

The State of Agentic Marketing

How AI Agents Are Reshaping Marketing Operations

91%

AI adoption in marketing

57%

Scaling AI usage

~25%

Email + automation lead

AT A GLANCE

Key Findings

AI adoption in marketing has moved from experimentation to default behavior. The challenge is no longer whether teams will use AI, but whether they can scale it across real marketing operations workflows.

Adoption has crossed the mainstream threshold. The next gap is operational: governance, measurement, and workflow scale.

Who should read this

- Marketing ops leaders
- CMOs planning AI investment
- Lifecycle and demand-gen teams
- RevOps and martech architects

91%

of marketing teams now use AI, up from 63% year over year.

45%

use AI for campaign optimization, the top MOps-related use case.

57%

plan to increase AI use for content production and operations.

41%

can confidently prove ROI from AI initiatives.

Sources: Jasper State of AI in Marketing 2026; Customer.io State of the AI Marketer 2025

FROM TOOLS TO AGENTS

What Agentic Marketing Actually Means

Not chatbots or content generators — systems that perceive data, reason through context, and execute multi-step workflows with minimal human intervention. The shift is from AI that helps you work faster to AI that does the work.

THE EVOLUTION

 Rules-Based Automation 1999-2022

Eloqua 1999 · Marketo 2006 · HubSpot 2006 · Pardot 2006

If-then workflows, drip campaigns, and manual segment building. Every trigger, branch, and campaign path had to be designed by hand.

 AI Copilots 2022-2024

Einstein GPT 2023 · HubSpot AI Assistant 2023 · Marketo GenAI 2023

AI helped teams draft copy, suggest subject lines, and recommend segments. But MOps still had to build, configure, QA, and launch every workflow manually.

 AI Agents 2024-Present

Agentforce 2024 · Marketing Cloud Next 2025 · HubSpot Breeze 2025 · BrazeAI 2025 · [Conversion Agents 2026](#)

Marketers define the goal and guardrails. Agents build the campaign, select the audience, orchestrate the journey, and prepare the work for approval.

THE EFFICIENCY MULTIPLIER

Agents don't just make individual tasks faster. They change the operating model. MOps teams move from building every workflow by hand to reviewing, approving, and improving agent-generated work.

Copilot vs. Agent in the MOps context

The distinction matters because it changes who does the work. A copilot helps you work faster inside the existing paradigm. An agent changes the paradigm.

COPILOT

AI suggests, human builds and executes every step.

Example. An AI copilot suggests a segment based on engagement data. The MOps person manually creates the list, builds the workflow, configures the email, and hits send.

AGENT

Human defines goal and guardrails; the agent builds the workflow and executes.

Example. An AI agent takes a goal — "drive upsell for customers approaching usage limits" — generates a brief, builds the audience from the data cloud, drafts content against brand guidelines, and sets up the journey with branching and send-time optimization. The MOps person reviews and approves — but didn't build any of it.

The honest state of play

Truly autonomous agents in B2B marketing ops are still early. But this is no longer theoretical. Real products already exist, and the operating model is shifting from building workflows to reviewing agent outputs.

Teams that start learning now will have an advantage over the next 12–18 months because orchestration, governance, and approval workflows take time to build.

WHY THIS FRAMING MATTERS


MOps readers are rightly skeptical of vendor hype. The point isn't that agents replace teams tomorrow — it's that the execution layer is already shifting, and the operating model is what compounds.

THE DATA

How MOps Teams Are Actually Using AI

Beyond content generation, the operational use cases dominate: segmentation, scoring, enrichment, and workflow automation. Here's where AI is actually being deployed inside MOps.


GRAPH 1 · TOP AI USE CASES IN MARKETING OPS



Sources: SurveyMonkey 2025; HubSpot 2025 & 2026; Loopex Digital 2026; ZoomInfo 2025.

The ROI Measurement Gap

AI investment is rising, but ROI confidence is falling. As AI moves from side experiments into core workflows, leadership wants proof tied to pipeline, revenue, and growth — not just hours saved.



