




The AI-Ready Marketing Organization

A Strategy-First Playbook
for B2B CMOs

By Drew Neisser &
CMO Huddles FOMO task force

FOREWORD

Why we started the FOMO Task Force

Marketing leaders are being asked to build the next version of their organizations while keeping the current version running. They face rising expectations, impatient boards, uneven AI capabilities, and teams trying to understand what all of this means for their careers. Buying another tool is often the easiest decision in that environment. Deciding what work should change, and why, is considerably harder.

We formed the CMO Huddles Future of Marketing Organizations task force to work through those harder decisions together. Across three rounds of discussions, a practitioner panel, and expert conversations, we explored how marketing organizations might evolve and what CMOs can responsibly do today. The most useful insights often came from the friction: between speed and judgment, productivity and growth, experimentation and operational discipline.

Christine Royston, former CMO of Wrike, captured the order of operations clearly: “To me, it’s strategy first, then how can I apply AI?” That is the organizing idea of this playbook. A stronger marketing organization begins with a choice about how the business will win, followed by a clear view of the work required to make that choice real.

Our point of view is directed primarily at CMOs leading established teams of roughly 25 to 250 marketers. These leaders cannot assume that a five-person startup’s operating model will transfer intact to a global organization with multiple products, regions, buying groups, and approval requirements. They need a path that creates room for experimentation while protecting the customer relationships and institutional knowledge that already make the business valuable.

DREW’S TAKE | GIVE THE ACCELERATOR A DESTINATION

AI is an accelerator. Strategy decides where you’re going.

We recommend reinvesting demonstrated capacity gains in customer understanding, experimentation, and business impact before defaulting to headcount reduction. Some companies will need to reduce costs, and some roles will change substantially. Those decisions deserve evidence about the work, its economics, and the capabilities the organization still needs.

This is a CMO Huddles point of view informed by practitioners and research, with room for disagreement and revision. We intend to revisit it every six months as the evidence improves. Our ambition is to give you a framework for making better decisions now, and a community with which to challenge those decisions as the future unfolds.

Drew Neisser
CMO Huddles

What Every CMO Should Know

A better organization begins with a strategic choice

An AI-ready marketing organization can deliberately combine people, knowledge, processes, and technology to improve business outcomes, and can show whether those choices are working. Readiness is an organizational capability. It includes the ability to choose where AI belongs, decline applications that fail the economic or quality test, and change course without dismantling the business each time a new model arrives.

The central recommendation is to choose a strategic outcome, redesign the work required to achieve it, and let demonstrated changes in that work inform roles and structure. A CMO pursuing stronger retention will make different investments from a CMO entering a new category. Both may use AI extensively, while retaining different human capabilities and organizing around different priorities.

IN THE HUDDLE | 2026 FIELD NOTE: FROM PRODUCTIVITY TO A STRATEGIC HYPOTHESIS

The task force repeatedly exposed a gap between faster activity and a stronger business. JD Dillon, CMO of Tigo Energy, described better product imagery and broader translation coverage, while observing, “I have not seen the game-changing that everybody else seems to have thus far.” Later, he described a proposed AI-supported installer information package tied to a specific loyalty-program growth objective. He explicitly said that example had not yet been implemented. The progression matters: a useful efficiency gain can lead to a more ambitious strategic hypothesis without becoming proof that the hypothesis has already succeeded.

RESEARCH SNAPSHOT (2026) | ADOPTION VERSUS INTEGRATION

Our CMO AI Maturity Benchmark study, conducted with Benchmarkit, makes the adoption-to-integration gap concrete: 98% of respondents use AI in at least one marketing function, while only 3% report orchestration across GTM with shared data, workflows, governance, and measurement.¹ Widespread use gives CMOs a starting point. Turning it into a dependable operating capability requires deliberate organizational work.

¹ CMO Huddles and Benchmarkit, Marketing AI Maturity Benchmark 2026, full draft. Adoption and GTM integration questions: n=176 valid responses. Percentages rounded; self-reported practices. Study link and methodology in Appendices A and C.

For a CMO with an established department, the practical challenge is continuity. The organization must keep serving buyers and customers while it learns to operate differently. We recommend a sequence of evidence gates: understand the work, redesign workflows, build shared intelligence, deploy AI, redesign roles, and then reconsider organizational boundaries. McKinsey found workflow redesign was the strongest of 25 tested attributes associated with self-reported earnings impact from generative AI.² The finding reinforces the value of changing how work gets done, while leaving each CMO responsible for testing the business case.

Your route through the playbook

Read Parts I and II to align your leadership team on the argument and principles. Use Part III to choose an operating model, plan the migration, and discuss the implications for each function. Part IV turns the argument into seven actions, while the appendices provide decision worksheets and a scorecard. Start with the CMO Huddles AI Maturity Calculator to assess your organization across strategy, workflows, talent, governance, and measurement, then use this playbook to address the gaps.

READING GUIDE

Why change? Chapters 1–3 connect strategy, work and economics.

What principles hold? Chapters 4–6 address judgment, leadership and the division of labor.

How do we evolve? Chapters 7–11 cover migration, structures, departments, people and governance.

Where do we start? Seven actions turn the argument into a leadership agenda.

What remains open? The conclusion identifies the evidence to watch; Appendices A–E retain research, contributors, methodology, decision worksheets and the scorecard.

The organizational models and staffing allocations are illustrative planning tools. They are deliberately presented alongside the conditions that would make them appropriate and the evidence needed to move further. They should support a discussion with your team and executive peers, rather than supply a headcount target detached from your strategy.

What We Believe

Strategy should give AI its purpose. Redesign work before redesigning people, and judge productivity by the business outcomes it enables. A better organization needs dependable shared knowledge and explicit human ownership of judgment.

2 McKinsey, The state of AI: How organizations are rewiring to capture value (2025). Survey of 1,491 participants, fielded July 2024; workflow redesign was the strongest of 25 tested associations with self-reported gen AI earnings impact. Correlation does not establish causation. Source link in Appendix A.

Governance and economics belong in the workflow from the start. Treat AI as an operating capability that must be maintained, evaluated and improved. These seven principles express the CMO Huddles point of view; Chapter 4 develops their practical implications.

Ten key findings

1. Strategic clarity determines where AI creates value. Christine Royston's strategy-first formulation places business intent ahead of tool selection. The best starting question is which customer or commercial outcome needs to improve, followed by which part of the work constrains it.

2. Task productivity and business performance need separate measures. More output can help the business, but it can also increase review queues and overwhelm customers. Julie Kaplan, CMO of Higher Logic, explained that productivity is often emphasized "because that's what we can measure," even when the underlying ambition is to compete more effectively.

3. Workflow redesign reveals the organizational problem. Paige O'Neill, CMO of Culture Amp, described decomposing marketing work into about 5,000 units and 36 scenarios. The exercise exposed duplication and collaboration breakdowns before it informed an automation roadmap; her estimate of automatable work was a forward-looking assessment.

4. Human judgment becomes a design requirement. Faster production increases the number of decisions an organization can face. Andrew Cox, CMO of Forrester, observed that AI "accelerates the number of decisions that need to be made," which makes decision authority and reviewer capacity central to the design.

5. Shared knowledge is an operational dependency. Dave Anderson, CMO of PointFive, explained, "I can only [exploit agents efficiently] because we built an intelligence layer." Reusable customer, product, and commercial knowledge makes multiple workflows more consistent, provided that ownership and access rules remain clear.

6. Agent proliferation creates management work. Experiments generate maintenance, access, quality, and cost obligations. Micheline Nijmeh, CMO of ThoughtSpot, brought that shift into focus: once many people can create agents, someone must connect and govern their work.

7. AI economics extend well beyond a subscription. Paige warned that "the bill is just now starting to come due" as individual use expands into continuous business workflows. Her concern was a hypothesis about scaling costs, not a finding that AI is universally more expensive than people.

8. Capability combinations matter more than fashionable titles. Heidi Bullock, CMO

of Tealium, argued for technical builders working with creative and innovative marketers. Business judgment, customer insight, technical execution, measurement, and change leadership all need explicit owners, even when one person covers several responsibilities.

9. Training needs a credible employment and growth story. People need time, support, and clarity about changing expectations. Heidi framed her responsibility as making employees “more marketable than when you first came in,” connecting adoption to career development.

10. Several organizational futures can be sensible. Some teams will retWain functional departments, others will integrate shared services, and some will organize more extensively around customer outcomes. The right choice depends on business complexity, available skills, infrastructure, and demonstrated performance, with no single chart established as the winner.

Five immediate recommendations

1. Choose one outcome and make it measurable

Select a meaningful priority such as improving qualified-opportunity progression, accelerating adoption, or reducing the cost of a completed customer onboarding. Establish the baseline, the accountable business owner, and the evidence that would justify

additional investment. A target without a measurement plan is still an aspiration.

2. Map one complete workflow

Follow the work across functional boundaries from trigger to accepted result. Include waiting, handoffs, rework, exceptions, and decisions, then remove unnecessary steps before introducing more technology. Amy King, CMO of Relias, noted that a demand-generation campaign could touch 30 people across functions.

3. Fund a bounded pilot and its operating costs

Give the pilot a defined audience, a spend ceiling, a quality standard, an owner, and a way to stop. Budget for integration, review, maintenance, and training alongside software consumption. Measure the entire workflow rather than the fastest step within it.

4. Establish shared knowledge and clear decision rights

Identify the authoritative sources for the pilot and who keeps them current. Specify which decisions AI can execute within rules, which outputs require review, and which exceptions need escalation. An intelligence layer begins with reliable information and ownership, even if its first implementation is modest.

5. Teach the team and redeploy the capacity deliberately

Use real work to develop AI fluency, evaluation skills, and the ability to improve processes. Decide in advance where

any recovered time will go, then track whether it actually moves into customer conversations, experiments, analysis, or reduced external spending. Avoid counting hours as financial savings when the cash expense remains unchanged.

The leadership conversation to have now

The CEO and CFO should understand what the pilot is intended to improve, what it costs to learn, and which claims remain unproven. The CIO or technology leader should understand the data and operating requirements. Sales, product, and customer success should help define the handoffs and outcomes that depend on their teams.

Rani Salehi, principal analyst at Forrester, urged leaders to establish “what must that organization deliver against those growth priorities.” That question gives the CMO a stronger basis for resource allocation than an arbitrary automation percentage. It also gives the executive team a way to distinguish progress from a collection of impressive demonstrations.

