



HUMAN-FIRST AI MARKETING PLAYBOOK

Scale faster with practical AI that handles the heavy lifting without risking the brand you've built.



Welcome!

Embed AI into your marketing without losing trust or brand integrity

If you're a business leader or run a marketing team, and you're using AI (or about to), this playbook is for you. It's designed to make life easier without your standards suffering.

Inside you'll find copy-and-paste tools you can use immediately. Get more done and grow your business without losing the trust you've worked so hard to build.



Contents

- 01 Why this matters now
- 04 Your trust settings for AI marketing
- 08 The AI publishing checklist
- 09 A simple sign-off model
- 10 The human-first AI pledge
- 11 Where AI fits in the marketing workflow
- 12 Prompting for consistent, credible output
- 15 Agents: the shared standard problem solver
- 17 The best agent builds for marketing
- 19 Getting the right foundations
- 20 Know your options for embedding AI ethically
- 24 Make this real in your business (next steps)

Why this matters now

If you're responsible for marketing right now, it can feel like this constant extra job that sits on your chest.

It's not "do some posting" anymore. You're expected to grow. Stay visible. Look confident. Sounds consistent. Keep standards high. Do it without hiring. Do it while the market is twitchy and buyers take longer. Do it while your attention is split across ops, people, finance, delivery, product. Argh!

Most leaders we speak to are overloaded. Marketing ends up living in your head, so you become the bottleneck without meaning to.

The channels don't help. Everything is louder. Everyone is posting and "thought leading". The stuff that used to work now works... sometimes. And because platforms reward volume and consistency, it can start to feel like you're feeding a machine rather than building an asset.

Then AI shows up

Relief. Finally - a way to plug the gap.

An affordable subscription and, just like that, problem solved. You can get "ten campaign ideas" in five seconds. A draft in minutes. Momentum, on demand.

So you try it.

Some of it is genuinely useful. It gets you moving instead of stuck.

But AI isn't the shortcut you think it is (yet)

Loads of teams are “using AI” now. That doesn't mean they're using it well. Or in a way that protects the brand.

The risk is that AI quietly becomes the person driving the marketing. Deciding what you sound like, what you claim. Filling in gaps you haven't had time to think through.

When content gets easier to produce, you pay in extra reviews and extra responsibility on the final approver.

If you don't decide what “good” looks like before you use it, the tool will happily fill the gaps for you. Sometimes it'll guess wrong in a way that's obvious.

More often it'll guess wrong in a way that looks fine, until it isn't. So before we get into tactics, prompts, agents, workflows, any of that, we need a backbone.





“If your marketing standards live inside someone’s prompt, you’re one rushed draft away from publishing something you can’t defend. As the person responsible, you’re the one who pays for that, not the tool.”

Your trust settings for AI marketing

These aren't settings you switch on inside an AI platform. They're the settings humans need to embody. Shared defaults that keep your marketing credible when output gets fast.

They stop AI from becoming a shortcut around judgement, or a way to publish things you can't stand behind. If your team remembers nothing else, remember these.

1

Strategy needs a human owner

AI can support decisions, but it can't make them.

Strategy is the bit where you choose:

- Who you're for
- Why you (and not the other option)
- What you stand for
- What you're willing to say
- What you're not willing to say
- Where you're going to focus
- What you'll ignore

AI is useful inside that work. It can help you explore options, pressure-test a message, summarise research and draft variations for different channels. But if you ask it to make decisions, you're basically outsourcing leadership to a prediction machine.

A simple test: if the output would change your direction, it needs a human owner.

2

AI isn't a fact-checker

AI can feed you confident nonsense. It's how these tools work. They're built to produce plausible language. So treat output like a draft, not a source.

If AI gives you facts, stats, "industry benchmarks", customer behaviours, legal claims, competitor details, even quotes... you need receipts. Links. Primary sources. Evidence you'd be happy to show if someone challenged you. Also, don't trust the receipt on face value either. Click it. Check it's reputable and says what the content claims it says.