HOW TEAMS MEASURE AI ROI

METHOD	SHARE
FTE hours saved	57%
Reduced vendor / agency spend	43%
Shortened campaign launch cycles	38%
Time saved in compliance / brand review	34%
Lift in campaign conversion / engagement	29%

THE TELL


Most teams still measure inputs, not outcomes. 57% track FTE hours saved, while only 29% track conversion or engagement lift. The teams proving ROI are the ones connecting AI workflows to pipeline, revenue, and business outcomes.

Source: Jasper State of AI in Marketing 2026 (n=1,400). Enterprise performance via Worth.com. High-maturity orgs achieve 61% ROI provability vs. 41% overall.

Year-over-Year AI Adoption Acceleration

Across multiple adoption signals, AI use is accelerating. Marketing and sales are seeing one of the sharpest increases of any business function.

● GenAI in marketing activities ● Business AI adoption ● Company-provided AI licenses ● Full-scale AI deployment



METRIC	EARLIER	LATER	CHANGE	SOURCE
GenAI in marketing activities	7.0% (2024)	15.1% (2025)	+116% YoY	CMO Survey (Duke/Deloitte/AMA)
Business AI adoption	29% (2021)	76% (2025)	+162% in 4 yrs	IBM Global AI Adoption Index
Company-provided AI licenses	71% (2024)	82% (2025)	+11 pts	Marketing AI Institute
Full-scale AI deployment	35% (2024)	46% (2025)	+11 pts	IBM 2026 (n=8,584)

Sources: Marketing AI Institute, 2025 State of Marketing AI Report; Jasper, State of AI in Marketing 2026.

Four Workflows

Transforming Marketing Ops

01 - Lead Scoring

BEFORE: RULES-BASED

Static point systems require manual tuning. Rules go stale, and sales stops trusting the score.

AFTER: AI/ML PREDICTIVE

AI models learn from real conversion data, update continuously, and improve scoring without manual maintenance.

METRIC	RULES-BASED	LATER	SOURCE
Accuracy	30–70%	80–90%	Twenty One Twelve, Reform.app
Sales acceptance improvement	Baseline	+35%	Forrester 2024
MQL→SQL conversion (manual)	35%	55% (AI)	Optifai, n=939 companies
False MQLs	Baseline	–40%	Data-Mania (citing Clari)

 REAL EXAMPLE

People Data Labs (Conversion customer) tried 6sense, RollWorks, and Marketo's native scoring — three different tools, three attempts. None got sales buy-in because the models were black boxes. They switched to Conversion's natural-language scoring where criteria are written in plain English.

02 – Email Personalization

BEFORE: RULES-BASED

Clone emails across dozens of audience segments. Hostfully was doing this across 100+ segments in Pardot, writing a prompt in ChatGPT, sending code to a developer, then manually dropping HTML into the builder.

AFTER: AI-DYNAMIC

One AI-dynamic template with content blocks that adapt copy, CTAs, and offers per recipient in real time.

REVENUE CONCENTRATION

2% of emails drive 37% of revenue

METRIC	VALUE	SOURCE
Automated flows: percent of volume	2%	Omnisend 2025 (~24B emails)
Automated flows: percent of revenue	37-41%	Omnisend / Klaviyo
Revenue per recipient (flows vs. campaigns)	18x higher	Klaviyo 2026 Benchmarks
Segmented vs. unsegmented revenue lift	760%	DMA
Send-time optimization impact	+5-10% opens/clicks	Backlinko
Dynamic content blocks	+10-40% higher click rates	Digital Marketing Knight

i

REAL EXAMPLE

Hostfully replaced 100+ Pardot segment clones with one AI-powered template in Conversion. HockeyStack built persona-based nurtures on Conversion and saw 41% CTR increase with 60% less campaign build time.

03 – Audience Building

BEFORE: STATIC LIST

Manually maintain static lists based on demographic filters. Lists decay 22–28% annually as contacts change jobs, emails bounce, and phone numbers rotate.

