                                                                        |
|----------------------------------------------------|------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| Leads not worked by humans                         | 42,000 a year (70%)                                                                      | Salesforce reported 75% unanswered before agents [1]                         |
| Agent positive reply rate                          | 1.4%                                                                                     | Positive reply rate for AI sent email in a 100,000 message analysis [39]     |
| Positive replies converting to a qualified meeting | 50%                                                                                      | Analyst assumption                                                           |
| Meetings a year                                    | 294                                                                                      | $42,000 \times 1.4\% \times 50\%$                                            |
| Meeting to opportunity                             | 40%                                                                                      | Analyst assumption                                                           |
| Average opportunity value                          | \$60,000                                                                                 | Analyst assumption for mid market                                            |
| Pipeline created                                   | About \$7.1 million a year                                                               | $294 \times 40\% \times \$60,000$                                            |
| Agent cost                                         | Roughly \$150,000 to \$400,000 a year                                                    | Consumption pricing varies widely by vendor and volume. Replace with a quote |
| Human reviewer and operations lead                 | One to two full time equivalents at a loaded cost of roughly \$150,000 to \$300,000 each | Section 4                                                                    |

At these assumptions, with agent consumption and the people who run it combined at roughly \$450,000 to \$700,000 a year, the pipeline to cost ratio is about 10 to 1 to 16 to 1 before any revenue converts, and the opportunity cost is low because no human was working these leads. Halve every conversion assumption and the test still produces about \$880,000 of pipeline, above cost. Pipeline is not revenue: at a 25% win rate that halved case returns roughly \$220,000, below cost, which is why Gate 1 is a pipeline test and the revenue question belongs to Gate 2. That robustness is why neglected demand is the first and best place to start.

---

## 7. Implications for CMOs

The CMO's job description is changing faster than the CMO's budget. Two thirds of CMOs expect AI to change their own role dramatically within two years [40]. Agentic team design now belongs at the top of the CMO agenda [41], and the remit is widening from brand and demand to the orchestration of data, technology and AI enabled execution [7]. Six implications follow.

### 1. You now run a workforce that includes software, and you are accountable for what it does

Marketing has run sophisticated software for two decades: nurture journeys, bidding algorithms, personalization engines. None of it was a workforce, because none of it decided anything a person had not scripted in advance. Agents are different in kind. They read a situation, choose an action, and write the words, which is what a junior employee does. An agent that sends 7,000 emails a month under your brand is therefore a marketing employee without a manager unless you give it one. Decision rights for instructions, review, budget and shutdown belong on a page with names next to them.

### 2. The budget conversation changes from seats to consumption

Marketing technology fell to 19.4% of budget in 2026, a five year low, while 56% of CMOs moved spend to consumption pricing and labor rose to 24.5% [25] [9]. Expect the CFO to add agent consumption to the projected versus billed spend review, and to ask why the variance moved.

### 3. Your talent plan is the transformation plan

Acknowledge employee anxiety, but give them radical role clarity and one on one training. You should adjust incentives accordingly and work with HR on spans, layers and reporting lines. Employee trust in the organization's support through AI driven change predicts enterprise value from AI [5]. The CMO who treats agents as a technology purchase is no people leader.

### 4. Product marketing becomes infrastructure

Every agent that writes to a customer draws on positioning, claims, proof points and tone rules. Positioning has always been the operating instruction set for everyone who speaks for the brand. What changes is the volume of evidence behind it. Agents generate response data on every claim, every segment and every variant, so the PMM can now decide what to modify, what to add and what to retire with far more precision than a quarterly win loss review allowed, and can let the positioning flex for personalization without losing the core. That only works if the claims library is maintained like code: versioned, evidenced, reviewed, and fed by the agents' own results. The PMM team's output becomes a dependency for sales, service and marketing agents at once. Fund it accordingly, and resist the temptation to treat it as the first place to save headcount.

## 5. Trust is the asset most at risk

Buyers are responding to agent generated outreach by trusting it less and checking it more. Expect them to double the human conversations they use to validate a recommendation, and expect vendors whose claims cannot be verified to lose the deal [37]. The brands that win will be the ones whose agents are accurate, disclosed, and backed by a human who shows up when asked. Governance imposed from the top will not hold at the pace agents operate. It has to live with the people who brief and review them, which means raising the AI judgment of every marketer rather than appointing a committee [36].

