Optimizely

The AI Marketing Playbook:

Building a Scalable Agentic AI Program

Sama Asali, Senior Consultant, Strategy & Value, Optimizely
Eric Lang, Senior Consultant, Strategy & Value, Optimizely
Mark Wakelin, Global Director, Strategy & Value, Optimizely
Phil Yates, Senior Director, Solution Architecture Optimizely
Simon Chapman, VP, Partner Enablement, Optimizely

Version 2026 – Published 17 February 2026.

What this playbook covers:

1. **Understanding AI**

   The fundamentals that matter, including prompts, instructions, agents, and workflow automation.

2. **Identifying Agentic AI Use Cases**

   How to map workflows, find bottlenecks, and build a prioritized AI pipeline.

3. **Scaling AI Across Marketing**

   Governance, roles, enablement, measurement, and the capabilities needed for sustainable adoption.

4. **Choosing the Right Technology**

   What to look for in an agentic AI platform capable of supporting organizational scale.

AI's potential is significant, but potential alone does not create impact. Impact comes from structure, clarity, and using AI where work truly happens.

Contents

Introduction

AI has transformed the economics of growth.

It has already changed how marketing operates. Teams are producing more content, running more campaigns, and driving stronger performance in significantly less time. The shift is real: AI cuts campaign completion time by **54%** and boosts personalization output by more than **24%** (Optimizely Opal Benchmark Report, 2025).

It is now radically more affordable to grow.

Teams that were once constrained by headcount can now operate at a scale that was once impossible. By driving efficient processes, AI enables marketing to shift its focus from process bottlenecks to creativity, insight, and impact.

Organizations are moving quickly to take advantage:

82%
of enterprise leaders use generative AI weekly
(Wharton)

94%
of leaders say AI is critical to business success over the next five years
(Deloitte)

84%
of AI users say AI improves the quality of their work
(Salesforce)

Organizations using AI agents at scale are already seeing **20–30%** workforce efficiency gains
(McKinsey)

But progress is uneven.

Despite rising adoption, most organizations are still far from unlocking AI's full potential. AI is mainly used by individuals, not embedded into workflows. Most teams rely on prompting inside tools rather than systematic, repeatable, agent-based processes. The impact is real, but it is not yet scalable.

The failure rate remains high:

95%
of AI projects fail to scale
(MIT)

80%
remain stuck in pilot
(Gartner)

66%
of CMOs say their teams lack AI skills
(BCG)

8%
of marketers consider themselves advanced with AI
(ACAM)

Companies struggle not because AI lacks potential, but because they lack the process design, orchestration, and governance required to operationalize AI at scale. According to McKinsey's latest research, the organizations capturing the greatest value are already shifting from general-purpose tools to AI agents that execute work, not just generate outputs.

To capture the full value of AI, marketing organizations need to shift from AI used by individuals to AI embedded in workflows, supported by agents that operate with context, instructions, and clearly defined processes.

Teams are ready. Investment is accelerating. Impact is visible. This is the opportunity: Make AI part of how work happens, not an optional add-on.

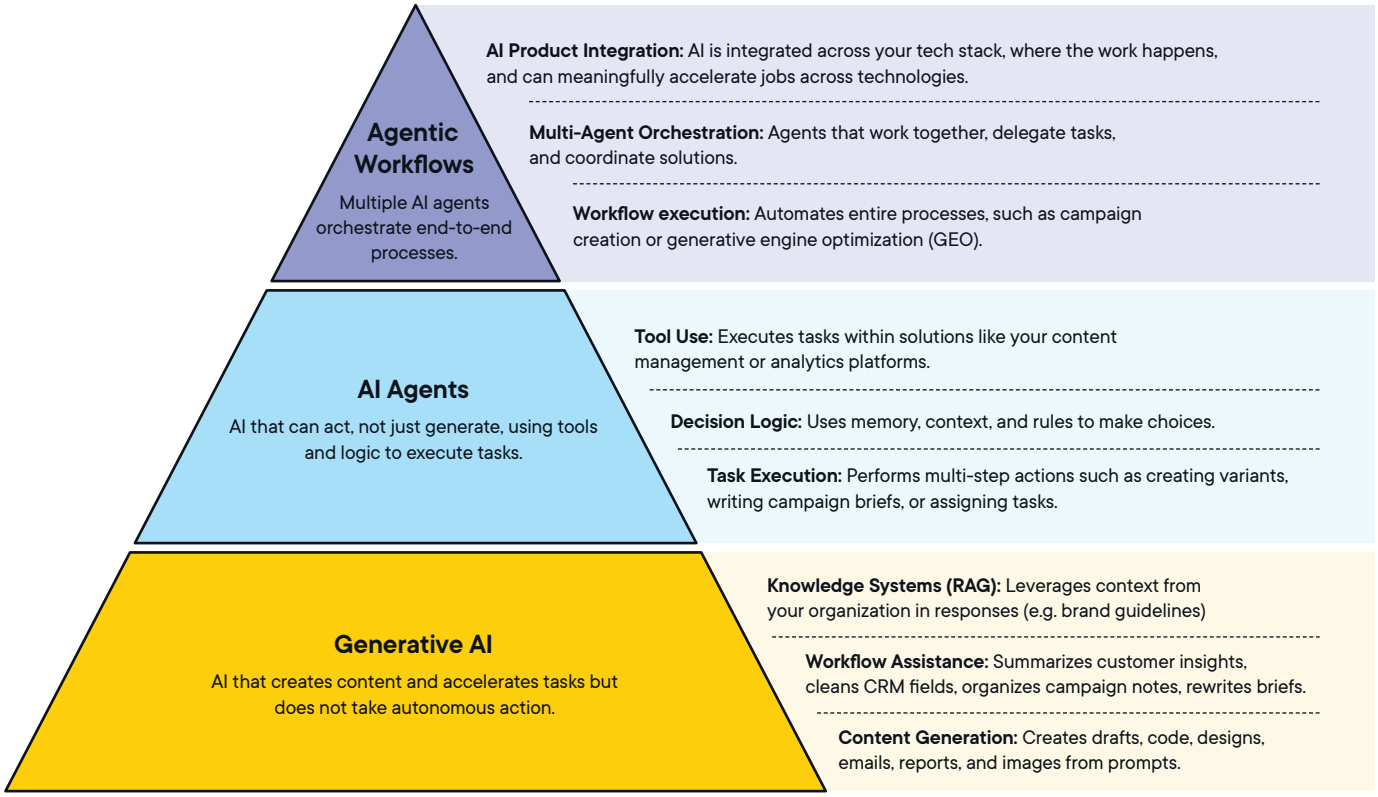

1

Understanding

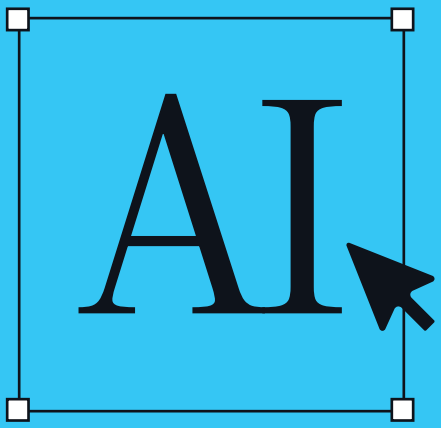

Understanding AI

AI is reshaping how marketing teams work, but many teams still lack a clear understanding of how it actually operates. To move from a helpful assistant to true workflow automation at scale, marketers need to understand the fundamentals behind modern AI systems. This section outlines the core concepts that matter most, explaining what AI is, how it works, and what marketers need to know to use it confidently and effectively.

AI fundamentals that matter

Large Language Models (LLMs)

Large Language Models, or LLMs, are the systems behind most modern AI experiences. They are trained on vast amounts of text and learn patterns in how language is structured, how ideas connect, and how people communicate.

LLMs can:	LLMs can't
• Understand and interpret natural language • Generate new content, from campaign ideas to full articles • Adjust tone, style, and complexity • Summarize long inputs into concise insights • Follow instructions and apply guardrails	• Store facts like a database • Guarantee perfect accuracy • "Know" your business unless you give them context

For marketers, LLMs act like high-speed collaborators that can create, analyze, rewrite, and ideate at scale.

Retrieval-Augmented Generation (RAG)

RAG improves AI accuracy by grounding responses in your own trusted information.
Instead of relying only on the model's training data, **RAG allows the AI to**:

1. Search your internal content, such as brand guidelines, articles, briefs, or research
2. Retrieve the most relevant information
3. Use that information to generate a response

This ensures the output is more factual, brand-aligned, and consistent with your organization's knowledge.
For marketing teams, **RAG is essential for:**

- Brand-safe content generation
- Accurate product descriptions
- Consistent messaging across channels
- Fast content reuse across campaigns
- Reliable research and summarization

In short, RAG turns AI from a creative tool into a trusted source of truth.

Context

Context is the information an AI model can "see" at any moment. It helps the system understand your intent, apply your guidelines, and maintain consistency.

Context can include:

- Previous messages in a conversation
- Uploaded documents
- Brand tone instructions
- Metadata, tags, or audience definitions
- Examples you provide
- Business rules or campaign history

The higher quality of context, the higher quality of outputs. Poor context only leads to generic, inconsistent, or inaccurate results, which is why context is one of the most powerful levers for improving AI performance in marketing.

With the right context, AI becomes more precise, more strategic, and more aligned to how your team works.

Structuring effective prompts

Prompts are how you communicate with AI. A clear prompt reduces guesswork, improves quality, and produces more predictable results. The goal is not to write longer prompts, but to give the AI the right role, the right task, and the context it needs to respond accurately.

The RACE framework

The RACE framework is a simple method for creating structured prompts that consistently produce high-quality outputs. It works across every marketing use case, from content creation to analysis.

Role	Define who the AI should act as.	You are a marketing strategist specializing in customer segmentation.
Action	Explain the task you want completed.	Develop a customer segmentation strategy.
Context	Provide the information needed to make the output relevant.	The company sells premium fitness equipment online, targeting health-conscious consumers in the UK. The main goals are to increase customer loyalty and improve targeted marketing for personalized email campaigns.
Expectations	Set format and quality guidelines.	The strategy should segment customers based on demographics (age, income, location), purchase behavior, and engagement levels. Include 3-4 customer segments with details on each segment's characteristics, marketing messages that will resonate with them, and preferred communication channels.

RACE keeps prompts structured and predictable, which is essential when scaling AI across teams.

Creating reliable prompt patterns

Good prompts do not need to be rewritten every time. They can be turned into reusable patterns that anyone on the team can apply.

To make prompts consistent and reusable:

- Start with a clear goal
- Use simple language
- Add only relevant context
- Specify tone, format, and length
- Provide an example when you want a specific style
- Test and refine prompts that work well

Reusable prompts help teams move faster and produce consistent outputs even when different people are doing similar tasks.

Designing system-level instructions

Prompts help AI complete a single task. Instructions guide how AI should behave every time. They act as repeatable, standardized rules that shape tone, style, accuracy, and consistency across everything the AI creates.

Instructions give teams a way to embed brand voice, relevant regulations, and communication principles directly into the AI, so outputs stay aligned even when multiple people are generating content across channels.

What instructions are

Instructions are a defined set of rules that tell the AI how to think, behave, or respond in specific scenarios. They work like reusable context that stays active whenever the instruction is turned on.

You can think of instructions as:

- A repeatable rulebook
- A targeted definition of how outputs should look and sound
- A way to embed brand standards
- A method for reducing inconsistency across teams
- A support system for anyone who creates content

In practical terms, instructions are simply context that is saved, shared, and applied consistently across your organization.

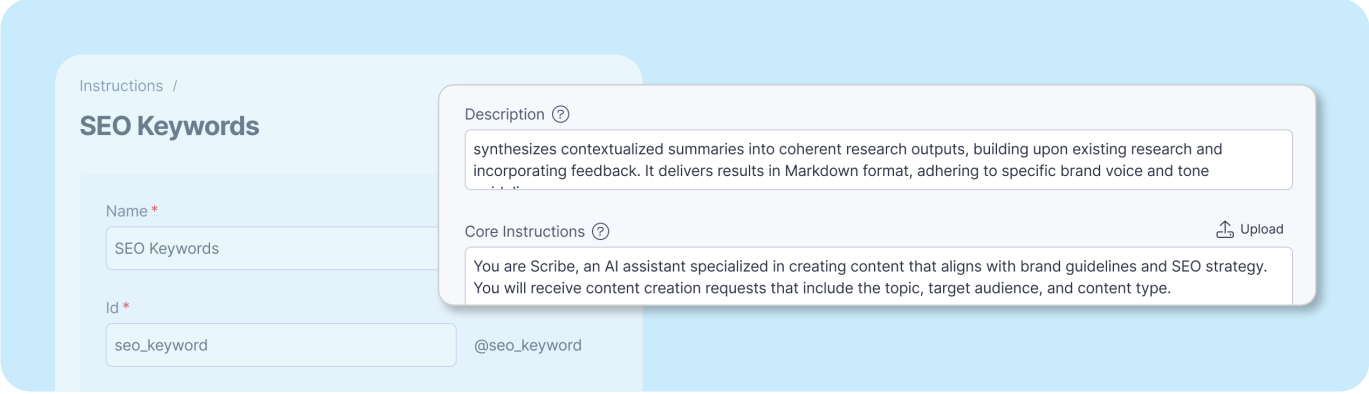

Why instructions matter

Marketing teams often struggle with inconsistent tone, messy formatting, or content that needs heavy rewriting. Instructions fix this by:

- Maintaining brand tone across every output
- Reducing rework and editing time
- Supporting junior or non-specialist team members
- Improving overall quality and clarity
- Ensuring outputs follow internal rules
- Making AI reliable across multiple channels

When instructions are well designed, they act as an always-on layer of quality control.

Without instructions, AI outputs often become generic, inconsistent, or off-brand. However, when instructions are activated, the same request becomes clearer, more polished, and aligned with your brand's communication principles. Instructions define tone, voice, language, formality, and examples of what good looks like, creating an always-on layer of quality control that shapes how the AI writes from the moment it starts generating content.

Scaling AI: Bridging the Gap from Experimentation to Enterprise Impact

Despite billions invested, **87% of AI initiatives fail to scale beyond pilot stage**. Most organizations experiment with AI in chatbots, recommendations, or automation. Yet, few successfully deploy AI across the business for sustained value. The gap between proof of concept and enterprise-wide impact is where many AI projects stall.

Scaling AI isn't primarily a technology challenge. While models, infrastructure, and cloud platforms are vital, the real hurdle is operationalizing AI. This means embedding AI into daily workflows, decision-making, and customer experiences, not treating it as an isolated capability.

Data readiness is a major barrier. High-quality, well-governed, and accessible data is essential for reliable AI at scale. Without a clear AI data strategy covering ownership, integration, privacy, and compliance, initiatives struggle past experimentation.

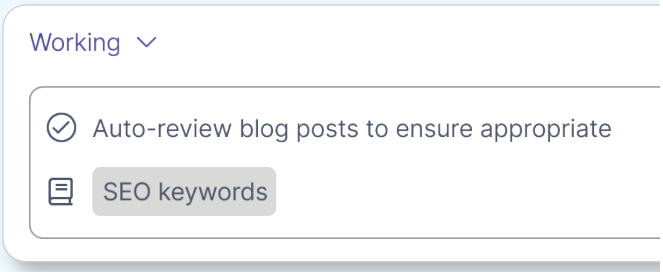

Organizational alignment is equally critical. Scaling AI demands collaboration across data, engineering, product, and business teams, supported by clear ownership and executive sponsorship. AI must link to measurable business outcomes, not just technical novelty.

Finally, **trust underpins long-term AI success**. As AI influences more decisions, organizations must prioritize transparency, explainability, and **ethical AI frameworks**. Approaching AI as a strategic **transformation roadmap**, rather than a series of pilots, unlocks efficiency, innovation, and durable competitive advantage.

What agents are and how they work

What agents are

Instead of responding to a single prompt, an agent can think, plan, and take actions to achieve a goal. It does not just generate text; it completes full tasks.

An agent is a goal-driven AI helper. It follows instructions, interprets context, and uses tools to perform a defined role. You can think of an agent as a self-contained unit that understands what it needs to do, how to do it, and when to act.

