

CMO · HEAD OF DIGITAL / MARKETING LEAD

AI Governance and Marketing Policy

What AI may touch, what it may not, who verifies, and what we disclose. A policy you can put in front of a legal team rather than a set of good intentions.

THE DECISION THIS TEMPLATE PROTECTS

Where AI is permitted in this function, and who remains accountable for the output.

ROLE	COURSE PLACEMENT	REVISION
CMO / Head of Digital	Base Module · AI and intelligence systems primer	Version 1.0 Current as of 19 August 2026

Write the permission tiers first; everything else follows from them. The principle that makes the policy coherent is that accountability never transfers to a tool. Whoever publishes is accountable, regardless of what produced the draft.

BEFORE YOU START, YOU NEED

- An honest inventory of how AI is currently being used, including uses nobody has approved.
- Legal or compliance contact.
- The client or customer data your team touches.

01 Permission tiers

TIER	WHAT IT COVERS	APPROVAL	VERIFICATION
Green: use freely	Drafting, ideation, summarising internal material, code assistance	None	Normal review
Amber: use with verification	Research, data analysis, customer-facing copy, creative	Team lead aware	Named verifier signs off

TIER	WHAT IT COVERS	APPROVAL	VERIFICATION
Red: not without explicit approval	Anything with personal data, regulated claims, synthetic likeness, automated customer communication	Named approver	Documented
Black: never	Fabricating evidence, synthetic testimonials, impersonating a real person, presenting AI as a human	None	N/A

The black tier is short and absolute. Everything on it is the kind of decision that ends careers and it should never be a matter of judgment in the moment.

02 Data rules

DATA TYPE	MAY IT GO INTO AN AI TOOL?	WHICH TOOLS	CONDITIONS
Public information			
Internal non-personal			
Customer personal data			
Employee data			
Client confidential material			
Unpublished financial information			

Check retention and training settings for every approved tool, and record who checked and when. Defaults change without notice.

03 Accountability

The load-bearing principle: accountability does not transfer to a tool.

- Every published output has a named accountable human.
- The verification record is completed before publication, not reconstructed after a complaint.
- 'The model produced it' is not an explanation and will not be accepted internally.
- Anyone may decline to sign off on output they cannot verify, without consequence.

04 Disclosure

CONTEXT	WHAT WE DISCLOSE	DECIDED BY
Published content		
Imagery and creative		
Customer-facing automated conversation		
Research and reports		
Work delivered to clients		

Set the standard once, at policy level. Deciding disclosure case by case means deciding it by convenience, and the pattern becomes visible to anyone paying attention.

05 Review and enforcement

ITEM	OWNER	CADENCE
Approved tool list and their settings		
Policy review		
Spot check of verification records		
Incident log		

WHAT HAPPENS WHEN THE POLICY IS BREACHED

JUDGMENT PROMPTS

1. Which tier is your team already operating outside, without approval?
2. If a customer asked whether a piece of your content was AI-generated, what would you say, and is it the same answer you would give a journalist?
3. Who checked the retention settings on your approved tools, and when?
4. Is there anything on the black list that somebody in your organisation would currently see as a grey area?

WHAT THIS TEMPLATE WILL NOT DO FOR YOU

- It is not legal advice. Take the completed policy to somebody qualified before adopting it.
- It will go stale. Tool capabilities and settings change constantly. Check the review date on this page.
- It cannot address model bias or output quality, only accountability and permitted use.
- A policy nobody enforces is worse than none, because it creates documented negligence.

THE DEFENSIBLE OUTPUT

An adopted policy with four permission tiers, data rules, a named accountable human per output, and a disclosure standard set at policy level.

AFactor | An online-first institution for career transformation and future industry leadership.

This template is published open. Copy it, adapt it, use it on client work. If you pass it on, keep this line on it.

H7 · Version 1.0 · Current as of 19 August 2026 · Review cycle: every 90 days