



Scenario 1 of 10

The intake summary

A caseworker pastes a new client's intake notes into a free AI chat tool to get a quick summary before a meeting. The notes include the client's name, immigration status and a safety concern.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 2 of 10

The donor thank-you

The development lead uses an approved AI tool to draft 40 thank-you letters. They add each donor's first name and gift amount, then send them all at once.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 3 of 10

The grant draft

The executive director asks AI to write the 'impact' section of a grant application. It produces strong numbers: '85% of participants found stable housing.' Nobody is sure where the figure came from.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 4 of 10

Translating a flyer

A volunteer uses AI to translate a program flyer into Arabic and French. Nobody on the team reads Arabic. The flyer goes up in the community centre.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 5 of 10

The note-taker at a case meeting

A partner agency joins a case-planning meeting and turns on an AI note-taker. The meeting covers a family's safety plan.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 6 of 10

Volunteer screening

A company offers a tool that uses AI to score volunteer applications and rank the top 20. It would save hours each intake round.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 7 of 10

The payment email

The finance lead gets an urgent email that looks like it's from the executive director: 'I'm in a meeting. Please pay this invoice today, \$4,800.' Then a voicemail in the executive director's voice says the same thing.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 8 of 10

The youth program photo

Comms wants a bright photo of kids in your youth program for social media. Someone suggests using AI to make one instead of asking families for consent.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 9 of 10

The feature nobody switched on

Your case-management system sends a notice: 'New! AI assistant now summarizes every client file.' It's already turned on.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?



Scenario 10 of 10

The board member's shortcut

A board member pastes the draft budget and a staff salary list into their personal ChatGPT to 'double-check the numbers' before the finance meeting.

Talk about it: Would we let AI do this? What could go wrong? What rule should we have?

Answer notes for the person leading the session

Read these after the group has talked. They are a starting point. Your team's answer can be stricter.

1. The intake summary

Red light. This is personal and sensitive information, and it went into a personal AI account. Stop, and tell the AI lead. If summaries would really help, the Systems version shows a safer way: an approved automation with a written privacy check, inside your own systems.

See the policy: Sections 6, 7 and 8

2. The donor thank-you

Mostly fine, with care. First names and gift amounts are personal information, so leave them out of the AI tool and add them after. A person reads every letter before it goes. Add the disclosure line to your email signature.

See the policy: Sections 2, 7 and 12

3. The grant draft

Never submit numbers you can't trace. AI can invent facts that sound right. Use AI to shape the writing, but every figure must come from your own records, and a person checks each one.

See the policy: Sections 2 and 6

4. Translating a flyer

Check before posting. A fluent person must review anything AI translates for the people you serve. A wrong word on a flyer can send someone to the wrong place, or tell them they don't qualify.

See the policy: Section 11

5. The note-taker at a case meeting

No. Politely ask them to turn it off, and take notes by hand. Meetings with private or safety information don't use AI note-takers. For other meetings, everyone must agree first.

See the policy: Section 7

6. Volunteer screening

AI can help sort, but a person decides. Before using anything like this, the AI lead checks it for bias: test it on past applications and look at whether some groups are ranked lower. If you can't see why it ranks people the way it does, don't use it for decisions.

See the policy: Sections 6 and 9

7. The payment email

Call back first, on a number you already know. AI can now fake emails and voices convincingly. Never act on a request for money, personal information or a password until you've confirmed it directly.

See the policy: Section 15 (Systems) or 14 (Starter)

8. The youth program photo

Don't show AI-made people as if they're real participants. Either get consent for a real photo, or use an image that's clearly labelled 'Image made with AI' and doesn't pretend to be your program.

See the policy: Sections 6 and 12

9. The feature nobody switched on

Treat it like a new tool. Turn it off until the AI lead has sent the Vendor Question Sheet and, if it reads personal information, done a Privacy Check. Then decide.

See the policy: Sections 4 and 13

10. The board member's shortcut

The policy covers board members too. Salaries are personal information, and a personal account isn't approved. Thank them for the care, share the Staff Card, and have the AI lead suggest an approved way to check the numbers.

See the policy: Sections 1, 4 and 7