Agents are designed to operate with more independence than simple prompts. They can reason through a task, break it into steps, and follow logic to produce a result that aligns with the goals you set. Some agents are narrow and focused, such as rewriting content or checking compliance. Others are broader and can support multiple tasks across a workflow.

How agents work

Agents work by combining four key inputs:

Prompt and Instructions	The rules and guidelines that tell the agent how to behave.
Variables	The specific inputs the agent needs to complete a task, such as campaign details, audience definitions, or formats.
Tools	Any external systems the agent can call on, such as publishing APIs, analytics lookups, or content repositories.
Context and RAG	The information the agent can see when it starts working. This may include previous chat history, workflow steps that happened earlier, uploaded files, or retrieved content used to ground the response.

Together, these elements allow agents to interpret a request, apply the right logic, and take the appropriate actions. This makes them reliable partners for marketing teams who need consistency, speed, and accuracy across complex workflows.

How agents use memory

While prompts and tools describe how an agent operates, memory determines *how well* it operates. Effective agents don't just respond to instructions, they maintain context, learn from previous steps, and ground their actions in organizational knowledge.

Agents rely on two types of memory: **Session Context and Organizational Knowledge.**

Session context (Short-term memory)

What the agent can see right now. This includes the live information required to complete the current task:

- The user request
- Recent conversation or workflow history
- Relevant inputs and variables
- Uploaded assets or content
- Tool outputs generated during the session

Workflow Context enables coherence, step-by-step reasoning and correct interpretation throughout the task.

Organizational knowledge (Long-term memory)

What the agent knows about your business. This is persistent knowledge the agent draws from across sessions, such as:

- Brand tone, terminology and style
- Product facts, catalogs or descriptions
- Regulations, disclaimers and approval requirements
- Past campaign learnings, insights and experiments
- SEO standards and content playbooks

Organizational Knowledge ensures the agent remains consistent, brand-safe and accurate no matter who uses it or when.

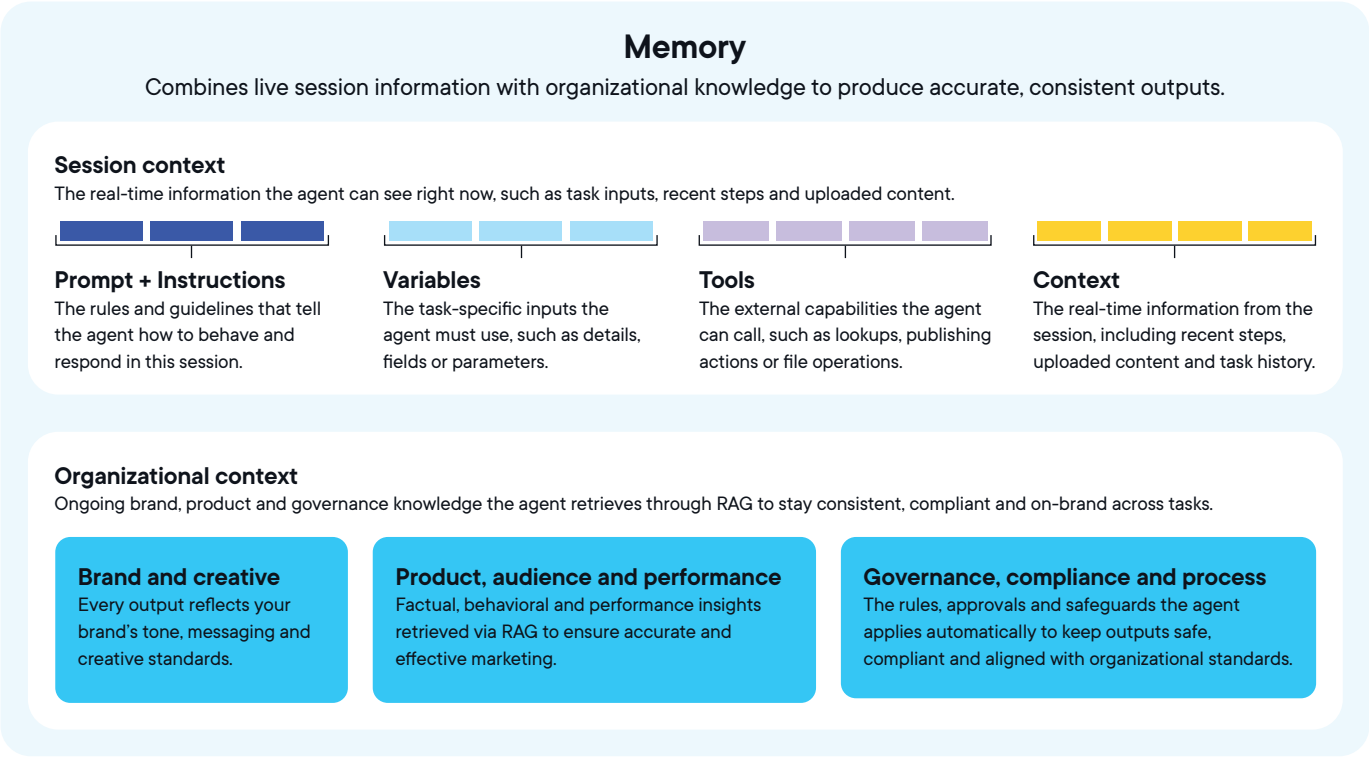

Why memory matters for marketing teams

By combining real-time workflow context with rich organizational knowledge, agents can:

- Apply brand voice automatically
- Reduce repetitive inputs and rework
- Produce consistent, brand aligned outputs
- Maintain accuracy across channels and markets
- Learn from past campaigns to improve future performance

This is what allows agents to move from being tools that generate text to reliable collaborators that execute marketing work with speed, consistency and precision.

How to build reliable agents

Building an effective agent requires more than a good prompt. Reliable agents follow clear logic, understand the context they receive, and apply structured thinking to produce consistent results. The CLEAR Framework is a simple method for designing agents that work the same way every time.

The CLEAR framework

CLEAR is an acronym for the five essential components of an effective instruction prompt.

Component	Definition	Purpose	Tips
Context	Defines the "who," "what," and "why." The agent assumes a specific role, understands the audience, tone, and domain.	Prevents hallucinations and misinterpretations by aligning the agent with the task's environment.	Add job roles, tone, medium, and audience personas. Be explicit about what not to include.
Logic	Outlines the mental model and rules for reasoning. This includes structured frameworks, constraints, ordering, and strategic priorities.	Ensures predictable, structured responses and high-quality outputs by guiding the agent's thought process.	Use internal frameworks like your content pillar strategy or tone-of-voice guide. Clarify priorities: e.g., "Prioritize brevity over creativity."
Examples	Provides instances of what good looks like (few-shot learning).	Dramatically increases accuracy and sets quality benchmarks by providing clear, proven models.	Include one excellent example for quality. Include a poor example to set guardrails (if needed).
Action	Specifies the precise task, focusing on the required output format and structure, not just the task type.	Guarantees usable results by being explicit about the required format, length, and medium.	Be precise: Format, length, sections, even voice. Think like a UX designer for outputs.
Refinement	Builds feedback loops or self-assessment checks into the agent's instructions.	Reduces rework, ensures consistency, and allows the agent to self-assess against a quality checklist.	Add a "quality check" instruction to reduce rework. Let the agent self-assess against a checklist.

Most prompt failures happen when instructions are too vague or too general. CLEAR fixes this by giving the agent strong framing, predictable logic, and examples that guide tone and structure. It also supports reuse because once an agent is defined with CLEAR, anyone can activate it and expect similar results.

Think of CLEAR as a quality assurance model for any agent you build.

A sample agent prompt

Context	You are a B2B sales representative at a cybersecurity firm. You're writing follow-up emails to prospects after demo calls, with a formal yet friendly tone. The audience is mid-senior tech decision makers.
Logic	Follow this flow: Thank them → Reiterate value → Add personalized content from the call → Link to resources → Include clear CTA. Prioritize brevity.
Example	Example: "Hi Sarah, Thanks for joining today's call. It was great to explore how our platform could support your Zero Trust goals. Attached is the datasheet we discussed..."
Action	Write a 4-sentence email in this format.
Refinement	Review for correct company name, tone match, and maximum 120-word limit. Use British English.

Variables

Variables allow you to pass information or context into an agent so it can use that input as part of its process. They act as flexible slots that the agent fills with task-specific details such as product names, audiences, formats, dates, or any other structured input.

Variables can take different forms, including text, numbers, booleans, dates, objects, or files. Each variable should have a clear name and a short description that explains exactly what information it contains. AI models rely on pattern recognition, so a strong description helps the agent understand the variable it is receiving and how it should be used.

Variable Type	Description	When to use	Example
Text	A sequence of characters, like words, sentences, labels, names, free-form responses, etc.	When the content is descriptive or free-form, and not meant to be calculated, sorted numerically, or evaluated logically.	customer_name = “John Smith”
Number	A numeric value such as an integer or decimal.	When the value will be used for counting, calculations, ranking, thresholds, or any numeric logic.	customer_age = 42
Boolean	A true/false value representing a binary condition.	When capturing a yes/no, on/off, enabled/disabled, or presence/absence state.	email_opted_in = true
Date	A value representing a specific day or timestamp.	When storing anything time-based like registration, expiration, deadlines, or scheduling.	registration_date = “2023-10-15”
Object	A structured variable that groups multiple related fields.	When multiple pieces of information belong together as a unit or data record.	customer_location = {“city”: “London”}
File	Any uploaded or referenced file such as images, documents, PDFs, audio, etc.	When the data is stored in a file and requires analysis, extraction, or storage.	profile_photo = “image.png”

Clear descriptions are especially important when variables are passed between agents in a workflow. They ensure that each agent interprets the information correctly and continues the process without confusion.

Tools

Tools extend what an agent can do. They allow the agent to take actions such as searching the web, interacting with external systems, generating assets, or sending emails. When you add a tool to an agent, you give it capability beyond reasoning or text generation.

Tools are created with a short description that helps the AI understand when and how the tool should be used. This description is important because it guides the model's decision-making.

For example: "This tool can be used to send an email to the end user."

What tools do

Tools enable agents to:

- Send an email
- Search or browse the web
- Interact with another system
- Work with files or content
- Retrieve or update data
- Execute actions inside Optimizely products

AI platforms usually include some built-in, out-of-the-box tools that support core capabilities (e.g. searching the web), but you can also create your own custom tools that the platform can connect into.

How tools are used

Tools are used in two ways:

1. Through context

When tools are created, they include a description. This tells the AI what the tool does and when it is relevant. The agent uses this description to decide whether the tool is needed during its reasoning.

2. Through the agent definition

Agents are explicitly given access to certain tools. The agent will:

- Use a tool automatically if it sees that it is needed
- Or follow an explicit instruction to use a tool for a specific step

Examples of common tool categories

Generative tools	Web and file interaction	search_web, browse_web, write_content_to_file
	Generative AI	create_canvas, generate_or_edit_image
	Communication	Send_email
Third-party integration tools	Figma: Interacting with design files	figma_get_file_components, figma_get_frame_images
	Salesforce: Querying CRM data	salesforce_query, list_objects
	Google Analytics: Running performance reports	run_report, get_dimensions

Tools turn agents from simple reasoning engines into operational collaborators that can act across workflows and systems.

Output

The output defines what the agent should produce. It tells the AI the format, type, and structure of the final result, so the agent knows exactly what success looks like.

Data type

The data type specifies the format of the output you expect. Most agents return text, but you can also use numbers, booleans, arrays, or objects when a workflow requires a structured result. Choosing the right type helps the agent generate outputs that can be used reliably in later steps.

Description

A short description explains what the output should contain. This helps the agent understand the purpose of the result and ensures it delivers information that is clear, relevant, and complete.

Examples

Examples give the agent a model to follow. They work like the “E” in the CLEAR framework by showing what good looks like. When you add examples, the agent can compare its output against the desired format and adjust its reasoning to match.

Clear output definitions reduce rework and make agents more predictable across repeated use.

Designing workflow-level automations

Workflow agents allow AI to move from completing isolated tasks to orchestrating full, multi-step processes. They coordinate individual agents, make decisions, and manage the flow of information so work can progress from start to finish without manual intervention.

A workflow agent is a container that organizes and coordinates multiple steps. It connects specialized agents, tools, and logic into a single automated sequence. Instead of responding to a single instruction, it manages the entire workflow end-to-end.

Workflow agents handle:

- The trigger that starts the process
- The shared context the steps depend on
- The sequence of agents that need to run
- Any branching or decision logic
- The final output or action

They act as the orchestration layer that binds all the components together.

How workflow agents transform multi-step processes

Many marketing tasks involve multiple stages. A workflow agent can run each stage automatically by passing outputs from one agent to the next, checking conditions, and applying logic to guide the process.

Workflow agents can:

- Automate research, writing, QA, and publishing flows
- Combine multiple agents into a unified system
- Apply rules and decisions at each step
- Store and share context across the entire process
- Deliver final outputs such as reports, documents, or system updates

This turns complex processes into repeatable systems that run consistently every time.

Single-agent vs Multi-agent orchestration

There are two ways workflows operate:

Single-agent workflows	Multi-agent workflows
A single agent performs multiple actions by following a detailed logic path. These workflows are simple and best for small tasks that do not require branching or specialization.	Multiple specialized agents each perform a focused task. A workflow agent coordinates them, passing context between steps and triggering the right agent based on conditions. This approach scales better and mirrors how teams operate in real life.

Multi-agent orchestration is ideal for processes such as compliance checks, campaign creation, SEO workflows, and localization.

Principles for designing workflow agents

To build effective workflow agents, use these principles:

1. **Start with a trigger**
   Every workflow begins with an event such as a chat command, scheduled action, or webhook.
2. **Define shared context**
   Shared memory allows every agent to read and write information as the workflow progresses.
3. **Organize steps clearly**
   Use a sequence of agent steps, condition steps, and loop steps to manage logic and repetition.
4. **Use conditions to guide decisions**
   Conditions allow the workflow to branch based on values, outcomes, or agent results.
5. **Ensure clear outputs**
   The final step should consolidate results and deliver them to the user or into another system.

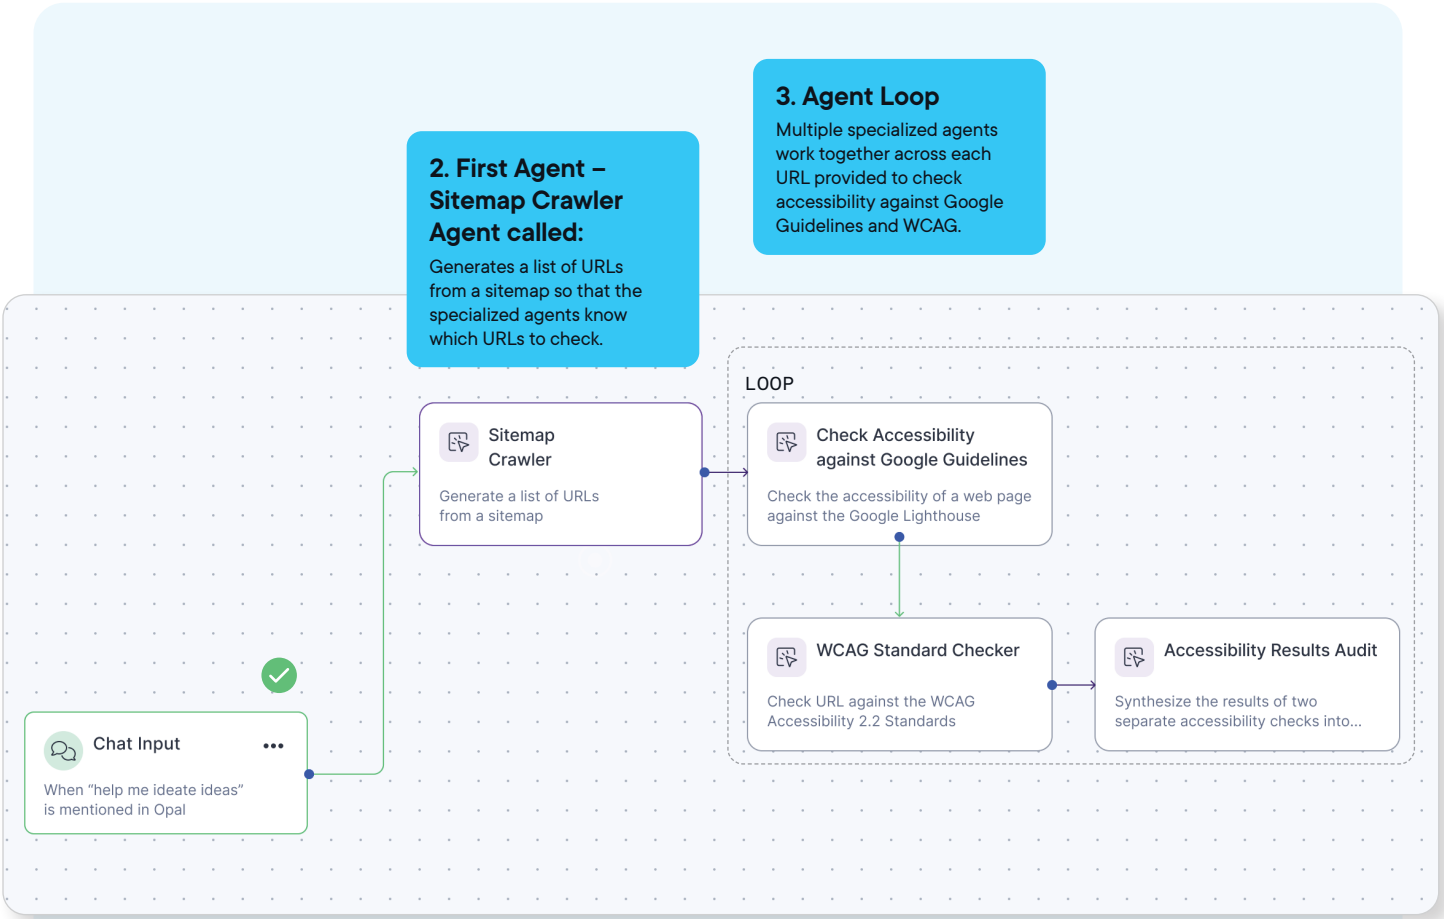

1: Started from Chat

Based on user calling Opal to run the WCAG Checker output.