What We’re Watching: Three predictions to test

1. Shared operating capabilities will become more visible

We expect more established marketing organizations to formalize responsibility for AI-enabled workflows, knowledge quality, and economic oversight. Some

will do this inside marketing operations or RevOps; others will share the capability with IT. The prediction will be supported if these responsibilities gain durable staffing, funding, and service commitments, rather than remaining volunteer work.

2. Production roles will broaden while judgment bottlenecks become clearer

As drafting and assembly become easier, more marketers may operate across formats and channels. That creates a need for stronger editorial direction, customer understanding, and evaluation, particularly where a small number of senior people must approve a growing volume of work. This prediction will weaken if review effort and coordination costs consume the expected capacity gains.

3. Organizational change will follow different speeds within the same company

Content production may evolve faster than enterprise relationship building or complex pricing decisions. We expect mixed operating models, with some workflows highly automated and others remaining predominantly human. A company that deliberately retains expert-led work can still be AI-ready if the choice is grounded in outcomes and economics.

What We’re Watching: What remains unresolved

The source conversations do not establish a universal percentage of work that

should be automated, a minimum human headcount, or an inevitable timetable for replacing functional teams. They also leave open how much incremental growth can be attributed to AI once product changes, market conditions, and sales execution are taken into account. Vendor capabilities and operating costs will continue to affect those answers.

RESEARCH SNAPSHOT (2026) ACTIVITY VERSUS IMPACT

The benchmark also reveals how far impact travels. Personal productivity and content volume or speed are each cited by 60% of respondents, compared with 26% citing pipeline contribution and 17% citing revenue or customer outcomes.³ These self-reported findings do not establish causation, but they give CMOs a useful challenge: show how faster work contributes to the business outcome that justified the investment.

The sequence to remember


Strategy → **Work** → **AI** → **Org chart**
first **second** **third** **last**


Use the fuller framework in Chapter 2 to turn that sequence into decisions. Move forward when the evidence supports it, revisit earlier decisions when the business changes, and protect the ability to learn throughout the transition.

3 CMO Huddles and Benchmarkit, Marketing AI Maturity Benchmark 2026, full draft. Reported impact question: n=173; multiple selections permitted. These are respondents' reported outcomes, not causal estimates. Study link in Appendix A.

PART I

WHY MUST MARKETING CHANGE?





CHAPTER 1

Why faster work is not enough

A marketing organization can become faster without becoming easier to manage. AI can accelerate research, drafts, design variants, and analysis while leaving the underlying approval chain untouched. When that happens, production grows faster than the organization's ability to select, coordinate, and learn from what it produces.

IN THE HUDDLE 2026 FIELD NOTE: SPEED AND INSTITUTIONAL KNOWLEDGE

The task force's opening conversations contained useful, concrete successes. JD Dillon described AI-renderWed product images and expanded translation coverage. Andrew Cox described a faster creative process, then raised a concern about the knowledge lost when roles went unfilled: "they understand the clients, they understand the needs of the market." Saving production time does not automatically preserve the expertise that makes production useful.'

The pressure is economic and organizational

Marketing budgets averaged 7.7% of company revenue in Gartner's 2025 CMO Spend Survey, and 59% of respondents said they lacked enough budget to execute their strategy.⁴ That pressure makes the economics of organizational change part of the CMO's job. A credible AI plan should explain both the cost of the new workflow and the business value it is expected to create.

AI changes the set of feasible choices inside those constraints. A marketer can test more variants, a product marketer can prepare more tailored sales materials, and a team can extract patterns from conversations it previously lacked time to read. Each expanded possibility still needs a reason to exist, a quality standard, and a route into an actual customer or business decision.

The organizational question becomes whether the department can convert lower effort per task into better outcomes across the system. If every function optimizes its own activity, the handoffs may remain slow and the customer may receive more disconnected messages. If the work is redesigned around a shared objective, AI may remove a constraint that previously limited the business.

IN THE HUDDLE | 2026 FIELD NOTE: WHEN SPEED CONSUMES THINKING TIME

Andrew warned that "AI is compressing the thinking time, the time to sit back and stare out the window." Creative judgment needs room to develop a point of view, test an assumption, or notice something unexpected in a customer conversation. A CMO who spends every recovered hour on additional production may remove precisely the capacity needed to make the work distinctive.

We recommend reviewing the whole flow of value before assuming another productivity gain is the best next investment. Track how long work waits, how often it returns for correction, whether customers use it, and whether a commercial decision changes. Those measures will reveal whether the department needs better tooling, fewer handoffs, clearer priorities, or more expert attention.

What makes this different from an ordinary tool rollout

Generative AI can participate in loosely specified language and knowledge tasks that previously required a person to interpret the brief. That increases its reach across functions, while making its reliability dependent on context and evaluation. Familiar software discipline

4 Gartner, 2025 CMO Spend Survey. Base: 402 marketing leaders, predominantly at companies exceeding \$1 billion in revenue. These averages provide context, not a budget prescription. Source link in Appendix A.

still applies, but the division of work between people and systems deserves a more deliberate design.

BCG's 2024 survey of 1,000 executives across 59 countries found that only 26% of companies had developed capabilities to move beyond proofs of concept and generate tangible AI value. Its analysis placed most implementation challenges in people and processes.⁵ For this playbook, the implication is practical: a deployment plan should include organizational work as explicitly as it includes technology.


Key Takeaways

Faster production can leave coordination unchanged. Measure waiting, rework and customer use, and protect the thinking time that makes the work distinctive.

Questions for Your Team

- 1. Where does work wait after production is complete?**
- 2. Which bottleneck affects the customer or business most?**
- 3. Where should recovered time go?**

⁵ BCG, AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value. Survey of 1,000 executives across 59 countries. Source link in Appendix A.



CHAPTER 2

What should strategy ask of AI?

“AI-first” sounds decisive, but it can allow a means to become the objective. A team can maximize tool usage, agent counts, or automated output while remaining uncertain about the advantage it is trying to create. Strategy-first gives the organization a basis for choosing among the many things AI makes possible.

IN THE HUDDLE | 2026 FIELD NOTE: CHRISTINE ROYSTON ON STRATEGY FIRST

Christine Royston made the distinction in direct terms: “AI first means AI is... is, like, the first part, it’s before you even think about the strategy.” Her alternative was explicit: “To me, it’s strategy first, then how can I apply AI?” The operational value of that sequence is that it forces tradeoffs. A team prioritizing enterprise trust may preserve a human interaction that a team optimizing a simple self-service purchase would automate.

01 / THE STRATEGY-FIRST DECISION SEQUENCE

1	Strategy	Where will the business win?
2	Business outcomes	What measurable result must improve?
3	Work	What must happen to deliver that result?
4	Workflows	How should tasks, decisions and handoffs connect?
5	Organization	Which capabilities and decision rights are required?
7	Org chart	Which reporting lines support the proven design?

Revisit earlier decisions when testing changes the evidence

Work backward from an advantage

Consider an illustrative company whose strategic ambition is to become the easiest supplier in its category to adopt. Its outcome might be a shorter time from purchase to the customer's first verified result. The work includes diagnosing the customer's needs, providing accurate guidance, coordinating implementation, and recognizing when the customer needs a person.

The workflow then makes that work executable: receive the signal, retrieve the correct product information, tailor guidance, review consequential claims, deliver support, and measure the result. AI might help with retrieval and preparation while a customer-facing specialist owns sensitive exceptions. The staffing and reporting implications become clearer only after the new workflow operates reliably.

This example does not require a promise to automate a particular percentage of marketing. It requires a credible

explanation of how the intervention improves adoption and what evidence would disprove the hypothesis. A tool that increases content volume without improving the adoption experience would be a weak fit for that strategy.

Productivity can serve strategy

Amy King cautioned against treating productivity and business improvement as mutually exclusive choices: "I think it's a false choice to pick one way or the other." Her point was to deploy AI in the best way to reach the business objective. Efficiency can be strategically important when it releases cash, removes a bottleneck, or allows the organization to serve customers it otherwise could not support.

The distinction is in the chain of reasoning. If a team saves time, specify whether that time is redeployed, whether service improves, whether an external expense falls, or whether the organization can avoid an otherwise necessary hire. Without that next step, time saved

remains an operational observation rather than a demonstrated business benefit.

**IN THE HUDDLE | 2026 FIELD NOTE:
GIVE AI A DESTINATION**

Heidi Bullock described choosing better customer satisfaction, understanding calls, and bringing data together before deciding where to apply AI: “Then you’re like, now I know where to point my AI.” A focused destination makes it easier to select a pilot, reject distractions, and explain the investment to employees.

The full framework is a decision sequence, not a waterfall project. Early testing may show that the necessary data are missing or that the proposed economics do not work. Return to the earlier choices when that happens, keeping the strategic objective visible as the implementation changes.

Key Takeaways


A strategic advantage should lead to a measurable outcome, a work redesign and a bounded AI application. Revisit earlier choices when testing exposes a weak assumption.

Questions for Your Team

- 1. Which advantage are we trying to create?**
- 2. What workflow constrains it?**
- 3. What result would justify further investment?**

**DECISION POINT | CHOOSE THE OUT-
COME BEFORE THE TOOL**

Agree on the customer or commercial result, its baseline, the accountable owner and the workflow expected to influence it. Record what would make the pilot worth expanding and what evidence would change your mind.



CHAPTER 3

What does an accepted result really cost?

The relevant unit of cost is the accepted outcome of a workflow. A generated draft is an intermediate product. If it requires extensive correction, fails approval, or never reaches a useful audience, a low generation price says little about its economic value.

Paige O'Neill distinguished individual productivity tools from systems "supposed to be running in perpetuity." That shift changes the budget conversation. Continuous workflows can accumulate consumption, integration, monitoring, maintenance, and review costs long after the initial demonstration has impressed the team.

Build a complete cost model

For each workflow, capture setup and integration, allocated licenses, model consumption, data retrieval and storage, monitoring, maintenance, human review, rework, and incident handling. Use the same scope for the existing process so that the comparison is fair. Avoid charging every shared platform dollar to a pilot while ignoring the existing team's shared costs, or doing the reverse.

A useful calculation is: cost per accepted output equals the workflow's total cost for the period divided by the number of outputs that meet the agreed quality standard. For a commercial outcome, measure cost per qualified opportunity or completed onboarding with an explicitly defined denominator. Keep accounting cost, cash spending, and the opportunity value of recovered employee time separate.

