

From Frameworks to Implementation

Putting the frameworks to work in your organisation

TODAY'S AGENDA

- 1 Welcome Back**
What you told us and what changed

- 2 Building Your AI Policy**
The one-page template, data classification, shadow AI

- 3 Working effectively with data**
Live data pipeline, donor analysis, and AI-assisted CRM prototyping

- 4 Board Governance & Ethics**
Trustee responsibilities and ethical boundaries

- 5 Tools & Next Steps**
Copilot vs Claude, pricing tiers, specialist tools, and your action plan



THE FIVE FRAMEWORKS

RACE Role · Action · Context · Expectation
Structures your prompts so AI knows who you are and what you need.

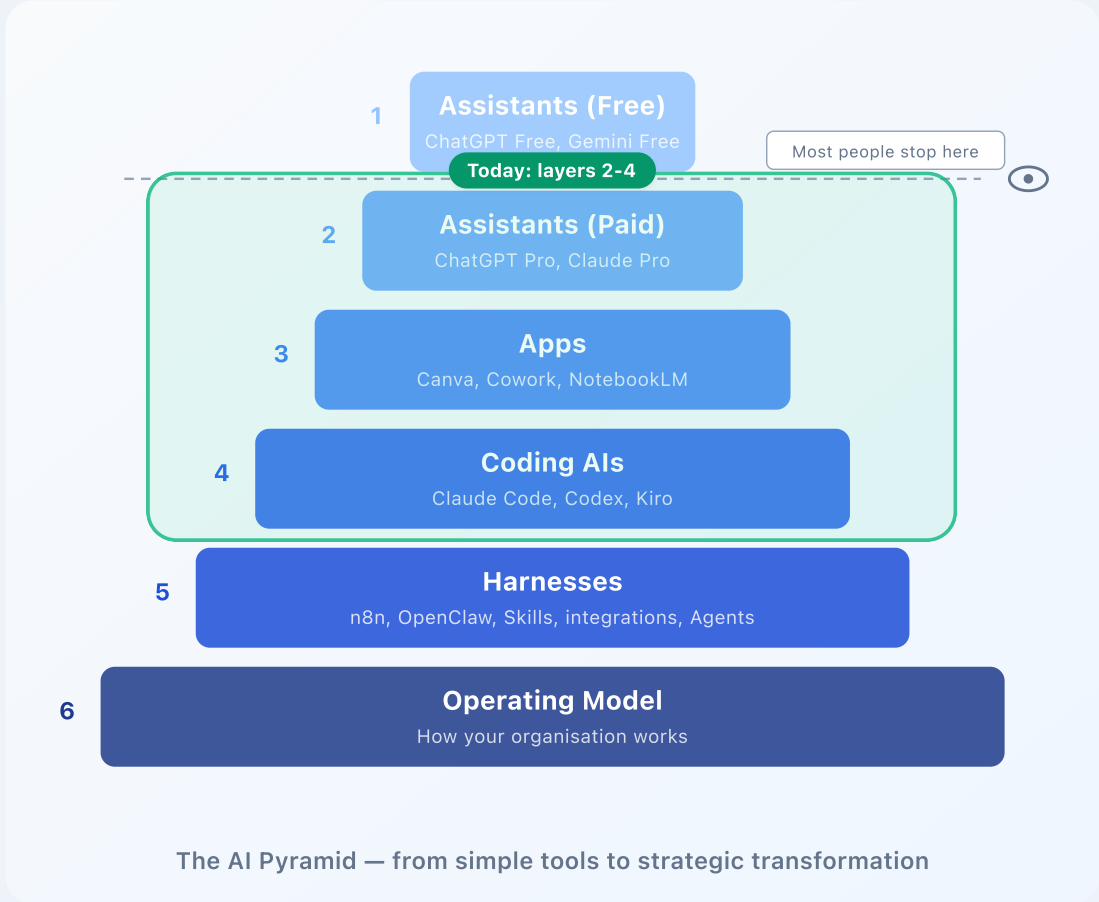
GUARD Ground · Underwrite · Align · Resist · Defend
Your five safety boundaries for daily AI use.