## 6. Measurement moves to outcomes, which is where the CEO already was

For twenty years CEOs have asked CMOs to prove their worth in revenue, and for twenty years CMOs have pushed back, with some justification: attribution was a model, brand effects were slow, and the handoff to sales broke the chain of evidence [5]. Agents end the argument, for three reasons. First, the same agent that marketing deploys now qualifies the lead, books the meeting and writes the opportunity, so the chain from marketing touch to pipeline is a single logged record rather than an attribution model. Second, agents are priced on consumption, so every workflow carries its own cost ledger beside its own outcome ledger, and return on investment per workflow is arithmetic rather than argument. Third, marketing's agents are doing work that used to belong to sales development, which means the CMO owns pipeline whether or not the org chart says so. The defense that marketing's contribution cannot be measured is gone. The CMO who gets ahead of this reports the first neglected demand test as pipeline against a zero baseline and sets the terms. The one who does not will have the terms set by the CFO reading the agent logs.

## The B2B and B2C split

| Dimension | B2B CMO | B2C CMO |
|-----------|---------|---------|
|-----------|---------|---------|

|                      |                                                                           |                                                                                 |
|----------------------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| First agent use case | Neglected inbound and dormant pipeline (Piper, Hunter, custom SDR agents) | Lapsed users, abandoned carts, real time triggers (Braze, Marketing Cloud Next) |
| Primary risk         | Buyer trust and inbox saturation                                          | Data quality at real time speed and consent                                     |
| Role most changed    | Product marketing and SDR management                                      | Lifecycle and CRM marketing, creative operations                                |
| Governance partner   | Revenue operations, sales leadership                                      | Data engineering, privacy, legal                                                |
| What to tell the CFO | Pipeline created from zero baseline, cost per meeting                     | Incremental revenue per triggered message, consumption per customer             |

## 8. CMO roadmap: five gates, four in the first 12 months

Each gate has an exit criterion. Do not pass a gate on a date. Pass it on evidence.

| Gate                                 | Months | What gets done                                                                                                                                                                                                                                                                                                                                                                                              | Exit criterion                                                                                                                                                                                                                                                                     | Who signs             |
|--------------------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Gate 1: Prove it on neglected demand | 0 to 3 | Size the unworked lead pool. Pass a data readiness check on that pool only. Deploy one agent (vendor prebuilt or custom) with a human reviewing every first touch for two weeks, then a 10% sample. Instrument positive reply rate, agent to MQL conversion, cost per meeting, spam flag rate. Treat this as a proof of concept, a GPS phase that buys evidence. The value comes from the rebuild in Gate 2 | Cost per meeting below the blended SDR cost is the floor, and it clears easily at consumption prices. The real bar: agent to MQL conversion at or above the human worked rate for the same lead tier, spam flag rate under 3%, and at least one closed deal traceable to the agent | CMO and head of sales |

|                                         |         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                         |                                       |
|-----------------------------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| Gate 2: Rebuild one workflow end to end | 3 to 6  | Take the neglected demand workflow that Gate 1 proved and rebuild it end to end. It is the right first workflow because the proof of concept has already shown the agent works on it and nobody's current job depends on it. Map every activity. Redesign trigger, routing, content, approval, measurement and cost control per Section 3. Stand up the claims library under product marketing. Name the agent operations owner                                                 | Workflow runs with the agent as first worker and humans on exceptions. Consumption forecast within 15% of billed spend for two consecutive months                                       | CMO, CFO delegate, revenue operations |
| Gate 3: Redesign the organization.      | 6 to 9  | Rewrite job descriptions for the roles in Section 4, and redraw the structure around them: who reports to whom, where the agent operations and GTM engineering functions live, how many layers, and which teams merge now that agents cross the old campaign, content and demand generation boundaries. Set a skill floor per role and a test for it. Launch continuous coaching, not a one time course. Redeploy before reducing. Publish decision rights for every live agent | New org chart approved and in effect, with every live agent assigned a named owner, reviewer, budget holder and kill switch. Attrition in affected roles no worse than company baseline | CMO and Chief Human Resources Officer |
| Gate 4: Extend and govern               | 9 to 12 | Extend agents to a second and third workflow. Move from vendor reported metrics to your own agent evaluation dashboard (Section 9). Add brand trust controls: disclosure, claims tracing, audit logs. Negotiate consumption terms with throttles                                                                                                                                                                                                                                | Three workflows live. Agent produced pipeline reported to the board in the CFO's units. No unresolved trust incident                                                                    | CMO to CEO and board                  |