3

Draft with AI, finish with a human

Your brand voice is made of specifics. It comes from:

- Lived customer context
- The proof you can stand behind
- Past decisions
- The language you choose
- The claims you avoid

That's what gives voice its credibility and makes it unique.

Use AI to get the first draft on the page, then keep a human in the loop to add the bits AI can't earn, like your opinion, your tone, the real-world context behind the claim and the slightly rough edges that make it sound like you rather than a bot.

4

Sensitive information needs guardrails

Customer info, personal data, sensitive commercial details, internal performance numbers... anything you'd be uncomfortable seeing shared outside the business needs extra care with AI tools.

Business and enterprise AI accounts do tend to offer stronger security and governance features, which helps reduce risk. But it doesn't remove it.

The aim is to make "safe use" easy and repeatable. Be clear on:

- what information is okay to use
- what needs redacting first
- which tools are approved
- which settings should be switched on

If you're unsure, bring IT or your data lead into the decision early, then write it down so the whole team is working from the same rules.

5

Speed must not break standards

If you double your output, you naturally increase the chances of avoidable mistakes, like:

- slipping in a claim you can't prove
- drifting tone across channels because different people (or prompts) are steering it
- publishing something that reads "fine" but doesn't feel like you
- creating a small trust wobble that takes weeks to smooth over

Speed only counts if your standards travel with the work. Build lightweight checks that happen every time (use the AI publishing checklist in this guide).

This is why values matter more than prompts. Prompts will change. Tools will change. Pressure will change. Values are the steady part, the shared line your team can hold even when things are moving fast.



The AI publishing checklist

Truth and claims

- Read every factual claim and it still makes sense
- Stats/benchmarks/quotes have sources
- Clicked sources: reputable, current and support the claim
- No invented numbers, dates, examples, results
- No guarantees or risky promises

Proof and evidence

- Claims are backed by evidence we can stand behind
- Evidence is usable if challenged
- Attribution is clean (named/linked as appropriate)

Tone

- Sounds like us
- Removed vague “could be anyone” lines
- Headline, opening, CTA match stance and confidence

Strategy fit

- Audience is clear in one line
- One main message
- Fits current priorities
- CTA matches intent

