



AI for marketing and go-to-market

Improving planning, content operations, audience insight, and campaign learning.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 08 / EMBER & CLAY

EXECUTIVE WORKING EDITION - 08 OCT 2026

08

AT A GLANCE

Improving planning, content operations, audience insight, and campaign learning.



Marketing teams are under pressure to produce more content, in more channels and languages, with the same headcount. Generative AI makes first drafts cheap, which shifts the bottleneck to strategy, brand judgment, approvals, and learning from results. This playbook explains where AI genuinely helps across planning, content operations, localization, and research, how to protect brand and claims accuracy, and how to measure whether campaigns improve rather than merely multiply.

WHO IT IS FOR

CMOs, heads of marketing and product marketing, content and demand generation leads, brand managers, and marketing operations teams planning how AI fits their go-to-market work.

TIME TO USE

60 minutes to read; one working session to apply



What you will have at the end

- A clear view of which marketing tasks benefit from AI and which need tighter control
- A content operating model with ownership, approvals, and a brand and claims check
- A measurement plan that separates output volume from campaign effectiveness
- A 90-day plan with a pilot, a test set, and a decision point

Contents

01	From cheaper content to better go-to-market	03
02	Use cases, fit, and boundaries	04
03	How generative tools fit into a marketing workflow	06
04	Roles, workflow, and approvals	07
05	Brand, claims, intellectual property, and data	09
06	Baseline, metrics, and learning loops	10
07	A staged rollout for marketing teams	12
08	Common failures, an illustrative scenario, and vendor questions	13
09	A first quarter that produces evidence	15
10	Terms used in this publication	16



EXECUTIVE BRIEF

From cheaper content to better go-to-market

01



AI lowers the cost of producing words and variants. It does not lower the cost of knowing what to say, to whom, and why they should believe it.

Marketing was one of the first functions to adopt generative AI because so much of the work is language: briefs, headlines, emails, landing pages, social posts, product descriptions, scripts, and reports. The early pattern is predictable. Output goes up quickly, quality becomes uneven, and the brand starts to sound like everyone else. Teams that gain lasting value use AI to speed up preparation, adaptation, and analysis, while keeping positioning, creative direction, and claims firmly in human hands.



Go-to-market is wider than content. It includes market research, segmentation, positioning, launch planning, sales enablement, pricing communication, and campaign learning. AI can help with each, mostly by summarizing large volumes of material, generating options to react to, and spotting patterns in feedback. The decisions about which segment to pursue and what promise to make remain leadership calls.



Where marketing teams see value first

- Content planning: topic clusters, editorial calendars, and briefs built from audience questions and search data
- Content operations: repurposing a core asset into formats, adapting to channels, and maintaining style consistency
- Localization: first-pass translation and cultural adaptation reviewed by native-language specialists
- Research and insight: synthesizing interviews, reviews, support tickets, and survey comments
- Campaign operations: variant testing, reporting drafts, and post-campaign learning summaries



THE RISK THAT MATTERS MOST

Volume without distinction. If every competitor can generate the same competent article in minutes, competent is no longer an advantage. Differentiation comes from original data, real customer stories, sharp positioning, and a point of view only your company can offer.

WORKING QUESTION

If AI doubled your content output tomorrow, which results would improve and which would simply be diluted?

WHERE IT FITS

Use cases, fit, and boundaries

02

Match AI to tasks where a draft is a useful starting point and a skilled editor can verify the result quickly.

The best marketing uses share a pattern: there is a clear source of truth, a defined format, and a skilled person who can judge quality in seconds. Repurposing a webinar transcript into a summary, a set of social posts, and an email is a good example. The facts already exist, the formats are known, and an editor can compare the draft to the source. The weaker uses ask the model to originate facts, such as market statistics, customer claims, or competitor comparisons.

USE CASE	FIT	CONTROL NEEDED
Repurposing owned content across formats	Strong	Editor checks against the source and brand voice guide
Audience research synthesis from interviews and reviews	Strong	Keep source quotes traceable; do not present themes as statistics

USE CASE	FIT	CONTROL NEEDED
Campaign brief and editorial calendar drafts	Strong	Strategist owns priorities; AI proposes options
Localization and translation first pass	Moderate	Native reviewer; glossary of approved terms and product names
Ad copy and subject line variants	Moderate	Claims and legal check; test against a control, not by taste
Competitor comparisons and market statistics	Weak	Verify every figure with a primary source before use
Testimonials, case studies, endorsements	Avoid	Never generate; use only real, approved customer material



What not to automate

- Positioning and messaging strategy, which require judgment about trade-offs
- Any statistic, quote, or customer story that the AI supplied rather than a person verified
- Regulated claims, such as health, financial, or environmental statements, without specialist review
- Customer-facing replies on social channels where tone and context carry risk
- Creative work where originality is the point, for example a brand campaign concept



DISCLOSURE AND TRUST

Some jurisdictions and platforms have rules or expectations about labeling AI-generated or synthetic content, especially images, video, and endorsements. Check current requirements for your markets and channels, and decide your own policy for when to disclose.