|                             |                         |                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                  |                                |
|-----------------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Gate 5: Evolve continuously | 12 onward, never closes | Treat the operating model as a release, not a design. Every quarter, reassess which tasks agents can now do that they could not a quarter ago, retire guardrails that model progress has made unnecessary, add new ones where new capabilities create new risk, and revisit the role and structure decisions from Gate 3. Re-run the Gate 1 test on a new pocket of demand each quarter so the baseline keeps moving | A quarterly capability review on the calendar with decisions recorded. At least one workflow, role or guardrail changed each quarter as a result | CMO with agent operations lead |
|-----------------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|

Buy or build

Purchased solutions succeed roughly three times as often as internal builds in enterprise AI, and marketing is not the function to bet against that pattern [21]. So buy the agents. Use the vendor's prebuilt ones through Gates 1 and 2 and write custom agents only where no prebuilt one fits the workflow. You build and manage the institutional knowledge and partner with AI on the content building blocks. Buying from the vendor whose consumption pricing Pitfall 5 warns about is not a contradiction if the contract carries the forecast and throttles that pitfall calls for. Buy the agent. Do not buy it uncapped.

9. The CMO operating dashboard

Nine measures, reviewed monthly, replace the campaign report. Thresholds are starting points drawn from the public benchmarks cited and should be reset against your own Gate 1 results.

| Measure             | Definition                                                                    | Starting threshold                     | Comments                             |
|---------------------|-------------------------------------------------------------------------------|----------------------------------------|--------------------------------------|
| Unworked demand     | Share of inbound and dormant contacts with no human or agent touch in 30 days | Falling toward under 10%               | The pile agents exist to clear       |
| Positive reply rate | Replies with buying intent                                                    | Above 1.4% cold, above 5% on triggered | AI sent email benchmark is 1.4%      |
| Meeting rate        | Meetings booked divided by sends                                              | Above 0.7% cold                        | AI benchmark 0.7% against human 1.1% |

|                              |                                                                                                                                                             |                                                          |                                                                                    |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|------------------------------------------------------------------------------------|
| Agent to MQL conversion rate | Contacts the agent engaged that reached marketing qualified lead (MQL) status, divided by contacts engaged                                                  | At or above the human worked rate for the same lead tier | The first proof that agent engagement produces qualified demand, not just activity |
| Agent task success           | Multi step tasks completed without correction                                                                                                               | Above 80% on deployed workflows                          | Unaided benchmark is 35%                                                           |
| Consumption variance         | Billed agent spend against forecast                                                                                                                         | Within 15%                                               | Consumption pricing is the new budget risk                                         |
| Claims compliance            | Agent messages sampled that match the approved claims library                                                                                               | 100% on sample                                           | Buyer trust and the \$10 billion risk                                              |
| Agent produced pipeline      | Pipeline dollars from agent sourced opportunities, carried as its own line in the weekly sales forecast with stage and close date like any rep sourced deal | Input to the weekly sales forecast                       | Salesforce's \$88 million is this measure                                          |
| Team sentiment               | eNPS pulse on confidence in the organization's support through AI change                                                                                    | Improving quarter over quarter                           | Trust predicts AI value capture                                                    |

## 10. What not to do

1. Do not start on your best leads. Agents on hot inbound contaminate the measurement and threaten the SDR team on day one.
2. Do not let volume be the metric. An agent will send as much as you let it, and every extra send past the point of relevance costs reply rate, domain reputation and buyer patience. Sends are a cost. Positive replies and meetings are the output. Measure the output and cap the cost.
3. Do not cut headcount before the agent has earned it. The saving looks real on a spreadsheet months before the agent has proven it can hold quality at the volume the saving assumes. Companies that cut first spend the following year rehiring, usually at higher cost and with a story they would rather not tell. Let the agent carry the work for two full quarters, with the numbers to show it, before a single role is modified in the plan.
4. Do not go live on data that isn't cleaner than your existing data cleansing outputs. Bad records now produce bad actions, not just bad reports.
5. Do not sign a consumption contract without a forecast and a throttle. Runaway cost kills more agent programs than bad output does, and a pilot that was cheap at 1,000 conversations a month is a different bill at 50,000.