AFTER: AI-DYNAMIC

AI dynamically builds and refreshes segments based on behavioral signals and engagement patterns instead of manually maintaining static lists.

LIST DECAY

Static lists lose 25% of value annually

METRIC	VALUE	SOURCE
Annual email list decay	22-28%	HubSpot + Bloomreach
B2B contacts experiencing change within 12 months	70.8%	Landbase/ZeroBounce
Business email addresses changing annuall	37.3%	Landbase/ZeroBounce
Cost of poor data quality (US annually)	\$3.1 trillion	Landbase
Behavioral segmentation repeat purchase rates	9x higher	Envive 2026
Behavior-based vs. scheduled emails	3x higher engagement	InboxAlly 2026
Segmented vs. unsegmented click rates	+100.95%	Mailchimp ~9M recipients

REAL EXAMPLE

Conversion's magic enrich node uses AI to classify messy CRM job titles into clean persona buckets automatically — solving the foundational data quality problem that makes segmentation fail. 50+ variants of the same role type get normalized without manual effort.

04 - Cross-sell / Upsell

BEFORE: VIBES-BASED

Trigger expansion campaigns based on static CRM fields like renewal date or company size. Vibes-based upsell where AEs reach out when it feels right.

AFTER: PRODUCT-SIGNAL

Connect product usage data (feature adoption, usage limits, expansion signals) into marketing workflows so expansion triggers are based on actual behavior.

EXPANSION REVENUE

40% of late-stage ARR comes from expansion

METRIC	VALUE	SOURCE
Expansion ARR as % of new ARR	40%	Benchmarkit 2024-2025
PQL conversion to paying customer	25-30%	Tunguz/Redpoint + SiriusDecisions
MQL conversion to paying customer	~2%	Same
PQLs likelihood to convert vs. MQLs	8x higher	Accenture via Amplitude
Customers adopting 5+ features churn reduction	60-80% lower	Gainsight/ChurnZero
Top-quartile B2B SaaS NRR	113%	McKinsey Nov 2025 n=55 companies

i

REAL EXAMPLE

Conversion's native warehouse integration (Snowflake/BigQuery) + AI SQL editor lets marketing teams trigger expansion campaigns based on real-time product usage data, replacing vibes-based expansion.

Product signals indicating expansion readiness:

- 1. Approaching usage/feature limits (seats, storage, API calls at 80%+ capacity)
- 2. High feature adoption breadth (using advanced features beyond core)
- 3. Team size/seat growth (inviting new users)
- 4. Integration usage (connecting with other tools = stickiness)
- 5. Increasing session frequency/depth

Omnisend, 2025 Ecommerce Marketing Report; Klaviyo, 2026 Email Marketing Benchmarks; HubSpot, Database Decay; Landbase, Data Freshness Statistics; Benchmarkit; McKinsey, Net Revenue Retention Advantage; Conversion blog.

The New Marketing Ops Org Chart

Agents absorb the execution layer — workflow building, campaign QA, data hygiene, reporting. That reshapes the team. Drawing on MKT1's B2B marketing-org framework (Emily Kramer), the org splits into three sub-functions plus producers.

BEFORE (TRADITIONAL 20+ PERSON ORG):

- Content & Brand, Growth, Product Marketing as three siloed sub-functions
- Dedicated specialist headcount for MOps, lifecycle, demand gen, analytics
- Producers needed to coordinate across silos and reduce handoff friction
- Ops buried inside growth marketing as a support function
- Vertical career ladders (junior SEO → SEO manager → SEO director)

AFTER (GEN MARKETER ORG):

- Same three sub-functions, staffed by **Gen Marketers** — AI-fluent generalists who own campaigns end-to-end
- Agents absorb execution: workflow building, QA, data hygiene, reporting
- The Producer role dissolves — Gen Marketers reduce handoffs by working campaign-by-campaign
- MOps evolves from support function into the **orchestration layer** across all three sub-functions
- Specialist expertise shifts to contractors when needed. Teams of 10 can cover work that used to require 20+.