Illustrative monthly cost	Existing workflow	AI-supported workflow
Production or orchestration labor	\$12,000	\$3,000
Human review and rework	\$3,000	\$4,500
Software and consumption	\$1,000	\$2,000
Allocated integration and maintenance	\$0	\$2,500
Total economic cost	\$16,000	\$12,000
Accepted outputs	100	120
Cost per accepted output	\$160	\$100

These figures are a worked example, not a benchmark or a participant result. They assume equal acceptance standards and include an allocated setup cost in the AI-supported workflow. If salaried labor stays in place, the modeled reduction is capacity value until the business redeploys that time productively or reduces actual spending; it is not automatically \$4,000 of monthly cash savings.

Optimize the workflow before choosing the model

Use the least expensive method that meets the task's quality and risk requirements. A rule, a database lookup, a reusable template, or a smaller model may be sufficient for a narrow task. Reserve more capable models and deeper review for work that benefits from them, and retest quality when the configuration changes.

Agents add another dimension because one request can trigger multiple actions, retries, and sub-tasks. Set limits on activity, require escalation for exceptions, and measure the cost of a successful run rather

than a single model call. A declining price per call can coexist with a rising total bill if usage grows faster than unit cost falls.

IN THE HUDDLE | 2026 FIELD NOTE: COUNT WHAT THE SYSTEM ACTUALLY COSTS

Dave Anderson described regularly reviewing “our costs of what we spent,” “what we’ve published,” and “how long is it taking.” That is a useful operating habit even for teams using much simpler workflows. The point is to connect spend, accepted work, and cycle time in one conversation.


We use marketing FinOps here to mean shared financial and operational ownership of AI-enabled marketing work. Marketing defines value, Finance validates the economic treatment, and the technical owner explains consumption and maintenance. Review cost per accepted result, quality, utilization, and business contribution together before expanding the workload.

Key Takeaways

Compare complete workflows at the same acceptance standard. Report cost per accepted result, and separate recovered capacity from spending that actually disappears.


Questions for Your Team

- 1. What costs are missing from our comparison?**
- 2. How much review and rework does each accepted result require?**
- 3. Where will saved capacity create value?**



PART II

WHAT PRINCIPLES HOLD?



CHAPTER 4

Seven principles for better decisions

1. Start with strategy

Name the market or customer advantage the organization is trying to create. Translate it into an outcome the marketing team can influence and a set of tradeoffs it must make. Christine's strategy-first framing is useful because it asks leaders to decide what matters before allocating AI effort.

For example, improving retention may justify investment in customer listening and adoption guidance ahead of another

acquisition-content workflow. That is a strategic choice to test, not a universal recommendation that retention should always come first. The company's economics and customer needs determine the priority.

2. Redesign work before redesigning people

Understand what happens between the request and the result, including

informal coordination and hidden expertise. Rani Salehi warned against “treating AI as a headcount strategy before really redesigning those roles and responsibilities.” An existing job title is a poor proxy for the work that will remain valuable.

Paige’s detailed work inventory illustrates one route into that analysis. A smaller team can begin with a single consequential workflow, documenting tasks and handoffs at enough detail to reveal the constraint. The purpose is a better decision, not the largest possible inventory.

3. Optimize for business outcomes before productivity

Keep efficiency measures, but connect them to the intended benefit. Julie Kaplan’s account of wanting to compete faster shows why productivity can be a necessary intermediate measure. The question is whether the increased capacity helps produce a result customers and the business value.

If a pilot reduces preparation time, specify what the team will do with the recovered hours and when its effect can reasonably be assessed. A long sales cycle may delay revenue evidence; that makes early indicators useful, provided they are labeled as early indicators. It does not turn those indicators into realized revenue.

4. Build knowledge before building agents

Give the workflow dependable information, with owners, permissions, and a refresh process. Dave’s shared-intelligence example suggests why a common source of customer and product context can support many outputs. It also raises a management obligation: shared information must be reliable enough for the decisions made from it.

A practical first version can be a tightly controlled collection of approved messages, product facts, customer evidence, and definitions. Test retrieval against questions the team actually needs to answer. Expand after identifying which missing or conflicting facts cause failures.

5. Humans own judgment

Give people explicit responsibility for business priorities, consequential claims, exceptions, and relationships. Avoid assuming that adding an approval step will automatically make a system reliable. The reviewer must have the expertise, source access, authority, and time to challenge the output.

MIT Sloan’s report on a meta-analysis of 106 experiments found that human-AI combinations performed worse on average than the better of humans or AI alone, although creative tasks showed more promising collaboration.⁶ This supports

6 MIT Sloan, Humans and AI: Do they work better together or alone? (2024). Meta-analysis of 370 results from 106 experiments; performance varies by task. Source link in Appendix A.

testing the division of labor for each workflow. It does not justify removing human accountability from consequential business decisions.

6. Govern AI like any strategic capability

Assign owners to the systems the organization depends on. Heidi Bullock proposed agreement on models, API access, team roadmaps, and participation from RevOps, IT, and legal. Those decisions become more important when an experiment begins to affect customers or shared data.

RESEARCH SNAPSHOT (2026) | GOVERNANCE IN PRACTICE

Only 9% of benchmark respondents reported governance embedded in AI workflows; 49% had either no defined governance or basic guidance that varied by team or use case.⁷ CMOs should establish a clear route from a safe experiment to an approved service. Specify what employees can try independently, what needs technical review, and what cannot go live without business approval, with named owners and predictable decisions.

7. Treat AI as an operating model

RESEARCH SNAPSHOT (2026) | OPERATING MODELS AND DATA

Budget for the ongoing work of maintaining knowledge, evaluating performance, supporting users, and retiring obsolete automations. Forrester emphasizes operating-model change, job redesign, and aligned business data as requirements for agentic enterprise applications.⁸⁸ Our benchmark gives that last requirement particular urgency: 53% of respondents named data quality or access among their top barriers to scaling AI.⁹⁹ The intelligence layer needs a business owner and a maintenance budget.

Amy King asked the operational questions plainly: “Where’s everything going to get stored?” and “who’s managing those agents?” A CMO should be able to answer them for every workflow the business relies on. If the answers depend on one enthusiastic employee’s memory, the capability is still fragile.

IN THE HUDDLE | 2026 FIELD NOTE: A WORKFLOW HAS MORE THAN ONE OWNER OF WORK

Andrew Cox described a standard demand campaign with 13 documented steps involving ten people. Removing a few production tasks would leave much of that coordination intact. A stronger redesign starts by asking which handoffs are necessary, who can make the cross-functional decision, and what information lets the next person act confidently.

7 CMO Huddles and Benchmarkit, Marketing AI Maturity Benchmark 2026, full draft. Governance question: n=174; single response. Percentages rounded. Study link in Appendix A.

8 Forrester, The Agentic Business Fabric Is How AI Will Transform Enterprise Applications (2025). Analyst perspective on operating models, jobs, and business data. Source link in Appendix A.

9 CMO Huddles and Benchmarkit, Marketing AI Maturity Benchmark 2026, full draft. Barriers question: n=176; respondents selected up to three barriers. Study link in Appendix A.

Together, the seven principles create a standard for an investment review. A proposal should explain the outcome, the work, the knowledge requirements, the human decisions, the operating owner, and the economics. A missing answer is a design task to resolve before scale.

Key Takeaways


The seven principles connect strategic intent with dependable execution. Knowledge, judgment, ownership and economics need explicit design before a workflow scales.

Questions for Your Team

- 1. Which principle is weakest in our current approach?**
- 2. Who owns that gap?**
- 3. What evidence would show it has been resolved?**

DECISION POINT | USE THE PRINCIPLES AS AN INVESTMENT TEST

Before funding the next workflow, ask for its intended outcome, work design, knowledge sources, human decisions, operating owner and full economics. Give every missing answer an owner before approving scale.



CHAPTER 5

What must the CMO own?

The CMO's work shifts toward designing the conditions under which good decisions and good execution can happen repeatedly. That includes selecting priorities, defining the marketing contribution to growth, allocating resources, developing leaders, and making sure the organization can learn from its results. AI can assist that work, but responsibility for the choices remains with the executive team.

Rani Salehi described wanting transformation, strategic planning, operations, and finances close together as a CMO's source of decision support. The organizational implication is that AI

oversight should sit near the decisions about where marketing invests. A technically impressive workflow with no connection to those decisions is easy to build and hard to justify.

Lead through decision architecture

Identify the recurring decisions that slow marketing down. These might include which customer problem deserves attention, which claims the company can support, which launch needs a full campaign, or when a regional request warrants an exception. Assign decision owners and clear escalation rules so that every disagreement does not become a CMO review.

The CMO should also define what evidence is sufficient for each kind of decision. A reversible creative test can tolerate more uncertainty than a claim about a customer's financial results. This distinction helps the organization move quickly where it can, while preserving care where a mistake would damage trust.

Protect the work that creates differentiation

Julie Kaplan described the enduring value of “Human connection, human relationships, human judgment” informed by conversations and changes in the marketplace. A department that automates production but loses access to customers can become efficient at repeating outdated assumptions. Senior marketing leaders should preserve direct customer exposure and make those insights available to the team.

That means protecting time for interviews, market interpretation, creative development, and alignment with sales and product. It also means making deliberate decisions about launch cadence. When product teams can ship more quickly, marketing still has to decide what customers can absorb and which changes deserve a coherent market story.

IN THE HUDDLE | 2026 FIELD NOTE: MODEL LEARNING IN PUBLIC

Heidi Bullock described demonstrating a new tool in a team meeting: “people saw me, they saw me make mistakes,” followed by “I’m learning this too.” A leader’s willingness to learn openly can make experimentation less threatening. It is most credible when accompanied by training time, practical support, and clear standards for live work.


A useful leadership rhythm has three levels. Weekly operating reviews surface exceptions, quality failures, and bottlenecks. Monthly value reviews connect spend and recovered capacity to business progress, while quarterly organization reviews examine whether role boundaries and reporting lines still fit the work. Each meeting should end with a decision, not just an update on AI activity.

Key Takeaways

The CMO sets priorities, decision rights and the conditions for learning. Protect customer exposure and make operating, value and organizational reviews lead to decisions.

Questions for Your Team

- 1. Which decisions unnecessarily reach the CMO?**
- 2. Where are we losing direct customer insight?**
- 3. What will our next value review decide?**



CHAPTER 6

How should humans and AI divide the work?

