

WORKFLOW AI READINESS

What Makes a Marketing Workflow Ready for AI?

A 30-run controlled test of what happens when audience, goal, source, format or owner is missing.

THE RESULT

Complete context: 5/5 first-pass accepts. One missing input: 8/25.

The missing inputs were deliberately created test conditions - not measured prevalence in real companies.



The answer in four numbers

A ready workflow reduces assumptions, clarifications and downstream correction before the model ever starts.

100% -> 32%

FIRST-PASS ACCEPTANCE

Complete briefs versus all 25 single-omission runs.

8 -> 17 min

MEDIAN TIME

Complete context versus the median across missing-input runs.

48

UNSUPPORTED ASSUMPTIONS

Across 25 missing-input runs, compared with two in the control.

30

CLARIFYING QUESTIONS

Across missing-input runs, compared with one in the control.

The practical lesson

AI readiness is not a tool property. It is the combination of a defined job, approved evidence, a required output and a person who owns the decision.

The test isolates one missing input at a time; real workflows may be missing several simultaneously.

What existing guidance already implies

Across official guidance, the recurring design principles are explicit instructions, documented context, appropriate evaluation and escalation when required information is missing.

CLEAR INPUTS

INSTRUCTIONS REDUCE AMBIGUITY

OpenAI's agent guide says high-quality instructions improve decision-making and recommends explicit actions and branches for incomplete information. [1]

CONTEXT OF USE

DOCUMENT THE EXPECTED SETTING

NIST recommends defining acceptable context, assumptions, limitations, value and the operational environment. [2]

SIMPLE FIRST

ADD COMPLEXITY ONLY WHEN NEEDED

Anthropic recommends simple, composable patterns and evaluation before more complex agentic systems. [3]

The gap this test addresses

Guidance identifies the ingredients. This controlled test estimates the operational penalty when one ingredient is deliberately removed from a marketing workflow.

The 30-run stress test

Five common B2B marketing tasks were run under one complete condition and five single-omission conditions. Model, task, approved facts and scoring rule stayed constant unless that item was the tested omission.

Condition	What changed	Runs
Complete	All required inputs supplied	5
No audience	Buyer role and knowledge level removed	5
No goal	Business outcome and intended decision removed	5
No source	Approved factual evidence removed	5
No format	Length, channel and structure removed	5
No owner	Approval and escalation owner removed	5

Four outcome measures

QUESTIONS

Explicit requests for missing information.

ASSUMPTIONS

Unapproved choices introduced by the system.

FIRST-PASS ACCEPT





Meets all criteria before correction.

TIME TO ACCEPT

Active minutes through clarification and revision.

First-pass acceptance collapsed when evidence was absent

Complete context passed in all five tasks. Missing an approved source produced zero first-pass accepts; missing a goal produced one.




8 of 25 incomplete briefs passed first time

This is a 32% first-pass acceptance rate, versus 100% for complete briefs in the same five tasks.

Missing context more than doubled time to acceptance

The median increased from eight minutes with complete context to 17 minutes across the 25 incomplete runs. Missing sources took longest.





2.1x slower with one missing input

The comparison is 17 minutes versus 8 minutes. It includes clarification and revision through final acceptance.

The system filled empty fields with guesses



Across the 25 incomplete runs, the output introduced 48 unapproved assumptions. Missing sources created the largest assumption load.



What counted as an assumption
A material audience, claim, objective, format or approval choice not supplied by the brief and not supported by the approved context.

Clarification is a feature - if the workflow allows it

Incomplete conditions generated 30 questions. Missing goals generated the most. A workflow that forbids questions converts uncertainty into silent assumptions.



Readiness rule

If a required field is missing, the system should ask, pause or escalate. Proceeding should be an explicit risk decision, not the default.

Not all missing inputs have the same cost

Source and goal omissions caused the largest combined damage across acceptance, assumptions and time. Format was the least harmful - but still measurable.

Missing input	1st-pass accept	Assumptions	Questions	Median time	Severity
Approved source	0%	14	7	24 min	Critical
Goal	20%	11	8	21 min	Critical
Audience	40%	9	6	17 min	High
Owner	40%	8	5	16 min	High
Format	60%	6	4	13 min	Moderate

Priority order

Lock evidence and goal first, then audience and owner. Format can often be repaired cheaply; an invented fact or objective can invalidate the entire output.

Analytical tasks were least tolerant of missing context

Insight memos and landing-page sections had no first-pass accepts across the five omission conditions. Social posts passed four of five.

Task	Passes across 5 omissions	Pass rate	Average time	Primary sensitivity
Blog brief	1 / 5	20%	19.4 min	Source + goal
Campaign email	3 / 5	60%	16.2 min	Goal
Insight memo	0 / 5	0%	21.8 min	Source + owner
Social post	4 / 5	80%	14.0 min	Source
Landing section	0 / 5	0%	18.2 min	Goal + audience

The design implication

Readiness should be assessed at the task level. A workflow may be ready for formatting and social adaptation while remaining unready for evidence-heavy synthesis.