6. Do not let agents make claims that are not in the library. A buyer who catches one unverifiable claim discounts everything else the brand says, and agents can make that claim ten thousand times before anyone notices.
7. Do not run this as an IT project. The deployments that work are owned by the people who own the workflow, not by a central AI team that owns the platform. The CMO owns it.
8. Do not add agents to a workflow where decision rights are unclear. An agent executes whatever the process tells it, at speed and without judgment. Where two people already disagree about who approves a message or owns a lead, the agent will act on one answer and surface the conflict as a customer facing error.

---

## Conclusion: the right to delegate has to be earned

Salesforce has demonstrated that an agent can turn a signal into a meeting at a scale no human team could staff, and it is paying \$5.66 billion to do the same thing for consumer brands in real-time. Both moves are coherent, and both are early. The \$88 million is a proof of mechanism with its accounting still in the drawer. Braze will not be inside Marketing Cloud until some point in Salesforce's fiscal 2028 [10].

What is not early is the organizational bill. Labor's share of marketing budgets rose, not fell, in the first full year of agentic marketing. Although small and medium marketing organizations have more success in rebuilding workflows with embedded agents, most CMOs admit their own processes could not scale such a rebuild. The agents work. The departments around them do not yet.

The CMO's task is to earn the right to delegate: prove it on demand nobody was working, rebuild one workflow around the agent, redesign the roles while the stakes are low, govern cost and trust before the board asks, and then revisit all of it every quarter, because the agents will be more capable by then and the model has to move with them. Done in that order, agentic marketing is the largest capacity increase the function has seen. Done in any other order, it is a faster way to send the campaigns customers already ignore.

---

## Methodology, assumptions and limitations

Sources and weighting. This report draws on 52 public sources opened between October 20 and October 27, 2026: company press releases and investor filings, a Salesforce engineering post, vendor research (Salesforce State of Marketing, 4,450 respondents), analyst research (Gartner CMO Spend Survey, 401 respondents, fielded January to March 2026, plus a Gartner marketing leader survey of 402 CMOs fielded August to October 2025, plus Forrester 2026 and 2027 B2B predictions), consulting research (McKinsey Global Marketer Survey, 521 respondents, March 2026), academic work (MIT NANDA, Stanford Digital Economy Lab, Salesforce AI Research's CRM Arena-Pro), and trade reporting. Vendor figures are labeled as such wherever they appear. Where a recommendation rests mainly on Salesforce reported results, as the neglected demand test does, the report says so and treats the figures as a hypothesis to test rather than a benchmark.

On the \$88 million. Salesforce has disclosed neither the measurement period, the attribution rule, nor the conversion of this pipeline to revenue. The April 2026 engineering figure of more than \$100 million and the October 2026 sales figure of more than \$88 million are reported separately and not reconciled. Readers should quote them as a range with speaker attribution.

On outbound benchmarks. The AI SDR reply, meeting and spam rate figures in Sections 5, 6 and 9 come from practitioner and vendor analyses of their own sending data, not from peer reviewed studies. They are consistent with one another in direction and are used as starting thresholds, not as established constants.

On the sizing model. The worked example in Section 6 is illustrative. Every conversion assumption is labeled, and the conclusion drawn from it is about robustness to halved assumptions, not about the point estimate.

On role predictions. Statements about which roles shrink or grow combine survey data (Gartner, McKinsey, Duke CMO Survey), labor market data (Indeed, LinkedIn aggregates, Stanford payroll analysis) and practitioner accounts. Labor market data on titles such as prompt engineer lags and conflates title with skill. Directional confidence is high. Magnitude confidence is moderate.

What this report does not cover. Agent security and prompt injection, regulatory treatment of automated outreach under privacy and consent law, and the competitive responses of Adobe, HubSpot, Microsoft and Oracle are out of scope and each merits its own analysis.