The practical unit of design is a task within a workflow, with a clear owner of the result. An entire function rarely belongs in one column. Research includes both gathering information and deciding what evidence means, while content includes both routine assembly and the original judgment that gives a brand a point of view.

The matrix below is a starting allocation to test. “AI” means execution within defined boundaries and monitoring, rather than independent ownership of the function. “Shared” work needs an explicit handoff or review rule so that responsibility does not disappear between a person and a system.

Work	Humans own	Shared work	AI can execute within rules
Research	Questions, source quality, interpretation	Synthesis and hypothesis testing	Retrieval, tagging, first-pass summaries
Content	Point of view, factual claims, final standard	Drafting and editorial revision	Derivatives from approved material
Campaigns	Audience, offer, budget, success criteria	Planning and optimization	Assembly and scheduled actions
Analytics	Metric meaning and causal interpretation	Anomaly investigation	Routine reporting and checks
Planning	Priorities and tradeoffs	Scenario analysis	Input collection and comparisons
Creative	Concept, taste, distinctiveness	Exploration and selection	Resizing and template variations
Customer	Trust, empathy, consequential commitments	Tailored guidance	Approved information retrieval
Operations	Architecture and accountability	Workflow design and testing	Validated routing and updates
Governance	Policy, permissions, exceptions	Quality and risk evaluation	Logging and rule-based alerts
Decisions	Strategy and high-stakes judgment	Reversible recommendations	Low-risk actions within limits

Design the review, then test it

A reviewer needs the original brief, supporting sources, an acceptance standard, and the ability to reject the result. Measure the time and rework involved. A nominal approval step that adds little scrutiny can create confidence without improving quality.

Research with 758 BCG consultants illustrates the importance of task fit. Harvard Business School's summary of the study reports faster completion and higher-quality work on tasks within the tested AI capability boundary, while emphasizing that performance varies across tasks.¹⁰ A successful test on drafting should not be treated as evidence that the same system can safely interpret a pricing exception or resolve a disputed customer claim.

10 Harvard Business School AI Institute, Navigating the Jagged Technological Frontier (2023). Field research involving 758 BCG consultants. Results depend on whether tasks fall within the tested AI capability boundary. Source link in Appendix A.

Keep accountability visible

Julie Kaplan asked what happens when a company can launch a new product each month: “Can your customers absorb that?” The capacity to produce a launch package does not answer the customer question. Someone still has to decide which audience needs the release, what the product changes for them, and how it fits the wider story.


For each workflow, write down who can authorize publication, spending, customer contact, and data changes. Also specify the fallback when the system fails or the reviewer is unavailable. The matrix becomes operational only when the team can apply those decisions during a real exception.

Key Takeaways

Allocate tasks within workflows and test the handoff between people and AI. Reviewers need expertise, evidence, time and authority, with a fallback when the system fails.


Questions for Your Team

- 1. Which decisions must a person own?**
- 2. Can reviewers independently check the evidence?**
- 3. Who acts when the system or reviewer is unavailable?**



PART III

HOW SHOULD THE ORGANIZATION EVOLVE?



CHAPTER 7

When has the organization earned the next step?

Transformation should advance when the organization has earned the next step. A calendar provides review points, while the evidence determines whether to expand, revise, or pause. Teams can run stages in parallel on different workflows, provided they do not skip the dependencies that make a particular deployment reliable.

IN THE HUDDLE | 2026 FIELD NOTE: AMBITION WITH A MANAGEABLE SCOPE

Amy King's recommendation was to "define the thing that move the needle, but are doable," after considering infrastructure and economics. That combination of ambition and practicality is the foundation of this sequence. Begin with a workflow whose improvement matters and whose boundaries the team can understand.

1 Understand the work

Recognized baseline
Named workflow owner

2 Redesign workflows

Fewer avoidable handoffs
Acceptance criteria & fallback

3 Build shared intelligence

Traceable, reliable sources
Working update process

4 Deploy AI

Repeatable quality
Acceptable full-workflow cost

5 Redesign roles

Credible workloads
Training + backup coverage

6 Redesign the organization

Sustained results
Clear cross-team authority

Evidence determines when to advance. Calendars set review points.

Stage 1 | Understand the work

The objective is to establish how value is produced today. Follow one workflow from trigger to accepted result, observing actual work alongside documented process. Capture human effort, elapsed time, waiting, systems, handoffs, exceptions, and the decisions that change the outcome.

Ask which steps the customer or business actually needs and where the process loses time or quality. Common mistakes include treating job descriptions as complete inventories, overlooking informal coordination, and valuing only visible production. Success means a baseline the people doing the work recognize, with an accountable workflow owner and a measurable constraint.

Stage 2 | Redesign workflows

The objective is a simpler route to the same or a better result. Remove unnecessary steps, clarify decision rights, and improve access to information before deciding which tasks need AI. Julie Kaplan identified “orchestration issues” as a constraint in campaign execution, underscoring that the problem may sit between activities.

Ask which handoffs can disappear and where a person needs to intervene. Common mistakes include automating a broken process, preserving every approval, and making a pilot dependent on several unready systems. Success means a documented future workflow with explicit acceptance criteria, fewer avoidable handoffs, and a tested fallback.

Stage 3 | Build shared intelligence

The objective is to give the workflow trustworthy context. Identify approved product information, customer evidence, definitions, brand guidance, and account

data, then assign owners and access rules. Begin with the minimum reliable knowledge set for the pilot rather than waiting for a perfect enterprise repository.

Ask which source wins when information conflicts, how stale content is detected, and whether the system can show its supporting evidence. Common mistakes include copying everything into a repository without ownership and making sensitive information broadly available for convenience. Success means tested answers to representative questions, traceable sources, and a working update process.

Stage 4 | Deploy AI

Test reliable execution on representative work, including difficult cases. Compare the result with the current process using the same acceptance standard, bounded permissions, and manageable volume.

Ask whether the complete workflow improves, what each accepted result costs, and where failures occur. Avoid easy-only tests and omitted rejects; success requires repeatable quality, manageable exceptions, acceptable economics, and an operating owner.

Stage 5 | Redesign roles

Align responsibilities with demonstrated changes in the work, then identify training or hiring needs. Review capacity after the workflow stabilizes, including new maintenance and review work.

Ask which expertise remains scarce and where recovered time goes. Avoid equating saved hours with removable headcount; success means clear responsibilities, credible workloads, backup coverage, development plans, and sustained service quality.

Stage 6 | Redesign the organization

Change reporting lines, boundaries, or resource allocation where the new work requires it. Compare the three structures against the constraints observed in the pilot.

Ask whether shared services reduce duplication, pods can make tradeoffs, and expertise retains a home. Avoid symbolic reorganization and ambiguous authority; success means sustained results and coordination without unacceptable cost, risk, or overload.

Review checkpoints over the next year

Checkpoint	Evidence to review	Decision it supports
Around 3 months	Work baseline, pilot design, knowledge and ownership gaps	Whether the pilot is ready for controlled deployment
Around 6 months	Quality, cost, exceptions, adoption, early outcome signals	Whether to extend the workflow or revise it
Around 9 months	Recovered capacity, sustained workload, skill gaps	Whether roles or shared services should change
Around 12 months	Business contribution, reliability, cross-team performance	Whether broader organizational redesign is justified

These are illustrative management checkpoints, not forecasts of when a company will reach each stage. A regulated enterprise may spend longer on knowledge access and validation, while a well-prepared team may move through a narrow workflow more quickly. The decision criteria matter more than the month number.

Key Takeaways


The six stages are evidence gates, with dependencies that matter more than a fixed timetable. Role and structure changes should follow demonstrated changes in the work.

Questions for Your Team

- 1. Which stage has this workflow actually reached?**
- 2. What dependency is still unproven?**
- 3. What would justify advancing, holding or stopping?**

DECISION POINT | ADVANCE ON EVIDENCE

At the next review, choose whether to expand, revise, hold or stop the pilot. Document quality, cost, exceptions and the outcome signal, then name the evidence needed for the next stage.



CHAPTER 8

Which organizational model fits your business?


A CMO can use the three scenarios below to discuss how much organizational change the business actually needs. These are organizational choices, not maturity levels. Each can be AI-ready, with different advantages and risks; no single design is established as the best path for every company.

The illustrative staffing examples hold a 60-person marketing organization constant to make allocation choices visible. They include the CMO and assume access to shared IT, security, legal, and Finance support outside marketing; those costs still

belong in the economic model. These are planning examples, not labor benchmarks or a recommendation to hire a particular number of people.

Scenario 1 | The Optimizer

The Optimizer retains functional departments and introduces AI within well-defined workflows. It suits an organization with valuable specialization, limited integration capacity, or a need to protect delivery while improving performance. It can be a durable model when functional boundaries are already serving the customer well.



Shared company support: IT, security, legal and finance

An illustrative 60-person allocation is one CMO, 12 people in brand and content, 10 in product marketing, 17 in demand and field, 10 in customer and lifecycle, and 10 in operations and analytics. AI enablement responsibilities sit within the last group and are shared with functional champions. The allocation preserves specialist leadership while requiring explicit workflow owners across the functions.

The advantage is lower transition complexity and a clearer home for professional expertise. The limitation is that local gains can leave cross-functional friction unchanged, particularly when a campaign still waits in several queues. Evaluate the model by end-to-end cycle time and business performance, not by adoption rates within each department.

A plausible first review comes after one quarter of bounded pilots. If those pilots reveal the same knowledge, production, or engineering needs across several functions, the CMO has a reason to consider more shared capability. If they do not, adding a new organizational layer may create unnecessary overhead.

IN THE HUDDLE | 2026 FIELD NOTE: MICRO-CHANGES CAN BE A LEGITIMATE STRATEGY


Julie Kaplan described making “micro-changes as needed” while keeping a relatively traditional team structure. Her example is useful because it resists the assumption that serious AI adoption must begin with a dramatic reorganization. The question is whether the structure continues to support the company’s priorities as the work changes.

Scenario 2 | The Integrator

The Integrator brings duplicated capabilities into shared services while maintaining business-facing teams. A shared content operation, knowledge capability, and workflow engineering group serve multiple demand, product, and customer priorities. This model suits an established organization where repeated work and disconnected systems have become material constraints.

potentially with a hub-and-spoke arrangement. That suggests a practical design question: which work benefits from a common system, and which work needs to remain close to the customer or product expert?