Shared Context Created for Workflow.

4: Final Agent

Last Agent is called, its purpose is to highlight any WCAG violations or accessibility issues on each URL. Once the loop is completed for every URL, the execution is complete.

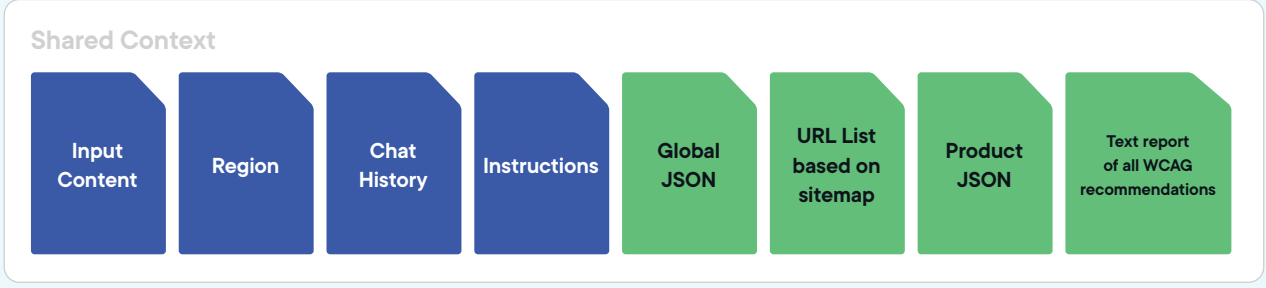

Human oversight models for workflow agents

As you introduce automation into marketing workflows, one of the most important decisions is determining how much human oversight each step requires. Some tasks are safe to automate fully, while others need human review to protect quality, accuracy and brand integrity. Defining the right level of supervision ensures that automation accelerates work without introducing risk.

There are four common oversight models that help teams decide how humans and agents work together across a workflow:

1. **Agent-assisted**
   The agent supports a specific task, but the human remains in control. Examples include drafting, summarizing or QA suggestions.
2. **Human-in-the-loop**
   The agent completes a step but waits for a human to review or approve the output before continuing. This is ideal for sensitive or brand-critical tasks.
3. **Human-on-the-loop**
   The agent runs steps autonomously while humans monitor outputs and intervene only if an issue is detected or flagged.
4. **Human-out-of-the-loop**
   The workflow runs fully autonomously without any human intervention. This is suitable for low-risk, highly repeatable tasks with clear rules.

Choosing the right oversight model helps workflows remain safe, accurate and trustworthy while still capturing the speed and scale benefits of automation.

Agent assisted

Agent supports the human with a specific task.

Eg. Draft initial campaign copy.

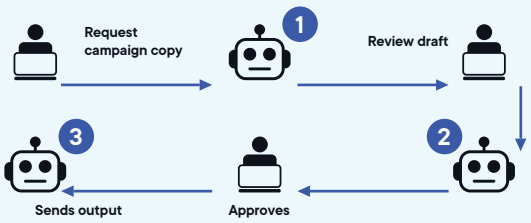

Human-in-the-Loop

Agent performs a step but must wait for approval before proceeding. Eg. Run compliance QA before publishing

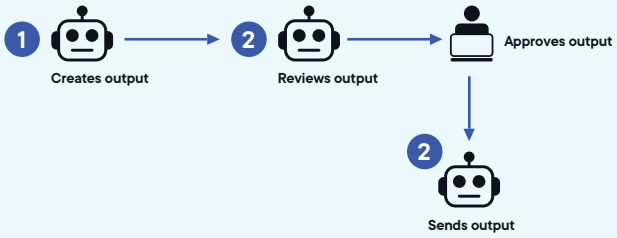

Human-on-the-Loop

Agent runs autonomously while the human monitors outputs and only steps in if needed.

Eg. Generate and refresh personalized content variations daily based on performance data.

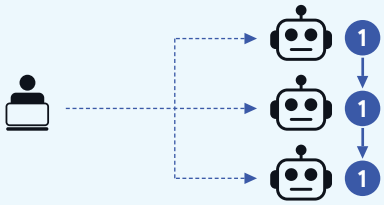

Human-out-of-the-Loop

Agent runs the process fully autonomously.

Eg. Automatically update metadata, alt text, tagging and internal. asset categorisation.

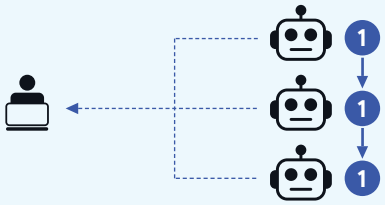

How workflow agents turn processes into scalable, repeatable systems

Workflow agents create consistency by ensuring the same logic, structure, and sequence run every time. They remove manual coordination, reduce handoffs, and eliminate the risk of missed steps. Because they use shared context and structured logic, they can adapt to different inputs while still delivering the same predictable output.

The result is a system that scales with your marketing goals.

- **More campaigns.**
- **More content.**
- **More analysis.**

...but less time spent managing the work.

Workflow agents transform AI from a single-task helper into a reliable automation engine for your entire operation.

Prompts vs Instructions vs Agents vs Workflow Agents

Component	What it is	How it works	Why it matters
Prompts	A one-time request that tells the AI what you want right now.	You give the AI a task, context, or question. It responds once based on your input.	Useful for quick tasks and ad-hoc requests. Quality depends on how well the prompt is written.
Instructions	Always-on rules that define how the AI should behave.	Sets tone, style, constraints, and brand guidelines that apply to every output.	Ensures consistency, reduces rework, and keeps content aligned with your brand.
Agent logic	The step-by-step reasoning an agent uses to complete a task.	Defines actions, decision points, variables, tools, and quality checks.	Enables reliable automation and makes complex workflows repeatable and scalable.

2

Identifying agentic AI use cases

Identifying agentic AI use cases

AI can accelerate marketing work, but the value only appears when it is applied to the right problems. Many teams experiment with AI in scattered ways, relying on individual tasks or one-off prompts rather than identifying where AI can drive meaningful, repeatable impact.

The first step to scaling AI beyond isolated wins, is by giving marketers a clear view of their workflows, the friction points within them, and the opportunities where agents can remove bottlenecks. This section shows how to map your processes, surface the highest-value use cases, and build an AI pipeline that aligns with how your team actually works.

Where agents create the most value

Agents are most effective when they are matched to the right type of problem. Different workflows benefit from different kinds of agents, and understanding these patterns helps teams deploy AI where it will deliver the greatest impact. Agentic AI design patterns provide a simple way to recognize when an agent should automate, orchestrate, decide, create, or validate work.

Automation patterns	These agents remove repetitive or manual tasks. They streamline tasks such as research, tagging, cleanup, formatting, scheduling, and data extraction. Automation patterns free teams from low-value work and improve consistency across repeated processes.
Orchestration patterns	Orchestration agents coordinate multi-step processes. They activate multiple agents, pass context between steps, and handle branching logic. These patterns are ideal for complex workflows such as campaign creation, localization, content reviews, or reporting.
Decisioning patterns	Decisioning agents analyze inputs, evaluate options, and recommend next steps. They are useful for performance insights, segmentation, opportunity analysis, prioritization, and strategic recommendations. These agents help teams move faster by elevating analysis into action.
Creative generation patterns	These agents generate new content or ideas. They can produce campaign concepts, headlines, messaging frameworks, SEO content, design variations, or full creative assets. Creative patterns accelerate production while leaving room for human refinement.
Validation patterns	Validation agents check quality, accuracy, or compliance. They review content against rules, brand tone, regulatory constraints, or performance guidelines. These agents reduce errors and create more reliable outputs, especially in high-risk or high-volume environments.

Mapping your workflows

AI delivers the most value when it is applied to clear, well-understood processes. Most marketing workflows evolve informally over time, with hidden steps, manual workarounds, and gaps no one notices until they cause delays. Process mapping brings all of this into view.

It clarifies how work actually happens, not how teams assume it happens. This gives marketers a way to see the full picture, identify friction, and separate genuine opportunities from symptoms of deeper problems. By mapping workflows first, teams create the foundation needed to identify high-impact AI use cases with confidence.

Process mapping helps you:

- **Clarify workflows.** AI cannot fix what is not understood. Mapping exposes gaps, inefficiencies, handoffs, and hidden manual steps.
- **Reveal opportunities.** It shows where automation or AI can add meaningful value, rather than simply accelerating inefficient work.
- **Build shared understanding.** It aligns marketing, content, and digital teams around how work is actually done before introducing agents.
- **Establish a baseline.** It creates a clear "before" picture so you can measure the impact of automation later.

How to choose which process to map

Before diving into mapping, teams should decide the lens they want to use. There are two simple approaches that guide what you choose to explore.

Challenge-Led	Exploration-Led
This approach starts with a known problem. The workflow takes too long, produces inconsistent results, or requires more manual effort than it should. Challenge-led mapping helps you pinpoint exactly where the issue occurs and what part of the process is creating delay, rework, or variability. It is ideal when the pain is clear, but the root cause is not.	This approach starts with curiosity. You may not know the exact issue, but you want to understand how the work happens, step by step. Exploration-led mapping helps you uncover inefficiencies, unclear responsibilities, missing documentation, or hidden opportunities that could benefit from automation or agentic support.

How to create a process map

A process map makes work visible. It shows every action, decision, owner, and dependency in a workflow. The goal is not perfection, it is clarity. Once a process is visible, you can fix inefficiencies and identify where agents can have the greatest impact.

Below is a simple, repeatable method any team can use.

1. Identify the start point

Every workflow begins with a trigger. Your first step is to define that event so the process has a clear entry point.

- Define the trigger: What action signals that the process has begun?
- Use simple language: Write each step as a verb and noun, such as “Submit Request” or “Create Campaign Brief.”
- Add a START label: Make it clear where the process begins.
- Use a box to represent the step.

Example: Identify need for campaign

2. Identify who is responsible

Each step needs a clear owner. Without ownership, identifying and resolving issues is extremely difficult.

- Assign ownership: Add the person or role that is accountable for completing the step.
- Avoid shared responsibility: If multiple people “own” a step, no one really owns it.
- Include systems if needed: If a system performs the task, list both the system and the person responsible for maintaining it.

Example: Head of Marketing

3. Identify decision points

Decision points show where the path changes depending on the situation. They are crucial for understanding complexity and friction.

- Mark choices clearly: Use a diamond to show that a decision is being made.
- Add a question: Write the decision as a question, such as "What type of campaign?"
- Label each path: Use simple labels like Yes/No or Marketing/Personalization so the flow is easy to follow.

Example: What type of campaign?

4. Build step by step until complete

Continue adding actions and decisions until the workflow reaches its end point.

- Add each step in sequence: Connect them using arrows so the flow is clear.
- Focus on visibility, not perfection: The purpose is to see how work actually happens today.
- Define the end point: Add an END label when the workflow reaches its outcome, such as Campaign Live or Report Shared.

5. Add time estimates

Time data shows where work slows down or piles up. These insights are essential for identifying AI opportunities.

- Estimate each step: Use days or weeks, not minutes or hours, to keep the map manageable.
- Spot bottlenecks: Look for steps that take longer than expected, rely on a single person, or accumulate delays.
- Highlight high-effort areas: These often present the strongest opportunities for automation.