How AI Reshapes Each Sub-function

SUB-FUNCTION	BEFORE AI	AFTER AI	WHAT CHANGES
Content & Brand (Fuel)	Team of writers, designers, video producers	AI generates first drafts, variations, and asset adaptations	Role shifts from creation to curation and direction
Growth Marketing (Engine)	Manual workflow building, campaign ops, reporting	Agents handle execution; GTM engineers emerge as technical implementers	Role shifts from execution to architecture
Product Marketing (Foundation)	Research-heavy, competitive analysis, positioning	AI accelerates research and synthesis	Most AI-resilient — requires deep judgment
Producers	Coordinate across silos, manage campaigns	Become MORE critical as AI multiplies volume needing coordination	Role merges with AI orchestration

Source: MKT1 Newsletter — "Building B2B Startup Marketing Orgs from 1 to 25+" (Aug 2024); "Introducing the Gen Marketer" (Sep 2025).

Supplementary: Marketing Team Budget Trends

METRIC	VALUE	SOURCE
Marketing budgets as % of revenue	7.7% (flat YoY)	Gartner 2025 CMO Spend, n=402
CMOs planning to reduce labor spending	39%	Gartner 2025
CMOs planning to cut agency budgets	39%	Gartner 2025
GenAI enabling agency budget reduction	22%	Gartner 2025

CAUTIONARY TALE

The Klarna case study

- Marketing team **cut** from 200 → ~100employees
- External agency spend cut 25%, **saving \$4M**
- AI "Copy Assistant" **handles** 80% ofcopywriting
- **Image production** costs cut \$6M

BUT THEN:

Klarna began rehiring human agents after quality declined. CEO: "We focused too much on efficiency and cost. The result was lower quality."

Sources: Gartner 2025 CMO Spend Survey; Klarna press (2025); Jasper 2026; Bloomberry GTM engineering analysis.

Common Pitfalls and How to Avoid Them

A note on the data: dedicated MOps failure-rate studies don't exist yet, so these stats draw from the closest adjacent domains — B2B marketing, B2C engagement, and enterprise AI deployment. The patterns are consistent: governance, data quality, trust, and measurement are universal failure points regardless of vertical.

01 - Governance

Hallucination liability

AI generates customer-facing content — emails, landing pages, chatbot responses — that contains fabricated claims. Air Canada learned this the hard way in 2024 when a chatbot invented a refund policy that didn't exist. The court held the company liable and rejected the argument that the chatbot was a separate entity. For MOps teams running AI-assisted lifecycle campaigns, the implication is direct: if an agent drafts an email that makes a promise your product can't keep, your company owns that promise.

Brand drift

AI outputs are technically correct but tonally off — wrong voice, wrong positioning, wrong level of nuance for the audience. This is harder to catch than hallucinations because nothing looks obviously broken. It just slowly erodes the consistency that took years to build. Without a style guide wired into the agent's workflow, every piece of generated content is a coin flip.

Compliance exposure

The FTC's 2024 sweep against deceptive AI claims and AI-generated fake reviews set a clear precedent for marketing. Any AI-generated testimonials, case study summaries, or review content in lifecycle campaigns now face the same scrutiny as manually written claims. Most MOps teams haven't updated their review processes to account for this.

THE COMMON THREAD:

the problem isn't the AI. It's the absence of a governance layer — who reviews what, what the agent is and isn't allowed to generate, and where a human has to sign off before something goes live. Teams that treat this as a "we'll figure it out later" problem are the ones most likely to figure it out in front of legal counsel.

02 – Data Pipeline Failures

The most common failure point — and it has nothing to do with model quality. Teams deploy agents before deciding what data the agent needs, how it gets piped in, and whether it's clean enough to acton.

