



Tofu Ads

Scaling AI adoption in marketing.

A tactical guide for leaders.

How B2B marketing teams move from scattered pilots to repeatable, measurable results.





What's inside.

Everyone's using AI. Almost nobody can prove it's working. 3

01 Where your team actually is 5
(no, really)

02 Why adoption stalls 7
(it's not the robots)

03 Pick the right use cases 9
(not all 47 of them)

04 Build the foundation 13
(the boring part that makes everything work)

05 Who owns AI, and who does what now 16

06 Why your team is secretly using ChatGPT 19

07 Measure what matters 23
(hint: not just hours saved)

08 Your 90-day plan 25

09 Your buyers are using AI too 27

10 Toolkit 28

Sources 31



START HERE

Everyone's using AI. Almost nobody can prove it's working.

Here's a stat that should make every marketing leader a little sweaty. In one year, AI use among marketers jumped from 63% to 91%. Over the same year, the share who can actually prove ROI *dropped*, from 49% to 41% ([Jasper, 2026](#)).

More usage. Less proof. That's not the trend line you want to show your board.

To be fair, AI has delivered something. Teams are drafting faster, summarizing faster and shipping more stuff than ever. But "we made more stuff" is not a line item your CFO gets excited about. The bar moved from hours saved to pipeline, and most of us are still clearing the old bar.

I've talked to more than 50 marketing leaders this year, and when I ask how they're measured, the answer is always the same: pipeline, pipeline, pipeline. Nobody has "number of AI-generated blog posts" on their scorecard.

The gap, in numbers

- **We bought the Peloton. We haven't started riding it.** CMOs now put 15.3% of their marketing budget into AI, but only 30% say they're ready to scale it. Seventy percent say their processes aren't mature enough ([Gartner 2026 CMO Spend Survey](#)).
- **Everyone's using it for content. Few see results.** About 9 in 10 B2B marketers use AI to create content. Fewer than 4 in 10 say it has improved performance ([Content Marketing Institute](#)).
- **B2C is beating us. Again.** Only 20% of B2B CMOs report significant, measurable revenue impact from AI, versus 31% of B2C CMOs ([BCG, June 2026](#)).

And the teams that figure it out are pulling away fast. The 7.5% of companies BCG calls "future-built" deliver 2.4x the revenue growth of laggards ([BCG, Sept 2026](#)). AI-ready marketing teams also get more budget: 8.9% of company revenue versus a 7.8% average (Gartner). Turns out "we can prove it works" is a pretty good budget strategy.

The (un)sexy truth

Scaling AI is not a tooling problem. It's an operating model problem.

BCG's rule of thumb: about 10% of the work is the algorithms, 20% is technology and data, and 70% is people and process. McKinsey's [2026 State of AI](#) found nearly three-quarters of AI high performers have fundamentally redesigned their workflows, versus about a quarter of everyone else.

Translation: the winners didn't find a magic tool. They changed how their teams work. Less exciting to tweet about. Much more effective.

What's in this guide

This is the guide I wish I could hand every marketing leader who asks me, "Okay, but what do I actually *do*?" It covers:

- Where your team really stands (be honest)
- Which use cases to pick, and which to ignore
- The foundation that keeps AI on brand and out of legal's inbox
- How roles, training and team habits need to change
- How people actually *feel* about using AI at work, and what to do about it
- How to measure impact in a way your CFO will believe
- A 90-day plan to go from pilots to proof

Every section ends with **Do this next**: things you can assign on Monday. And every stat is linked, because if there's one thing marketers love more than a good stat, it's a good stat with no source.



Where your team actually is

(no, really)

Good news first: if your team is still "piloting," you are extremely normal.

The [Marketing AI Institute's 2025 report](#), the biggest marketing-specific dataset out there, found only 14% of marketing teams are actually scaling AI. About 46% are piloting. The rest are still figuring out what the acronyms mean. Salesforce's [2026 State of Marketing](#) tells the same story: 75% of marketing orgs use AI, 39% have fully integrated it, and only 13% use autonomous agents. (So if your LinkedIn feed has you convinced everyone else is running a fleet of agents, they're not. Breathe.)

Now the less-good news: where you are matters for your numbers. MIT CISR's [maturity research](#) found companies stuck in the experimenting and piloting stages grew *below* their industry average. Companies that had changed how work gets done grew 11 points *above* it.

The money shows up when pilots turn into ways of working. Not before.

The four stages

Stage	What it looks like	How you know you're here	What gets you to the next one
1. Experimenting	People try AI on their own, often with personal accounts	No approved tools, no policy, results are anecdotes from Slack	Approve tools, publish basic guardrails, name an owner
2. Individual productivity	Most people use approved tools to draft and research faster	Usage is high, quality is mixed, ROI is "we saved some hours"	Pick 3 to 5 team workflows and redesign them end to end
3. Team workflows	Specific workflows run on shared prompts, shared brand context and real review steps	Shared libraries, review tiers, before/after metrics	Connect workflows with shared data and measured business impact
4. Systematized and agentic	Connected workflows and agents run parts of campaigns while humans steer and approve	Pipeline and cost impact reported to the board	Keep expanding as trust and measurement allow

Most teams I talk to are in Stage 2: lots of activity, very little proof. That's the gap this guide is about closing.

A word on Stage 4, because everyone wants to skip to it. Gartner expects more than 40% of agentic AI projects to be canceled by 2027, and it warns that plenty of products are "agent-washed," meaning old features with a shiny new label. Agents are amazing at running workflows that already work. They are equally amazing at running broken ones, just faster.

The 2-minute self-assessment

Score each statement from 1 (lol, no) to 5 (absolutely). The full version is in the Toolkit.

1. We have approved AI tools and a written policy people actually know exists.
2. Someone specific owns AI adoption in marketing.
3. At least three workflows have been redesigned around AI, with shared prompts and review steps.
4. Our AI tools pull from a maintained source of brand voice, messaging and proof points.
5. We measured a baseline before rollout and can show before/after results.
6. Managers actively encourage AI use and talk openly about how they use it.
7. We can connect at least one AI workflow to pipeline, conversion or cost.