The benefits are reusable knowledge, more consistent standards, and less duplicate building. The risks include central queues, weaker local responsiveness, and confusion about who has the final say.



Shared company support: IT, security, legal and finance

An illustrative 60-person allocation is one CMO, 12 people in brand and portfolio strategy, 15 in acquisition and field, 12 in customer and lifecycle, 10 in content and creative operations, and 10 in intelligence, engineering, and measurement. Shared services need an intake process, service commitments, and a way to resolve competing priorities. Otherwise, consolidation merely relocates the bottleneck.

Rani Salehi described leaders questioning whether content resources scattered across teams should move toward a content strategy and operations organization,

Protect against those risks by giving each workflow a business owner while shared services own the quality and reliability of their contribution.

Allow several quarters to test service delivery and capacity. At six and nine months, inspect rework and responsiveness; revisit boundaries that add approvals without improving quality.

The minimum capability coverage

Every scenario needs ownership of strategy, customer insight, brand and product judgment, execution, operations,

measurement, and people development. Combine responsibilities in smaller teams and divide them as complexity warrants.

For teams of roughly 25–50, start with an operations lead, a hands-on builder who may be shared, and named knowledge and measurement responsibilities. For teams of roughly 51–100, dedicated workflow capacity and shared content operations may become worthwhile. For teams of roughly 101–250, regional and portfolio complexity may justify distributed practitioners with central standards; headcount alone does not determine the choice.


Scenario 3 | The Outcomer

This scenario organizes more work around customer outcomes, supported by shared intelligence, engineering, creative, and measurement capabilities. It suits a company with reliable data, mature technical support, leaders comfortable with cross-functional decisions, and evidence that its current boundaries impede results. It is the most demanding migration option in this playbook.

An illustrative 60-person allocation is one CMO, 10 people in market and portfolio strategy, 13 in an acquisition pod, 11 in an adoption and expansion pod, 10 in content and creative, and 15 in intelligence, engineering, and measurement. Pod members combine skills around a defined outcome, while shared specialists retain responsibility for standards and reusable infrastructure. People counts should follow the actual workload and complexity of the business, not this illustration.

Christine Royston described organizing pods around outcomes and envisioned humans working alongside agents across those pods. That is a plausible direction for established teams because it places accountability near a commercial objective. It requires clarity on who can reprioritize work when a pod's goal conflicts with a shared brand, product, or company requirement.

The advantages are shorter handoffs, direct learning, and flexible deployment of capabilities. Risks include inconsistent execution, key-person dependency, duplicated systems, and weaker



Shared company support: IT, security, legal and finance

professional development; maintain communities of practice, standards, and backup coverage.

**IN THE HUDDLE | 2026 FIELD NOTE:
WHAT DAVE'S FIVE-PERSON TEAM ESTABLISHES**

Dave Anderson described a five-person marketing team operating with a large bot layer and shared company intelligence. He explained, "You still need product marketers. It just means that these roles can blur a little bit." That example broadens the range of possible operating models, while its size and technical context limit its transferability to a 60-person or 200-person team.

Test this model in a bounded product, segment, or journey. A 9–12 month checkpoint can assess broader potential, while enterprise migration may take longer and require further validation.

Choosing among the three


Choose the Optimizer when specialization and continuity dominate, the Integrator when duplicated capability and handoffs constrain performance, and the outcome-based model when cross-functional accountability is the central problem and the foundations are ready. A company may use all three in different parts of marketing. The choice should be revisited as the evidence changes.

Key Takeaways

Optimizer, Integrator and Outcomer are choices, not a maturity ranking. Match the model to the constraint, and treat the 60-person allocations as illustrations rather than staffing targets.

Questions for Your Team

- 1. Is our main constraint specialization, duplication or cross-team authority?**
- 2. Are the foundations ready for the model we want?**
- 3. Where would a different model make more sense?**



CHAPTER 9

How might each department evolve?

The near-term possibilities below concern roughly three to six months; longer-term options concern nine to twelve months and beyond, after evidence gates are met. These are scenarios to test against your business conditions.

Demand generation

Demand teams coordinate audiences, offers, media, routing, and follow-up. Near term, simplify one campaign and assemble approved variations; longer term, consider an acquisition pod with shared content and engineering support.

Humans retain responsibility for audience strategy, offer quality, budget, qualification standards, and interpretation of results. AI can assist with account research, asset assembly, routing, and bounded optimization. Amy King's description of campaigns touching 30 people is a reminder to measure the whole system, including sales acceptance, rather than just asset production.

Brand

Today, brand teams often own visual consistency and corporate narrative while

answering requests from many functions. Near term, they can turn approved principles, examples, and claims into reusable guidance. Longer term, brand leadership may work more closely with product and customer insight to maintain a coherent narrative across increasingly personalized experiences.

Humans own distinctiveness, taste, reputation, and decisions about what the company should say. AI can produce variations and check work against defined standards, but those standards need continuing editorial judgment.

IN THE HUDDLE | CURRENT EXAMPLE (2026): BRAND JUDGMENT

Dave Anderson captured the distinction: “I wouldn’t get Claude to do your brand; it would be a very, very bad brand.” His observation reinforces the need for a human point of view when using a tool to produce brand work.

Product marketing

Today, product marketing links customer needs, positioning, launches, and sales enablement. Near term, AI can reduce information gathering and first-draft preparation, freeing time for competitive interpretation and customer conversations. Longer term, product marketers may directly shape more digital experiences and guide self-service enablement built on approved information.

Humans own positioning, pricing judgment, market choices, and the translation of features into customer value. AI can retrieve product facts, summarize conversations, and prepare materials for review. Rani’s proposal for a more unified brand narrative across product and solutions supports tighter collaboration, without implying that specialist product knowledge becomes unnecessary.

Field marketing

Today, field teams combine logistics with relationships, local market knowledge, and sales alignment. Near term, AI can help with preparation, invitations, segmentation, follow-up drafts, and analysis. Longer term, field marketers may operate within segment or account pods while drawing on centralized production support.

Humans retain host judgment, partner relationships, situational awareness, and the quality of the live experience. AI can reduce administrative effort, but physical venues, travel, and hospitality still have costs. Amy noted that “events will still cost a lot of money, regardless of what agents do for us.”

RevOps and marketing operations

Today, operations teams maintain systems, data flows, campaign execution, and reporting. Near term, they can inventory AI workflows, establish

approved access, instrument costs, and help business owners test reliability. Longer term, they may own a shared service for workflow engineering, monitoring, and technical coordination with IT.

Humans own architecture, permission design, prioritization, incident response, and the meaning of operational controls. AI can handle bounded data checks, documentation, routine updates, and monitoring support. Heidi's description of the go-to-market engineer as a "modern ops person" emphasizes the continuity of operational responsibility beneath a changing toolset.

Content

Today, content teams may be spread across brand, demand, product, and regional groups. Near term, consolidate the source material and remove repetitive drafting and formatting steps. Longer term, a shared content operation may manage reusable narratives, evidence, and modular assets for several customer journeys.

Humans own editorial priorities, original reporting, source verification, and the point of view. AI can adapt approved material across formats and audiences, with review proportional to the change. In the panel, one CMO shared, "It's really the workflows that matter to me the most these days," moving attention from isolated agent creation to connected work.

Lifecycle marketing

Today, lifecycle teams often operate campaigns around stages, triggers, and segments that were defined at different times by different functions. Near term, validate a small set of meaningful customer signals and improve the response to one of them. Longer term, lifecycle work may sit inside a pod accountable for adoption, expansion, or renewal.

Humans own the customer journey, contact policy, offer, and escalation standards. AI can help recognize signals, prepare appropriate messages, and execute approved actions. Amy's discussion of product bundles and customer outcomes suggests that lifecycle planning should follow the customer's progress rather than the company's internal product catalog.

Research and market intelligence

Today, valuable customer and market information is often dispersed across interviews, call recordings, documents, and individual memory. Near term, AI can make that evidence easier to retrieve and compare. Longer term, the function can maintain a continuous learning system that helps teams test market assumptions and update priorities.

Humans own research questions, sampling, source credibility, and interpretation. AI can assist with

transcription, thematic coding, and finding patterns, while analysts test for missing or contradictory evidence. Faster synthesis should create more time to investigate uncertainty, rather than make a plausible summary feel more certain than the underlying data warrant.

Creative

Today, creative teams balance original concepts with a high volume of adaptations and production requests. In the near term, AI can reduce repetitive production and expand exploration within a clear brand direction. Longer term, smaller production groups may support broader creative leadership, with capacity redirected toward stronger ideas and better testing.

Humans own concept selection, cultural judgment, emotional relevance, and final quality. AI can help with variations, mockups, and format conversion. Andrew's concern about compressed thinking time remains central: preserve room for creative development instead of assuming that faster rendering means original thinking should occur at the same speed.

Analytics

Today, analytics teams spend substantial effort reconciling data and producing recurring reports. Near term, automate reliable preparation and routine checks while making definitions explicit. Longer term, analysts may work more closely with

workflow owners to design experiments, investigate exceptions, and connect operating changes to business performance.

Humans own metric definitions, causal reasoning, uncertainty, and decisions about what to test next. AI can prepare summaries and identify anomalies, but an explanation generated from observational data is still a hypothesis. Amy anticipated more need for "assessment and analytics and performance" skills as teams evaluate the systems doing the work.

Customer marketing

Today, customer marketing connects advocacy, community, adoption, and expansion, often with incomplete access to product and customer-success information. Near term, improve evidence collection, identify relevant customer stories, and support one adoption or advocacy workflow. Longer term, customer marketing may become a central partner in an outcome pod spanning marketing, product, sales, and customer success.

Humans own relationships, consent, advocacy judgment, and the interpretation of what success means to a customer. AI can surface potential stories, prepare interview prompts, and help package approved evidence. Julie's proposed workflow for case studies and testimonials was also a learning opportunity: collecting those signals could support other uses across the business.

The boundary that matters across departments

A function should keep the work that benefits from its expertise and share work that benefits from a common system. That requires an explicit service relationship, not simply a decision to centralize. Define what the requesting team supplies, what the shared service delivers, who approves the result, and how competing requests are prioritized.


Review the boundaries when a bottleneck persists, a capability becomes reusable, or a strategic priority changes. Do not infer that a department should disappear because several of its production tasks become automatable. Its enduring contribution may be judgment, relationships, or a body of knowledge the rest of the organization depends on.