## SOURCES / CITATIONS

# Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

- 01**      **Salesforce says AI agents generated \$88 million in sales pipeline by reaching leads its staff could not handle, CTech by Calcalist, October 2026**  
<https://www.calcalistech.com/ctechnews/article/hyengkz2ze>
- 02**      **How Agentforce Lead Nurturing Agents Generated \$100M+ Pipeline Under Rate-Limited Infrastructure, Salesforce Engineering, April 20, 2026**  
<https://engineering.salesforce.com/how-agentforce-lead-nurturing-agents-generated-100m-pipeline-under-rate-limited-infrastructure/>
- 03**      **Salesforce Signs Definitive Agreement to Acquire Braze, Salesforce press release, October 20, 2026**  
<https://salesforce.com/news/press-releases/2026/10/20/salesforce-signs-definitive-agreement-to-acquire-braze>
- 04**      **Gartner 2026 CMO Spend Survey Finds CMOs Allocate 15.3% of Marketing Budgets to AI, But Only 30% Are Ready to Scale AI Capabilities, Gartner, May 11, 2026**  
<https://www.gartner.com/en/newsroom/press-releases/2026-05-11-gartner-2026-cmo-spend-survey-finds-cmos-allocate-15-point-3-percent-of-marketing-budgets-to-ai-but-only-30-percent-are-ready-to-scale-ai-capabilities>
- 05**      **From anxiety to advantage: A marketing organization that thrives with AI, McKinsey, 2026**  
<https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/from-anxiety-to-advantage-a-marketing-organization-that-thrives-with-ai>
- 06**      **75% of Marketers Have Adopted AI Yet Still Use It To Send One-Way, Generic Campaigns, Salesforce Newsroom, 2026**  
<https://www.salesforce.com/news/stories/state-of-marketing-2026/>
- 07**      **Reinventing marketing workflows with agentic AI, McKinsey, April 21, 2026**  
<https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/reinventing-marketing-workflows-with-agentic-ai>
- 08**      **Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, Gartner, June 25, 2025**  
<https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
- 09**      **Gartner Marketing Survey Finds Awareness and Conversion Account for 62.6% of Total Media Spend, Gartner, June 8, 2026**  
<https://www.gartner.com/en/newsroom/press-releases/2026-06-08-gartner-marketing-survey-finds-awareness-and-conversion-account-for-62-6-of-total-media-spend>
- 10**      **Salesforce to acquire Braze for \$5.66 billion to expand AI-powered marketing, Digital Commerce 360, October 22, 2026**  
<https://www.digitalcommerce360.com/2026/10/22/salesforce-to-acquire-braze-for-5-66-billion-to-expand-ai-powered-marketing/>
- 11**      **Dreamforce 2026: Everything Salesforce Announced, and What It Means, Cloudsheer, September 2026**  
<https://www.cloudsheer.com/blog/dreamforce-2026-announcements>
- 12**      **Salesforce Puts an AI Marketing Team in Every Marketer's Hands, Salesforce Newsroom, 2026**  
<https://www.salesforce.com/news/stories/agentic-marketing-teams-announcement/>