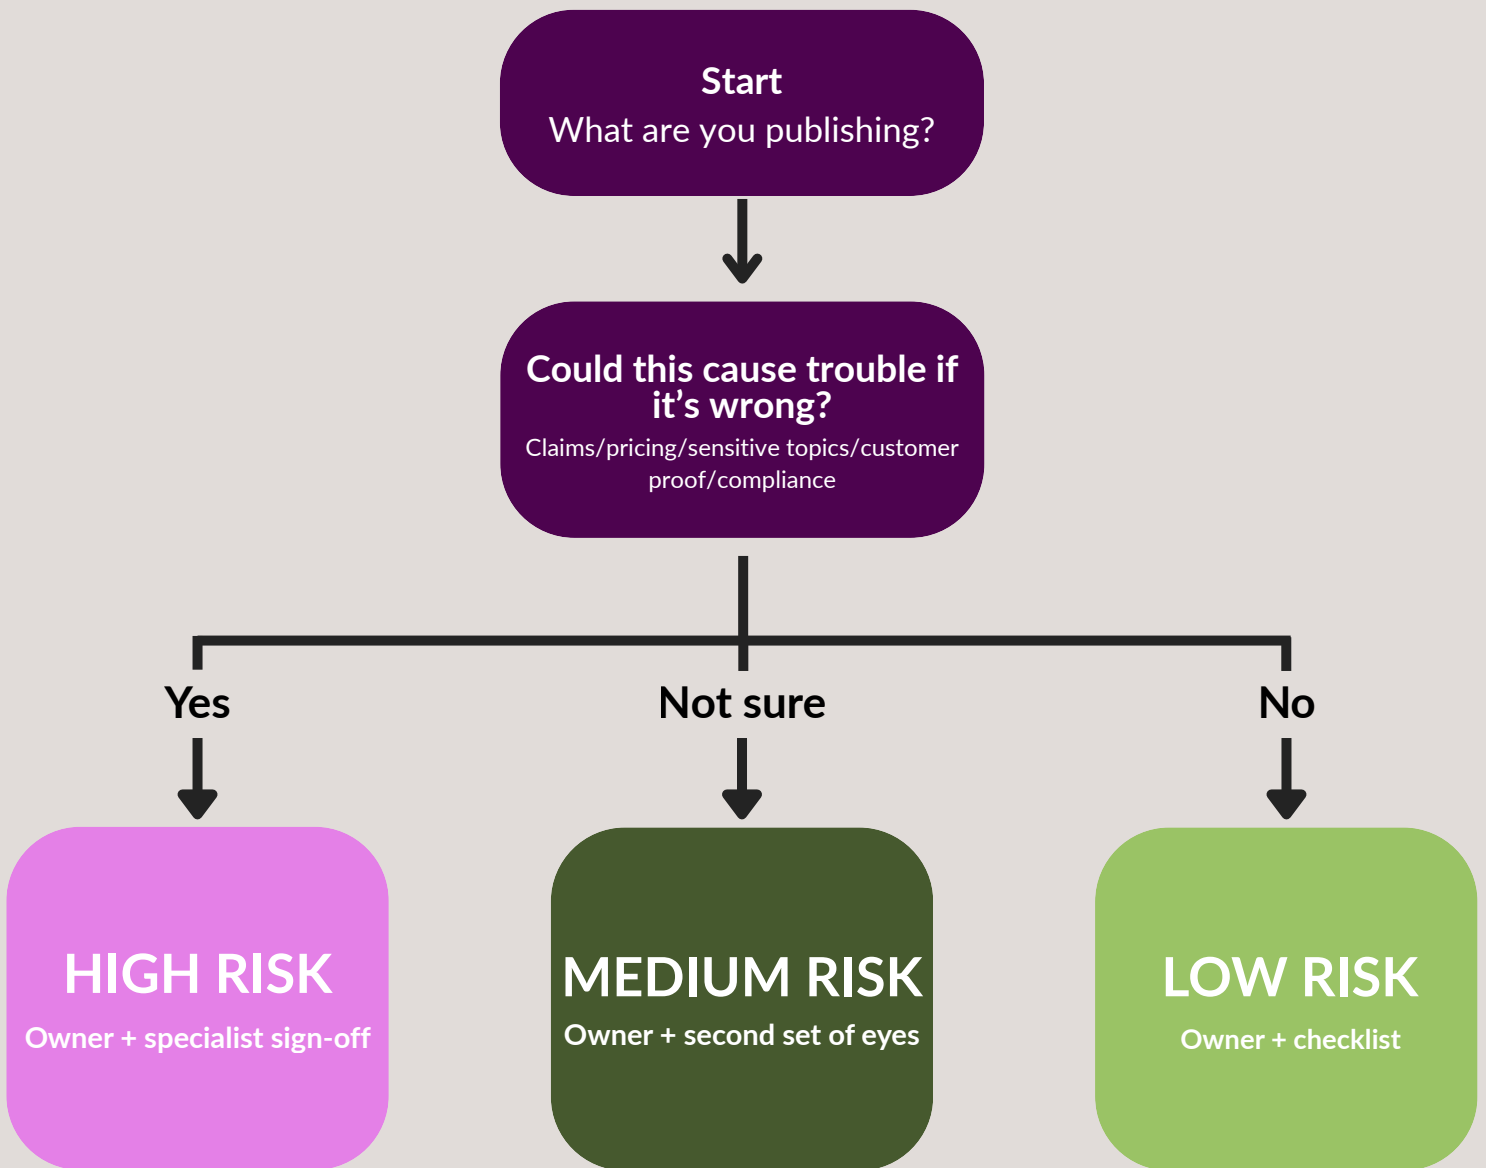
Permission and safety

- No sensitive or confidential info
- Customer references approved
- Compliance reviewed if needed
- Approved tools/settings used

