Prompt Engineering Checklist 2026

A one-page pre-flight check for any prompt going into production — from BestPrompt.art

BEFORE YOU WRITE IT

- ■ Confirmed whether the model has a native reasoning / thinking mode
- ■ Named the exact audience and desired output format
- ■ Decided zero-shot vs. few-shot based on task type, not habit
- ■ Checked whether this task actually needs examples at all

STRUCTURE

- ■ Role is light-touch, not an elaborate persona
- ■ XML tags used only if the prompt is genuinely complex
- ■ Constraints section: format, length, audience, tone all explicit
- ■ "Say I don't know" permission line included

BEFORE YOU SHIP IT

- ■ Ran the same prompt 5–10 times to check output consistency
- ■ Traced every quantitative claim in a template to a real source
- ■ Voice sample included if brand tone matching matters
- ■ Considered whether this belongs in a context file, not just a prompt

SECURITY & SCALE

- ■ Reviewed for prompt injection if the model reads external content
- ■ Tool permissions scoped to least privilege, not broad access
- ■ Documented in a shared prompt library with known failure modes
- ■ Re-tested after any model upgrade — don't assume portability

Companion resource to “Prompt Engineering in 2026: The Complete Guide” — bestprompt.art. Free to use and share internally. Last updated August 20, 2026.