CRAFT Collect · Refine · Analyse · Fact-check · Transform
Turns raw data into insight using AI as your analyst.

Triple T Trust · Training · Tooling
The three organisational pillars of sustainable AI adoption.

REACH™ Recognition · Engagement · Acumen · Competency · Habit
Your personal AI leadership maturity journey.

THE AI PYRAMID



TAKEAWAY Five frameworks, one vocabulary. Today we move from understanding them to applying them.

Session 1 Feedback

WELCOME BACK

WHAT YOU TOLD US

29 Written feedback responses	82% Rated the session 5/5
16 Email addresses for AILR trial	9 Requested an AI policy template

SESSION 2 TOPIC REQUESTS



BUILDING ON, NOT OVER

THE STANDOUT FEEDBACK
"Excellent presentation Mal you took some fear out of AI, loved the frameworks used, thank you." - Session 1 attendee, 13 May 2026

SESSION 2 ASSUMES FLUENCY
We build on all five frameworks today, not over them. If anything feels unclear, the Session 1 deck is available at the link shared in the chat.

WHAT CHANGES TODAY
Session 1 gave you the frameworks. Session 2 puts them to work on real problems your organisations face: writing policy, processing data, governing AI at board level, and choosing tools with confidence.

SESSION 1 RECORDING & RESOURCES
thirdsectorlab.co.uk: AI Leadership Readiness for Charity Leaders
YouTube: Session 1 recording

TAKEAWAY Today's session was designed from the ground up using your responses. Every section maps directly to what you asked for.

Another poll before we begin

WELCOME BACK

WHAT WE'RE LOOKING AT

No wrong answers. Your responses help calibrate the depth and pace of today's session.

How do you feel about AI right now?

- Word cloud to capture the emotional temperature in the room, helpful to see if sentiment has shifted since Session 1.


Do you have an AI Policy at your charity?

- Establishes the baseline governance position, directly relevant to what we cover in Section 2 today.

Which of these tools do you use for data analysis?

- Helps calibrate the live data pipeline demo in Section 3, we'll tailor examples to the tools already in use.

SLIDO #3485150



If the wall stays blank, your browser may block embedding — open the same Slido link in a window or share your screen instead.

TAKEAWAY *We start from real answers in the room so the next section lands at the right depth for everyone here.*

Building Your AI Policy

The number one request. A live, worked policy template you can take back today.

Why a Policy

The One-Page Template

Data Classification

Shadow AI to Agreed Systems

Why Every Charity Needs an AI Policy Now

THE PROBLEM

43%

Cite data privacy as top AI barrier (CDSR 2025)

22%

Have reviewed governance for trustee oversight of AI

CDSR 2025, P.15

"Data privacy, GDPR and security concerns (43%) are up from 31% last year. This rises to 56% of large charities, who say this is the biggest barrier they face."

WHY YOU NEED AN AI POLICY

Most charities are already using AI. Staff are using free tools with organisational content without guidance, without boundaries, without visibility. This is **shadow AI** and it carries real governance and reputational risk. An AI policy is not a compliance exercise: it is part of your organisation's infrastructure.

WHAT A POLICY DOES

- Gives staff clear permission and boundaries, removing anxiety and guesswork
- Creates accountability for which tools are approved and under what conditions
- Shows trustees that adoption is managed, not accidental
- Earns trust to move faster, not a brake on adoption

AI POLICY MATURITY STAGES

WHICH STAGE DESCRIBES YOUR CHARITY? (SLIDO POLL)

Stage 1

No policy

Staff use whatever works, no visibility

Stage 2

Informal norms

Some shared understanding, nothing written

Stage 3

Documented policy

Written guidance, approved tools, data classification

Stage 4

Reviewed & ratified

Trustees signed off, review cycle running

CDSR 2025, P.14

Only 22% of charities have reviewed governance for trustee oversight of AI. Most are at Stage 1 or 2.

THE GOOD NEWS

A policy does not have to be long to be effective. The goal is clarity, not compliance theatre.

TAKEAWAY A one-page AI policy is the single most impactful governance document a charity can produce in 2026.