Publishable

- I'd put my name on it
- It sounds like us
- We can defend it if challenged
- Our best customer would recognise us
- Named person signed off

A sign-off model built for AI-supported work



High risk: pricing, guarantees, regulated topics, performance claims, customer proof

Medium risk: blogs, landing pages, newsletters, ads

Low risk: simple social updates and promos

The human-first AI pledge

A one-page promise your team can share and use as a final check before anything goes out (no, we're not asking you to stand in a circle and read it aloud every morning... that would be hilarious to watch, though).

We use AI to assist, not replace
Humans own strategy and judgement

We don't publish claims without proof
No source = no claim

We treat data like a duty of care
If we wouldn't email it to a stranger, we don't paste it into a tool

We protect our voice
AI can draft, but humans add the context and truth that make it us

We don't let speed lower the bar
More output only counts if standards stay the same (or get better)



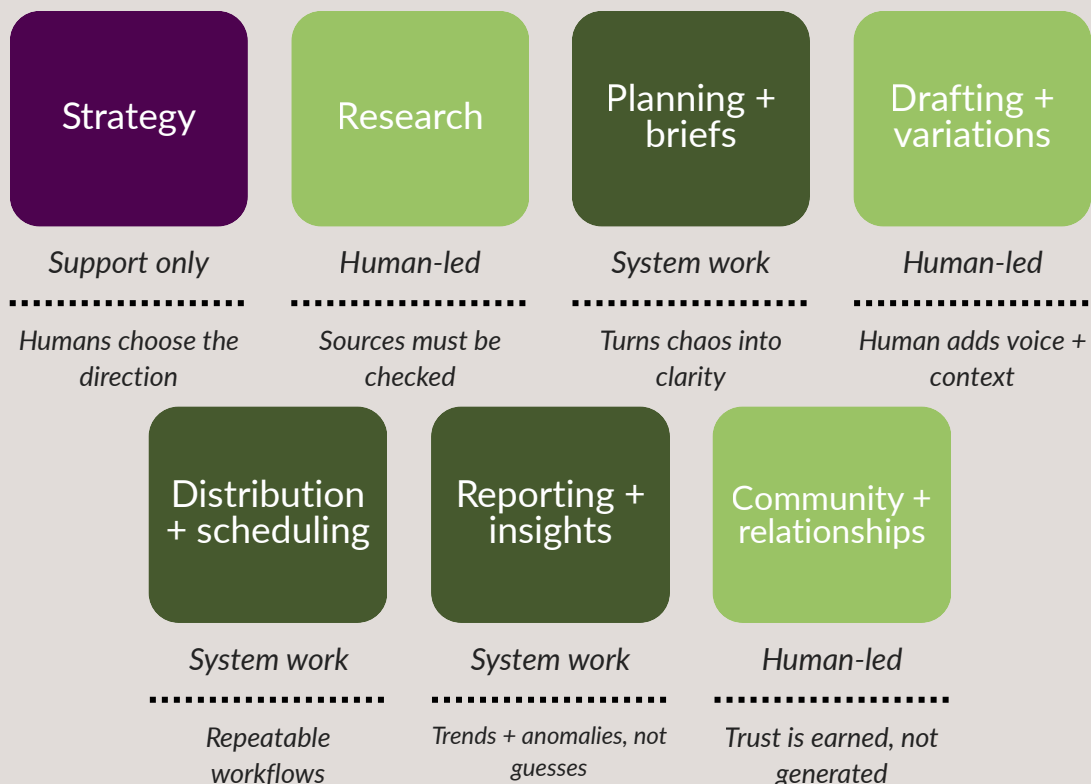
Where AI fits in the marketing workflow

Most teams get into trouble with AI because they use it everywhere in the same way. They treat it like a general replacement for thinking, rather than a tool that's strong in specific parts of the marketing lifecycle.

