

Should We Use This AI Tool?

A simple checklist for non-profits.

You do not need to know everything about AI to make a sensible decision about a tool.

Start with the work, the people affected and the level of risk. Then decide whether the tool is worth testing, needs more thought or is not right for this job.

What problem are we solving?

Describe the problem in one sentence.

Is it a real pressure on the team? Could a clearer process, training, a template or a non-AI tool solve it just as well? AI should add enough value to justify the time, money and oversight it needs.

Who benefits, and who could be affected?

Think about staff, volunteers, funders, service users, donors and partners.

If the tool makes a mistake, who could be harmed or treated unfairly? Take extra care where it might affect someone's safety, rights, access to support, employment, funding or reputation.

What information will it use?

List the information the tool will need. This could include documents, emails, records, personal data, special category data or confidential material.

Only put information into a tool when you have a clear reason, the right safeguards and permission to do so. A public AI tool is not the place for personal, confidential, identifying or safeguarding information.

What does the provider do with our information?

Find out what happens to your data, prompts and outputs.

Ask whether the provider uses information to train its models, where information is stored, how long it is kept, who can access it, and whether your organisation can control those settings.

What will the tool be allowed to do?

Be specific about whether the tool will only generate a draft or whether it will connect to systems, search files, update records, send messages or take actions.

Keep permissions proportionate to the task. Do not give a tool broad access simply because it is easier.

What human checking is needed?

Decide who checks the output, what they need to check and when a person must make the final decision.

The more a tool can affect people, money, services or your organisation's reputation, the more meaningful the human involvement needs to be.

What will it really cost?

Look beyond the subscription price. Consider staff time, training, set-up, support, monitoring, accessibility, reviewing outputs and the work involved if the tool goes wrong.

A tool that creates more checking than it saves may not be the right answer.

Can we get out again?

Ask whether you can export your information, stop using the tool, remove access and move to another system without causing a problem.

You should not be locked into a tool you no longer trust or need.

How will we test it?

Start with one low-risk use. Use a small group, a clear time limit and a simple way to decide whether it has helped.

Keep a note of what worked, what did not, what mistakes it made and what needs to change before you go further.

Make a decision.

Proceed when the purpose is clear, the benefits are proportionate and the risks are understood and managed.

Pilot when the tool may help, but you still need evidence about how it works in your charity.

Pause when you cannot answer the questions clearly, the risks are too high, or AI is not actually the right answer.

This guide is a practical starting point, not legal, data protection or procurement advice. Seek specialist advice where your organisation's circumstances require it.