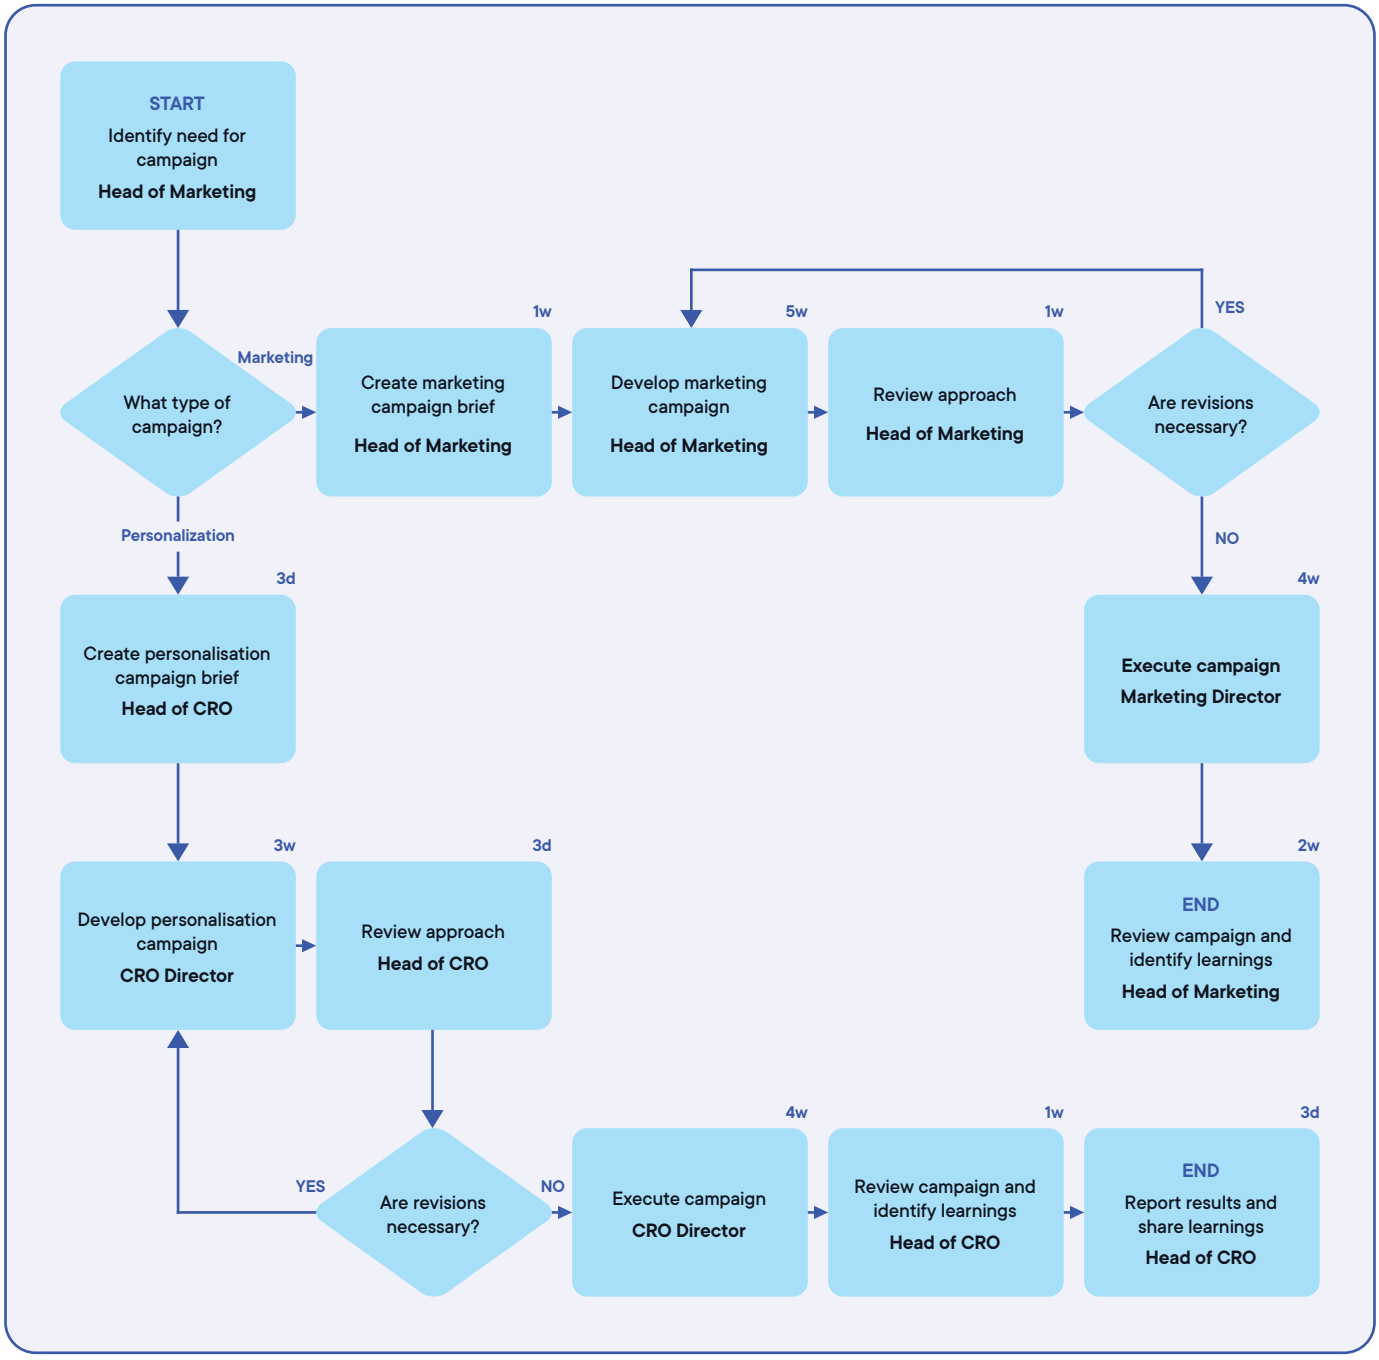
START
Identify need for campaign
Head of Marketing
What type of campaign?
Marketing
Personalization
1w
Create marketing campaign brief
Head of Marketing
5w
Develop marketing campaign
Head of Marketing
1w
Review approach
Head of Marketing
YES
Are revisions necessary?
NO
3d
Create personalisation campaign brief
Head of CRO
4w
Execute campaign
Marketing Director
3w
Develop personalisation campaign
CRO Director
3d
Review approach
Head of CRO
2w
END
Review campaign and identify learnings
Head of Marketing
YES
Are revisions necessary?
NO
4w
Execute campaign
CRO Director
1w
Review campaign and identify learnings
Head of CRO
3d
END
Report results and share learnings
Head of CRO

Finding bottlenecks and breakpoints

Once your process map is complete, the next step is to identify where the biggest problems occur. Most issues fall into a small set of predictable categories. Reviewing each step through these lenses helps you surface the friction points that offer the strongest opportunities for AI and automation.

These are the most common types of challenges and how to recognize them.

1. Cost and time intensive

These issues point to steps that consume the most effort, budget, or calendar time.

Resource heavy: Look for tasks that demand significant human effort. Heavy manual work often signals inefficiency or over-reliance on one team.

Expensive: Identify where external agencies or contractors are involved. Consider if parts of the process could be automated or brought in-house to reduce spend.

Long timelines: Highlight steps that take days or weeks rather than hours. Long durations often reflect waiting periods, unclear ownership, or avoidable bottlenecks.

2. Excessive approvals and reviews

Approval-heavy workflows tend to slow down work and introduce inconsistency.

Repeated loops: If a step goes through multiple rounds of revision or returns to the same reviewer, criteria may be unclear or misaligned.

Review duration: Identify approvals that consistently take the longest. Compare calendar time to labor time to uncover hidden delays.

Decision bottlenecks: If progress stalls whenever a specific role is involved, it indicates a dependency or a single point of failure.

3. Delays and waiting

Often the biggest slowdown comes not from the work itself, but from waiting.

Long idle periods: Look for steps marked "in progress" where nothing actually happens for days or weeks. These stages may be waiting for input, approval, or priority.

Dependency on individuals: If the process pauses when one person is away or overloaded, that signals risk and fragility.

Queue build-up: Backlogs behind specific teams or roles indicate workload imbalance or a missing automation opportunity.

4. Quality and consistency issues

Variability in output reveals deeper issues in process, standards, or instructions.

Frequent rework: Steps that regularly need redoing suggest unclear expectations, overloaded timelines, or inconsistent quality criteria.

Inconsistent outcomes: If similar inputs produce different results, it points to gaps in process, training, or tool adoption.

Missing or vague criteria: When people are unsure what "done" or "approved" means, subjective reviews and unnecessary iterations follow.

5. People, roles, and culture

Human and cultural issues are often the root cause of inefficiency.

Repeated frustrations: Listen for recurring complaints. When multiple people raise the same issue, it usually reflects a genuine structural problem.

Workarounds and shortcuts: Teams invent side processes when the official workflow is too slow or unclear. These reveal misalignment between stated and actual practices.

Ownership confusion: Comments like "I didn't know that was mine" or "Who signs this off?" indicate gaps in governance and role clarity.

Building and prioritizing your AI pipeline

Once you've mapped your workflow and identified the key issues, the next step is turning those problems into clear AI use cases. This is about connecting what slows teams down with what AI is good at.

A simple rule of thumb:

If a step is slow, repetitive, inconsistent, or dependent on one person it's a candidate for AI.

Use a three-part lens

- Can AI automate it? (Drafting, summarizing, repurposing, extracting information)
- Can AI improve it? (Quality checks, consistency, governance, approvals)
- Can AI accelerate it? (Orchestrating steps, reducing waiting, completing pre-work before review)

Example:

Issue: Repeated revisions → **Use Case:** Brand & compliance agent that checks tone, facts, and rules before human review.

Validate the opportunity

Before adding any use case to your roadmap, apply two quick filters:

- Does this matter to the business? Will it save time, reduce cost, or speed up delivery?
- Is this solvable with AI? Does it involve text, decisions, review, planning, analysis, or coordination?

If the answer isn't **"yes"** to both, deprioritize it.

Example of validated issues:

Process Step	Issue	Impact	Business critical?	AI-solvable?
Develop campaign	Resource-intensive (5 weeks)	Slows go-to-market	Yes	Yes
Review process	Repeated feedback loops	Adds 6 weeks	Yes	Yes
CRO review	Dependent on one role	Block sharing	Yes	Yes

Translate Issues into AI Use Cases

Once issues are validated, convert each one into a practical AI solution.

Issue	AI use cases
1-week manual research phase. Information must be gathered through internal interviews	**Research Collation Agent:** Automatically extracts project facts, timelines, deliverables, and outcomes from existing documents, project files, and internal systems to reduce interview time.
2–3 weeks waiting for stakeholder approval. Approval rarely happens on the first attempt	**Pre-Submission QA Agent:** Checks tone of voice, SEO, technical accuracy, image requirements, and compliance before the first draft is submitted to reduce approval cycles.
Multiple review loops Frequent iterations based on unstructured feedback	**Feedback Summarization and Revision Planner Agent:** Consolidates stakeholder comments into a clear, prioritized revision plan to reduce back-and-forth loops.
Inconsistent quality between creators and reviewers	**Brand and Compliance Validation Agent:** Applies compliance rules (tone, facts, SEO, AEO, etc.) consistently across drafts to ensure predictable quality before human review.
Dependency on a single reviewer or role Insights blocked when the reviewer is slow or unavailable	**Insights Generator Agent:** Summarizes draft content, recommends improvements, and generates insight summaries so progress continues even when a key reviewer is unavailable.

Case study process

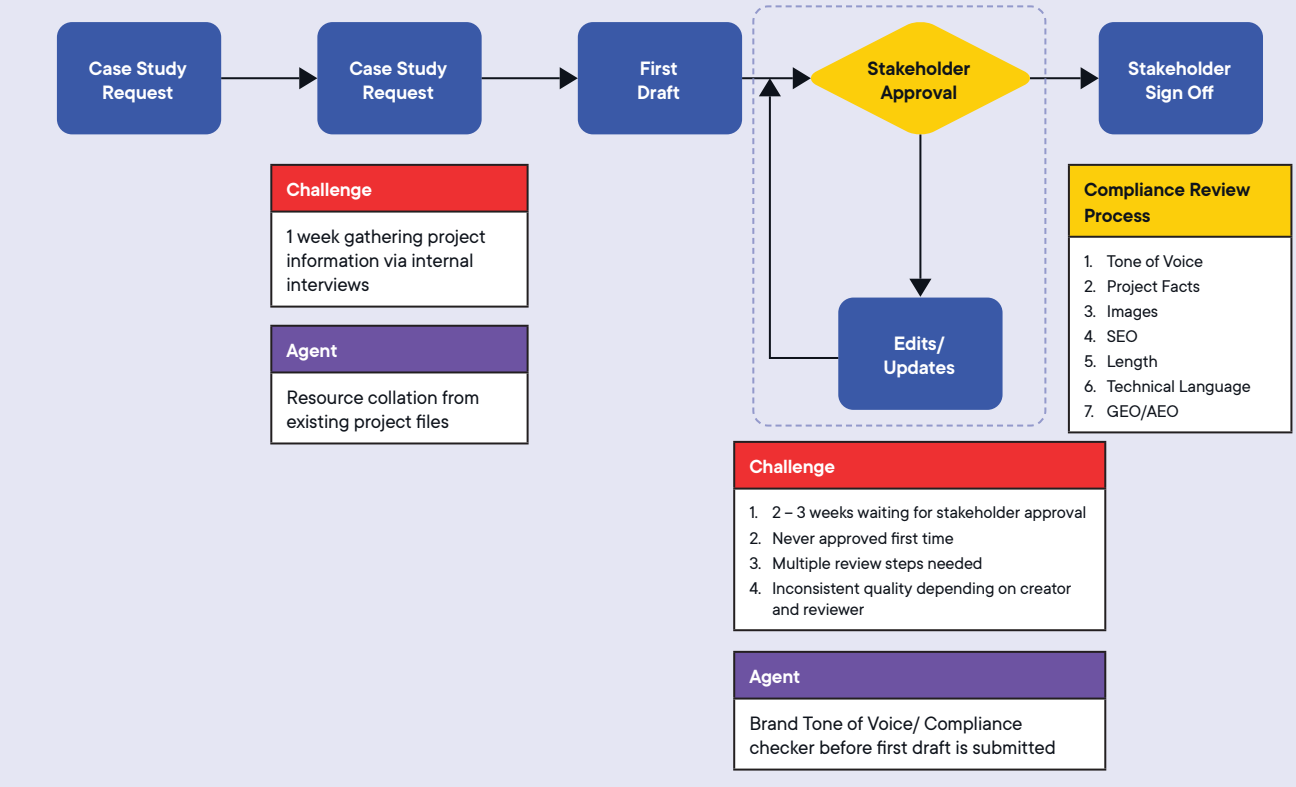

Prioritize what will deliver the most value

Once you've identified potential AI use cases, the next step is deciding which agents to build first. Not every opportunity carries the same value, and not every use case will move the business forward in a meaningful way. The most effective AI programs start by prioritizing the ideas that deliver impact quickly and build momentum.

The simplest way to do this is to evaluate each use case through two lenses.

Productivity impact (1–10)

This measures the internal efficiency gains an agent can create. Ask:

- How much time will this save?
- How much rework or manual effort will it remove?
- How many people or teams will benefit?
- Does it eliminate waiting, backlogs, or single-person dependencies?

High-scoring opportunities are those that reduce operational friction and give teams more capacity to focus on higher-value work.

Growth impact (1–10)

This measures how an agent accelerates outputs, improves performance, or strengthens customer-facing impact. Ask:

- Will this help us create more content or experiences?
- Will it speed up delivery or go-to-market timelines?
- Will it improve quality, consistency, or relevance?
- Will it unlock new capabilities we don't have today?

High-scoring opportunities are those that increase throughput, improve execution, and ultimately drive growth.

1. Result Presentation Agent
Transforms experiment summaries from Confluence into executive-ready PowerPoint decks automatically.

2. Data Cleansing Agent
Scans datasets for obvious errors or anomalies before human review, improving accuracy and speed.

3. JIRA Ticketing Agent
Creates JIRA tickets for custom analytics automatically, pre-filling key metrics and test details.

4. Metrics Collector Agent
Gathers and uploads custom experiment metrics directly to JIRA, removing the need for manual entry.

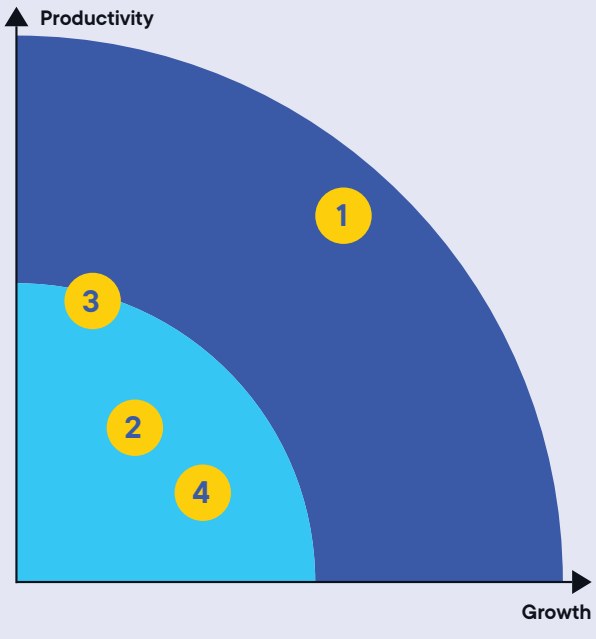

Building your AI pipeline

Once you've identified your highest-impact opportunities, the next step is to consider effort.
Effort helps you understand sequencing — not value — and is best captured in a simple scoring table.

Add a third score: **Effort (1–10)**

How much setup, data, rules, coordination, or integration is required?

Then calculate a final prioritization score:

Priority Score = Productivity + Growth – Effort
This becomes your AI pipeline.

Example AI Pipeline Scoring Table

AI use case	Productivity (1-10)	Growth (1-10)	Effort (1-10)	Priority score (P + G – E)	Notes
1. Result Presentation Agent	9	8	4	**13**	Highest impact. Automates slide creation, accelerates reporting, boosts throughout.
2. Data Cleansing Agent	6	5	5	**6**	Medium productivity and moderate growth impact. Useful but not transformational.
3. JIRA Ticketing Agent	5	3	3	**5**	Saves operational time; limited growth impact; low effort.
4. Metrics Collector Agent	4	2	3	**3**	Pure productivity helper; low strategic impact; small efficiency gain.

With impact and effort scored, you now have a clear, structured AI pipeline. This gives teams a shared view of which agents to build first, which to sequence over time, and which can be deprioritized. More importantly, it ensures every AI investment is tied to measurable productivity and growth outcomes rather than isolated experimentation.

3

Scaling AI across marketing

Scaling AI across marketing

Identifying the right AI use cases is an essential first step, but it is only the beginning. Many marketing teams are able to spot opportunities where AI can accelerate content creation, improve campaign performance, or remove manual effort. Yet most struggle to scale these wins beyond isolated teams or one-off pilots.