WORKING QUESTION

Which of your current content tasks have a clear source of truth that a draft could be checked against?



HOW IT WORKS

How generative tools fit into a marketing workflow

03

Quality comes from what you feed the model and how you constrain it, not from clever phrasing alone.

A generative model writes by predicting plausible text. Left alone, it produces the average of what it has seen, which is why unguided output sounds generic. To get on-brand, accurate output you provide three things: context (the audience, the goal, the source material), constraints (the voice, the format, the claims allowed), and examples (strong pieces your team approved). A brand guide turned into a structured instruction set, plus a library of approved messages, does more for quality than any amount of prompt trickery.

01 Define the brief

State the audience, the single message, the proof available, and the action you want. A weak brief produces a weak draft at speed.

02 Supply approved material

Provide product facts, positioning, proof points, and prior winning examples. Tell the model to use only these for factual claims.

03 Generate options

Ask for several distinct angles rather than one polished version, so that a strategist chooses direction before polishing begins.

04 Edit and verify

An editor checks facts, tone, and claims, rewrites the weak parts, and records what was changed and why.

05 Publish and learn

Tag the asset with how it was made, track performance, and feed the lessons back into the brief and the guide.

Why style guides matter more now

When humans wrote everything, brand voice lived in the heads of a few editors. With AI in the loop, it must be written down in a form a tool can follow: preferred and banned phrases,

sentence length, how formal to be with each audience, how to describe the product, and examples of good and bad copy. Teams often discover their brand guide is mostly logos and colors. Writing the voice section properly is a worthwhile project in itself, whether or not you use AI.



PLAIN-LANGUAGE WARNING

A model can produce a confident statistic, a fake quote, or a nonexistent study. Treat every number and attribution in a draft as unverified until a person has traced it to a source.

WORKING QUESTION

Could a new hire, or a tool, write in your brand voice from your written guidelines alone?



OPERATING MODEL

Roles, workflow, and approvals

04



A content system with AI needs a clear owner for voice, a clear owner for claims, and a fast route for approvals, or the speed gain disappears into rework.

When drafting becomes fast, review becomes the constraint. Teams that do not redesign approvals find that the legal or brand reviewer is now buried under three times the volume. The answer is not to skip review. It is to tier content by risk, so low-risk assets move quickly and high-risk assets get proper attention. A social caption repurposed from an approved blog post does not need the same scrutiny as a new product claim or a regulated advertisement.



ROLE	OWNS	DECIDES
Head of marketing	Strategy, risk appetite, budget	Which use cases and channels use AI

ROLE	OWNS	DECIDES
Brand and content lead	Voice guide, templates, editorial standards	What counts as on-brand and publishable
Product marketing	Positioning, proof points, claims register	What may be claimed about the product
Marketing operations	Tools, data, integrations, measurement	Access, tracking, reporting definitions
Legal and compliance reviewer	Claims, privacy, advertising rules	Escalation rules and prohibited claims
Editors and channel owners	Final quality of every asset	Publish, rework, or reject

Tier content by risk

- Tier 1, low risk: repurposed owned content and internal summaries; editor review only
- Tier 2, medium risk: new copy for known products; editor plus product marketing check
- Tier 3, high risk: new claims, comparisons, regulated topics, spokespeople; full legal and brand review
- Record the tier on each asset so reviewers can see what is expected of them
- Review the tiers every quarter, based on errors found and escapes after publication

Keep a claims register: a single list of approved statements about the product, with the evidence behind each and the date it was last checked. Drafting tools can be pointed at the register, and reviewers can check against it. This reduces the most common failure in AI-assisted marketing, which is a plausible claim that nobody can support.

WORKING QUESTION

If your reviewers received three times as many drafts next month, which tier would they be forced to rush?

RISKS AND CONTROLS

Brand, claims, intellectual property, and data

05

The main exposures are unsupported claims, ownership and copyright uncertainty, careless use of customer data, and sameness that erodes the brand.