STAT	DOMAIN	SOURCE
74% of companies haven't demonstrated tangible AI value — root cause is people/process, not technology	Enterprise CxOs, all industries, 59 countries	BCG Oct 2024, n=1,000
Over 40% of agentic AI projects will be canceled by end of 2027	Enterprise IT, all industries	Gartner Jun 2025, n=3,412
45% of data used by marketing teams is incomplete, inaccurate, or outdated	Marketing data broadly	Data Axle
Only 33% of marketers say they can activate their data effectively	Marketing analytics	Supermetrics 2026

ON IMPLEMENTATION

Setting up effective AI pipelines is genuinely hard. Connecting your CRM, warehouse, and engagement channels into a unified data layer that an agent can actually act on requires significant engineering effort when done in-house. Most MOps teams don't have dedicated data engineers — they're operators, not infrastructure builders. Building and maintaining custom integrations between Snowflake, Salesforce, and your MAP is a full-time job that compounds in complexity every time a schema changes or a new data source gets added.


This is where platform choice matters. An AI-native MAP with built-in warehouse connectors, native CRM sync, and embedded AI workflows collapses this complexity into configuration rather than engineering. The data pipeline, the scoring model, the enrichment logic, and the campaign execution all live in one system — which means the agent has access to the data it needs without middleware, custom ETL jobs, or engineering tickets. The difference between "we spent 3 months wiring up a pipeline" and "we were live in 2 weeks" is almost always a platform architecture question, not a team capability question.

03 – Measurement

Deploying an agent without defining what success looks like upfront means you'll never know whether it worked — and you won't be able to justify scaling it.

STAT	DOMAIN	SOURCE
Only 41% can prove AI ROI (down from 49% YoY)	B2B + B2C marketing professionals	Jasper 2026, n=1,400
57% measure AI ROI by FTE hours saved; only 29% measure conversion lift	B2B + B2C marketing professionals	Jasper 2026

Measurement framework that works (three layers):



The teams that measure at layer 3 are who can prove ROI. Everyone else is stuck at layer 1.

Sources: Jasper 2026; Braze 2026 CER; BCG; Gartner; IAB; Data Axle; Supermetrics 2026.

Your First 90 Days with Agentic Marketing

5-step action plan for MOps teams:

- ✓ Audit your martech stack for agent-readiness
- ✓ Pick one high-impact workflow to pilot (lead scoring or data enrichment)
- ✓ Define governance guardrails before scaling
- ✓ Restructure ops roles around orchestration, not manual execution
- ✓ Measure what matters — outcome-based, not activity-based

Phased Implementation

PHASE	DAYS	FOCUS
Foundation	1–30	Audit data sources, set 3–5 baseline KPIs, pick 2–3 pilot workflows, establish governance rules
Operationalize	31–60	Run pilots, connect AI to CRM/MAP/ad platforms, measure against baselines
Optimize & scale	61–90	Refine, expand across channels, leverage advanced features, plan next 90-day cycle

Lowest-Risk Workflows to Pilot

Tier 1 (start here):

Email subject line testing, content first drafts, reporting automation, data enrichment/cleanup

Tier 2 (phase 2):

Lead scoring, customer segmentation, send-time optimization

Pilot Cautionary:

Most successful AI implementations started with pilot projects. Companies that skip pilots face higher failure rates and higher costs.

Time-to-Value Benchmarks

TIMEFRAME	EXPECTED RESULTS
0–6 weeks	Quick wins; content time drops from 8–10 hrs to <2 hrs per piece
1–3 months	5–12% campaign improvement
3–6 months	20–30% campaign ROI improvement
6–12 months	18–35% gains from compounding optimization

Sources: Humane Labs, AI Pilot Project Guide; Madgicx, 15,000+ Campaign Benchmarks; Conversion.ai implementation framework.

THE HONEST TRUTH

Why The Best Teams Are Already Moving

Here's what most reports won't tell you: knowing what to do is not the hard part. The playbook is straightforward. The hard part is doing it on a stack that wasn't designed for any of this.