Mostly 1s and 2s: Stage 1. Mostly 3s: Stage 2. Mostly 4s: Stage 3. Mostly 5s: Stage 4, and also please call me, I'd love to have you on the podcast.

Do this next

- **Have your leadership team score it separately, then compare.** If your VP of Demand Gen says 4 and your head of content says 1, congratulations, you just found your first problem.
- **Set one goal for the next 12 months:** get three workflows to Stage 3, with measured results.
- **Re-score every quarter** and share the trend with the whole team.



Why adoption stalls

(it's not the robots)

When AI adoption stalls, it's almost never because the tech didn't work. It's because people don't know how to use it, legal became a bottleneck, the data is a mess, or nobody measured anything before starting. Good news: every one of these has a fix.

Nobody actually taught anyone anything

Lack of training has been marketers' #1 AI barrier for **five years running**, cited by 62% in the [Marketing AI Institute's 2025 survey](#). Only about a third of marketing teams provide AI training. We handed people a Ferrari, skipped driver's ed, and then wondered why everyone keeps it in the garage.

It's also an emotional problem, not just a skills one. Gartner found 80% of CMOs say staff fear and anxiety block experimentation. People who don't feel confident with AI get anxious about it, and anxious people don't practice. Fun loop.

Legal and brand review becomes the bottleneck

AI makes it easy to produce 10x the content. It does not make your legal team 10x bigger. In [Jasper's 2026 survey](#), brand, legal and compliance review was the #1 thing stopping teams from scaling, and governance blockers rose 3.4x in a single year.

The fix isn't less review. It's smarter review: pre-approved claims, templates, and clear rules for what needs whose sign-off (more in Section 4).

Your data is held together with duct tape

Only 26% of marketers are fully satisfied with how well their data is connected (Salesforce). Gartner [predicts](#) organizations will abandon 60% of AI projects that don't have AI-ready data. Personalization at scale with a dirty CRM just means personalizing the wrong thing to the wrong person, very efficiently.

Pilots that never graduate

MIT's NANDA initiative studied hundreds of enterprise AI efforts for its [2025 "GenAI Divide" report](#). Of organizations that looked at enterprise AI tools, 60% evaluated them, 20% piloted, and just **5%** reached production. (It's interview-based and calls itself directional, but if you've ever run a pilot with no end date, this will feel extremely familiar.)

One finding every marketing leader should know: tools bought from specialist vendors made it to production about **twice as often** as tools built internally. Your team probably doesn't need to build its own content engine from scratch. I promise.

Shadow AI

Your team is already using AI. Maybe not the tools you approved, but they're using it.

MIT found only 40% of companies had bought an official LLM subscription, while workers at more than 90% of companies use personal AI tools for work. The [KPMG and University of Melbourne global study](#) of 48,000 people found nearly half had used AI in ways that break company policy, including uploading sensitive data to public tools.

Yes, that's a risk. It's also a giant flashing signal that your team *wants* this. The fastest way to kill shadow AI is to give people approved tools that are at least as good as the ones they found on their own.

Treating AI like an IT project

BCG's [10-20-70 rule](#): 10% of the effort is algorithms, 20% is tech and data, and 70% is people and process. If your AI plan is 90% "which tool should we buy," you've optimized the smallest slice of the pie.

Do this next

- **Approve tools fast.** Get an approved-tools list out within 30 days, even if it's short. Shadow use drops when good options exist.
- **No metric, no pilot.** Every pilot gets a success metric and a kill date before it starts.
- **Build review into the workflow.** Create a library of pre-approved claims and proof points, and define which content needs which level of review.
- **Buy before you build** for most marketing workflows. Save custom builds for what's truly unique to you.
- **Fund training first.** If you only have budget for one thing, spend it on hands-on enablement (see Section 5).



SECTION 03

Pick the right use cases

(not all 47 of them)

The teams getting real results didn't try to "use AI for content." They picked a handful of narrow workflows, tied each one to a metric they already report, and went deep. Bain found B2B growth winners run about **4.5** AI use cases in production. Laggards run 3.3. The difference isn't that winners do everything. It's that they finish things.

Most teams do the opposite. In a [2025 Adobe study](#), 42% of senior B2B execs said scaling stalled because they couldn't identify high-impact use cases. Meanwhile, everyone is using AI to write blog posts. Hmm.

How to pick: value, feasibility, risk

Score every candidate on three questions:

- 1. Value: Does it move a number you already report?** Pipeline, conversion, cost per opportunity, cycle time, agency spend. If you can't name the metric, it's not a use case. It's a hobby.
- 2. Feasibility: Can you actually do it?** Do you have the data, the tool and a human who owns the workflow?
- 3. Risk: What happens if it's wrong?** An internal research summary that's a little off is fine. A customer quote that's invented is a very bad day.

Plot them on a simple value vs. feasibility 2x2 (Gartner's favorite shape) and start in the high-value, high-feasibility corner. [Bain's advice](#) is to start at the front of the funnel and to redesign the process, not just automate the old one. There's a worksheet in the Toolkit.

Where the receipts are

The most convincing results come from workflows with a clear before and after:

Company	Workflow	What they report
Adobe	Email and campaign production	Email creation time down 96%; one onboarding campaign went from 27 days to 1, with click-through up 84%
Lumen Technologies	B2B campaign launch with brand-trained models	Launch time cut from 25 days to 9
IBM	Internal creative assistant	7,000+ drafts in year one; client stories drafted in 5 days instead of 10
Salesforce	AI nurture agent on low-score leads	14,000 "meh" leads turned into 84 opportunities and 12 closed deals in 30 days
Ramp	AI analysis of sales calls	Found LinkedIn came up ~3x more than Meta, the opposite of what attribution said

These are self-reported, so take them as direction, not gospel. But notice the pattern: a specific workflow, a measured starting point, and a business outcome. Not one of them says "we wrote more blog posts."