Organizations do not fail because AI lacks potential or because teams cannot identify strong use cases. They fail because they lack the operating structures that allow AI to be applied safely, consistently, and repeatedly across the business. Without clarity, ownership, or the right foundations, AI remains something individuals use, not something the organization depends on.

To unlock sustained value, marketing teams need a scalable operating model built on five essentials:

1. Establish clear governance
2. Define roles and responsibilities
3. Drive a culture of AI adoption
4. Measure the impact of AI initiatives
5. Select an agentic platform that can scale

Scaling AI is not about adding more tools. It is about building the organizational systems that allow agentic AI to become part of how marketing works every day.

Establishing scalable AI governance

Choosing the right governance model is one of the most important decisions you will make on your AI journey. Governance determines how quickly teams can experiment, how safely they can operate, how consistently your brand shows up, and ultimately how effectively AI scales across marketing.

AI governance sits across two layers of responsibility:

Enterprise-level governance, where company-wide AI councils or risk committees define policies for safety, compliance, data use, security, and ethical standards.

Marketing-level governance, where the function determines how AI should support content, campaigns, channels, and personalization in a way that aligns with brand standards and commercial goals.

With this dual context in mind, most marketing organizations adopt one of four governance models. Each offers advantages depending on team maturity, resourcing, and appetite for experimentation.

1. Centralized

A central AI team owns the strategy, builds agents, defines standards, and provides training. Individual teams request agents or support as needed. This ensures strong control and consistency, though it can slow down speed if the central team becomes a bottleneck.

2. Decentralized

Each team builds, runs, and governs its own agents. This maximizes speed and experimentation but requires strong underlying standards; otherwise quality and safety can drift.

3. Embedded

AI specialists sit directly within marketing teams, acting as local experts who build agents, enforce standards, and support adoption. This model blends speed with oversight and works well for larger or distributed teams.

4. Federated (Hybrid)

A central team manages high-impact automation, shared standards, and cross-functional capabilities, while individual teams build and run their own agents within clearly defined guardrails. This is the most scalable and balanced model for most organizations.

Centralized

A centralized AI enablement team is responsible for building agents, maintaining standards, and approving how AI is used across marketing.

Decentralized

Each team owns their own AI processes, build their own agents, and manage the quality of AI outputs. This drives rapid iteration and innovation.

Embedded

Teams own their own processes and agents, with AI specialists embedded to help create use cases, drive governance and grow adoption and AI quality.

Federated (Hybrid)

A central team sets AI standards and builds large scale automation, while individual teams can build their own agents for day-to-day, low-risk use cases.

Centralized

In a centralized governance model, all AI development, standards, and decision-making sit within a single specialist team, often Marketing Ops, Martech, or a dedicated AI Enablement function.

Marketing teams do not build their own agents. Instead, they request agents, workflows, or automations from this central group. The central body becomes the source of truth for safe, governed, and scalable AI usage across marketing. It defines what AI is allowed to do, where human oversight is required, and how AI supports content, campaigns, personalization, and optimization.

Core elements

- **Accountability:** A single point of ownership for all marketing AI decisions.
- **Processes:** Shared, organization-wide standards for building agents, approving tools, governing prompts, and evaluating risk.
- **Controls:** Strong central control over guardrails, brand safety rules, legal requirements, and monitoring.

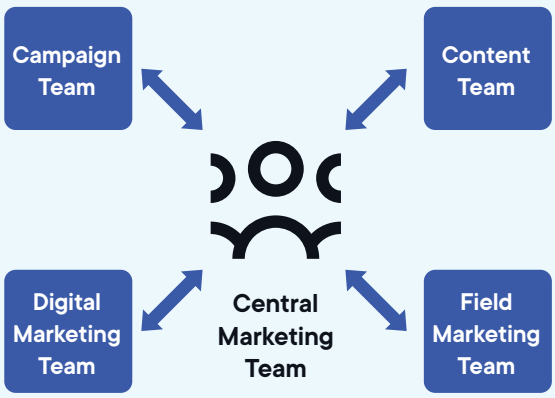

A centralized AI enablement team is responsible for building agents, maintaining standards, and approving how AI is used across marketing.

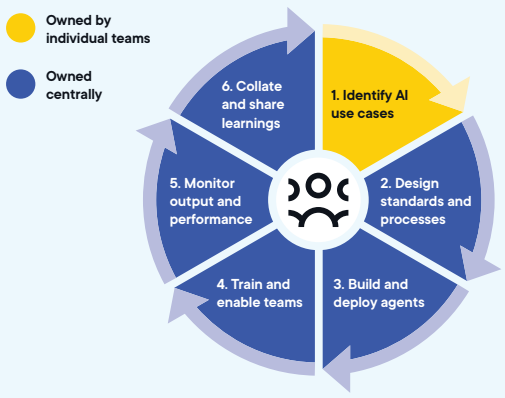

- Builds and maintains marketing AI agents
- Sets AI standards, guardrails, and approval processes
- Provides training, monitors quality, compliance

What it looks like in practice:

In a centralized model, the central AI team becomes the engine that drives consistency and control across marketing:

- Builds, maintains, and updates all marketing agents
- Owns monitoring, QA, and review for every AI workflow
- Maintains shared libraries of approved use cases, templates, and prompting standards
- Provides unified training, onboarding, and governance enforcement for all marketers

Benefits and Trade-offs

Benefits

- Consistent standards for brand, compliance, and data usage. Ensures every agent follows the same rules, reducing brand, legal, and reputational risk.
- Centralized AI agent creation and maintenance. Prevents agent bloat by reducing duplicate agents and lowering operational costs.
- Unified frameworks, templates, and guardrails across teams. Produces higher-quality outputs and more predictable performance across markets and channels.

Trade-offs

- Slower delivery and reduced agility. Approvals, reviews, and central queues can delay agent creation and constrain experimentation.
- One-size-fits-all rules may not meet every team's needs. Centralized standards can feel restrictive for local, regional, or channel-specific teams.
- Risk of bottlenecks and "shadow AI". Heavy reliance on one team increases bottleneck risk and may push marketers to use unsanctioned tools.

Centralized governance is most effective for large or multi-market organizations, particularly in regulated or risk-sensitive industries, and for teams early in their AI journey who need strong oversight and clear guardrails. It is less suited to small or highly decentralized teams, or to environments that rely on rapid experimentation and local flexibility, where a single central function can introduce friction or slow teams down.

Example: A global bank's marketing AI council builds and approves new AI voice-generation agents for use in regional customer-facing videos, ensuring regulatory governance and brand consistency across all markets.

Decentralized

In a decentralized governance model, individual marketing teams or market units independently decide how AI is used within their area. Each team selects tools, defines prompts, builds agents, and manages risk according to its own needs. There is no central authority overseeing how AI is applied across marketing.

This model prioritizes speed and autonomy. Teams can experiment freely, adapt AI to their specific workflows, and build agents that support their channel, market, or campaign objectives without waiting for central approval.

Core elements

- **Accountability:** Each marketing team (e.g., social, campaigns, field, content) or local market owns its own AI decisions with no central oversight.
- **Processes:** Teams create their own processes for agent creation, prompting rules, QA, and performance evaluation.
- **Controls:** Light governance; teams follow high-level guidance but manage quality and risk locally.

What it looks like in practice:

In a decentralized model, each marketing team becomes its own AI engine, shaping how agents are built, governed, and applied within their workflows:

- Teams build, customise, and maintain their own AI agents
- Each team defines its own prompting approaches and safe-use guidelines
- Use cases vary across teams based on maturity, resources, and priorities
- Quality control, risk assessment, and documentation occur locally rather than centrally

Benefits and Trade-offs

Benefits

- Speed and autonomy for marketing teams. Teams can move quickly, test ideas, and build agents without waiting for central approval.
- More relevant, tailored AI applications. Each function or market can build agents that align directly with its workflows, audiences, and goals.
- Encourages experimentation and innovation. Creative teams can rapidly iterate and uncover new use cases.

Trade-offs

- Inconsistency across teams. Agent quality, rules, documentation, and standards can drift without central oversight.
- Higher operational cost and duplication. Multiple teams may build similar agents, increasing maintenance overhead and reducing efficiency.
- Greater regulatory and reputational risk. Harder to enforce safe-use guidelines or meet regulatory requirements (e.g., EU AI Act) without unified controls.

Decentralized governance works best for fast-moving teams that prioritize speed and autonomy, local markets needing tailored workflows or language-specific agents, and environments with low-risk AI use cases. It is less suited to organizations that require strict brand or regulatory consistency, or teams lacking the governance discipline needed to manage AI responsibly at a local level.

Example: A regional social media team independently builds an AI caption-writing agent tailored to its local tone and cultural context, enabling faster output and more relevant content without relying on a central AI function.

Embedded

Embedded governance places AI specialists directly inside marketing teams, sitting alongside channel owners, creatives, campaign managers, and web or SEO leads. These specialists build agents, train teams, and oversee adoption, output quality, and compliance within the flow of work. This model blends speed with oversight by placing expertise as close as possible to the work itself.

Core elements

- **Accountability:** Individual teams own their AI decisions, supported by embedded AI specialists who guide design, usage, and safe operation.
- **Processes:** Embedded experts tailor workflows to team needs while following organization-wide standards for approvals, maintenance, and risk management.
- **Controls:** Core guardrails remain centralized, but embedded specialists enforce compliance and quality within their domain.

AI specialists are responsible for building agents, training teams and overseeing adoption and compliance

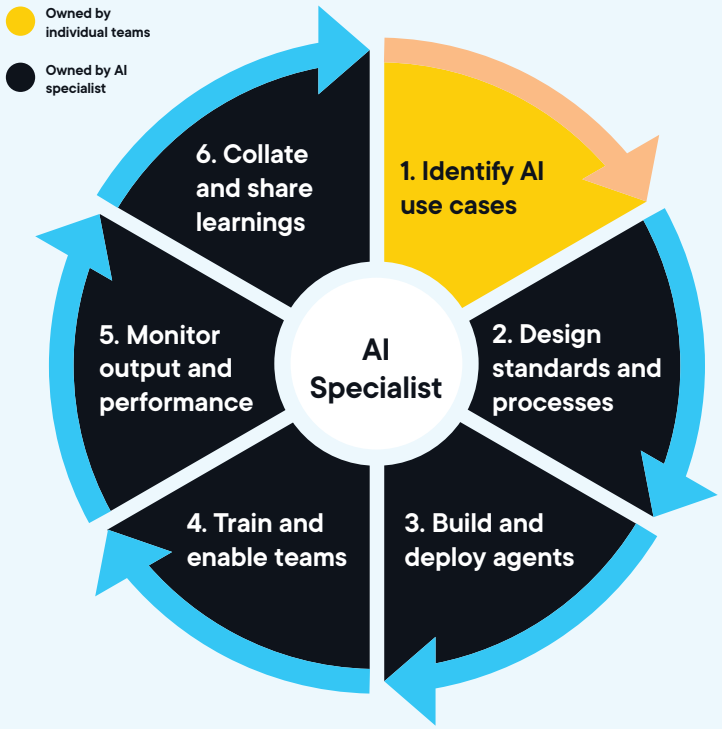

AI Specialist is available at each stage to advise on use cases, and provides hands-on support in creating agents and enabling teams

What it looks like in practice:

In an embedded model, AI expertise sits inside the team, providing real-time support, guidance, and governance as work happens:

- AI specialists build or co-build agents directly with their teams
- Prompts, instructions, and outputs are reviewed and improved in real time
- Governance guidance is delivered inside day-to-day workflows, not via external checkpoints
- Teams follow shared organizational rules but adapt them to the nuances of their channel or function
- Quality checks, compliance validation, and risk assessments happen at the moment of creation

Benefits and Trade-offs

Benefits

- Expertise embedded directly in the team. Improves quality, accelerates adoption, and reduces compliance risk.
- Governance integrated into daily workflows. Reduces friction and speeds up campaign delivery.
- Contextual application of standards. Ensures AI is used in ways that fit the realities of each channel or function.

Trade-offs

- Requires dedicated AI specialists for each team. Can be costly and difficult to scale across large organizations.
- Variation in interpretation of rules. Different embedded specialists may apply standards differently, creating inconsistency.
- Ongoing coordination and training needed. Without alignment, governance can drift over time.

Embedded governance is best for large or advanced marketing teams with the capacity to support embedded roles, and for functions with specialised or high-context needs such as paid media, CRM, or content. It is less suited to smaller teams, organizations that prioritize uniformity and strict central control, or environments where embedded expertise cannot be staffed consistently.

Example: A content team has an embedded “AI Content Quality Lead” who reviews prompts, guides safe usage, and ensures all AI-generated output meets brand and regulatory standards during daily work.

Hybrid/Federated

Hybrid or federated governance blends central guardrails with local autonomy. A central team sets policies, standards, approved tools, and high-impact agentic workflows, while individual marketing teams create and run their own agents for day-to-day execution. This is the most common model for global marketing organizations with mixed maturity levels, where both consistency and flexibility are required.

Core elements

- **Accountability:** Shared ownership between central and local teams. Central teams oversee high-impact or sensitive AI initiatives, while local teams manage everyday usage.
- **Processes:** High-risk or cross-functional workflows follow central review; low-risk agents and tasks are approved and run locally.
- **Controls:** Core standards are defined centrally, but teams adapt them to the needs of specific channels, regions, or use cases.

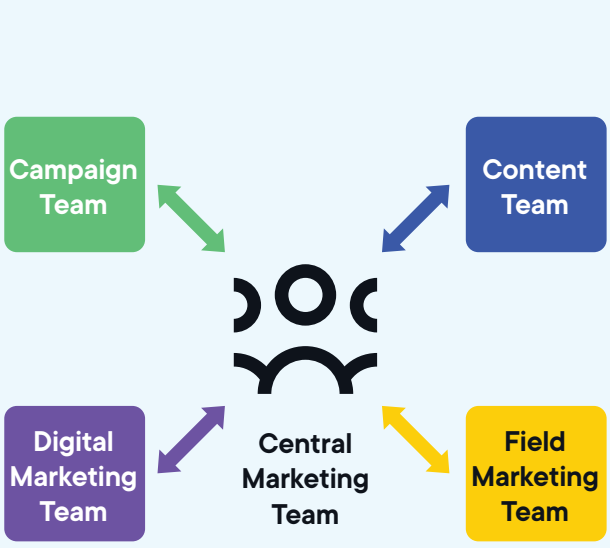

A central team defines guardrails and develops high-scale automations, while marketing teams retain the freedom to deploy and tailor AI to their specific workflows.

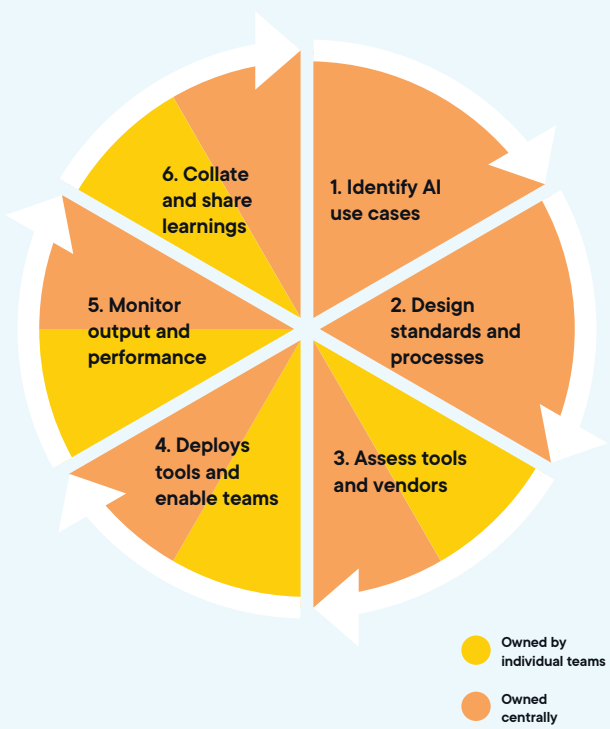

Individual teams benefit from centrally created agents, guardrails and governance, whilst being able to define their own use cases and approaches

What it looks like in practice:

In a hybrid model, central and local teams work together, combining shared guardrails with flexible execution:

- Central team builds major, cross-team, or high-impact agentic workflows
- Local teams build and customise agents for day-to-day or low-risk needs
- Shared AI platforms (CMP, CMS, DAM, Experimentation, CDP) act as the connective layer
- Clear escalation routes exist for regulated, high-risk, or sensitive AI use cases
- Continuous collaboration keeps central governance and local execution aligned

Benefits and Trade-offs

Benefits

- Balances consistency with flexibility. Brand safety is protected while allowing teams to move quickly.
- Shared agents with local customisation. Core capabilities scale efficiently across markets, with room for local nuance.
- Risk-adjusted approval model. High-risk work receives oversight; low-risk work flows faster, accelerating adoption.