Marketing is public. A mistake does not stay inside a spreadsheet; it appears in a campaign, an ad, or a post that can be screenshotted. The controls below are practical and mostly cheap. The key is to apply them before volume grows, not after the first embarrassing error.

RISK	EXAMPLE	CONTROL
Unsupported claims	Invented statistic or superlative in ad	Claims register, source trace, tiered review
Intellectual property	Generated text or images that resemble protected work; unclear ownership	Policy on tools and outputs; check vendor terms and current law with counsel
Customer data	Customer lists or call transcripts entered into public tools	Approved tools, data classification, no personal data in unapproved systems
Privacy and consent	Targeting or personalization beyond what consent allows	Check data protection and marketing rules; limit sensitive inference
Brand sameness	Content that reads like generic voice copy	Voice guide, original data, distinct point of view, editorial bar
Bias and exclusion	Stereotyped imagery or audience assumptions	Review checklist and diverse reviewers on sensitive campaigns

Pre-publication check for AI-assisted assets

- Every statistic, quote, and named source has been traced to a primary source
- Product claims match the claims register
- The voice matches the guide and does not sound like a template

- Any image or video is cleared for use and labeled where required
- No customer or personal data appears that should not be public
- A named person has approved the asset and the tier is recorded



Rules about advertising standards, endorsements, synthetic media, copyright, and privacy vary by country and are evolving. This playbook cannot tell you what is lawful in your markets. Ask your counsel, check current regulator guidance, and use general frameworks such as the NIST AI Risk Management Framework or ISO/IEC 42001 as a way to organize internal policy, not as a substitute for legal review.

WORKING QUESTION

Who in your team can say with certainty which images and text in your last campaign were AI-generated?



MEASURING VALUE

Baseline, metrics, and learning loops

06



Output counts are easy and misleading. Measure cycle time and cost for efficiency, and measure results by comparing against a control for effectiveness.

Before using AI widely, record how long common deliverables take now, how many review rounds they need, and what they cost to produce. Then record performance of the content itself: engagement quality, conversion, and contribution to pipeline. Without this baseline you can say you produced more, but you cannot say you did better. Be wary of vanity measures such as page views for AI-generated articles, which can rise while qualified leads do not.



METRIC	HOW TO MEASURE	WATCH FOR
Production cycle time	Brief to publish, by asset type, before and after	Review queues that absorb the saving

METRIC	HOW TO MEASURE	WATCH FOR
Edit depth	Share of draft changed by editors	Very high edit depth means weak inputs; near zero may mean weak review
Error escapes	Factual or claim corrections after publication	The most important quality signal
Conversion and qualified leads	Landing page or campaign conversion versus control	Volume gains with flat conversion
Share of voice and branded search	Brand mentions and branded queries over time	Slow signals; track quarterly
Cost per asset and per qualified lead	Fully loaded team and tool cost	Ignoring review and rework time
Audience feedback	Reply sentiment, unsubscribe and complaint rates	Rising unsubscribes after personalization

Test creative the right way

AI makes it easy to generate twenty subject lines, and that tempts teams to choose by taste. Instead, run controlled tests: split the audience, change one thing, and decide the winner in advance by a metric tied to business results, such as qualified replies rather than opens. Mail privacy features and bots can distort open rates, so lean on downstream actions. Keep a learning log of what worked, for whom, and why, because the main gain from fast variant creation is faster learning, not more variants.

For quality, build a test set of 30 to 50 briefs with approved reference outputs and a rubric covering accuracy, voice, usefulness, and claims safety. Score new tools and prompt changes against it before they reach production. Thresholds here are illustrative: for instance, zero unsupported claims and most drafts needing only light edits. Set your own with your team and keep them stable so that you can compare over time.

WORKING QUESTION

What decision about next quarter's campaigns would change if you could see edit depth and error escapes by channel?

IMPLEMENTATION GUIDE

A staged rollout for marketing teams

07

Begin with internal and low-risk work, write down the voice and the claims, and widen the scope as quality holds.

01 Write the foundations

Turn the brand guide into a usable voice document and create the claims register. These two assets drive quality for every later step.

02 Pick two workflows

For example, repurposing long-form content and synthesizing customer feedback. Both have clear sources and fast review.

03 Set tiers and approvals

Agree the risk tiers, name reviewers, and set expected turnaround times so that review does not become a queue.

04 Train the team on real briefs

Practice with live projects. Teach people to supply context, ask for options, challenge claims, and record edits.

05 Measure against the baseline

Track cycle time, edit depth, error escapes, and conversion for the pilot work versus similar past work or a control group.

06 Expand deliberately

Add localization, ad variants, and personalization one at a time, with a fresh test set and review rules for each.