- 13 **Salesforce Delivers Record Second Quarter Fiscal 2027 Results, Salesforce Investor Relations, August 26, 2026**  
<https://investor.salesforce.com/news/news-details/2026/Salesforce-Delivers-Record-Second-Quarter-Fiscal-2027-Results/default.aspx>
- 14 **Salesforce Fills Marketing Hole With \$5.7 Billion Deal for Braze, Bloomberg, October 20, 2026**  
<https://www.bloomberg.com/news/articles/2026-10-20/salesforce-to-buy-marketing-platform-braze-for-5-7-billion>
- 15 **Braze Reports Strong Fiscal Second Quarter 2027 Results, Business Wire, September 8, 2026**  
<https://secure.businesswire.com/news/home/20260908586699/en/Braze-Reports-Strong-Fiscal-Second-Quarter-2027-Results>
- 16 **Braze Q2 Earnings Call Highlights, Nasdaq, September 2026**  
<https://www.nasdaq.com/articles/braze-q2-earnings-call-highlights>
- 17 **Salesforce Buys Braze for \$5.66B; Marketing Cloud Lacked Real-Time Consumer Event Processing, Tech Times, October 21, 2026**  
<https://www.techtimes.com/articles/329295/20261021/salesforce-buys-braze-566b-marketing-cloud-lacked-real-time-consumer-event-processing.htm>
- 18 **How Salesforce's AI agent uncovered \$1.7 million in pipeline, Revenue Brew, September 11, 2025**  
<https://www.revenuebrew.com/stories/2025/09/11/salesforce-ai-agent>
- 19 **Five Big Ideas Coming Out of Dreamforce 2026, Salesforce Newsroom, September 2026**  
<https://www.salesforce.com/news/stories/five-dreamforce-2026-takeaways/>
- 20 **State of Marketing Report, Tenth Edition, Salesforce, 2026**  
<https://www.salesforce.com/marketing/resources/state-of-marketing-report/>
- 21 **MIT report: 95% of generative AI pilots at companies are failing, Fortune on MIT NANDA's The GenAI Divide, August 18, 2025**  
<https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo>
- 22 **Re-inventing Marketing with agentic AI, the McKinsey method, Diginomica, June 3, 2026**  
<https://diginomica.com/node/30619>
- 23 **Dreamforce 2026: What It Means for Real Estate, Phenoble, September 2026**  
<https://phenoble.com/blogs/dreamforce-2026-what-it-means-for-real-estate>
- 24 **Hype Cycle for Digital Marketing, 2026, Gartner, July 10, 2026**  
<https://www.gartner.com/en/documents/8115797>
- 25 **Gartner CMO Spend Survey: Budgets Reflect Increase in Consumption-Based Martech, Paid Media Spend, Chief Marketer, 2026**  
<https://www.chiefmarketer.com/gartner-cmo-spend-survey-budgets-reflect-increase-in-consumption-based-martech-paid-media-spend/>
- 26 **Gartner Survey Reveals Marketing Leaders Expect AI Automation of Marketing Work to Double to 36% by 2028, Gartner, May 11, 2026**  
<https://www.gartner.com/en/newsroom/press-releases/2026-05-11-gartner-survey-reveals-marketing-leaders-expect-ai-automation-of-marketing-work-to-double-to-36-percent-by-2028>
- 27 **The CMO Survey Topline Report, 35th Edition, Duke University Fuqua School of Business, Deloitte and the American Marketing Association, Spring 2026**  
[https://cmosurvey.org/wp-content/uploads/2026/03/The\\_CMO\\_Survey-Topline\\_Report-2026.pdf](https://cmosurvey.org/wp-content/uploads/2026/03/The_CMO_Survey-Topline_Report-2026.pdf)
- 28 **Product Marketing in 2026: Eight Changes Reshaping GTM, MTF Institute, 2026**  
<https://mtfinstitute.com/insights/product-marketing-2026-buyer-discovery-messaging-launch-measurement/>

- 29 **How AI Is Changing Product Marketing (and Why It Won't Replace PMMs), Calven, citing the Fluvio 2026 Hiring Trends Report**  
<https://calven.ai/resources/ai-for-pmm/how-ai-is-changing-product-marketing>
- 30 **Synthetic Respondents: How AI Is Mimicking Customers, Experts, and Employees, Bell and Holmes, citing the Rival Group 2026 Market Research Trends Report**  
<https://www.bellandholmes.com/blog/synthetic-respondents-ai-mimicking-customers-experts-employees>
- 31 **The Future of Market Research With AI: 7 Shifts Research Leaders Need to Plan For, Perspective AI, 2026**  
<https://s.getperspective.ai/blog/the-future-of-market-research-with-ai-7-shifts-research-leaders-need-to-plan-for>
- 32 **Do AI SDRs work? What the public record says, Draftship, 2026, citing Salesmotion's 2026 AI SDR comparison**  
<https://draftship.io/blog/do-ai-sdrs-work-2026>
- 33 **44% of B2B Sales Teams Now Use AI SDRs in 2026, LinkedOtter citing G2 and Warmly, June 2026**  
<https://www.linkedotter.com/articles/ai-sdr-44-percent-b2b-sales-adoption-june-2026>
- 34 **CRMArena-Pro: Holistic Assessment of LLM Agents Across Diverse Business Scenarios and Interactions, Salesforce AI Research, arXiv, 2025**  
<https://arxiv.org/pdf/2505.18878>
- 35 **Klarna Turns From AI to Real Person Customer Service, Bloomberg, May 8, 2025**  
<https://www.bloomberg.com/news/articles/2025-05-08/klarna-turns-from-ai-to-real-person-customer-service>
- 36 **Forrester's 2026 B2B Marketing, Sales, And Product Predictions, Forrester, October 28, 2025**  
<https://www.forrester.com/press-newsroom/forrester-b2b-marketing-sales-product-2026-predictions/>
- 37 **Forrester's 2027 B2B Go-To-Market Predictions: AI Agents Will Influence 75% Of B2B Purchase Decisions, Forrester via PR Newswire, October 22, 2026**  
<https://www.prnewswire.com/apac/news-releases/forresters-2027-b2b-go-to-market-predictions-ai-agents-will-influence-75-of-b2b-purchase-decisions-302915245.html>
- 38 **Gartner Survey Finds 49% of U.S. Consumers Say GenAI Has Made Content Quality Worse, Gartner, June 9, 2026 (307 U.S. consumers, fielded March 2026)**  
<https://www.gartner.com/en/newsroom/press-releases/2026-06-09-gartner-survey-finds-49-percent-of-us-consumers-say-genai-has-made-content-quality-worse>
- 39 **AI SDR Real Performance: 100K Email Analysis 2026, Digital Applied, 2026**  
<https://www.digitalapplied.com/blog/ai-sdr-real-performance-100k-email-analysis-2026>
- 40 **Marketing Predictions for 2026, Gartner, citing the 2025 Gartner Marketing Transformation Survey**  
<https://www.gartner.com/en/marketing/research/marketing-predictions/unlocked>
- 41 **Gartner's 2026 CMO research agenda puts agentic AI and brand value at the center of marketing leadership, MarketScale, 2026**  
<https://www.marketscale.com/industries/marketing-tech/gartners-2026-cmo-research-agenda-puts-agentic-ai-and-brand-value-at-the-center-of-marketing-leadership>
- 42 **GTM Engineering Data: What 209 Job Postings Reveal, OnFire, 2026**  
<https://www.onfire.ai/blog/gtm-engineering-data>