Where I'd look first

- **Personalization and ABM.** This is the biggest gap between what's possible and what most teams do. McKinsey's [2026 B2B research](#) found market-share leaders are 4x more likely to run one-to-one personalization. Yet most teams only personalize email, and "personalization" there often means `{FirstName}` (CMI). Full disclosure: this is literally what we build at Tofu, so I'm biased. But the data was saying this before we were.
- **Paid creative.** LinkedIn says campaigns with 5+ ad variants get 20%+ higher click-through than single-ad campaigns. AI makes variants nearly free. Your job shifts from making ads to picking winners.
- **AI search visibility.** Half of software buyers now start research in an AI chatbot (Section 9). If ChatGPT doesn't know you exist, a growing share of buyers won't either.
- **RFPs and proposals.** In [Loopio's 2025 report](#), 68% of proposal teams used AI, writing time per RFP fell from 30 to 25 hours, and win rates ticked up.
- **Marketing ops.** Call mining, reporting, data cleanup, lead routing. The most (un)sexy work on this list, and often the fastest payback.

Software companies vs. services firms

Nobody has studied this split directly, but the evidence points to different starting lines.

Software companies usually win with scale: personalized account pages, ad variant testing, AI search content and nurture agents. Your product and proof points are already well documented, which makes AI output easier to keep accurate.

Services firms (consultancies, agencies, IT and professional services) win on expertise and relationships. Hinge's [2026 research](#) found referrals and human outreach still drive nearly two-thirds of new business, and firms' top priority has shifted from creating content to promoting their experts. So use AI to make your experts go further, not to replace them: turn one partner's webinar into a month of content, draft RFP responses from past wins, personalize outreach around real client outcomes. Bain found IT services firms only see big win-rate gains when they redesign the sales process, not just automate it.

Plays to steal

- 1. The webinar multiplier.** One expert webinar becomes 12 to 15 assets: a blog post, a few LinkedIn posts for the speaker, an email, ad variants and an FAQ page. The expert only reviews claims and quotes, not every comma.
- 2. The five-variant rule.** Every paid campaign launches with at least five AI-generated variants across message and visual. Kill losers weekly. Feed winning angles into the next brief. (Purpose-built tools like Tofu's Ads Agent can generate on-brand variants in each platform's formats, but the discipline matters more than the tool.)
- 3. Tier-1 account pages.** Generate a page for each top account from CRM data, call notes and relevant case studies. Sales approves before launch. Measure meetings booked against a holdout group.
- 4. The AI search sprint.** List the 50 questions buyers ask AI about your category. Publish clear, structured answers with original data. Track how often AI tools cite you each month.
- 5. The RFP copilot.** Build a governed library of past answers. AI drafts, the bid lead edits. Track hours per RFP and win rate.
- 6. The leftovers agent.** Run an AI nurture agent on leads below your MQL threshold for 30 days. Count the opportunities. You may be sitting on more pipeline than you think.
- 7. Call mining for budget calls.** Have AI tag sales calls for "how did you hear about us," compare it to your attribution model, and test before moving a dollar.

Do this next

- List 10 to 15 candidate use cases and score them on value, feasibility and risk.
- Pick three to five. For each one, name the owner, the metric and the baseline.
- Make at least one a personalization or paid creative workflow. That's where the gap between leaders and everyone else is widest.



Build the foundation

(the boring part that makes everything work)

AI is only as good as what you feed it. Garbage in, very confident garbage out. Four things decide whether your AI sounds like you, uses accurate data and stays out of trouble: brand context, data, stack and governance.

Teach AI who you are, once

Off-brand AI content isn't a hypothetical. In an [Optimizely study](#) of 2,003 B2B marketing leaders, one in four admitted they regularly publish content they *know* is off-brand because volume pressure wins. Same research: 76% spend three or more hours a week fixing AI output. That's a lot of hours spent editing the thing that was supposed to save you hours.

Gartner's [advice](#) is blunt: marketers have to actively teach AI their brand.

The usual failure mode? Every person and every tool carries a different, slightly stale version of your brand in their prompts. Your SDR's version says you're "the leading platform." Your content lead's version is from two rebrands ago. The fix is one maintained **brand context layer** that every AI tool pulls from:

- Voice and tone, with examples of good and bad copy
- Positioning and messaging by persona and segment
- Your ICP and buyer personas
- Approved claims, stats, proof points and customer quotes
- Your best past content
- Banned phrases, competitor rules and legal no-nos

Give it an owner, update it on a schedule, and feed editors' corrections back in so the same mistake doesn't show up 400 times. This is also the biggest difference between marketing-built AI platforms and a generic chat window: the good ones apply your brand context every single time, instead of trusting each person to paste it in. (At Tofu, this is literally step one of every customer rollout. Nothing good happens until the AI knows who you are.)

Fix the data your use cases need, not everything

Only 26% of marketers are fully satisfied with how their data connects ([Salesforce 2026](#)), and the average team juggles seven data sources. Gartner predicts organizations will abandon 60% of AI projects that lack AI-ready data.

Don't try to boil the CRM ocean. For each use case from Section 3, list the fields it depends on (industry, account tier, persona, stage, recent engagement) and clean those first. Nothing says "we personalized this just for you" like getting the company name wrong.

Prune your stack before you add to it

Gartner says marketing teams use only [49% of their martech capabilities](#). So half the stack is basically a very expensive bookmark folder. And 45% of martech leaders say the AI agents their existing vendors bolted on fall short of what was promised ([Gartner](#)).

My stack rules:

- **Audit before you buy.** Cut what nobody uses before adding anything new.
- **Buy for specialized workflows.** Remember MIT: vendor tools reached production about twice as often as homegrown builds.
- **Pick tools that connect.** AI that reads your CRM, brand context and content library beats AI that lives on an island.
- **Make vendors show receipts.** Ask for customer results against a baseline. (15 questions to ask are in the Toolkit.)