Key Takeaways

Departments retain judgment and expertise while sharing reusable production and infrastructure. Change boundaries when evidence shows a persistent constraint, with clear service commitments.

Questions for Your Team

- 1. Which work must stay close to a specialist or customer?**
- 2. Which capability should be shared?**
- 3. Who resolves competing priorities across that boundary?**



CHAPTER 10

How will your people grow with the work?

AI readiness is built through the way people work, learn, and make decisions together. A team needs both practical fluency with tools and the confidence to challenge a result that looks polished. It also needs a credible account of how responsibilities will change and how employees can develop the capabilities that remain valuable.

IN THE HUDDLE | 2026 FIELD NOTE: FROM FEAR TO EXPERIMENTATION

Paige O'Neill described a journey from "fear to suspicion to now excitement," supported by training, collaboration, and experimentation. That sequence cannot be reduced to a tool-access announcement. It requires leadership time and a working environment in which people can admit what they do not understand.

Hire for capability combinations

Look for evidence of business judgment, learning ability, hands-on problem solving, and the ability to evaluate work. For a technical role, ask a candidate to explain how a workflow could fail, how it would be maintained, and how its value would be measured. For a creative or product role, ask what the candidate would reject in an AI-generated output and why.

Heidi Bullock said, "I want people who can be hands-on problem solvers," while recognizing that the tools themselves will change. A useful hiring exercise therefore tests a real work problem and the reasoning behind a choice, rather than familiarity with one interface. Give candidates equivalent access and

conditions so the exercise measures the intended capability.

Make training part of work

Use three levels of development. Everyone needs practical fluency, source checking, and an understanding of approved data use. Workflow owners need process design, acceptance criteria, and measurement skills, while builders need deeper technical capability, testing, monitoring, and maintenance practices.

A useful learning cycle begins with a real problem, includes a supported attempt, and ends with a review of both the output and the method. Protect time for that cycle and make successful patterns available to others. Otherwise, the organization will reward hidden individual shortcuts without developing shared capability.

Harvard Business Review's "Reskilling in the Age of AI" addresses the organizational challenge of developing new skills.¹¹ HBS's account of the research notes interviews with leaders at about 40 organizations investing in reskilling.¹² That broader work is consistent with treating capability development as a management responsibility rather than leaving it to employee initiative alone.

Recognize uncertainty without lowering standards

Julie Kaplan described employees struggling because "Their role used to be prescribed, and now it's not prescribed." Resistance can reflect uncertainty, insufficient support, or reasonable concerns about quality and workload. Diagnose the cause before treating every hesitation as unwillingness to learn.

Managers should explain the expected capabilities, provide access and time, and review progress using actual work. Psychological safety supports candid discussion of errors; it does not remove responsibility for meeting standards. Employees should be able to report a failure early without being encouraged to hide it behind a confident AI-generated answer.

Develop careers beyond production volume

Create paths for deep expertise and for orchestration across functions. A writer may develop into an editor who shapes the narrative and evidence system, while an operations practitioner may become a workflow designer. Those changes require real development opportunities and authority, rather than new titles attached to old workloads.

¹¹ Tamayo et al., Reskilling in the Age of AI, Harvard Business Review (2023). Source link in Appendix A.

¹² Harvard Business School, Reskilling in the Age of AI Wins 2023 HBR Prize (2024). Institutional account describes interviews with leaders at about 40 organizations. Source link in Appendix A.

|Preserve entry-level learning deliberately. If routine drafting and analysis disappear from junior roles, create supervised opportunities to learn how good work is made and evaluated. Reviewing an answer without understanding the underlying task is a weak foundation for future judgment.

Change performance reviews with the work

Evaluate contribution to outcomes, quality, collaboration, learning, and improvements that others can reuse. Do not make prompt counts, generated asset volume, or number of agents the principal measure of performance. Such measures can reward unnecessary activity and obscure the value of stopping low-quality work.

Ask employees what they improved, what evidence supports the improvement, which failures changed their approach, and what support they still need. Include workload and review burden in the discussion. A team that produces more while becoming exhausted may be consuming its future capacity to innovate.

Be honest about headcount

DREW'S TAKE | REINVEST BEFORE YOU REDUCE

Our recommendation is to test reinvestment before defaulting to reduction. That does not imply every company can preserve every role or that all automation gains should be retained within marketing. It does mean leaders should distinguish a financial decision to reduce cost from an operational finding that a new workflow can reliably absorb the work.


Avoid promises the organization cannot keep. Explain what is known, what decisions remain open, and how skills, performance, business needs, and demonstrated capacity will inform those decisions. Honest uncertainty is more credible than declaring that nobody will be affected while quietly removing the work that gave their role meaning.

Key Takeaways

Develop fluency, workflow ownership and technical depth through real work. Keep career expectations credible, preserve junior learning and test reinvestment before assuming capacity gains justify cuts.

Questions for Your Team

- 1. What will each role need to learn?**
- 2. Where will employees practice with support?**
- 3. How will we know recovered capacity improved the business?**



CHAPTER 11

What makes an AI workflow ready to scale?

When people rely on an AI workflow, ownership, access, action boundaries, and maintenance become operating requirements. Governance makes those requirements part of daily management.

IN THE HUDDLE | 2026 FIELD NOTE: PUBLIC WORK NEEDS OPERATING CONTROLS

Heidi Bullock emphasized that marketing's outputs reach the public and that customer data require careful testing. That is a useful correction to comparisons with internal experiments. A workflow that drafts privately and one that contacts customers automatically need different controls, even if they use the same underlying model.

Establish a minimum operating record

Register each live workflow's purpose, owners, sources, permissions, version, acceptance criteria, cost limits, monitoring, fallback, and review date. Include dependencies and backup coverage; retire work whose purpose has ended or whose cost exceeds its value.

Dave explained, "I built an org chart for the agents too so I could work out what they were doing." Visibility into the work matters more than the count.

Match controls to the consequence

Workflow exposure	Typical controls to establish	Accountable decision
Internal preparation using approved information	Source access, output checks, clear limitations	Team owner decides acceptable use
External content or customer guidance	Factual and brand review, evidence, approved release	Business owner authorizes publication
Actions affecting systems, spend, or customers	Narrow permissions, limits, logs, exception escalation, rollback	Authorized owner approves action boundaries
Sensitive data or consequential claims	Security and legal review appropriate to the case, restricted access, expert validation	Relevant corporate control owner approves conditions

This is an operating framework, not a substitute for company policy or case-specific legal advice. Involve the appropriate experts for personal data, intellectual property, contractual restrictions, and regulated claims. The business owner still needs to understand the practical limits, rather than treating expert review as a blanket transfer of responsibility.

Evaluate the complete chain

Test inputs, retrieval, outputs, downstream actions, and the human review process. Include incomplete information, conflicting sources, unusual requests, and failure cases. Record which failures can be corrected automatically, which must be escalated, and which should stop the workflow.

Do not assume a model upgrade is an unconditional improvement. Recheck the cases that matter to the business, compare accepted-output cost, and preserve a fallback where possible. A system is

ready to scale when the organization understands both its useful range and its failure behavior.

Review value at three levels

At the task level, measure effort, quality, and accepted output. At the workflow level, measure elapsed time, handoffs, exceptions, cost, and delivery reliability. At the business level, measure the outcome the investment was intended to influence, using a comparison that accounts for other material changes.

For a campaign pilot, for example, faster production is a task result, shorter launch time is a workflow result, and improved qualified-opportunity progression is a potential business result. The third cannot be inferred from the first two. Track it separately and be explicit about whether the evidence is experimental, observational, or anecdotal.

Give the CFO a useful comparison

Compare the proposed workflow with credible alternatives: retain the current process, simplify it without AI, buy a service, use an existing platform, or build. Include setup, operating costs, review, maintenance, and the value of the best alternative use of the team's time. Report sensitivity to volume, failure rates, and consumption rather than presenting a single optimistic case.

When an outcome is revenue-related, avoid counting pipeline as revenue or attributing all growth to the workflow. When an outcome is efficiency-related, separate actual expense reduction from recovered capacity. Finance can help validate that distinction before the team presents a result to the board.

Set scale and stop criteria before launch

Agree on the minimum quality, maximum cost per accepted result, acceptable exception rate, and evidence of benefit required to expand. Specify who can stop the workflow and how the existing process will cover the gap. A pilot that fails its test can still create value by preventing a larger unproductive investment.

The same discipline applies to buying and building. Paige described a “buy first mentality” while recognizing that both approaches have a place. A custom

workflow is attractive when it serves a distinctive need and the company can maintain it; a purchased capability is attractive when it meets the requirement more reliably and economically. The answer can change as the market and the team's capabilities evolve.

IN THE HUDDLE | 2026 FIELD NOTE: ECONOMICS REMAIN UNCERTAIN

Amy King compared the current uncertainty to earlier cloud adoption: “It's hard to know what it's going to cost in a few years.” That uncertainty strengthens the case for flexible commitments, transparent usage, and regular economic review. It does not justify postponing all experimentation or assuming future price declines will rescue a weak investment.

The governance objective is a portfolio of useful, maintainable workflows with accountable owners. Expansion should follow evidence that the system works at the next level of volume and complexity. The number of live agents is an inventory measure, not the definition of progress.

Key Takeaways


Governance combines accountable ownership, bounded actions and full-workflow economics. Evaluate task, workflow and business results separately, using agreed scale and stop criteria.

Questions for Your Team

- 1. Who owns each live workflow and its fallback?**
- 2. Are we measuring accepted results and complete costs?**
- 3. What evidence would make us scale, revise or retire it?**

DECISION POINT | MAKE THE SCALE DECISION EXPLICIT

Agree the quality floor, cost ceiling, acceptable exception rate and evidence of benefit before launch. Name the person who can stop the workflow and the fallback that protects customers.



PART IV

WHAT SHOULD YOU DO MONDAY MORNING?

Action 1 | Choose one strategic outcome

Ask the leadership team to identify the most valuable customer or commercial improvement marketing can influence. Be specific about the audience, the constraint, and why the improvement matters now. The priority might be adoption, deal progression, market entry, or a lower cost of serving a defined need.

Write an outcome charter

State the business objective, the workflow expected to influence it, the baseline, the accountable owner, and the intended measurement period. Name one primary outcome and a small number of quality or customer guardrails. Explain the mechanism that should connect the workflow change to the result.