Most MOps teams that try to implement agentic workflows in-house hit the same wall. They spend months wiring up **data pipelines**, writing **custom integrations**, debugging **sync failures**, and maintaining brittle connections between CRM, warehouse, and MAP before a single agent runs a single workflow.

This is why the highest-performing teams aren't **building from scratch**. They're choosing platforms where the **AI, data layer, and execution engine** are one system. Not bolted together. Not routed through middleware. **Native**.

The enterprises seeing the fastest time-to-value share three things: a MAP that connects natively to **CRM and warehouse data**, AI embedded in the **workflow engine**, and a deployment model that lets the ops team focus on strategy from day one.

The question is whether your current platform can get you there without a **six-month infrastructure project**.



Conversion

See how teams go from zero to live agentic workflows in under 4 weeks → conversion.ai

Methodology & About Conversion

Lowest-Risk Workflows to Pilot

CAPABILITY	TRADITIONAL AUTOMATION	AI-ASSISTED (COPILOT)	AI-NATIVE (AGENT)
Campaign creation	Built by hand — every trigger, email, segment manual	AI helps write copy faster; structure still manual	Define goal and guardrails; AI builds, selects audience, executes
Audience selection	Manual segmentation on fixed demographics/rules	AI suggests segments; human configures logic	AI selects and continuously refines based on behavior
Lead scoring	Fixed points-based rules	AI provides predictive scores; human reviews	Behavioral + firmographic + engagement; auto-adjusting
Email personalization	Insert first name; basic conditional blocks	AI generates content variations; human selects	Real-time individualization across behavior, history, context
Timing optimization	Scheduled batch sends	Send-time optimization suggested	Automatically optimizes time, channel, frequency per individual
Post-campaign learning	Manual review; standalone campaigns	Some insights surfaced automatically	Feedback loop closes automatically; learnings compound

Source: ActiveCampaign, "AI-Native Marketing Explained," Feb 2026

SOURCE	TYPE	SAMPLE	COMPANY DOMAIN
Jasper State of AI 2026	Vendor-sponsored (3rd-party conducted)	n=1,400	Marketing professionals, mid-market to enterprise
McKinsey State of AI 2025	Tier-1 research	n=1,993, 105 countries	All industries and functions globally
Braze 2026 CER	Vendor research	2,200 execs + 4,000 consumers	References to B2C brands, \$10M+ revenue, 15 markets
Marketing AI Institute 2025	Industry body	n=1,882	Marketing professionals across industries
BCG "Where's the Value in AI?"	Tier-1 consulting	n=1,000 CxOs, 59 countries	Enterprise CxOs across industries globally
Gartner CMO Spend 2025	Tier-1 analyst	n=402 CMOs	Enterprise \$1B+ revenue, NA/Europe
Customer.io State of AI Marketer	Vendor research	n=500+	Lifecycle/automation marketers, SMB to mid-market
DemandScience 2026	Industry vendor	n=750	B2B performance marketing leaders
Bloombergy GTM Engineer Study	Independent analyst	1,000+ postings	B2B SaaS job postings, US-focused
Omnisend 2025	Vendor (1st-party data)	~24B emails	References to Ecommerce/ DTC, predominantly B2C
CoSchedule 2025	Industry body	n=1,005	Marketing professionals across industries
DoubleVerify 2025	Vendor research	n=1,970	References to Digital advertisers and media buyers

Full source URLs available on request. Stats reproduced from the underlyingresearch; where a stat comes from outside B2B MOPs, its domain is labeled incontext.

Conversion

Conversion is an AI-native B2B marketing automation platform built for Salesforce-first enterprises with data warehouses.

It brings CRM sync, warehouse integration, enrichment workflows, campaign execution, and AI agents into one system, helping MOps teams move from manual workflow building to agentic campaign orchestration.

Built for teams that need:

- ✓ Native CRM sync
- ✓ Warehouse integration
- ✓ Embedded AI workflows
- ✓ Human approval before anything ships
- ✓ Faster campaign execution without custom infrastructure

[BOOK A DEMO →](#)