Guardrails that make you faster

I know "governance" sounds like where good ideas go to die. Done right, it's the opposite: it makes teams faster, because people know what they can do without asking permission.

Most teams have a policy by now. Few have guardrails that work. In a [2025 survey](#) of marketing and C-suite leaders, only 33% said their AI content guardrails are strong and consistently applied, and 80% still rely on spot checks.

The secret weapon is **review tiers**, so heavy review only happens where it matters:

Tier	Content	Who reviews
1	Internal drafts, research, summaries	You
2	Public content: blogs, social, emails, ads	A named editor checks facts, claims and brand
3	High-risk content: stats, customer quotes, legal or competitive claims, synthetic people or voices	Legal or compliance signs off

A few rules that apply to basically every B2B marketer (summaries, not legal advice, please don't @ me):

- **No fake reviews or testimonials.** The FTC's [rule](#) bans AI-generated reviews from people who don't exist. The FTC is also going after unsupported AI performance claims, so if your product page says "10x productivity," have the receipts.
- **The EU wants disclosure.** EU AI Act transparency rules [took effect in August 2026](#). Deepfakes and chatbots must be disclosed. AI text on matters of public interest must be disclosed unless a human editorially reviewed it.
- **Copyright needs a human.** The [US Copyright Office](#) says purely AI-generated content can't be copyrighted. Prompts alone don't count.
- **Disclose what's material.** The IAB's [framework](#) says label synthetic people, voices and imagery, but not routine AI-assisted copy.

Notice the theme? A named human who reviews and approves public-facing work solves brand, legal and copyright all at once. Build that approval step into every workflow, agents included. Nothing goes live without a human saying yes.

Do this next

- **Build your brand context layer** and name its owner. Start with voice, messaging, ICP and 20 approved proof points.
- **Clean the data fields** your chosen use cases actually depend on.
- **Run a 30-day tool usage audit** and cancel what nobody uses. Your finance team will send you flowers.
- **Publish your three review tiers** and an approved-tools list. There's a policy starter in the Toolkit.



Who owns AI, and who does what now

If I could only give you one stat from this whole guide, it'd be this one. McKinsey tested 25 factors to see what actually drives bottom-line impact from AI, and **workflow redesign** came out on top. Not the model. Not the vendor. Changing how the work gets done.

And that only happens when a senior leader owns it.

CMOs, this one's on you

Gartner **warns** that CMOs who hand AI off to IT, an agency or "the ops person" keep it stuck doing small tasks. Here's the part I find kind of hilarious: 65% of CMOs expect AI to dramatically change their role, but only 32% think their own skills need to change much. Gartner predicts that by 2027, low AI literacy will be one of the top three reasons big-company CMOs get replaced. So, maybe block some time on the calendar.

Ownership shows up in results, too. KPMG's **Q2 2026 AI Pulse** found organizations with a clearly accountable owner were 3.5x as likely to report real AI ROI (14% vs. 4%). Yet only 29% of marketing teams have an AI council and just 25% have a roadmap.

A simple model that works:

- **The CMO** is accountable and visibly uses AI. (Yes, visibly. More on why in Section 6.)
- **An AI product owner**, often in marketing ops, runs adoption with at least half their time. They own the roadmap, the brand context layer and the metrics.
- **A small council** (marketing, RevOps, IT or security, legal, plus one champion per team) meets monthly to approve tools, review results and decide what scales.

The new roles are builders

Marketing job descriptions are changing fast. A **CXL analysis** found the share of US marketing postings requiring AI skills jumped from 15% to 40% in just two quarters of 2026. Tools like Clay, n8n and Claude Code now show up in job requirements. For marketing roles. Wild.

Three roles worth knowing:

- **GTM engineer.** Builds automated workflows across data, outreach and campaigns. Postings grew about 200% in 2025, though the total number is still small.
- **Content engineer.** Builds the content machine (briefs, refreshes, publishing) with human checkpoints and AI search targets built in.
- **AI product owner.** Runs agent-assisted workflows and watches their performance. BCG found leading CMOs creating this role along with AI governance owners and AI search teams.

You don't need all three tomorrow. But before your next producer hire, ask: would a builder multiply the output of the people I already have?

Five hours of training is the magic number

BCG's [AI at Work research](#) found people who get **at least five hours** of training, plus in-person coaching, are dramatically more likely to become regular AI users. Only a third say they've been properly trained. And [Section's 2026 research](#) found just 5.5% of US knowledge workers are truly AI-proficient. 5.5%! Workers whose managers *require* AI use score 1.5x higher.

What good training actually looks like:

- **Hands-on, on real work.** Run the first five hours in person, built around each team's actual deliverables. Not "10 prompts that will change your life."
- **A skills floor for every role.** Zapier's [AI fluency rubric](#) has four levels (Unacceptable, Capable, Adoptive, Transformative) with marketing-specific examples, and "Capable" is the minimum for every hire. Steal it.
- **Champions.** One per team, with two to four protected hours a week, early access to new tools and ownership of the team's shared prompt library.
- **Little and often.** A weekly 15-minute tool drop or demo keeps the energy going after the launch buzz fades (and it always fades).

Redesign workflows. Don't just sprinkle AI on them.

Adding AI to each step of an old process saves minutes. Redesigning the process saves days. BCG's [2026 research](#) found companies that reshape workflows end to end are 24 points more likely to see measurable business improvement. Clear strategy adds about 25 points of impact. Better tools alone? About 5. Among B2B SaaS marketing leaders in a [Wynter study](#), the 23% reporting positive ROI had built specific, repeatable workflows instead of winging it.

How to redesign a workflow: map every step and handoff, ask which steps only exist because the work used to be slow, then rebuild it as AI plus human checkpoints. Measure cycle time and output before and after.

One warning. Marketing saves more time with AI than any other function: BCG found 60% of regular AI users in marketing save about **a full workday a week**. But two-thirds of workers get no guidance on what to do with that time. Decide where saved hours go before they quietly turn into more meetings.