Launch readiness

- Voice document and claims register exist and have owners
- Approved tools list and data rules are shared with everyone, including agencies
- Risk tiers and review turnaround targets are agreed
- Baseline cycle times and quality measures are recorded

- Agency and freelancer contracts say whether and how AI may be used
- There is a defined route for correcting published errors quickly

Do not forget agencies and partners

Much marketing output comes from agencies and freelancers. Update briefs and contracts to state whether AI use is allowed, what must be disclosed to you, which tools may handle your data, and who is responsible for claims and rights. Ask for the same traceability you expect internally. Otherwise AI-generated material may enter your campaigns with no record of how it was made or checked.



WORKING QUESTION

Which two workflows in your team have the clearest source material and the quickest review loop?



PITFALLS AND BUYING

Common failures, an illustrative scenario, and vendor questions

08



Most disappointments trace to skipping the voice and claims groundwork, judging by volume, or buying tools before choosing workflows.

- Publishing at scale before the voice guide exists, producing content that sounds like everyone else
- Counting assets produced as success while qualified leads stay flat
- Using AI for statistics, comparisons, or testimonials without source checks
- Letting review queues grow until reviewers approve without reading
- Allowing every channel owner to choose their own tool, so data and rights are unmanaged



ILLUSTRATIVE SCENARIO, NOT A REAL COMPANY

A mid-sized services firm triples its blog output with AI. Traffic rises, leads do not, and a prospect points out a cited study that does not exist. The firm halves output, adds a source trace and an editor, and invests in original survey data of its own. Output falls; lead quality and credibility recover.

Questions to ask any vendor

Vendor diligence

- Is our content, brand material, or customer data used to train models for others
- Which models and sub-processors does the product use, and how are changes communicated
- Who owns outputs, and what indemnities or warranties exist about intellectual property
- Can the tool be limited to our approved material and claims
- What controls exist for brand voice, approvals, audit trails, and user permissions
- How does it handle image, video, or voice generation, and how is synthetic content labeled
- How do we export our content and data if we leave

Run a bake-off using your own briefs, your own voice guide, and a hard set of cases that include missing facts and risky claims. Observe not only the best output but how often the tool refuses to guess and asks for the missing information. A tool that fills every gap with plausible text is a liability in a public-facing function.

WORKING QUESTION

Which of these failure patterns would be hardest for your team to notice before a customer did?

30 / 60 / 90

A first quarter that produces evidence

09



Build the voice and claims foundations first, prove value on low-risk workflows, and then widen the scope with measured quality.

First 30 days: ground it

- Convert the brand guide into a usable voice document
- Create the claims register with owners and review dates
- Choose two low-risk workflows and record their baseline cycle times
- Publish approved tools, data rules, and the disclosure policy



Days 31 to 60: test it

- Run the pilot with risk tiers and named reviewers
- Score outputs against a test set of 30 to 50 briefs
- Track edit depth and error escapes weekly
- Update agency briefs and contracts to cover AI use



Days 61 to 90: decide

- Compare pilot results with the baseline and a control
- Decide which workflows to extend, such as localization or ad variants
- Set a quarterly review of tiers, tools, and the claims register
- Start controlled creative tests tied to qualified outcomes



GLOSSARY

Terms used in this publication

10

TERM	MEANING
Claims register	A governed list of approved statements about products or services, each linked to its supporting evidence and review date.
Content operations	The people, process, and tools used to plan, produce, review, publish, and maintain content at scale.
Edit depth	The proportion of an AI draft that an editor changes before publication, used as a signal of draft quality.
Go-to-market	The plan for how a company brings an offering to a defined audience, covering positioning, pricing, channels, and sales support.
Hallucination	A fluent but unsupported or false statement produced by a language model.
Localization	Adapting content to a language, culture, and market, beyond literal translation.
Positioning	The deliberate choice of how an offering is understood relative to alternatives for a particular audience.
Risk tier	A category that sets the level of review an asset needs based on the consequences of an error.
Segmentation	Dividing a market into groups with shared needs or behaviors so that offers and messages can be tailored.
Synthetic media	Images, audio, or video generated or heavily altered by AI, which may carry labeling expectations.
Test set	A fixed group of briefs with approved reference outputs used to compare tool quality over time.
Voice guide	A written description of how the brand sounds, with examples, used by people and tools to keep content consistent.

A note on use

Free assessments are available at halden.company/assessments, and you can reach the team at halden.company/contact. Figures here are illustrative, and nothing in this playbook is legal advice.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

