

Applying the Practical Delegation Test

A short guide for deciding when to use AI, how far to delegate, and what still needs human judgement

What this guide is for

AI can often produce plausible work very quickly. That can be useful, but speed is not the same as dependable work.

Before you hand a task to AI, or before you pass an AI-generated output to someone else, you need to check four things: whether the task is clear enough, whether the result can be reviewed properly, what judgement still belongs to a person, and where any remaining uncertainty will go.

The Practical Delegation Test is a quick way to make that decision. It is not a scoring model, and it does not replace the formal rules or approval routes that apply to your work. It is a practical check to use at the point of work, especially when someone else may rely on the output.

Use it when you are deciding whether to ask AI to do a real piece of work rather than simply help you think.

Common examples include:

- draft or rewrite a message, report, policy, user story, meeting note, or brief;
- analyse data, summarise evidence, or compare options;
- write or change code, tests, queries, pipelines, or configuration;
- classify, triage, prioritise, or route work;
- generate a plan, checklist, decision note, or set of recommendations.

You do not need to run the full test for every tiny autocomplete suggestion or low-risk wording change. Use it when a wrong answer, weak assumption, poor handoff, or misleadingly polished output would create real work or real risk.

If AI is doing several steps in sequence, apply the test at the meaningful handoff points rather than only checking the final output.

The four questions

Ask these four questions before deciding how far to delegate the task.

Question	What you are checking	If the answer is weak
Can I describe the outcome clearly?	Whether you know what good looks like	Use AI to explore or clarify, not to execute
Can I verify the result cheaply enough?	Whether checking the answer is realistic and proportionate	Add checks, reduce delegation, or redesign the task
What judgement must remain human?	Which decisions, trade-offs, risks, or commitments cannot be handed off	Keep the decision and acceptance visibly human
Where does the uncertainty go?	Whether doubt is being handled properly or pushed downstream	Stop, name the uncertainty, and fix the handoff

The aim is not to delegate as much as possible. The aim is to delegate only where the combined human-and-AI workflow can produce a result that someone can responsibly use.

How to apply the test

1. Can I describe the outcome clearly?

Write down the actual task in one sentence. Avoid vague requests such as "look into this", "make this better", or "produce something useful". A task is easier to delegate when the output, audience, constraints, and purpose are visible.

Useful prompts:

- What output do I want?
- Who will use it?
- What would make it acceptable?
- What would make it unusable?
- What facts, policies, formats, or exclusions matter?

If you cannot answer those questions, AI may still help, but its role should be exploratory. Ask it to help clarify options, surface questions, or organise your thinking before you ask it to produce something final.

2. Can I verify the result cheaply enough?

Before generating the output, decide how you will know whether it is good enough. This is the most commonly missed step. "Cheaply enough" does not mean minimal effort; it means the checking cost is visible, proportionate to what is at stake, and worth paying. Delegation may still be worthwhile when review takes serious effort, especially if the task would have taken much longer to complete manually. The problem is when the checking burden is hidden, unrealistic, or needs expertise you do not have.

Useful prompts:

- Can I test the output?
- Can I compare it with a known source, standard, example, or expected result?
- Would a likely error be obvious, or could it be hidden inside a plausible answer?
- Who is qualified to review it?
- If it is wrong, how easy is it to fix later?

Examples of stronger verification include running tests, checking a query result against source data, comparing a summary with the original document, using a checklist, reviewing against policy, or asking a subject-matter expert to check a clearly bounded point.

3. What judgement must remain human?

AI can help produce and organise work, but it should not silently make the decision underneath it. A person still needs to decide what the organisation believes, what risk is acceptable, and what commitments can legitimately be made.

Before using the output, name the human judgement that remains.

Useful prompts:

- What decision am I still responsible for?
- What risk or trade-off cannot be delegated?
- What commitment would this output make if someone acted on it?
- Who has authority to accept, send, publish, merge, approve, or rely on it?
- Would I be comfortable explaining how I checked it?

This matters even for ordinary work. A customer response may look well written but still make the wrong commitment. A code change may compile but introduce the wrong assumption. A summary may be fluent but omit the one caveat that changes the decision.

4. Where does the uncertainty go?

If you are not sure whether the result is correct, do not hide that uncertainty under polished wording. Decide who needs to know, who needs to check, and whether the work should stop until the uncertainty is resolved.

Useful prompts:

- What am I still unsure about?
- Have I named that uncertainty for the next person?
- Am I passing work onward because it is ready, or because it looks plausible?
- Will a reviewer have enough time, context, and authority to challenge it?
- Is this a clean handoff or just a shifted burden?

Good delegation makes remaining uncertainty visible. Poor delegation sends uncertainty downstream and leaves someone else to discover it under pressure.

Better handoff	Weak handoff
Review is expected, scoped, and resourced	Review is implied, improvised, or left for the recipient to discover
Remaining uncertainty is named clearly	Remaining uncertainty is hidden by polish or confidence

Choose the right level of delegation

After the four questions, choose one of four levels.

Level	Use when	Your responsibility
Direct delegation	The task is clear, checking is simple, consequences are low or reversible, and little judgement is being handed off	Specify the task, check the result, and accept the outcome
Checkpointed delegation	AI can do substantial work, but assumptions or transitions need review	Review checkpoints, test assumptions, and decide what survives
Human-led with AI support	The task contains ambiguity, judgement, sensitivity, or costly verification	Keep control of substance and use AI as support
Do not delegate	The task is too unclear, too sensitive, too hard to verify, or likely to push uncertainty onto others	Clarify, redesign, escalate, or do the work directly

For example, asking AI to draft unit tests for a small, well-understood function may be direct delegation if you can run and inspect the tests. Asking AI to draft a response to a complaint should usually be human-led with AI support, because the facts, tone, commitments, and policy position still need human judgement. Asking AI to rank employees for redundancy should not be delegated as a decision, even if AI could produce a ranked list.

Work that needs extra care

Be cautious when the task involves:


- customers, residents, patients, employees, or vulnerable people;
- legal, regulatory, HR, safeguarding, finance, security, or safety consequences;
- decisions that affect access to services, support, money, employment, or rights;
- confidential or sensitive information;
- public statements, complaints, disciplinary matters, or contractual commitments;
- code, data, or automation that could affect production systems;
- summaries that busy decision-makers may rely on without reading the source.

These tasks may still use AI, but usually as support rather than as a delegate. Keep a human clearly responsible for facts, judgement, approval, and the final decision to use the output.

A one-minute version

Before using AI for a real task, ask:

1. Can I describe the outcome clearly?
2. Can I verify the result cheaply enough?
3. What judgement must remain human?
4. Where does the uncertainty go?



Use the test to make delegation clear enough to check, own, and hand off responsibly.

If the answers are strong, delegation may be appropriate. If one answer is weak, reduce the level of delegation and add checks. If several answers are weak, use AI only for exploration or do not delegate the task.

The practical test is simple: do not ask AI to produce work faster than you can responsibly understand, check, and stand behind it.