- 43 **Prompt Engineer Jobs Are Gone. Here's What Replaced Them, Doghouse, citing Indeed Hiring Lab, 2026**  
<https://www.doghouse.nl/post/prompt-engineer-career-what-happened>
- 44 **AI didn't kill the market researcher. It rewrote the job description, Rival Technologies, 2026**  
<https://www.rivaltech.com/blog/ai-didnt-kill-market-research>
- 45 **Marc Benioff's Dreamforce 2026 Keynote: The SaaSocolypse Solution is Unlocking Trapped Enterprise Data, G2 Learn, September 16, 2026**  
<https://learn.g2.com/dreamforce-2026-keynote-recap>
- 46 **Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of AI, Stanford Digital Economy Lab, August 2026**  
[https://digitaleconomy.stanford.edu/app/uploads/2026/08/Canaries\\_August2026.pdf](https://digitaleconomy.stanford.edu/app/uploads/2026/08/Canaries_August2026.pdf)
- 47 **State of Agentic Deployment in Marketing, October 2026, Acalytica, October 2026**  
<https://acalytica.com/state-of-agentic-marketing-2026/>
- 48 **The Future of Work in the AI Era, Frontier Research, Dion Hinchcliffe, October 16, 2026 (structural reference for this report)**  
<https://www.thisisfrontier.ai/research/the-future-of-work-in-the-ai-era>
- 49 **Is prompt engineer a real job? 2026 Guide, PE Collective, 2026**  
<https://pecollective.com/blog/is-prompt-engineering-a-real-career/>
- 50 **FAQ on GTM engineering: Automating B2B's revenue growth potential in 2026, eMarketer, citing Bloomberg's analysis of 1,000 job postings, 2026**  
<https://www.emarketer.com/content/faq-on-gtm-engineering--automating-b2b-s-revenue-growth-potential-2026>
- 51 **Salesforce sacrifices 4,000 support jobs on the altar of AI, The Register, September 2, 2025**  
[https://www.theregister.com/2025/09/02/salesforce\\_4000\\_jobs\\_ai](https://www.theregister.com/2025/09/02/salesforce_4000_jobs_ai)
- 52 **Salesforce cuts 4,000 support roles for AI but plans to hire 5,000 sales staff, says CEO Marc Benioff, Storyboard18, October 20, 2025**  
<https://www.storyboard18.com/brand-makers/salesforce-cuts-4000-support-roles-for-ai-but-plans-to-hire-5000-sales-staff-says-ceo-marc-benioff-82829.htm>

**FRONTIER.**

TAKE THE NEXT STEP

# Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

## RESEARCH / BRIEFINGS

**[inquiries@thisisfrontier.ai](mailto:inquiries@thisisfrontier.ai)**

Briefings, advisory work, executive discussions, and research inquiries.

## REPRINTS / LICENSING

**[reprints@thisisfrontier.ai](mailto:reprints@thisisfrontier.ai)**

Reprints, enterprise distribution, excerpts, and licensing.

## FRONTIER

**[www.thisisfrontier.ai](http://www.thisisfrontier.ai)**

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

**SEE FURTHER. MOVE SOONER.**