A missing source changes the job

The same instruction can look complete while asking the system to invent its evidence base.

SOURCE REMOVED

“Write a landing-page proof section showing why coordinated AI agents outperform a general chatbot.”

Observed behavior: the output supplied plausible percentages and generalized customer outcomes that were not authorized.



SOURCE RESTORED

“Use only the attached 96-position task test. State the sample, report 22/24 versus 8/24, and label the test as protocol-based.”

Observed behavior: the output used bounded evidence and preserved the study limitation.

The model did not become smarter. The workflow gave it an approved evidence boundary.

The six-field AI readiness check

A workflow is ready to test when all six fields are explicit enough that a system can act, stop and be evaluated.



AUDIENCE

Who receives it? What do they already know?



GOAL

What business decision or behavior should change?



SOURCE

Which facts are approved, current and traceable?



FORMAT

Channel, length, structure and technical constraints.



OWNER

Who approves, rejects or resolves ambiguity?



ACCEPTANCE

What observable conditions mean the work is done?

Stop rule

If source, goal or owner is missing for a consequential task, pause before generation or route to a named human.

How to make a workflow AI-ready

Readiness is a design exercise, not a prompt-writing trick. Fix the information and decision architecture before adding more agents.

MAKE FIELDS REQUIRED

Prevent execution when critical context is blank.

USE APPROVED SOURCE PACKS

Version evidence and record retrieval dates.

ADD ASK / PAUSE / STOP

Define behavior for incomplete or conflicting inputs.

NAME THE DECISION OWNER

Separate content production from business approval.

TEST ONE OMISSION AT A TIME

Use controlled failure tests before deployment.

MEASURE ACCEPTANCE

Track questions, assumptions, retries and time to accepted work.

Readiness metric

First-pass acceptance under realistic missing-input tests - not demo success on the perfect prompt.

The 30-run result register

Each cell shows first-pass result and time to final acceptance. P = pass; F = fail. The complete condition passed in all five tasks.

Condition	Blog brief	Campaign email	Insight memo	Social post	Landing section
Complete	P / 9m	P / 7m	P / 10m	P / 6m	P / 8m
No audience	F / 18m	P / 15m	F / 20m	P / 13m	F / 17m
No goal	F / 22m	F / 20m	F / 25m	P / 16m	F / 21m
No source	F / 26m	F / 21m	F / 28m	F / 19m	F / 24m
No format	P / 14m	P / 11m	F / 17m	P / 10m	F / 13m
No owner	F / 17m	P / 14m	F / 19m	P / 12m	F / 16m

Condition totals

Condition	Questions	Assumptions	Passes	Median
Complete	1	2	5/5	8 min
No audience	6	9	2/5	17 min
No goal	8	11	1/5	21 min
No source	7	14	0/5	24 min
No format	4	6	3/5	13 min
No owner	5	8	2/5	16 min

Limit: one author coded the outputs. The test set is small and synthetic; omissions were deliberately created and are not prevalence estimates.

Sources, scope and reuse notes

Reviewed in September 2026. External sources frame the method; all 30-run outcomes are original Teamulate analysis from the controlled test described here.

1. [OpenAI. A Practical Guide to Building AI Agents.](#)
2. [NIST. Artificial Intelligence Risk Management Framework: Generative AI Profile. NIST AI 600-1 \(2024\).](#)
3. [Anthropic. Building Effective Agents \(2024\).](#)
4. [Microsoft. 2026 Work Trend Index: Agents, human agency, and opportunity.](#)
5. [Lee, H.-P. et al. The Impact of Generative AI on Critical Thinking. CHI 2025.](#)
6. [Dell'Acqua, F. et al. Navigating the Jagged Technological Frontier. SSRN 4573321 \(2023\).](#)
7. [Noy, S. and Zhang, W. Experimental Evidence on the Productivity Effects of Generative AI. Science \(2023\).](#)
8. [Teamulate. The Hidden Cost of Fixing AI Marketing Work. Research Topic 07 \(2026\).](#)

Definitions

First-pass acceptance: all task criteria met before clarification or correction. Assumption: a material choice not supplied or supported.
Time to acceptance: active minutes through clarification, revision and final approval.

Interpretation

The experiment estimates the effect of these designed omissions in this test - not their frequency in real organizations.

Reuse notes

When quoting results, preserve the five-task sample, single-author coding and deliberately created omission conditions. Cite original publishers for external evidence.

Suggested citation: Teamulate (2026). What Makes a Marketing Workflow Ready for AI? Evidence Report.