Make headcount decisions on purpose

AI is already changing team size, mostly quietly. In Wynter's research, 47% of B2B SaaS marketing leaders cut, reduced or didn't backfill roles because of AI in the past year, mostly through attrition. Content, design and junior roles are seen as most exposed. Stanford [research](#) found employment for 22 to 25 year olds in the most AI-exposed jobs fell 13% relative to others.

Here's the long-term problem: today's junior marketers are tomorrow's VPs. If AI eats every entry-level task, who are you promoting in five years? Decide deliberately how people learn the craft.

Do this next

- **Name the owners.** The accountable exec, an AI product owner, and the first monthly council meeting on the calendar.
- **Set a skills floor** for each role and add one AI question to performance reviews.
- **Book five hands-on training hours per person** this quarter. Managers go first.
- **Redesign one workflow end to end** this quarter, with before and after metrics.
- **Add one question to every headcount request:** what did we try with AI, and where will the saved time go?



Why your team is secretly using ChatGPT

Tools and training get people started. Whether they *keep* going comes down to stuff that never shows up on a dashboard: Is it safe to admit I use AI? Will my colleagues trust my work? Does my manager actually want me doing this?

This is my favorite section, because it's where most AI rollouts quietly die.

Half your team is hiding their AI use

In the [KPMG and University of Melbourne study](#), 57% of employees said they hide their AI use or pass off AI work as their own. Slack's [Workforce Index](#) found 48% of desk workers would be uncomfortable telling their manager they used AI. Why? It feels like cheating, and they're worried about looking lazy or less competent.

Here's the twist: they're not wrong. A [2025 study in PNAS](#) ran four experiments with more than 4,400 people. Workers who used AI were rated *lazier and less competent* than workers who got the exact same help from a person. Ouch.

But (and this is the important part) the penalty mostly disappeared when the person doing the judging used AI often themselves. Managers who used AI a lot actually *preferred* job candidates who used it.

So the stigma isn't about AI. It's about whether the people in charge use it. **Leaders using AI out loud is the fastest way to kill the shame.** And it pays: Slack found workers comfortable sharing their AI use are 67% more likely to use it at all.

The anxiety is real (and marketing feels it)

Pew found 52% of US workers are [more worried than hopeful](#) about AI at work. BCG found job worry is *higher* where AI redesign goes deepest (46% vs. 34%), and higher among managers than frontline staff. For writers and designers, it can feel like a threat to who they are, not just what they do.

There's a nuance I love here. A [2026 study of 403 professional writers](#) found writers who saw AI as a collaborator were more productive and happier. But the ones who *also* kept a little competitive edge ("I can still do this better than the machine") were the only ones who kept their skills sharp. Best stance: use the AI, and keep your craft.

Calibrate trust, don't max it out

You don't want your team trusting AI *more*. You want them trusting it the right amount, for the right tasks.

The landmark proof is a [2023 Harvard and BCG experiment](#) with 758 consultants. On tasks well suited to AI, people using GPT-4 finished 12% more tasks, 25% faster, at 40%+ higher quality. Amazing. On a task just outside what AI does well, they were **19 points less likely** to get the right answer than people working without it. The researchers called this the "jagged frontier": AI is brilliant at some things and confidently, cheerfully wrong at others that look almost identical.

And over-trust is already everywhere. KPMG found 66% of employees use AI output without checking it, and 56% have made mistakes at work because of it. Microsoft and Carnegie Mellon researchers found the more people trust AI, the less critical thinking they apply. Which is... exactly backwards.

Workslop is a hidden tax on your team

You've received this. You may have sent it. It's the 6-page "strategy doc" that looks polished and says absolutely nothing.

Researchers at Stanford and BetterUp named it "[workslop](#)": AI-generated work that looks finished but isn't. 40% of US desk workers got some in the past month. Each one took nearly **two hours** to deal with, and about half of recipients thought less of the sender afterward. Managers get hit the hardest.

For marketing, where AI can 10x output overnight, this is a real risk. More drafts aren't progress if every one needs a colleague to fix it.

AI changes how teams work together

(mostly for the better)

The best collaboration research is a [2025 field experiment at Procter & Gamble](#) with 776 professionals. One person working with AI produced ideas as good as two-person teams without it. Teams *with* AI were the most likely to produce top-tier ideas.

Two more findings I think every marketing leader should know:

- **AI broke down silos.** Without AI, R&D people pitched technical ideas and commercial people pitched commercial ones. With AI, both pitched balanced ideas. Imagine product marketing and demand gen actually agreeing on something.
- **People were happier.** AI users reported more positive emotions and less frustration. AI can make work feel better, not just faster.

One quieter shift: in [Anthropic's study of its own staff](#), one engineer said 80 to 90% of the questions they used to ask colleagues now go to AI first. Efficient! But it means junior people ask senior people less, and that's how mentoring quietly disappears. You'll need to rebuild those learning moments on purpose.

Managers are your biggest lever

If you only do one thing from this section, do this one. Gallup surveyed more than 19,000 US workers. Employees whose manager actively supports AI use are:

- **2.1x** as likely to use AI a few times a week or more
- **6.5x** as likely to find their company's AI tools useful
- **8.8x** as likely to say AI helps them do what they do best

And only 28% strongly agree their manager is supportive. That's a lot of upside sitting on the table.

Culture follows the same pattern. After AI rollouts, about a quarter of employees say culture got better and a quarter say it got worse ([Gallup, Aug 2026](#)). Where managers back AI, it's far more likely to improve. BCG found positive feelings about AI jump from 15% to 55% with strong leadership support.

The gap? Skills and example. Half of Fortune 500 CHROs doubt their managers can guide AI use, and Section found only about one in five managers showed their team their own AI use in the past month.