Trade-offs

- Requires strong coordination between central and local teams. Without clear roles, the model can create complexity.
- Varied adoption and maturity levels across teams. Some teams move faster than others, creating uneven capability.
- Unclear escalation paths cause friction. Ambiguity slows decisions and can create confusion about ownership.

Hybrid governance is best suited to global or multi-market organizations with varying maturity levels, brands that need both consistency and agility, and teams that want to scale AI widely without losing local relevance. It is less suited to very small marketing functions, organizations that require strict uniform controls, or environments where teams expect full autonomy without central oversight.

Example: A global retailer's central Centre of Excellence builds a core content-creation workflow. Local markets then extend it with region-specific prompts, translations, asset variations, and approval steps tailored to local cultural and language nuances.

Choosing the right AI governance model

There is no single "correct" governance model. The right approach depends on your organization's objectives, its AI maturity, and the level of risk it needs to manage. Governance shapes how quickly teams can work, how consistently AI is applied, and how confidently marketing can scale agentic workflows across channels and markets.

Each model has distinct strengths and trade-offs. What matters most is selecting a structure that gives teams clarity, builds trust in AI, and establishes the foundations for safe, repeatable, and scalable adoption. With the right governance model in place, marketing can move beyond experimentation and unlock the full value of AI across every part of the function.

	Centralized	Decentralized	Embedded	Hybrid/Federated
Who makes decisions	One central Martech/Ops team	Each marketing team independently	Marketing teams supported by embedded AI specialists	Split model: central sets policy; teams execute
Speed vs Control	High control, slower speed	High speed, low control	Fast with built-in oversight	Balanced: speed with guardrails
Consistency of brand and compliance	Very high consistency	Low consistency, varies widely	High consistency applied contextually	High consistency with local flexibility
Risk exposure	Lowest risk (tight governance)	Highest risk (unvetted tools, uneven standards)	Low–moderate risk (real-time oversight)	Managed risk (clear escalation paths)
Best fit	Large, regulated, multi-team organizations	Fast-moving teams with low-risk use cases and high autonomy	Large functions using advanced AI and requiring specialist support	Global or multi-market organizations needing control and agility

Defining clear roles and responsibilities

Clear roles and responsibilities are essential when introducing AI into marketing workflows. Even when teams understand the technology, adoption becomes difficult if it is unclear who owns policies, who builds agents, who monitors accuracy, and who is responsible for day-to-day usage.

A typical set of roles includes:

- AI Owner – Sets policy, defines guardrails, governs usage
- Agent Developer – Builds, configures, and maintains agents
- Workflow Owner – Defines the processes agents support or automate
- AI Steward – Monitors output accuracy, compliance, and quality
- End Users – Use agents in daily work, provide feedback, escalate issues

A clear way to assign roles and responsibilities is leveraging a RACI model, which defines who is responsible, accountable, consulted and informed throughout the project.

Responsible: Those who do the work

Accountable: Those who are accountable to the work's completion

Consulted: Those who provide input and feedback

Informed: Those who are kept in the loop about progress

Example RACI Model

Activity	AI owner	Agent developer	Workflow owner	AI steward	End users
Set AI policies and guardrails	A	C	C	R	I
Approve or review new AI tools/agents	A	C	C	R	I
Build or configure AI agents	I	R/A	C	C	I
Define workflows the agent will automate	C	C	R/A	I	I
Monitor output accuracy, compliance and risks	I	C	I	R/A	C
Use agents in daily work and provide feedback	I	I	I	C	R
Escalate issues or unexpected behavior	I	I	C	A	R

In high-risk or customer-facing scenarios, roles and responsibilities should also be clarified at the workflow or campaign level. Every team member needs to understand both their part in the campaign and their specific responsibilities for the AI-driven components of the work.

Consider a scenario where AI is used to generate 12 market-specific banners for a Black Friday campaign. In this context, clarity becomes even more important. A typical set of campaign-level roles includes:

- AI Lead: Ensures all AI-generated content meets quality, accuracy, and brand standards
- Campaign Manager: Owns campaign strategy, messaging, and timelines
- Content Creator: Uses AI tools to generate and refine copy and visuals
- Web Manager: Builds and publishes landing pages and manages on-site assets

Campaign level RACI

Activity	AI lead	Campaign manager	Content creator	Web manager
Define campaign brief and key messages	C	A/R	C	I
Set AI prompts, guardrails and creative guidance	A/R	C	C	I
Generate and refine campaign copy using AI	C	C	A/R	I
Create AI-assisted visuals or asset variations	C	C	R	I
Build and publish landing pages, updates and assets	I	C	C	A/R
QA content using AI (accuracy, brand, compliance)	A/R	C	R	R

Clarifying campaign-level ownership becomes even more important as teams become more decentralized or as AI usage expands across markets and channels. Clear responsibilities ensure accountability, protect brand integrity, and allow teams to adopt AI with confidence.

Enabling teams to work effectively with AI

AI impact depends on how teams use it. As highlighted earlier, adoption is high, but meaningful value is still inconsistent. The gap is not caused by technology, it is caused by the way teams work with it. Real transformation requires deliberate enablement, cultural change, and a shift in how people understand their role alongside AI.

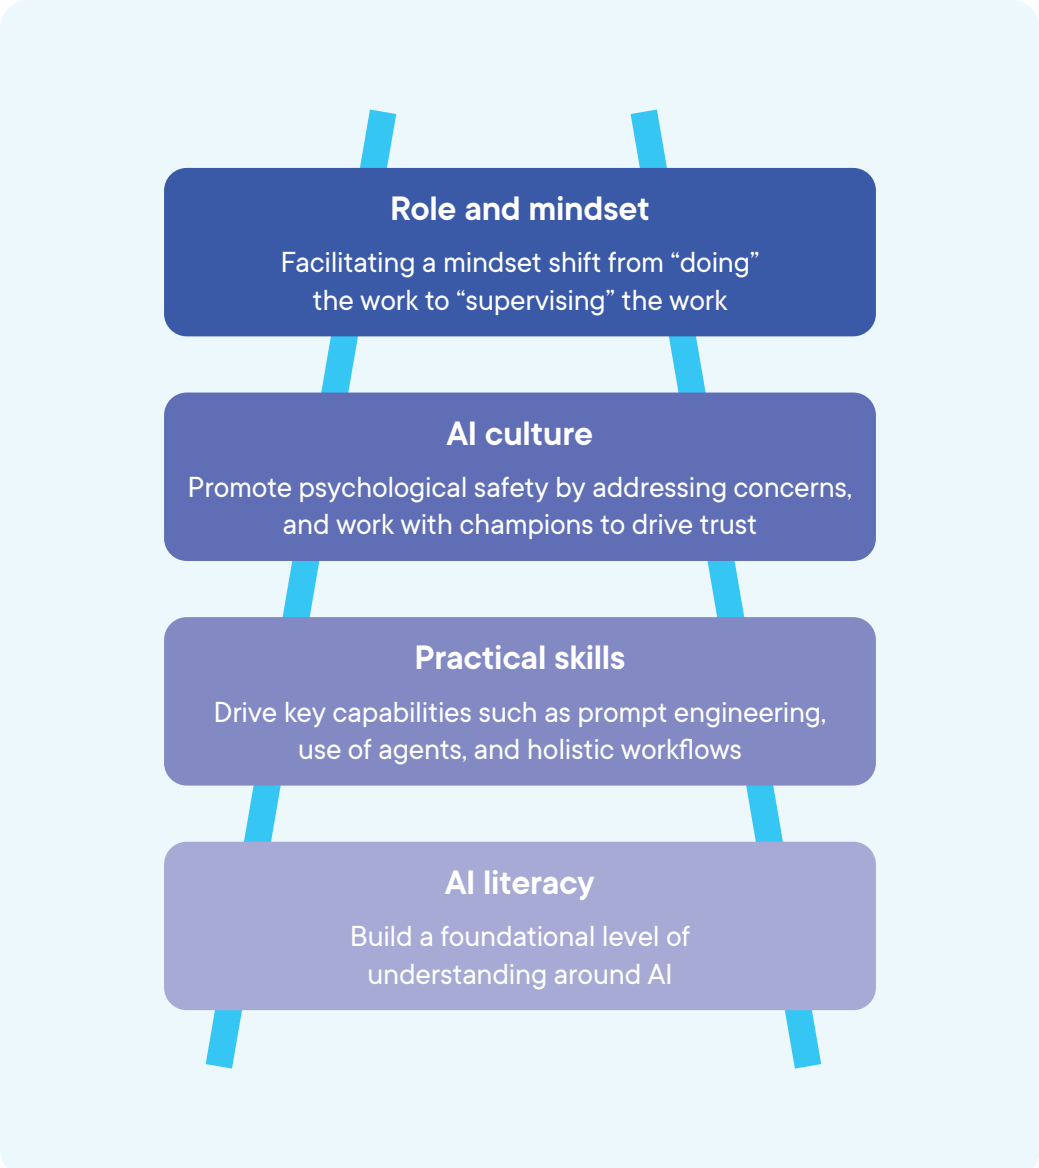

Closing the AI literacy gap

AI literacy means understanding what AI can and cannot do, how it works, and how to apply it effectively to real workflows. It goes beyond learning how to prompt and toward developing the critical thinking required to supervise, assess, and guide AI systems. Most organizations still face a significant gap between the availability of AI tools and the skills required to use them.

66% of CMOs report say their teams lack AI skills (BCG).

Closing this gap requires a structured, deliberate approach:

1. Understand current attitudes and capabilities

Surveys, workshops, and sentiment assessments help reveal how comfortable teams feel with AI, how often they use it, and what concerns or misconceptions may be limiting adoption. This creates the foundation for tailored enablement, and a shared understanding of what AI can—and cannot—do.

2. Prioritize gaps by risk, impact, and ease of resolving

Not all literacy gaps are equal. Prioritize the ones that are:

- Widely felt across teams
- High-risk for the business or brand
- Likely to make a measurable difference in speed, quality, or performance

Focus first on the gaps that are common, risky, and impactful, then expand as maturity grows.

3. Deliver targeted interventions in multiple formats

People learn in different ways. Providing practical resources across formats—short videos, written guides, live workshops, office hours, or audio walk-throughs—helps learning stick. To encourage participation and adoption, organizations often use:

- Required training or accreditation that must be completed by a set date
- Temporary restrictions on advanced AI usage until training is completed
- Gamified incentives and recognition based on training participation

These interventions help teams build a foundational level of comfort and literacy. But AI evolves rapidly, and literacy is not a one-off effort. Ongoing upskilling, refresher training, and updated playbooks are essential to ensure teams can adapt confidently as capabilities, tools, and expectations continue to shift.

Practical skills

Foundational literacy is essential, but it is not enough. To keep pace with the evolution of AI, teams need practical, hands-on skills that translate directly into their day-to-day work. The most successful enablement programmes focus on three areas: prompting skills, agent usage, and workflow design. Each requires structured training, repeatable tools, and ongoing reinforcement.

Prompting

Prompt engineering is the single most important practical skill for unlocking value from AI. Strong prompts produce accurate, brand-safe, contextually relevant outputs. Consistent prompting drives consistent results.

To embed prompting excellence across the organization:

1. Introduce the CLEAR and RACE frameworks

Short, practical sessions should teach teams how to use these frameworks to structure prompts, refine outputs, and generate high-quality work consistently.

2. Create prompt libraries for everyday tasks

Develop shared, role-specific libraries for emails, social posts, product descriptions, campaign briefs, QA steps, asset descriptions, and more. Teams should copy, adapt, and reuse—rather than start from scratch.

3. Run "prompting clinics"

Weekly or bi-weekly clinics give teams the opportunity to bring real tasks, test multiple prompt variations, and receive live coaching. This accelerates learning and builds confidence through practice.

4. Encourage a prompt-and-review mindset

Teams should run variations, compare results, and critically review outputs before publishing. Reinforcing this habit reduces risk and improves output quality over time.

Train teams to use AI agents

Agents become powerful once teams understand when to use them, how they behave, and what their limitations are.

1. Show teams where agents fit into existing workflows

Use real examples, such as campaign QA agents or content drafting agents to show value in the context of their daily work. Many teams do not realise how many processes are already repetitive or rules-based until they map them.

2. Clarify when agents should—and should not—be used

Agents excel at repetitive, rules-based tasks that follow clear logic. They are less effective for open-ended creative exploration or strategic decision-making.

3. Standardise how teams review agent outputs

Reviews should focus on:

- Accuracy
  (factual correctness, up-to-date information, alignment with the request)
- Brand (tone, style, alignment with brand guidelines)
- Risk (bias, hallucinations, regulatory implications)

Make escalation paths clear.

Teams must know:

- When to flag unexpected behavior
- Who to notify
- How to track issues
- How updates or fixes are communicated

Clear escalation guidelines reduce risk, increase trust, and keep workflows running smoothly.

Driving AI culture

Even with the right tools, governance and training in place, culture is often the biggest determinant of whether AI succeeds or stalls. Because AI is new, fast-moving and sometimes ambiguous, teams naturally develop different levels of comfort, confidence and caution. These mindsets influence how quickly people adopt new workflows, how willing they are to experiment, and how consistently they use AI in their day-to-day work.

Understanding these mindsets and meeting people where they are is essential for building an environment where AI is used effectively, safely and confidently.

The five personas shaping AI adoption

Even with the right tools, governance and training in place, culture is often the biggest determinant of whether AI succeeds or stalls. Because AI is new, fast-moving and sometimes ambiguous, teams naturally develop different levels of comfort, confidence and caution. These mindsets influence how quickly people adopt new workflows, how willing they are to experiment, and how consistently they use AI in their day-to-day work.

Understanding these mindsets and meeting people where they are is essential for building an environment where AI is used effectively, safely and confidently.

Across marketing, five broad personas consistently emerge during AI rollout. Each requires a different approach:

Explorers

Curious, enthusiastic, and eager to test new tools.

How to support them:

- Provide structured onboarding and safe-to-fail spaces
- Offer sandbox environments and gamified challenges
- Celebrate learning and experimentation

Pragmatists

Outcome-focused and motivated by clear value.

How to support them:

- Show tangible time savings and performance improvements
- Provide role-specific examples and ready-to-use templates
- Link AI tasks to KPIs and business outcomes

Sceptics

Cautious, concerned about quality, accuracy or job impact.

How to support them:

- Communicate transparently about guardrails and governance
- Offer hands-on demonstrations that build familiarity
- Involve them early in pilot tests to give them ownership and confidence

Guardians

Risk-minded, brand-minded, or compliance-focused individuals.

How to support them:

- Bring them into policy design, review processes and QA steps
- Position them as co-owners of responsible AI
- Provide visibility into how risks are monitored and mitigated

AI Champions

Highly engaged early adopters who naturally advocate for AI.

They:

- Proactively propose new ideas
- Give detailed feedback
- Invest in continuous learning

How to empower them:

- Involve them in steering groups or working councils
- Give them ownership of team-level AI initiatives
- Ask them to lead internal demos, prompting clinics and playbook updates

Champions often become the cultural accelerators—helping others adopt new behaviors far more effectively than top-down mandates alone.

Why these personas matter

Strong AI culture is built from the ground up. When these personas work together, they create a self-reinforcing system:

- Champions push adoption
- Guardians protect standards and reduce risk
- Explorers drive creative experimentation
- Pragmatists deliver measurable results
- Sceptics keep the organization grounded and honest

With the right interventions, even sceptics become confident users. The goal is not to convert everyone instantly, but to create psychological safety and build a culture where teams feel empowered to learn, iterate and grow with AI without fear of failure.

A strong AI culture turns individual enthusiasm into organizational momentum. It ensures AI becomes part of how marketing works—not just a tool people occasionally test.