An illustrative charter might test whether more relevant onboarding guidance helps a defined customer segment reach its first useful result sooner. It should also explain how the team will detect incorrect guidance, excessive contact, or additional support burden. The target and thresholds should come from your baseline and business requirements.

Resolve the resource choice

Agree what the team will stop, defer, or reduce to make room for the pilot. If the benefit is recovered capacity, identify

where it will be reinvested. If the benefit is cash savings, specify which expense will actually change.

How to know it worked

The CEO, CMO, and relevant commercial partner should be able to describe the same objective without mentioning a particular tool. They should also agree on what would make the pilot worth expanding. If the outcome cannot be explained without an automation percentage, the charter needs another pass.

Heidi Bullock's formulation, "now I know where to point my AI," is the test for this exercise. The destination should make the technology choice easier.

Action 2 | Map the work

Bring together the people who perform one strategically important workflow. Include the upstream requester and the downstream recipient so that the map reflects the whole process. Choose a manageable example, such as producing and activating a campaign for one segment, rather than trying to inventory the entire department in one meeting.

The first working session

Walk through a recently completed example, following the actual sequence of events. Record the trigger, inputs, tasks,

systems, decisions, waiting, rework, and acceptance standard. Ask what happened outside the documented process and which person supplied context that was not written down.

Use the map to distinguish necessary work from historical habit. A repeated approval may compensate for poor input quality, while a manual handoff may exist only because two teams use different definitions. Resolve the cause before deciding to automate the symptom.

What to leave with

Create a shared map, a baseline for effort and elapsed time, and a short list of constraints. Name one owner who can coordinate the workflow across teams. Record any missing evidence as a follow-up assignment with an accountable person. Bring your CMO Huddles AI Maturity Calculator results into the discussion to identify capability gaps that could affect the workflow and agree on which to address first.

How to know it worked

People doing the work should recognize the map, including its exceptions. The team should be able to explain where delay or rework occurs and why that matters to the business. The useful result is a clearer intervention, not an elaborate process diagram.

Andrew Cox's account of 13 campaign steps involving ten people provides a practical starting question: which handoffs earn their place? Apply that question to your own process before using another agent to move work between the same queues.

Action 3 | Create a workflow pilot

Select a workflow with a meaningful potential benefit, accessible information, a clear owner, and a reversible deployment path. Keep the first scope narrow enough that the team can observe failures and learn from them. A pilot should test an operating hypothesis rather than demonstrate every capability of a tool.

Design the comparison

Define the existing process and the proposed change, using the same audience and acceptance standard where possible. If a controlled comparison is feasible, plan it before deployment. If it is not, document the limits of a before-and-after comparison and record other changes that could affect the result.

Test representative cases, including difficult and incomplete inputs. Record both accepted and rejected outputs, human review time, rework, consumption, and elapsed time. Decide how long

the workflow must run to observe a meaningful result, taking the business cycle into account.

Prepare the operating conditions

Name the business and technical owners, set permissions and a spend ceiling, and establish the fallback. Give reviewers supporting evidence and authority to reject results. Tell affected teams how to report a problem and who will respond.

How to know it worked

The pilot should produce a defensible decision to scale, revise, or stop. A reliable workflow with acceptable economics and a credible outcome signal may justify expansion; faster generation alone does not. Keep the decision record so future teams can reuse what was learned.

Amy King put the principle plainly: “the proof has to come from some kind of experiment.” Build the pilot to generate that proof, including the possibility that the proposed approach is not worth pursuing.

Action 4 | Train everyone

Make AI fluency a supported expectation across the team, with different depth for different responsibilities. Everyone should understand approved use, source checking, and how to evaluate an output. Workflow owners and builders need additional skills

in process design, testing, economics, and maintenance.

Begin with real work

Ask each person to identify a recurring task that is slow, frustrating, or difficult to do well. Choose appropriate examples for a guided session, then compare the result with the existing approach. Discuss what improved, what failed, and what human judgment remained necessary.

Give employees protected time to practice and a place to ask questions. Have managers demonstrate their own learning, including mistakes. Share useful patterns after they have been tested, so the organization develops common capability rather than accumulating isolated tricks.

Make the career connection explicit

Describe how the role is evolving and which capabilities will matter. Offer opportunities to apply new skills to meaningful work, and distinguish a learning gap from a performance problem. Avoid asking people to become technical builders unless that is a real part of their role and workload.

How to know it worked

Evaluate representative work before and after training. Look for better judgment, fewer preventable errors,

and improvements that colleagues can reuse. Attendance and tool usage are participation measures, not sufficient evidence of capability.

Heidi Bullock's promise was to make employees "more marketable than when you first came in." Use that standard to assess whether training is genuinely expanding the team's ability to contribute.

Action 5 | Build your knowledge layer

Start with the information needed by the pilot. Gather the approved product facts, positioning, customer evidence, definitions, and process guidance that the team currently has to hunt down. Identify the authoritative source and owner for each category.

Make reliability visible

Record where information came from, who approved it, when it should be reviewed, and who may use it. Resolve contradictory definitions before connecting them to automated work. Keep sensitive information within the access boundaries already required by the business.

Build a small set of representative questions and check whether users can retrieve accurate answers with supporting evidence. Include questions that should remain unanswered because

the information is missing or uncertain. A system that admits it lacks evidence is more useful than one that invents a confident answer.

Plan for maintenance

Define how updates enter the knowledge set and how obsolete content is removed. Assign backup ownership and a review rhythm. Capture failures from the pilot so recurring gaps become improvements to shared information rather than repeated manual corrections.

How to know it worked

The workflow should use the same approved facts across users and outputs, and the team should be able to trace consequential claims. Measure the frequency of missing, stale, and conflicting information. Expand the knowledge set where those failures indicate a real need.

Dave called the shared intelligence layer "the only reason why this works." Treat dependable knowledge as operating infrastructure, with an owner and a maintenance commitment.

Action 6 | Measure business outcomes

RESEARCH SNAPSHOT (2026) | MEASUREMENT MATURITY

Only 6% of benchmark respondents connected AI-enabled workflows to business or capacity metrics, while 19% tracked only adoption or usage.¹³ Agree on the measurement design before the pilot starts. Separate task performance, workflow performance, and business contribution, with consistent definitions agreed across Finance, sales, customer success, and product wherever the metric crosses organizational boundaries.

Build a small scorecard

Track the primary business outcome, quality, end-to-end cycle time, cost per accepted result, and operational exceptions. Add a measure of recovered capacity and where it was used. Keep the scorecard focused enough to support decisions in a regular management meeting.

Use a credible comparison and label its limitations. A campaign launched during a major product release or market change may perform differently for reasons unrelated to AI. Record those conditions and avoid claiming causation from a convenient before-and-after story.

Translate efficiency carefully

Report actual spending reductions separately from the value of time saved. If the team retained its staff, describe the new capacity and the work it enabled. If the result is avoided hiring, explain the counterfactual and the evidence that additional staffing would otherwise have been needed.

How to know it worked

An executive unfamiliar with the technology should understand the benefit, the cost, and how confident the team is in the result. The scorecard should lead to a decision about scale, redesign, or retirement. More metrics are useful only if they improve that decision.

Grant Johnson, CMO of Chief Outsiders, returned to the importance of setting a baseline and tracking improvement in business metrics over time. Use that discipline to make the AI investment part of the same performance conversation as the rest of marketing.

Action 7 | Review your organization

Review the organization after the pilot has revealed how the work changes. Compare the responsibilities the workflow needs with the roles and reporting lines

13 CMO Huddles and Benchmarkit, Marketing AI Maturity Benchmark 2026, full draft. Measurement question: n=172; single response. Percentages rounded. Study link in Appendix A.

that exist today. Look for persistent handoffs, duplicated capability, overloaded reviewers, and knowledge that depends on one person.

Test the smallest useful change

Consider a clarified decision owner, a shared service, a revised role, or a temporary outcome pod before assuming a full reorganization is necessary. Define what the change should improve and how it will be evaluated. Preserve the functional expertise and customer relationships that make the work effective.

Use the three scenarios in Chapter 8 to discuss the options. Ask which model best fits your company's product complexity, customer relationships, technical support, and financial constraints. Make the choice about the work the company needs to perform, not the visual appeal of a new chart.

Plan the transition

Specify who owns the customer experience during the change, how employees will learn new responsibilities, and which services cannot be interrupted. Include backup coverage for technical and knowledge roles. Review the workload left behind when responsibilities move.

How to know it worked

The change should improve coordination, decision speed, quality, or business performance without unacceptable cost or overload. Assess those outcomes after a defined operating period. If the chart changes but the same bottlenecks remain, revisit the work rather than declaring the transformation complete.

DREW'S TAKE | LET THE WORK SHAPE THE CHART

The org chart is a consequence of your choices. Keep testing those choices as the evidence changes. An AI-ready organization remains capable of adapting after the first redesign.

What We're Watching | The questions for the next edition

The most useful outcome of the FOMO task force is a better sequence of decisions. Start with the advantage the business wants to create, understand the work that supports it, and redesign the workflow with clear ownership of knowledge, judgment, and economics. Let demonstrated changes in the work inform how roles and teams evolve.

The conversations also show why the work requires judgment. Paige described the pressure to transform an established organization, while Dave's smaller team illustrates a more expansive set of possibilities for shared intelligence and AI-supported production. Amy emphasized infrastructure and economics, Heidi emphasized architecture and capability, and Julie and JD kept bringing the discussion back to business purpose and human relationships. Those perspectives are complementary without being identical.

Several questions deserve more evidence. Which workflows deliver sustained business improvement after full

operating costs are counted? How does review quality hold up as output grows? What happens to employee learning when routine work is automated, and which organizational structures maintain both expertise and cross-functional speed? The answers may differ by company size, industry, and customer complexity.

Our intention is to revisit this playbook every six months. Future editions should incorporate longitudinal results, unsuccessful experiments, changes in economics, and feedback from teams that have tried these recommendations. A useful case should explain the baseline, the intervention, the operating conditions, the result, and what the team would do differently.

We invite the CMO Huddles community to challenge the framework and contribute that evidence. Share the workflow that improved, the investment that failed its test, or the organizational change that made a real difference to customers. The next edition will be stronger if it includes what proved difficult alongside what proved possible.

For now, choose a consequential outcome and one piece of work that stands in its way. Bring the right people into the conversation, test an improvement, and use what you learn to make the next decision. That is how the future marketing organization becomes a practical leadership task.

