

# Prompt Review Checklist

This is a pre-send check for a single prompt, run once in the seconds before you hit enter. It is not the prompt evaluation checklist, which judges outputs across many runs against a rubric. The review asks one question: does this prompt give the assistant enough role, context, constraints, examples, and a testable definition of done to return something you can actually use on the first try?

## Role and intent

- The prompt names the role the assistant should take, even if it is one sentence.
- The goal is stated as a concrete definition of done, not a vague ask.
- The single most important outcome is obvious from the first two lines.

## Context

- Every fact the model cannot infer is included: stack, file path, framework, data shape.
- The relevant constraint from the wider system is present, such as a schema that must stay stable.
- Background that does not change the answer has been cut so the key facts stand out.
- When a long file or log is pasted, the actual instruction sits at the end, below it.

## Constraints and format

- Limits are written as explicit rules, not left implied in your head.
- The exact output shape is named: a diff, a JSON object with set keys, a table, a numbered list.
- One short worked example is included when the format is easy to get wrong.
- Instructions say what to do rather than only what to avoid.

## Definition of done

- Success is something you could check or write a test for, not a mood like "make it good".
- The prompt tells the model to ask or flag assumptions when a detail is missing.
- You would be willing to bet the first response is usable; if not, find the one weak line and fix it.