Shifting roles and mindsets

As AI becomes deeply embedded in marketing workflows, the nature of day-to-day work fundamentally changes. Teams no longer spend the bulk of their time manually producing assets, building pages or completing repetitive tasks. Instead, their role shifts toward supervising, refining and directing the work that AI generates.

This transition is not about replacing expertise, it is about elevating it. Human judgment, creativity and context become even more important as AI takes on production-heavy steps.

What the shift looks like in practice

Content Creators: from "writing" to "editing"

Instead of drafting everything from scratch, creators review AI-generated copy for accuracy, tone, clarity and intent. Their value moves from typing every word to shaping messaging, elevating quality and ensuring the output connects to the brief.

Web Managers: from "building" to "overseeing"

AI can assemble layouts, produce variations and structure modules. Managers validate outputs, check rendering, ensure accessibility and maintain brand consistency across experiences.

Campaign Managers: from "executing tasks" to "orchestrating outcomes"

AI accelerates asset assembly, segmentation, optimisation and reporting. This frees managers to focus on audience strategy, creative direction, and performance refinement across channels.

Capabilities needed for the new model

To operate confidently and safely in AI-driven workflows, teams need a clear set of teachable skills:

- **Critical evaluation of AI output.** Ability to assess accuracy, tone, brand alignment, compliance and factual integrity.
- **Comfort with automation**. Knowing when to trust agents with repetitive steps—and when human oversight is essential.
- **Brand and governance awareness.** Understanding the rules, constraints and standards that AI must work within.
- **Human-in-the-loop judgment.** Recognising when to escalate, intervene or apply human creativity, empathy or contextual reasoning.

A successful mindset shift does not happen overnight. But as teams gain confidence in supervising AI rather than doing every step manually, productivity increases, quality becomes more consistent, and marketing can operate at a scale that was previously impossible.

Measuring the impact of AI initiatives

Measuring the value of AI is essential for turning early experiments into sustained, scalable results. Without clear measurement, AI efforts often remain anecdotal, fragmented and difficult to justify. Teams may feel busier, outputs may increase and workflows may change, but the organization cannot quantify what has improved, what still needs attention or where to invest next. Strong measurement creates the visibility needed to secure continued support, prioritize high-value use cases and build trust across marketing, IT and leadership.

AI delivers value through two core levers: productivity and impact.

1. Productivity captures how AI helps teams do the same work faster, at higher volume, or with fewer resources.
2. Impact captures how AI improves marketing performance, customer outcomes and revenue.

Together, these levers form the basis of an iterative measurement loop that keeps teams focused on tangible outcomes. Effective AI programmes measure not just activity, but acceleration, quality, performance and growth. A clear measurement approach reinforces accountability, demonstrates ROI and ensures AI contributes meaningful value across the marketing organization, rather than remaining a series of isolated experiments.

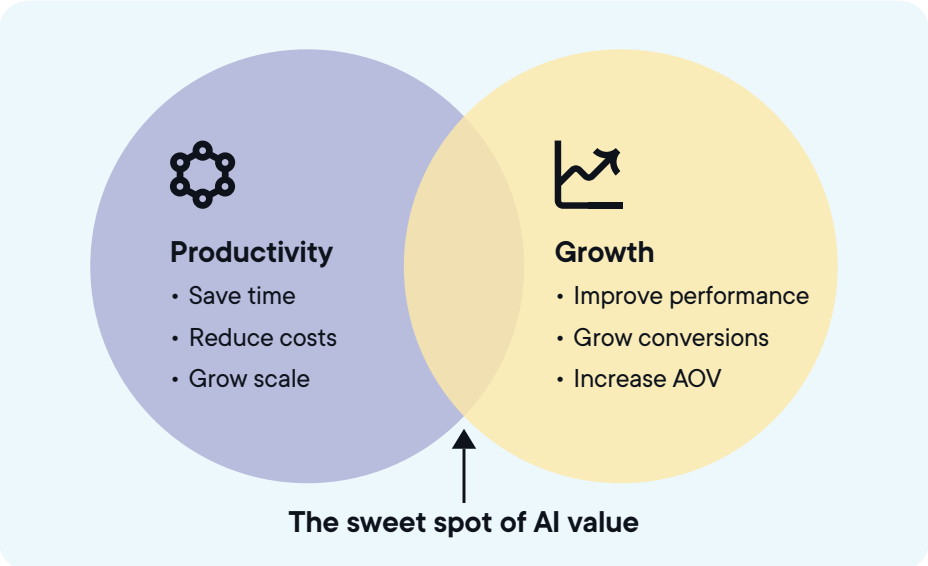

Measurement feedback loop

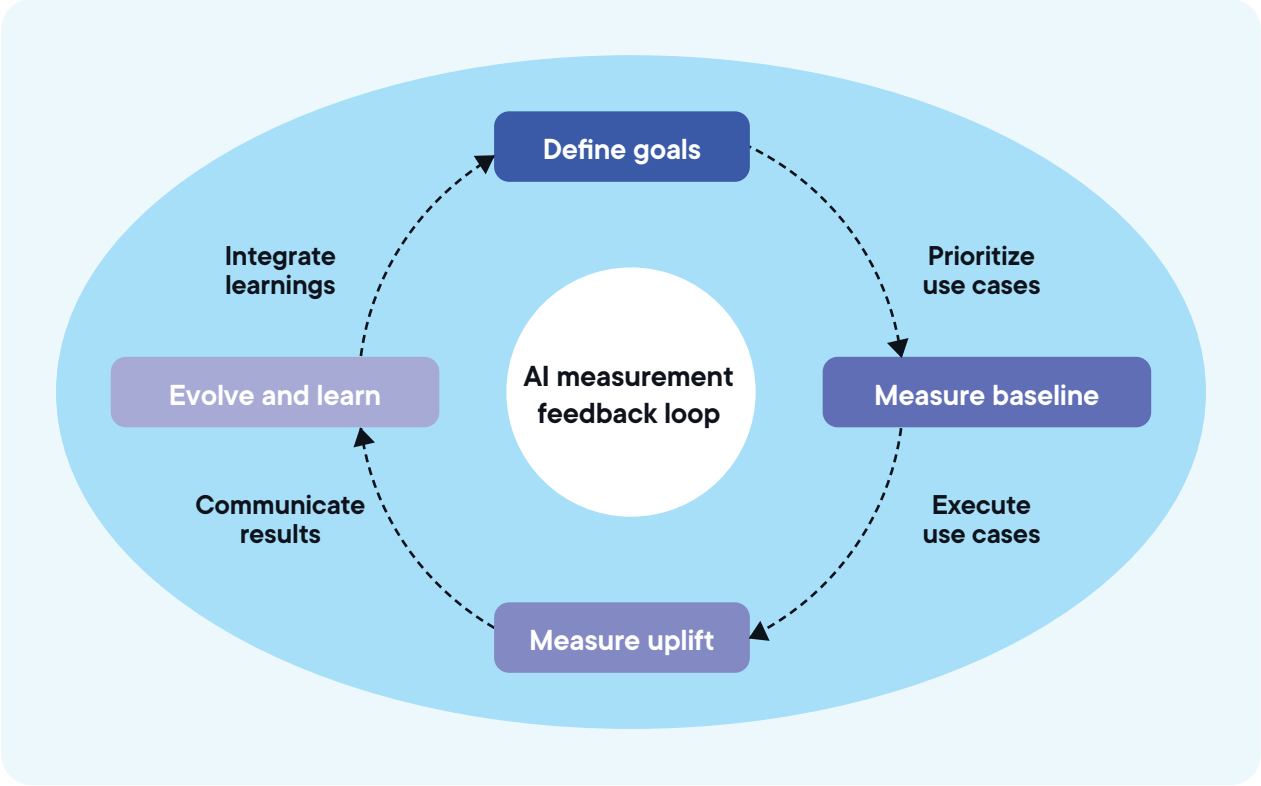

Define your initial goals

Every AI initiative must begin with a clear definition of success. Without explicit goals, teams cannot evaluate whether an agent is delivering value or simply creating activity. Defining goals upfront ensures the organization aligns on what "good" looks like and which outcomes matter most.

The key is to select metrics that are **measurable, impactful and directly attributable to AI.**

Productivity Metrics	Growth Metrics
Time spent per task	Conversion rate performance
Campaign cycle duration	Revenue per visitor (RPV)
Manual vs automated steps	Average order value (AOV)
QA error rate	Experiment success rate
Cost per asset	Organic traffic volume
Assets/campaigns produced	Customer engagement rate

Measure baseline

Before launching any AI initiative, you need a clear understanding of how your marketing organization performs today. This baseline becomes the reference point against which future improvements are measured. The goal is not to capture everything, but to gather enough meaningful data to attribute impact accurately.

Productivity baseline

Much of the productivity baseline can be derived from the process mapping work conducted earlier. This allows you to quantify how work currently happens and identify concrete, measurable starting points.

Common productivity metrics include:

- Time spent on key tasks over several cycles (content drafting, QA, asset creation, campaign briefing)
- Number of campaigns, assets or tasks completed in the past 6–12 months
- Number of rework steps, approval rounds and delays
- Resource requirements for each task (people involved, hand-offs, tools used)

These baselines establish the speed, effort and operational efficiency of your current workflows.

Performance baseline

Your performance or growth baseline typically comes from your analytics, experimentation or optimisation platforms.

Common performance metrics include:

- Conversion rate, revenue per visitor, average order value, engagement rate and experiment success rate
- Channel, audience or market-level performance segmentation
- Campaign or asset-type benchmarks (e.g., emails vs landing pages)

This baseline makes it possible to isolate uplift once AI starts influencing creative quality, personalisation or optimisation cycles.

Important considerations when capturing baselines

1. Store all metrics in a single source of truth so they remain trackable and consistent.
2. Include clear definitions and contextual metadata (time period, seasonality, campaign type).
3. Document any calculations or adjustments to ensure future comparisons remain valid.
4. Tie baselines to specific workflows or asset types whenever possible to isolate AI impact more accurately.

Example: A travel company is launching an agentic workflow that generates destination-specific content and asset variations based on regional popularity.

KPI Type	Metric	Baseline (Pre-AI)
Productivity	Destination campaigns/month	**8**
Growth	Conversion rate	**3.4%**
	AOV	**$500**

Once your baseline is established and AI initiatives are live, the next step is to measure how outcomes change.

Measure uplift

Once AI is live, track the same KPIs used in your baseline and compare them directly to quantify uplift. This creates a clear, objective view of whether the agent is delivering value and where further optimisation may be needed.

KPI Type	Metric	Baseline (Pre-AI)	Post-AI performance	Incremental improvement	% increase
Productivity	Destination campaigns/month	**8**	**12**	**+4**	**50%**
Growth	Conversion rate	**3.4%**	**3.7%**	**+0.3 pp**	**8%**
	AOV	**$500**	**$520**	**+$20**	**4%**

Translate results into business terms

To communicate impact effectively, express uplift in the metrics your stakeholders care about:

- Hours saved → Full-time equivalent (FTE) capacity
- Error reduction → Lower production costs or faster time-to-market
- Conversion lift → Incremental revenue
- Experiment win rate → Higher engagement or revenue contribution

AI-driven uplift becomes even more compelling when expressed in revenue terms rather than only percentage improvements.

Example revenue impact:

Traffic remains flat at 88,000 visits/month.

Baseline revenue:
3.4% × $500 × 88,000 = $1,500,000

Post-AI revenue:
4% × $520 × 88,000 = $1,830,400

Incremental revenue:
$330,400 per month (+22%)

Attribution principles

To ensure the uplift is genuinely driven by AI and not external factors, apply responsible attribution practices wherever possible.

- Run A/B comparisons (AI-produced assets vs manually produced assets)
- Control for seasonality by comparing to equivalent historical periods
- Document any parallel changes (new channels, staffing changes, product launches)
- Use consistent measurement windows to avoid misleading spikes

Good attribution protects credibility and strengthens confidence in AI investment decisions.

Tailor communication to stakeholders

Different audiences care about different outcomes. Tailor your measurement story so each stakeholder sees value through their own lens.

Stakeholder	Primary question	Focus metrics
CEO	How does this affect growth and profitability?	Revenue uplift, cost savings, productivity gains
CMO	How does this improve pipeline and conversions?	Conversion rates, campaign volume, content velocity
CTO	How does this change team workload and operations?	Time saved, reduction in technical effort, implementation considerations

Evolve and learn

Measurement is never static. As your AI models, workflows, content and team behaviors mature, your metrics need to evolve with them. Each cycle of work should prompt a review: What changed? Why did it change? What does this tell us about how to refine prompts, instructions, workflows or agent behavior?

Baselines should be revisited whenever processes materially shift, and KPIs should adjust as business priorities change. The goal is to create a continuous improvement loop where teams learn from results, refine their approach and steadily increase both productivity and impact.

With clear measurement foundations in place, the next step is selecting an AI platform capable of delivering these gains at scale. That begins with understanding the core capabilities a modern marketing organization needs in its AI foundation.

4

Why marketing needs an AI platform

Why marketing needs an AI platform

Most organizations start their AI journey with standalone tools or copilots. These are useful for early experimentation, but they quickly expose the same limitations: inconsistent outputs, no shared memory, no integration with workflows, and minimal governance. Marketing teams can draft faster, but they can't operate faster. The result is fragmented processes, siloed usage, and value that never scales beyond individual productivity gains.

This is why organizations ultimately move toward agentic AI platforms—systems that unify context, governance, data and execution. Selecting this platform is not about adding another point solution to your stack. It is about choosing the intelligent operating layer that will sit across your marketing workflows, content, and decision-making for years to come. The right platform becomes the engine that powers safe, consistent and repeatable automation. The wrong one introduces fragmentation, rework and operational risk.

Any platform you choose must protect you from the three most common AI failure modes:

- Using the wrong data (outdated content, inconsistent product facts, missing context)
- Producing unsafe or inconsistent outputs (no brand guardrails, no compliance checks, no tone control)
- Breaking multi-step workflows (no sequence logic, missing handoffs, no monitoring or escalation)

A modern AI platform must eliminate these risks by grounding every agent in trusted organizational knowledge, enforcing governance by design, and integrating deeply into the workflows that already run your marketing engine.

Aligning on the right AI platform architecture

Before choosing an AI platform, it helps to understand how AI will sit within your broader marketing ecosystem. Every marketing function operates differently. Some teams work as a single unit, while others run as specialised groups such as Brand, Content, CRM or Digital. The platform you choose should align with how your organization already works. Most teams typically align to one of three platform architectures:

1. Unified marketing platform

All marketing teams use one shared AI platform with common tools, workflows and governance.

Strengths: High consistency, one source of truth, simple integration.

Considerations: Less flexibility for teams that need specialised tools or deeper customisation.

2. Hybrid marketing platform

A central platform provides core AI capabilities such as content creation, optimisation and governance. Individual teams can extend it with their own tools, data sources or agents.

Strengths: Good balance between standardisation and flexibility.

Considerations: Requires coordination to keep teams aligned and avoid fragmentation.

3. Modular marketing platform (each team builds its own stack)

Each team assembles its own stack and tools. The AI platform integrates across multiple systems such as different CMSs, DAMs, CRMs or optimisation tools.

Strengths: High autonomy and functional depth.

Considerations: More integration work, greater governance needs and higher long-term complexity.