What shifts across teams

- **Sales.** Buyers still rate human reps 21 to 39 points above AI on understanding their needs and building confidence ([Gartner, May 2026](#)). AI should arm your reps, not replace them. This is where marketing and sales finally get to sing from the same songbook: marketing owns the approved messaging and proof points that power sales' AI outreach.
- **Agencies.** Gartner found 22% of CMOs say AI has cut their reliance on agencies, and 39% planned to trim agency budgets. Most B2B teams still want a hybrid model, but the scope is shifting: production moves in-house, agencies focus on strategy, big creative ideas and specialist channels, increasingly on output-based pricing.
- **Product marketing.** AI blurs the lines between functions (remember P&G). Expect more joint work on positioning, launches and customer research.

Do this next

1. **Make your AI use public.** The CMO and every VP share one real AI workflow a week: the prompt, the output and what they changed. Leaders going first is the antidote to the shame spiral.
2. **Normalize disclosure.** Add a neutral "how I used AI" line to briefs and review docs. Celebrate good use publicly. Never treat it as a gotcha.
3. **Run a monthly show-and-tell.** Three people, ten minutes each, one real workflow. Save each one in a shared library with an owner.
4. **Publish a frontier map.** Three lists: what AI does on its own, what AI drafts and a human checks, and what AI doesn't touch (think pricing, customer quotes, final positioning). Update it quarterly.
5. **Ban workstop.** Nothing AI-generated goes to a colleague without the sender's own edit pass, a fact check and a quick note on what AI did.
6. **Train managers first.** Hands-on hours for them before their teams, plus a 15-minute AI check-in for 1:1s: what did you try, what worked, what would you share?
7. **Protect the craft.** Position writers and designers as editors-in-chief and owners of voice and taste. Keep some flagship work deliberately human-led.
8. **Rebuild mentoring.** Pair juniors and seniors to review AI-assisted work together, so judgment still gets passed down.
9. **Re-scope agency contracts** around what you now do in-house versus what still needs outside expertise.
10. **Make it safe to fail.** Fund three to five small experiments a quarter, expect some to flop, and share what you learned. Out loud.



Measure what matters

(hint: not just hours saved)

The one rule: report two numbers, not one. A productivity number and a business number, both measured against a baseline. Teams that only report "hours saved" are losing the budget fight, and they don't even know it yet.

Why "hours saved" won't save you

Among marketers who track AI ROI, the most common metric is hours saved (57%). Only 29% track growth metrics like conversion or engagement lift (Jasper 2026). Meanwhile, Gartner is literally telling CFOs not to [mistake activity for value](#). Pilots launched, tools deployed and hours saved don't count unless they change a cost line or a decision.

You also need some patience. Deloitte's [AI ROI research](#) found most execs expect a typical AI use case to take two to four years to pay off, versus 7 to 12 months for other tech. Only 20% of enterprises see revenue gains from AI so far, while 66% see productivity gains ([Deloitte 2026](#)). Productivity shows up first. Revenue shows up later, and only if you're tracking it.

The metrics ladder

Measure on four rungs. The early ones show momentum. The later ones show money. Report all four, but know your board only really cares about the top one.

Rung	When	What to track
1. Adoption	Weeks 1 to 6	Weekly active users as a share of the team; training hours completed (target 5+); workflows with an approved playbook; approved vs. shadow tool use
2. Efficiency	Weeks 4 to 12	Hours saved per marketer per week; cycle time per asset; campaign launch time; cost per asset; agency spend avoided; review rounds per asset
3. Effectiveness	Months 2 to 6	Output at constant headcount; quality and rework rate vs. a human-only control; conversion lift from AI variants in A/B tests; share of target accounts getting personalized content
4. Business impact	Quarterly, to the board	Pipeline from AI-enabled programs vs. baseline; revenue per marketing FTE; cost per opportunity; total AI cost vs. realized gains (net ROI and payback)

How to measure without fooling yourself

- **Baseline first. Always.** Before you roll anything out, time three to five core workflows for two weeks. No baseline, no bragging.
- **Use a comparison group.** Keep a matched group working without AI, or roll out in waves and compare. IBM's [guidance](#) also says count rework, which can quietly eat your gains.
- **Report users AND everyone.** In a [randomized trial](#) of Microsoft Copilot across 56 companies, regular users cut email time by about three hours a week. Averaged across everyone who got access, it was 1.4 hours. Both are true. Your CFO will ask for the second one.
- **Discount your hours.** Forrester's ROI studies only count half of time saved as productive. Conservative, sure. But nothing builds credibility like underpromising.
- **Sanity-check vendor claims.** The St. Louis Fed [estimates](#) AI users save about 2.2 hours a week on average. If someone promises you 20, ask them to show their math.
- **Track what AI costs.** KPMG found orgs with strong visibility into AI spend were five times as likely to report real ROI (15% vs. 3%).

How to report it upstairs

One page. That's it.

1. **The headline:** the one business result that moved ("campaign launch time went from 25 days to 9").
2. **The ladder:** one line per rung, with trend and target.
3. **The cost:** total AI spend (licenses plus usage) vs. realized savings and gains.
4. **What we killed:** pilots you stopped and why. Nothing builds trust like showing you can say no.
5. **What's next:** the workflows you'll scale, plus one bigger bet.
6. **The caveats:** be honest about attribution. Boards trust leaders who flag uncertainty way more than leaders who don't.

Do this next

- **Pick one metric per rung** for each workflow you chose in Section 3.
- **Run a two-week baseline** before any new rollout.
- **Start an AI cost tracker** for licenses and usage. (Spreadsheet is fine. Really.)
- **Put a quarterly one-page AI readout** on the exec team's calendar.



Your 90-day plan

Ninety days is enough to get three to five workflows from "we're piloting" to "here are the numbers." MIT's NANDA research found top mid-market companies went from pilot to full rollout in about 90 days, while large enterprises took nine months or more. The difference wasn't budget. It was narrow scope, use cases picked by the people actually doing the work, and tools judged on business results.

Days 1 to 30: Get clear, get a baseline, get permission

Goal: decide what you're doing, measure where you're starting, and make it safe to begin.