Appendix A | Research Bibliography

The external sources below provide context and evidence for specific claims in the text. Their populations, dates, and methods differ; they are not combined into a single benchmark. Research findings are distinguished from the playbook's recommendations and illustrative planning examples.

CMO Huddles and Benchmarkit, **Marketing AI Maturity Benchmark 2026, full draft report**. Survey fielded July–August 2026; 194 captured responses and 153 complete, scoreable responses, with valid bases varying by question. Study source: [Marketing AI Maturity Benchmark 2026 - Full Draft Report](#).

McKinsey, **The state of AI: How organizations are rewiring to capture value**, March 2025. Survey of 1,491 participants in 101 countries, fielded in July 2024; workflow redesign and earnings findings are self-reported associations. [Read McKinsey research](#).

Gartner, **2025 CMO Spend Survey press release**, May 2025. Survey of 402 marketing leaders, largely at companies with more than \$1 billion in revenue; cited for budget context. [Read Gartner findings](#).

BCG, **AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value**, October 2024. Survey of 1,000 executives across 59 countries; cited for organizational capabilities and implementation challenges. [Read BCG findings](#).

MIT Sloan, **Humans and AI: Do they work better together or alone?**, October 2024. Institutional account of research by Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone; meta-analysis of 370 results from 106 experiments published between 2020 and mid-2023. [Read MIT Sloan findings](#).

Faram Medhora, Forrester, **The Agentic Business Fabric Is How AI Will Transform Enterprise Applications**, September 2025. Analyst perspective on operating-model, job, and business-data changes; used as a conceptual source, not a quantitative marketing benchmark. [Read Forrester perspective](#).

Harvard Business School AI Institute, **Navigating the Jagged Technological Frontier**, September 2023. Summary of field research involving 758 BCG consultants; illustrates variation in task fit and human-AI work design. [Read HBS study summary](#).

Jorge Tamayo, Leila Doumi, Sagar Goel, Orsolya Kovács-Ondrejko, and Raffaella Sadun, Harvard Business Review, **Reskilling in the Age of AI**, September–October 2023. Further reading on workforce development. [Read HBR article](#).

Harvard Business School, **Reskilling in the Age of AI Wins 2023 HBR Prize**, April 2024. Public institutional account of the HBR research, including its interviews with leaders at about 40 organizations. [Read HBS research account](#).

Appendix B | Task Force Members and Contributors

The task force brought together practitioners confronting different organizational conditions. Their perspectives helped shape this CMO Huddles point of view. Participation does not imply endorsement of every recommendation.

Task force contributor	Title and company
Amy King	CMO, Relias
Andrew Cox	CMO, Forrester
Christine Royston	Former CMO, Wrike
Grant Johnson	CMO, Chief Outsiders
Heidi Bullock	CMO, Tealium
JD Dillon	CMO, Tigo Energy
Julie Kaplan	CMO, Higher Logic
Kathie Johnson	Interim CMO, Digital.ai
Paige O’Neill	CMO, Culture Amp
Drew Neisser	Founder and CEO, CMO Huddles

Additional contributors include Micheline Nijmeh, CMO of ThoughtSpot; Rani Salehi, principal analyst at Forrester; and Dave Anderson, CMO of PointFive. Where speaker identity is uncertain, quotations are attributed to “one CMO.”

Appendix C | Methodology and Source Notes

This playbook combines practitioner perspectives from the FOMO task force and expert conversations with published research and the CMO Huddles X Benchmarkit CMO AI Maturity Benchmark study. Participant examples illustrate decisions and tradeoffs; they should not be treated as representative performance estimates.

CMO Huddles and Benchmarkit fielded the study among B2B marketing leaders in July and August 2026. The data file contains 194 captured responses, including 153 complete, scoreable responses used for maturity scores and tiers. Individual findings use the valid response base for each question, as specified in the footnotes. Percentages are rounded; multiple-response questions do not sum to 100%.

The study's 17 scored questions cover strategy, workflow, talent, governance, and measurement, using a weighted model. Its thresholds are diagnostic design choices, not externally validated industry standards. Findings describe respondents' self-reported practices and outcomes, and do not establish that AI caused a business result. Use the CMO Huddles AI Maturity Calculator to assess your organization with the 17-question model and review your dimension scores alongside peer benchmark context.

Appendix D | Decision Worksheets

Should I automate this?

Use this worksheet for one task within a named workflow. First confirm that the task should exist at all and that a simpler rule or process change would not meet the need. Then test whether automation improves the accepted result once review, failure, and maintenance are include.

Decision field	What to record
Purpose	Customer or business result the task supports
Current work	Trigger, inputs, output, owner, effort, waiting, and rework
Alternative	Eliminate, simplify, use rules, assist with AI, or automate
Quality	Acceptance standard and representative difficult cases
Human role	Decisions, review, exceptions, and authority to stop
Economics	Setup, consumption, review, maintenance, and expected volume
Test	Comparison, measurement period, and evidence needed
Decision	Proceed, revise, or retain current method, with rationale

Proceed when the task has a clear purpose, the output can be evaluated, and the operating conditions are supportable. Retain or redesign human work where the system cannot meet the required standard at an acceptable cost. Record the reason so the decision can be revisited when conditions change.

Should I hire?

Start with the work and capability gap rather than the vacant title. Consider whether the need is enduring, whether existing staff can develop the capability, and whether a shared or external resource would be more appropriate. Include workload created by AI systems, not only the work they remove.

Decision field	What to record
Unmet need	Outcome at risk and capability missing
Capacity evidence	Current workload, recovered capacity, and remaining constraint
Options	Develop, redeploy, share, contract, or hire
Role design	Judgment, expertise, collaboration, and technical responsibilities
Business case	Full cost and expected contribution relative to alternatives
Success review	Evidence the role is solving the intended problem

A hire is justified by an important, persistent need and a credible role design. An avoided hire is credible only when the organization can demonstrate how the work will be covered. Neither decision should depend solely on an assumed automation percentage.

Should I reorganize?

Look for a recurring organizational constraint that simpler changes have not resolved. A new structure should improve decisions, handoffs, service, or resource allocation in a way the team can observe. Identify what must remain stable during the transition.

Decision field	What to record
Trigger	Persistent friction or strategic change requiring action
Evidence	Where the present structure impedes the outcome
Smallest change	Decision rights, role revision, shared service, pod, or reporting line
Model	Optimizer, Integrator, outcome-based design, or a combination
Transition	Customer continuity, training, backup coverage, and ownership
Evaluation	Expected improvement, review period, and reversal conditions

Proceed when the change addresses a demonstrated constraint and the transition can be managed. If the proposed chart merely redistributes existing confusion, return to the workflow. A structure is useful when it makes the work and accountability clearer.

Should I buy?

Evaluate a purchased capability against the actual workflow and acceptance standard. Test relevant cases, understand data handling and integration requirements, and make ongoing economics visible. Ask who owns the service when the vendor changes a feature, price, or dependency.

Decision field	What to record
Requirement	Problem, audience, workflow, and required standard
Fit	Evidence from representative work, including failures
Total cost	License, usage, implementation, review, support, and exit
Dependencies	Data access, integration, internal skills, and vendor reliance
Ownership	Business owner and technical support responsibility
Exit	Exportability, fallback, contract flexibility, and retirement plan

Buy when the capability meets the need more reliably or economically than the alternatives. Keep the comparison specific to the organization's volume and complexity. A feature-rich product can still be a poor fit for the workflow that matters.

Should I build?

Build where the workflow serves a distinctive need, existing options are insufficient, and the organization can maintain the result. A working prototype is only the beginning of the investment. Compare the lifetime responsibility with the value of the advantage the custom system creates.

Decision field	What to record
Strategic reason	Why this need warrants a custom capability
Alternatives	Existing tools, purchased services, simpler process changes
Technical scope	Data, integrations, permissions, testing, and monitoring
Ownership	Builder, maintainer, business owner, and backup coverage
Economics	Setup, operation, review, maintenance, and change costs
Reuse	Which other workflows benefit from the same capability
Decision gates	Prototype evidence, production standard, scale, and stop rules

The strongest custom investments often create reusable capability, such as an approved knowledge service or a workflow that handles a distinctive customer requirement. The

weakest depend on one person, lack a support plan, or reproduce something already available at lower total cost. Revisit the build decision as both vendor offerings and internal capability evolve.

A common decision record

For any of the five decisions, record the chosen option, the alternatives considered, the evidence, the accountable owner, and the next review point. Include the condition that would change the decision. This makes the framework useful beyond the initial meeting and prevents assumptions from hardening into permanent policy.

Keep uncertainty visible in the record. A provisional decision can be the right one when the next step is reversible and designed to produce useful evidence. The objective is disciplined learning in service of the business.

Appendix E | AI-Ready Marketing Scorecard

Use this one-page scorecard in a monthly value review. Complete it for a priority workflow or a small portfolio with comparable outcomes, using consistent definitions. Keep one accountable owner for the business result and one for operational reliability.

Dimension	Measure and evidence	Review question
Strategic contribution	Primary outcome versus baseline and comparison	Is the intended business result improving?
Customer value	Adoption, satisfaction, relevance, or other defined signal	Is the customer experience better?
Quality	Acceptance rate, factual errors, rework	Does the output meet the same or a higher standard?
Workflow performance	End-to-end time, waiting, handoffs	Has the real constraint improved?
Economics	Full cost per accepted result and actual cash effect	Is the benefit worth the complete cost?
Capacity	Hours recovered and their demonstrated use	What valuable work became possible?
Knowledge	Missing, stale, conflicting, or unsupported information	Can the workflow rely on its sources?
Reliability	Exceptions, incidents, recovery, and owner coverage	Can we operate this dependably?
People	Capability, workload, learning, and review burden	Is the team becoming more capable?
Decision	Scale, revise, hold, or retire	What evidence supports the next step?

Record the baseline, current result, comparison method, accountable owner, and next action beside each relevant measure. Mark an item unknown when evidence is missing, and assign the work needed to resolve it. Do not manufacture a green status from a lack of observed problems.

Use three decision states: expand when agreed outcome, quality, and economic criteria are met; revise when the hypothesis remains plausible but a specific constraint needs work; stop or retire when the benefit no longer justifies the cost or the operating conditions cannot be met. A hold can be appropriate when the business cycle has not yet produced enough evidence.

Finish the review by deciding what to do, who owns it, and when the evidence will be reviewed again. An AI-ready marketing organization learns from its work and changes its choices accordingly.