A simple way to stay safe (and get real value) is to map your funnel and lifecycle, then decide where AI should be:

- supporting
- operating
- leading

The goal is to use AI in the places where it improves speed and consistency without taking over judgement or voice.



- Support only** = AI assists, humans decide
- System work** = AI can run repeatable steps
- Human-led** = Relationship + judgement stays human

Prompting for consistent, credible output

You don't need loads of prompts. You need one simple way of asking that keeps your drafts or research consistent and easy to stand behind, even when different people are using AI on different days.

The easiest way is to treat your prompt like a short brief. You're telling the tool: who it's helping, what you're making, what it can and can't say and what it must use as the truth.

A reusable prompt structure

Job → Background → Rules → What to use → How to sound → What to deliver

1 Job (what you want it to be)

Give it a clear role that matches the task.

Examples: “You’re a content strategist helping me write a LinkedIn post.” / “You’re helping me turn notes into a blog outline.”

2 Background (what is this and why)

- Who is this for?
- What are we trying to achieve?
- Where will it be used? (email, blog, LinkedIn, landing page)

3 Rules (the boundaries)

- What must stay true?
- What kind of claims are allowed?
- What topics are sensitive?
- What should it steer away from?

4 What to use (your source material)

Paste the things it’s allowed to use:

- your notes
- product details
- approved proof points
- links you trust
- customer language

5 What to deliver (so it’s easy to review)

Ask for:

- a clear headline/opening/body/CTA
- two versions (short and slightly longer)
- a list of anything that needs a source

A copy-paste “base prompt”

Use the below prompt, but ensure you edit it with the correct information before pressing submit.

AI Content Generator



Please type your prompt below...

“You’re a [insert job title]. Help me write an on-brand draft for [audience] with the goal of [what you want to happen], for use on [channel]. Keep it truthful and only use the information I provide below as the source of truth - don’t add facts, stats, quotes, or examples that aren’t included. If anything is unclear, ask me 1–2 questions or label it as a guess. Please give me Version A (short) and Version B (slightly longer), each with a clear headline, opening, body, and CTA, plus a bullet list of any lines that would need a source before we publish.”

Agents: the “shared standard” problem solver

Prompts are great, but you’re asking a lot from your team to use the same prompt every time and that can add more friction than it takes away.

Under pressure, people naturally tweak things. Someone forgets the truth rules or skips the source notes. Someone uses a different tool and the tone shifts. Before long, your standards are back living in people’s heads and the brand suffers.

Agents make life easier

Think of an agent as AI with a job description and house rules baked in. Instead of relying on every human to remember the same prompt structure every time, the agent shows up with the structure and guardrails already loaded.

Agents work best when they’re designed for a specific job. One agent per task. Avoid building a vague “marketing helper” that tries to do everything and ends up doing nothing consistently. Specific job = predictable output = easier review.





“Agents are not the boss. They’re the process. That’s why they solve the shared-standard problem so well.”

Agents for marketing activities



Brief Builder Agent

Turns messy notes into a clean brief (goal, audience, offer, proof, CTA and objections).

Support only



Source Check Agent

Flags claims that need proof, checks links and lists what's credible vs what needs replacing.

With receipts



On-Draft Brand Agent

Creates first drafts using your tone rules and approved proof points (and highlights any gaps).

System work



Repurposing Agent

Takes one source asset (blog/webinar/case study) and spins channel versions that keep the same messaging.

Support only



Tone Consistency Agent

Reviews content and points out where it stops sounding like you (then suggests tighter alternatives).

System work



Risk Flagging Agent

Spots risky phrasing (promises, guarantees, sensitive claims).