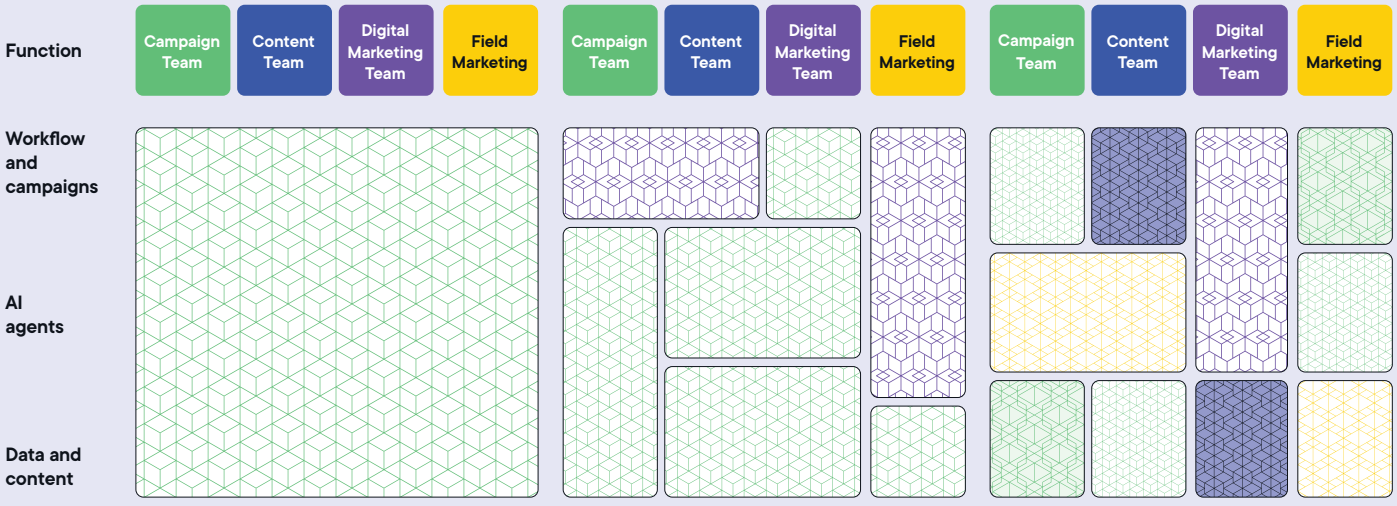

A single vendor powers all layers across marketing teams. This enables fast deployment and strong consistency, but limits customisation and team-specific needs.

A central platform provides the core capabilities, while teams extend with their own tools and data. Balances standardisation with flexibility across markets and channels.

Each marketing team assembles its own stack. Maximises specialisation but increases integration overhead, governance complexity, and long-term maintenance costs.

How to choose the right architecture

The best way to decide which architecture fits your organization is to anchor the decision in how your marketing teams actually operate today. These three questions provide a clear way to assess the right approach.

Governance and team structure: Are teams centrally coordinated or independently run?

- Centralized oversight often aligns well with unified or hybrid architectures
- More independent or regional teams often fit better with modular setups

Speed and ease of deployment: How quickly does AI need to be operational?

- Unified platforms deliver the fastest rollout
- Modular architectures take longer due to integrations and variability

Flexibility needs: How many tools, markets or systems will AI need to connect to?

- Hybrid architectures often strike the best balance between flexibility and simplicity

Choosing the right architecture creates the foundation for how AI will operate across your marketing function. With this in place, the next step is defining the capabilities your AI platform needs to support agentic workflows and deliver measurable value.

Evaluating AI platform capabilities

When evaluating AI platforms, it's easy to start with long feature lists or vendor comparisons. But the more effective approach is to focus on the capabilities that truly matter for how your teams work today and how agentic workflows will operate across your organization tomorrow.

A strong platform must do two things at once:

- Support the way your teams already plan, create, publish and optimise marketing work, and
- Provide the intelligence, governance and automation required to scale AI safely and reliably.

The most practical way to assess this is through diagnostic questions across five capability areas:

Category	Diagnostic question
1. Data Access and Context Integration	Does the platform reliably access and use the right data from your marketing tools so outputs are grounded in brand, product and performance context?
2. Governance, Ownership and Control	Does it support the permissions, auditability, controls and governance model your organization requires?
3. Workflow Integration and Collaboration	Can teams use AI directly within their existing tools and approval flows, without creating extra steps, new processes or operational friction?
4. Automation and Multi-Step Execution	Can it automate multi-step tasks using reusable agents, structured workflows and human checkpoints — and scale these across teams and markets?
5. Output Quality and Customisation	Can it consistently produce accurate, on-brand, channel-specific outputs and adapt to your tone, formats, regulatory rules and localisation needs?

This framing helps teams avoid chasing generic "AI features" and instead evaluate platforms on the capabilities that enable real, repeatable, enterprise-level impact.

Common AI constraints

Even with clear evaluation criteria, every marketing organization works within real constraints that shape what is possible and how quickly AI can scale. Understanding these constraints early helps teams avoid unrealistic expectations and choose a platform that fits both their goals and their environment.

Here are the five constraint themes that appear most often:

1. Legacy systems and integration limitations

Older CMS, CRM or DAM systems, limited APIs or long upgrade cycles can slow down platform adoption and restrict how deeply AI can integrate.

2. Fragmented or inconsistent data foundations

Multiple sources of truth, inconsistent taxonomies, missing metadata or unstructured content make it harder for AI to access accurate, usable information.

3. Workflow and process variability

Teams and regions often work in different ways, which can make automation difficult to standardise or scale across the organization.

4. Resource and capability gaps

Limited engineering support, scarce AI expertise or uneven levels of AI readiness can influence how quickly teams can adopt and maintain agentic workflows.

5. Governance and regulatory complexity

Evolving AI policies, regional regulations and unclear ownership models can slow implementation and increase the need for well-defined guardrails.

Balancing your evaluation criteria with your organizational constraints builds a realistic foundation for selecting a platform that fits your systems, your structure and your ambitions.

Balancing platform fit and system constraints

With your evaluation criteria and organizational constraints defined, the question becomes: which platforms can truly operate at scale? The answer lies in finding technologies that score highly on both platform fit and system compatibility — the capabilities that determine whether AI will enhance your workflows or become another operational burden.

Platforms in the top-right of the matrix are the ones capable of supporting agentic workflows at scale. They share five defining characteristics:

1. Integrated into real marketing workflows

AI only creates value when it fits into how work actually happens. The platform must connect directly to the systems that power your operations — CMS, CMP, experimentation, analytics and approval flows — so agents operate inside existing processes, not as disconnected add-ons.

2. Combines LLM intelligence with organizational context

Generic models cannot produce brand-safe, accurate or channel-ready outputs on their own. The platform needs access to your organization's knowledge: brand tone, product data, content libraries, experiment results, governance rules and localisation requirements.

3. Enforces governance and consistency by default

Reliability and safety cannot depend on individuals. The platform should automatically apply tone, terminology, disclaimers, permissions and approval logic so that safe, consistent output is the default.

4. Supports multi-step, agent-driven execution

Marketing work is a sequence — briefing, drafting, refinement, QA, localisation, publishing. The platform must support multi-step workflows where agents follow structured logic, hand off to humans when needed and deliver consistent results.

5. Scales through automation, not extra manual effort

As adoption grows, effort should decrease. The platform should offer reusable agents, shared instructions and automated workflows that compound efficiency over time.

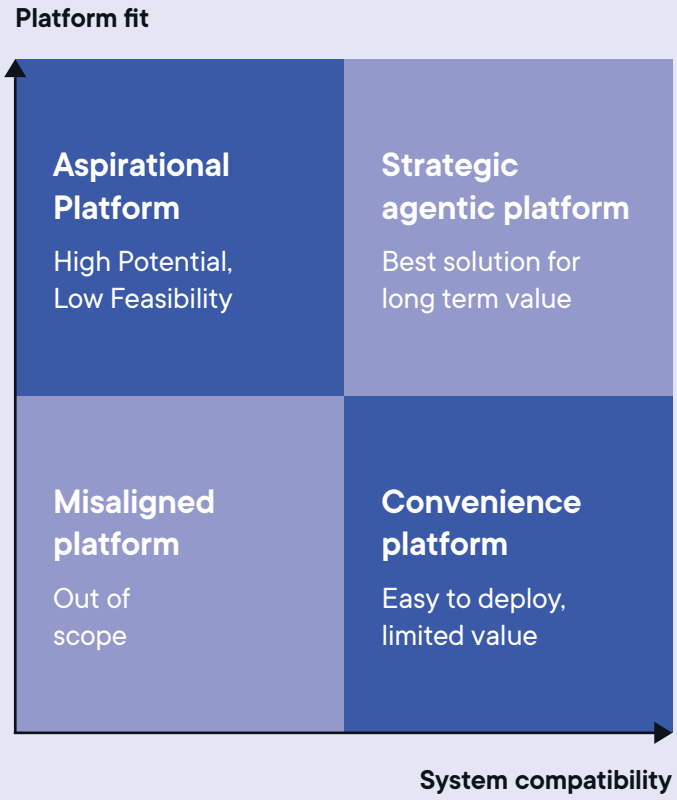

Platform Fit:
How well an AI platform aligns with your marketing organization's needs, including workflows, governance, data requirements and desired outcomes.

System Compatibility:
System compatibility is the platform's ability to work with your existing systems (and their constraints) today and evolve seamlessly as your architecture changes tomorrow.

A platform that delivers on these characteristics strengthens your entire marketing ecosystem by enhancing the way work already gets done. The stronger the alignment with your systems, data and governance, the faster adoption moves and the more value compounds.

5

Building a scalable agentic operating model

Building a scalable agentic operating model

Marketing is entering a new operating era — one where AI is no longer a tool used by individuals, but a coordinated system that shapes how work flows across the entire function. Early experimentation helps teams draft faster, but true transformation happens when AI becomes part of how marketing operates: when agents work together, share organizational knowledge, follow governance rules and orchestrate full end-to-end workflows.

This is where meaningful change occurs.

Not when teams generate content faster, but when they operate differently. A scalable agentic operating model depends on five foundations:

- Clear governance that defines how AI is used, who owns it and how risk is managed
- Defined roles and responsibilities that give teams clarity and confidence
- A strong culture of adoption built on literacy, practical skills and psychological safety
- Measurement frameworks that track productivity and performance outcomes
- An AI platform capable of unifying context, workflows, governance and automation

When these foundations are in place, AI becomes embedded in how marketing works; not an add-on, but an accelerator. Teams spend less time coordinating and more time applying strategy, creativity and judgment. Workflows become consistent. Quality becomes predictable. Output scales without adding operational drag.

This is the operating model leading organizations are moving toward:

Multiple agents, connected by a single platform, orchestrating work with speed and precision.

To support this model, a modern AI foundation must:

- Connect directly into the tools and workflows marketers use every day
- Ground every action in accurate brand, product and performance context
- Enforce governance automatically
- Support multi-step orchestration with appropriate human oversight
- Scale through automation rather than additional manual effort

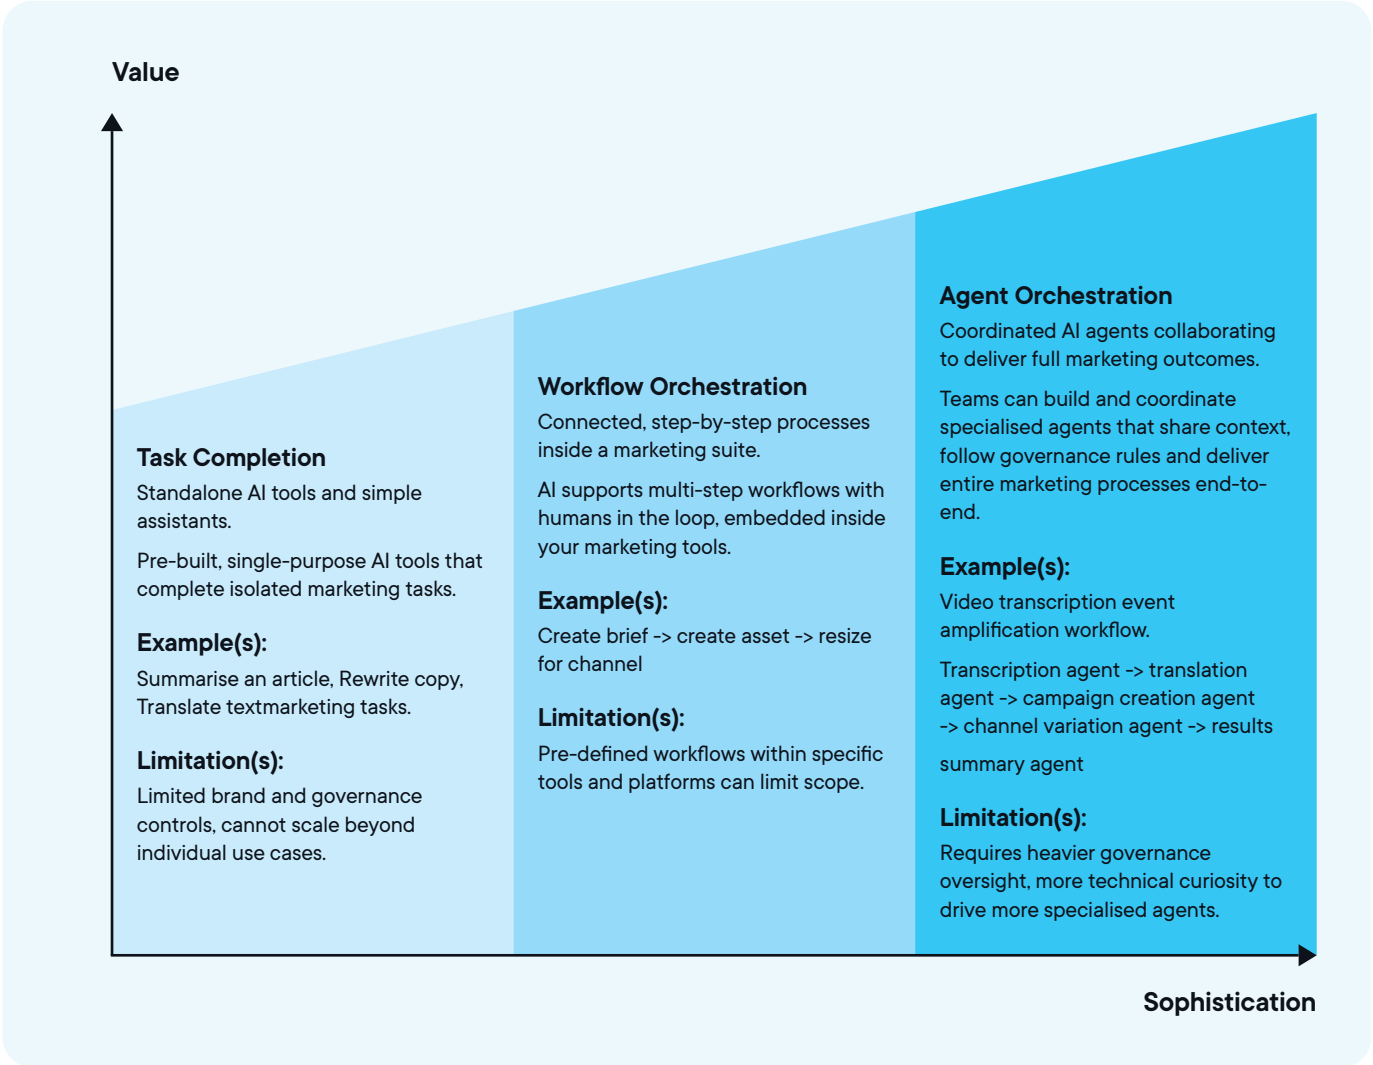

What makes Optimizely Opal different

Opal is built for this new operating era. It is not another AI feature or point solution; it is the agentic orchestration layer that unifies your marketing ecosystem. By aligning context, governance and workflow execution in one platform, Opal enables teams to automate confidently, collaborate effectively and deliver consistent, high-quality work at scale.

> “
>
> I think I’m in love...It’s an embedded, interconnected tool that improves content management, data analysis, experimentation, workflows—and much more.
>
> **- Digital Agency**

1. Context-aware by design

Opal uses retrieval-augmented generation integrated with your digital marketing suite. Agents operate with accurate organizational knowledge: brand tone, product details, past campaigns, experiment insights and regulatory rules. It improves continuously as more work flows through the system

2. Embedded in real workflows

Opal runs inside the tools marketers already use: CMP, CMS, experimentation, analytics and optimisation. There is no context switching. Agents handle drafting, grounding, QA logic, sequencing and execution inside existing processes.

3. Governed by default

Every agent and workflow inherits approved data sources, permissions, tone standards, governance rules and review checkpoints. Governance becomes system-level, consistent and automatic, not reliant on individuals.

4. Agentic orchestration

Opal coordinates multi-step marketing processes with structured logic, shared context and the right human touchpoints. Agents collaborate, escalate when needed and deliver predictable outcomes across channels, teams and markets.

With Opal, teams don't just use AI...they operate with AI.

The result is a marketing function that moves faster, delivers higher-quality work and scales with confidence. A function where agents collaborate, where humans guide and refine, and where the entire organization benefits from a shared, intelligent foundation that compounds value over time.

This is the operating model that will define the next decade of marketing.