- **Name the owners.** The accountable exec, the AI product owner and the council (marketing, RevOps, IT or security, legal).
- **Ship the guardrails.** A one-page AI policy, the three review tiers and an approved-tools list. Ask the team what they're already using, and treat the answers as demand, not confessions.
- **Crowdsource use cases.** Get 10 to 15 ideas from the people doing the work. Score them and pick three to five.
- **Baseline everything.** Two weeks of time, cost, output and quality on those workflows.
- **Pick tools.** Buy before you build.
- **Start the brand context layer.** Voice, messaging, ICP and approved proof points.
- **Recruit champions and book training.** One champion per team. Five hands-on hours per person, managers first.

You walk away with: a policy and review tiers, a use-case scorecard, a baseline sheet, a metrics ladder with targets, and a one-page charter per pilot with a success metric and a kill date.

Days 31 to 60: Pilot, redesign, build habits

Goal: prove the workflows work and start building the habits that make it stick.

- **Launch pilots** with one group per workflow. Redesign each one end to end, approvals included. Don't just bolt AI onto the old process.
- **Throw a (productive) party.** Lumen's marketing team ran a "Copilot Olympics": 40-minute team challenges where winning prompts went into a shared library. It grew from one team to all 250 marketers. Zapier calls its company-wide hackathon the single highest-leverage thing it did for adoption. People learn faster when it's fun. Who knew.

- **Start the Section 6 habits:** leaders sharing their AI use weekly, manager check-ins in 1:1s, monthly show-and-tell and the no-workstop rule.
- **Track it.** Adoption weekly, efficiency against baseline, and at least one A/B test for effectiveness.

You walk away with: a pilot dashboard, a before/after readout per workflow, updated playbooks, a shared prompt and workflow library, and a scale-or-kill call on every pilot.

Days 61 to 90: Scale, cut and show the money

Goal: scale what worked, kill what didn't, and show the business result.

- **Scale the winners.** Roll the one or two workflows that hit target out to every relevant team. Stop or rework the rest. BCG found AI leaders chase about half as many opportunities as their peers but scale more than twice as many. Focus wins.
- **Bank the gains.** Move freed-up agency spend and saved hours into named growth programs, so the value shows up somewhere visible. Unclaimed time evaporates.
- **Deliver the one-pager.** Your first exec readout using the metrics ladder, AI costs and honest caveats included.
- **Plan the next wave:** the next set of workflows, plus one bigger bet.

You walk away with: an exec impact report, scaled playbooks, an AI cost tracker and next quarter's use-case backlog.

Ways to blow it

- **Too many pilots.** Five well-run workflows beat twenty half-run ones. Every time.
- **No baseline.** You can't prove you improved from a starting line you never drew.
- **Tools first, people last.** Remember 10-20-70. Most of this is training, habits and process.
- **Skipping managers.** They're the strongest predictor of whether usage sticks (Section 6).
- **Letting saved time vanish.** Decide where it goes before it becomes three new recurring meetings.



Your buyers are using AI too

Plot twist: the biggest AI shift in marketing might not be happening inside your team at all. It's happening on the buyer's side.

- **Buyers start with AI now.** G2's 2026 research found 51% of B2B software buyers start their research in an AI chatbot instead of a search engine. And 69% switched away from the vendor they originally planned to buy because of what the chatbot told them ([via Demand Gen Report](#)). Let that one sink in.
- **The shortlist is set before you show up.** 6sense [found](#) 94% of buying groups rank vendors before they ever talk to one, and the early favorite wins about 80% of the time.
- **Clicks are drying up.** When Google shows an AI Overview, organic click-through on informational searches drops 61% ([Seer Interactive](#)). The silver lining: brands cited *inside* the Overview got 35% more clicks than brands that weren't.
- **The agents are coming (slowly).** Forrester [predicts](#) one in five B2B sellers will have to answer AI buyer agents with their own agents in 2026. Gartner expects AI-automated marketing work to roughly double, from 16% to 36%, by 2028.

So what do you do about it? Three things:

1. **Track your AI visibility.** How often do ChatGPT, Perplexity and Google's AI cite you for your category's big questions? Put it next to organic traffic on your dashboard.
2. **Make your proof machine-readable.** Clear product facts, use cases, pricing logic and customer evidence that AI tools can find and quote correctly. If the robots can't explain what you do, your buyers won't hear about it.
3. **Pilot one narrow agent.** Pick a workflow you've already proven with humans in the loop, give it clear limits and human approval, and measure it like any other pilot. Skip the "agentic transformation" keynote.

Here's where I'll leave you. The teams that win the next phase won't be the ones with the flashiest AI demos. They'll be the ones doing the (un)sexy work in this guide right now: clean context, redesigned workflows, trained people and honest measurement. Build agents on that foundation and they multiply your results. Build them on chaos and they multiply the chaos.

Now go assign something on Monday.



SECTION 10

Toolkit

Steal these. Seriously. Copy them, adapt them, put your logo on them. Each one is built on credible public frameworks, linked under each tool, so you don't have to start from a blank page.

A. AI maturity self-assessment

Score each statement from 1 (lol, no) to 5 (absolutely). Average each area. Anything below 3 gets fixed before you scale. No grading on a curve.

Area	Statement	Score (1 to 5)
Strategy	We have written AI goals tied to business metrics	
Strategy	A named executive owns AI adoption in marketing	
People	Every marketer has had 5+ hours of hands-on AI training	
People	Managers actively encourage AI use and share their own	
People	People openly discuss how they use AI	
Process	At least three workflows are redesigned around AI	
Process	We have clear review tiers for AI-assisted content	
Platform	We have approved tools, and shadow use is low	
Platform	Our AI tools draw on a maintained brand context layer	
Data	The data our use cases need is clean and connected	
Performance	We measured baselines and can show before/after results	
Performance	We can link at least one AI workflow to pipeline or cost	

Based on the [Trust Insights AI-Ready Marketing Strategy Kit](#) and the stage model in Section 1.