System work



FAQ/ Knowledge Base Agent

Drafts customer-facing answers using only approved internal info (and asks when it's missing).

Support only



Reporting Summary Agent

Turns performance data into a weekly narrative (what changed, what it might mean, what to do next) with clearly separated facts.

System work

Do you have the right foundations?

Building a personalised GPT or writing good prompts depends on having clear, well organised information.

If your information is messy or unclear, you can still get an answer. But it probably won't be complete. It might miss important details. And you'll need to double check it.

The better organised your inputs are, the better and more reliable the output will be.

The foundations you need aren't "*more prompts*". They're the things prompts depend on:

- A clear strategy so the tool knows what it's trying to achieve (and what it should ignore)
- A shared set of facts and proof points so it doesn't invent credibility
- A consistent voice and set of examples so it doesn't smooth you into generic
- Simple templates and structure so different people get similar outputs
- Clean, approved inputs (notes, research, customer language) so it's working from reality

This is why the strategy work has to come first. AI speeds up execution. It doesn't create direction. If you point it at a fuzzy strategy, it will help you produce fuzzy marketing faster.

A good rule of thumb: if you can't explain your offer and your audience clearly to a new hire, you're not ready to automate anything. Cement the fundamentals and your prompts (and agents) will get you better results.

Know your support options for embedding AI responsibly

If you want to use AI credibly in marketing, you've got a few routes. The right one depends on your confidence, capabilities, how much time you can spare and how risky it would be for your brand to get it wrong.

Option 1

Do it yourself

This works when your marketing is already in good shape. You know your audience, your tone of voice, your key messages and your values - and it's all written down and backed up by strong, consistent content on your website. Your team has already been experimenting with AI and knows how to use it well.

What you'll need next:

- A Business or Enterprise account on your chosen AI platform for better security
- Support from your IT provider (who is knowledgeable in AI)
- A dedicated person to set up and maintain your agents
- A proper sign-off process so nothing important slips through when things move fast



Option 2

Light support

Ideal when you've got the basics in place and a team who can run marketing, but you want help setting things up so AI use is consistent and safe. Think: getting your foundations organised, your standards written down, your team trained up and your workflows working before you scale output.

What you'll need next:

- A short set-up phase to pull your key documents into shape (brand voice, messages, proof points, templates)
- Agree your sign-off rules
- Choose approved tools and settings
- Set up a few task-specific agents



Option 3

Complete delivery

This is for when you can't DIY yet because the basics aren't fully nailed. Your strategy and messages aren't established, campaigns are more ad hoc than repeatable or your website needs updating. AI can help, but it will only amplify what's already there. So this option means getting the foundations clear first, then using AI to deliver faster without drifting off-brand or making claims you can't stand behind.

What you'll need next:

- A marketing partner (or in-house lead) who can own the work end-to-end and knows how to use AI without cutting corners
- Time from the people who hold the truth (founder/leadership, sales, delivery, product) to shape the strategy
- New or better materials (brand guidelines, website, decks, proposals, past campaigns, case studies, customer feedback)
- Agreement on your non-negotiables
- Someone to own it internally once it's running



Make this real in your business (next steps)

We hope this playbook has been genuinely helpful (and that it's taken a bit of weight off your shoulders). So, what's next?

Most of the businesses we work with need **Option 3**. Their brand has evolved, but their marketing hasn't kept up. In that situation, it's too early to bring in AI. We help you get your strategy right first, then we set up the tools so you can use AI safely and successfully.

We also work with teams at **Option 2**. Through 1:1 mentoring, team sessions, and practical AI-enabled tools, we help you build the confidence and capability to use AI without compromising credibility or brand quality.

Not sure which option fits you? We're happy to help you understand where you are and what the simplest next step looks like.



Build a marketing engine you can stand behind

If you want AI to make your marketing faster and better without losing your grip on quality, let's talk.



Get in touch

mpand.co

hello@mpand.co

+44 (0)161 478 6009

