

From AI pilot to accountable marketing workflow

Before a promising AI pilot becomes everyday practice, check who owns it, what it uses and how the output is reviewed. A working guide for research, writing, campaigns and reporting.

Reviewed: 4 October 2026

Define an operating model for AI

Start with one repeatable task and a baseline measured without AI. Record time spent on drafting, review and correction separately. Speed only becomes useful when the approved output meets the same quality standard. NIST's framework offers a governance structure, not a certification that a workflow is safe.

Baseline → Controlled pilot → Evaluation → Approved use

Input boundary

Approved source set and permitted data categories

Confirm access, retention and reuse terms before supplying confidential material

Evaluation

Representative briefs, factual errors and review time

Use the same test set when comparing models or prompts

Release control

Named approver and retained final version

Log model, prompt version, sources and reviewer decision

Application checklist

- Treat retrieved documents as evidence, not as instructions to the model.
- Define a stop condition for unsupported claims or exposure of restricted data.
- Expand only after a documented review of quality, cost and responsibility.

From AI pilot to accountable marketing workflow

Working checklist

- 1** Start with one clearly defined task. Agree on the source material, the expected result and who is responsible. Start with a task such as turning an approved hotel fact sheet into market specific draft copy. Keep publishing and pricing changes outside the pilot unless explicitly designed and authorised.

- 2** Agree on the sources everyone will use. List approved documents, dates, content owners and access restrictions. Mark facts that expire, such as rates and opening periods. The model must flag missing evidence rather than invent a plausible answer.

- 3** Set input boundaries. Decide which data may enter the approved tool and under which settings. Use anonymised or synthetic test inputs where possible. Confirm supplier terms, retention and permissions with the responsible teams before real client or guest data is introduced.

- 4** Define acceptance criteria. Score factual accuracy, completeness, brand voice, language quality and suitability for the audience. Set a zero tolerance rule for invented prices, services or claims. Keep a small fixed evaluation set so successive versions face the same challenge.

- 5** Separate drafting from approval. Assign a subject specialist to facts and an editor to language and positioning. Record the final approver. Escalate sensitive claims and public interest content for the relevant policy or legal review rather than relying on a generic AI disclaimer.

From AI pilot to accountable marketing workflow

Notes, owners and next actions.

6

Count the time the whole process takes. Compare drafting, checking, revisions and approval time with the previous workflow. Record errors and correction effort. Faster first drafts are not a productivity gain if review costs rise by more than the time saved.

7

Keep a record of the work and a way to continue without AI. Store the source version, prompt or instructions, tool version where available, reviewer changes and approved output. Define who can stop the workflow and how production continues manually after a failure.

8

Expand only after evaluation. Review a representative batch, including difficult cases and both Greek and English. Scale the tasks that meet quality and efficiency thresholds. Recheck when the model, source material or publishing context changes.

From AI pilot to accountable marketing workflow

Pilot: bilingual hotel campaign copy

Illustrative working example. Figures and decisions are not client results or market benchmarks.

Source and output

Input: approved room facts, seasonal dates, rate inclusions and brand voice. Output: three draft propositions for each market, with a source reference beside every factual claim. Unconfirmed facts remain questions for the hotel.

Acceptance

The reservations team checks and approves the facts. An editor rewrites each language for its audience. The campaign owner approves the final version and any disclosure requirements with the relevant specialist. Nothing publishes automatically during this pilot.

Value test

Compare how long it takes to produce approved copy and how often facts need correcting. Keep the new process only if it meets your quality standards and saves enough time to justify using it.

Further reading:

NIST / AI Risk Management Framework

European Commission • updated 6 August 2026

NIST • AI Risk Management Framework