B. AI governance policy starter

A practical marketing AI policy covers twelve things. Keep it to a few pages. If it's longer than your employee handbook, nobody will read it.

1. **Purpose, scope and owners.** Which teams, agencies and contractors it covers; the owner; review twice a year.
2. **Approved tools.** Approved, restricted and banned tools, plus a simple way to request new ones.
3. **Data handling.** Which kinds of data can go into which tools. No customer personal data or confidential financials in unapproved tools. Turn off model training on your data.
4. **Brand context.** Where the source of truth lives, who owns it and how often it's updated.
5. **Review tiers.** Self-review for internal drafts; a named editor for public content; legal sign-off for statistics, customer quotes, legal or comparative claims and synthetic people or voices.
6. **Claims and testimonials.** No AI-generated reviews, testimonials or personas presented as real. Every performance claim needs documented support.
7. **Disclosure.** Rules for labeling synthetic media and chatbots, aligned with the EU AI Act and IAB guidance.
8. **Intellectual property.** Meaningful human contribution on work you want to own; no prompting for living artists' styles or competitors' trademarks.
9. **Agencies and contractors.** The same rules, written into contracts, with disclosure of AI use.
10. **Training.** Required onboarding and role-specific training, with records kept.
11. **Incident response.** How to report an inaccurate, off-brand or leaked output, and who handles it.
12. **Measurement.** How adoption, quality and ROI are tracked.

Sources: [NIST Generative AI Profile](#), [IAB AI disclosure framework](#). This is a starting point, not legal advice; review with counsel.

C. Vendor evaluation questions

Ask every AI vendor these fifteen questions. (Yes, including us. We like these questions.)

1. Do you train or fine-tune models on our data or prompts? Is that off by default and written into the contract?
2. Where is our data processed and stored, how long is it kept, and can we delete it?
3. Which underlying models and subprocessors do you use, and will you tell us before they change?
4. Can you share a SOC 2 Type II report, ISO certifications and a completed [CSA AI security questionnaire](#)?
5. How does the product use our brand context (voice, claims, ICP, past content), and who can edit it?
6. How do you prevent and flag made-up facts, statistics or quotes? Can outputs cite their sources?

7. Which of our systems does it connect to (CRM, marketing automation, CMS, asset library), and which way does data flow?
8. What access controls, single sign-on and audit logs are available?
9. Can we build approval steps and review tiers into the product?
10. Do you offer IP indemnity, and what data were your models trained on?
11. Do you support content labeling or provenance standards?
12. How is AI use disclosed in any buyer-facing chatbot or agent?
13. How quickly will you notify us of a security incident or harmful output?
14. What results have customers like us achieved against a baseline, and can we speak to them?
15. How is pricing structured, and how do we avoid surprise usage bills?

For a fuller template, see MMA Global's [Marketing AI Vendor Evaluation Template](#).

D. Use-case prioritization worksheet

List your candidate use cases and score each from 1 (low) to 5 (high). Multiply value by feasibility, divide by risk, and start with your top three to five. Math! In marketing!

Use case	Metric it moves	Value (1 to 5)	Feasibility (1 to 5)	Risk (1 to 5)	Owner	Score
Example: ad variant testing	Cost per opportunity	4	5	2	Demand gen lead	10
Example: tier-1 account pages	Meetings booked	5	3	3	ABM lead	5

How to score:

- **Value:** How directly does it move a metric you already report, and how big is the opportunity?
- **Feasibility:** Do you have the data, the tool and an owner? Is the work repetitive and time-consuming?
- **Risk:** How much damage would a wrong output do? Customer-facing claims score high; internal drafts score low.

Adapted from Gartner's value vs. feasibility approach and the Trust Insights TRIPS scoring model.



REFERENCES

Sources

Key research cited in this guide. Vendor-published surveys are marked.

Marketing and CMO research

- Gartner 2026 CMO Spend Survey (May 2026)
- BCG CMO Survey: Making the agentic marketing transformation a reality (Jun 2026)
- Marketing AI Institute, 2025 State of Marketing AI Report
- Jasper, State of AI in Marketing 2026 (vendor)
- Salesforce, State of Marketing, 10th edition (vendor, 2026)
- Content Marketing Institute, B2B Trends for 2026
- Wynter, How B2B Marketing Actually Uses AI (May 2026)
- Hinge Research Institute, High Growth in Any Economy (Jan 2026)

Enterprise AI adoption and value

- McKinsey, The State of AI (2025 and 2026)
- BCG, AI value gap research (Sep 2026)
- BCG, AI at Work 2025 and 2026
- MIT NANDA, The GenAI Divide: State of AI in Business 2025 (interview-based; directional)
- Deloitte, AI ROI: The paradox of rising investment and elusive returns (2025)
- KPMG Global AI Pulse, Q2 2026

Human behavior and collaboration

- Reif, Larrick and Soll, Evidence of a social evaluation penalty for using AI, PNAS (2025)
- KPMG and University of Melbourne, Trust, Attitudes and Use of AI (2025)
- Gallup, Manager support drives employee AI adoption (2025)
- Dell'Acqua et al., Navigating the Jagged Technological Frontier (2023)
- Dell'Acqua et al., The Cybernetic Teammate, NBER (2025)
- Niederhoffer et al., AI-Generated Workstop Is Destroying Productivity, HBR (2025)
- Slack Workforce Index (2024)

Governance and regulation

- European Commission, Code of Practice on marking and labelling AI-generated content (2026)
- FTC, Final rule banning fake reviews and testimonials (2024)
- US Copyright Office, Copyright and AI, Part 2 (2025)
- NIST AI 600-1, Generative AI Profile

Buyer behavior and AI search

- 6sense, 2025 Buyer Experience Report
- Seer Interactive, AI Overviews impact on CTR (2025)
- Forrester, 2026 B2B predictions

This guide is educational and vendor-neutral. It is not legal advice; review governance and disclosure decisions with counsel.

Published by Tofu, the AI ad agent for B2B marketing teams.