stellar

AI Content Production
Workflow Readiness Checklist

Use this checklist to audit your team's AI content workflow before scaling or to identify where your current process is breaking down. Make a copy and work through it with your team.

1. Workflow Scope & Stages

If you haven't mapped the workflow, the rest of this checklist will produce false confidence.

- ☐ The workflow is documented *(planning → brief → draft → edit → QA → publish → distribution)*
- ☐ Each stage has an owner and a "definition of done"
- ☐ Handoffs are explicit *(what gets passed forward, in what format, where it lives)*
- ☐ You've identified the highest risk stages *(usually briefing, QA, and pre-publish)*
- ☐ Exceptions and conditions are defined

2. Roles & Decision Ownership

Roles that blur or collapse are one of the most common workflow failure points.

- ☐ It's clear who is responsible for deciding what gets drafted
- ☐ Editing and QA are defined as separate functions
- ☐ Ownership of each quality dimension is assigned to a specific role
- ☐ There is a defined escalation path when AI output doesn't meet standards
- ☐ Sign-off requirements are documented for each content type
- ☐ Role compression is intentional and trained for, not assumed
- ☐ AI operators are trained to use AI intentionally

3. Where AI Is Applied

Teams focused only on drafting are concentrating risk and underutilizing AI's value.

- ☐ AI supports article-level planning *(SERP analysis, content gaps, brief inputs)*
- ☐ AI is used for editing and optimization support *(clarity, concision, SEO alignment)*
- ☐ AI identifies visual opportunities within content
- ☐ AI assists with publishing requirements *(metadata, schema, linking, distribution tasks)*

4. Inputs & Briefs

The most predictable failure point: using a human brief as an AI input.

- ☐ Briefs are written specifically for AI, not adapted from human writer briefs
- ☐ Briefs include explicit constraints, not assumed context
- ☐ Relevant examples are attached or referenced in the brief
- ☐ Briefs specify what to include and what to avoid
- ☐ Tone and voice requirements are defined precisely *(not "make it engaging")*
- ☐ Structural requirements are documented at the brief stage, not discovered in review

5. Consistency Standards

Standardize where rigid consistency is required. Let it flex where it isn't.

- ☐ Brand voice requirements are precisely defined and documented
- ☐ High-consistency content types have structural templates *(e.g. location pages, product descriptions)*
- ☐ Review criteria are written down, not implied
- ☐ Outlines specify brand positions and priorities, not just structure
- ☐ You've identified which elements must be consistent vs. which can vary
- ☐ Over-standardization in low-risk areas has been evaluated and trimmed

6. Prompting Standards

Unstandardized prompting is a direct cause of output variability and brand drift.

- ☐ Shared prompt templates exist for recurring content types
- ☐ Prompts encode brand voice, not just task instructions
- ☐ Prompts are version-controlled or stored in a shared location
- ☐ Team members are not each writing their own prompts from scratch
- ☐ Prompt performance is reviewed when output quality drops

7. QA & Evaluation

QA validates against defined standards. Editing improves a piece. Conflating them creates problems at scale.

- ☐ QA standards are documented, not held in someone's head
- ☐ AI telltales are documented as a reference for evaluators
- ☐ Evaluators use a defined checklist or rubric during review
- ☐ Evaluators are trained to catch AI-specific failure modes *(factual drift, voice inconsistency, plausible-sounding errors)*
- ☐ Content does not move forward without passing a defined QA gate
- ☐ Automation bias is accounted for in review design *("looks good" is not a sufficient pass criterion)*

8. Human Oversight

AI doesn't remove humans from the workflow. It changes what they do.

- ☐ There are defined points where human judgment overrides AI output
- ☐ Accuracy and factual claims are verified by a human with domain knowledge
- ☐ Strategic alignment is assessed by someone with context, not just surface editing
- ☐ AI inaccuracy consequences are accounted for in your workflow design

9. Measurement & Feedback Loops

Without measurement, "standards" don't improve. They just get argued about.

- ☐ You track QA outcomes by content type, including pass rate and top failure reasons
- ☐ You track rework, including revision rounds and time-to-approve
- ☐ You track post-publish fixes, including corrections and compliance issues
- ☐ Recurring failures are categorized and used to update briefs, prompts, and QA criteria
- ☐ Workflow performance is reviewed on a defined cadence
- ☐ "Definition of done" criteria are revised when the same downstream issues repeat