THE SIX POLICY SECTIONS

- 1 Purpose**
Why this policy exists and what it covers
- 2 Approved Tools**
Which AI tools are sanctioned for which use cases (free vs paid tiers); AI notetakers require advance consent from all meeting participants
- 3 Data Classification**
What can and cannot be entered into AI tools (green / amber / red)
- 4 Verification Requirements**
What must be checked before AI output is used or shared
- 5 Trustee Visibility & Sign-Off**
Who approves new tools, who reviews incidents, and trustee oversight
- 6 Review Cycle**
When and how the policy is reviewed (recommend: every 6 months given pace of change)

WORKED EXAMPLE: BRIGHTHOPE

SECTION 2: APPROVED TOOLS

- Included** **Microsoft Copilot (M365)**
Grant drafting, internal comms; meeting notes only with advance consent from all participants
- Free** **ChatGPT free tier**
Personal skill development only, no organisational content
- Paid** **ChatGPT Plus / Claude Pro**
Grant applications and report analysis with anonymised data

SECTION 4: VERIFICATION REQUIREMENT

"All AI-drafted funding applications must be reviewed and edited by the named grant writer before submission."

AI NOTETAKERS: ADVANCE CONSENT REQUIRED

Any AI tool that records, transcribes, or summarises a meeting must be disclosed to all participants before it starts. Everyone must be informed in advance. Silence is not consent. Charity meetings often involve beneficiary disclosures, staff concerns, or sensitive funding conversations. Your policy must name this explicitly.

START HERE

This is a starting point, not a final document. Adapt it to your organisation.

THE THREE TIERS

● Green Share freely

- Public information, published reports, anonymised aggregates
- Synthetic or fictional data
- Your own published content

Can enter into any tool at any tier

● Amber Handle with care

- Aggregated data with appropriate consent
- Internal documents that do not name individuals
- De-identified case summaries

Paid personal tier minimum (£15–20/month), no free tools

● Red Never in a free tool

- Named beneficiaries, case notes, addresses, contact details
- Staff personal data, HR records
- Financial records identifying individuals, donor details

Enterprise/team tier only, or not at all

THE DECISION RULE

DECISION TREE

→ Does this data name a person or identify a beneficiary?

YES: RED. Do not enter into any AI tool unless on an enterprise plan with a DPA.

→ Does it contain aggregated organisational data?

YES: AMBER. Paid tier minimum.

→ Is it public, synthetic, or your own published content?

YES: GREEN. Any tool.

THE GUARD CONNECTION

This classification system is the operationalisation of GUARD's Defend boundary. The framework told you what to do. This tells you how.

WHEN IN DOUBT

Go up a tier. It is always safer to treat data as more sensitive than you think it is.

TAKEAWAY Green / Amber / Red. Three tiers, one rule: when in doubt, go up a tier.

GETTING TO STAGE 3 THIS MONTH

- 1 Audit which AI tools your team is currently using (a quick survey takes 10 minutes)
- 2 Decide on your approved tools list: start with what people already use
- 3 Apply the Green/Amber/Red classification to your three most common data types
- 4 Write a one-paragraph purpose statement
- 5 Share a draft with your senior leadership team before the end of June

THE TRUSTEE MOMENT

Boards do not need to approve every prompt. They need to know a policy exists, who owns it, and when it will be reviewed.

LIVE DEMO MOMENT 1

We will try this live: create a first-draft AI policy for BrightHope, then review what the AI gets right, what it misses, and what a trustee-ready version would need.

WORKSHOP: BRIGHTHOPE AI POLICY

RACE PROMPT: TRY THIS NOW

[R] You are a charity governance specialist with expertise in AI policy for small UK charities.

[A] Write a one-page AI policy for BrightHope.

[C] BrightHope is a small youth employment charity with 8 staff. They currently use Microsoft Copilot (M365, included in their subscription) for meeting notes and email drafting, ChatGPT free tier for personal skill development only, and ChatGPT Plus / Claude Pro for grant applications and report analysis with anonymised data. They use Granola as their AI notetaker but require that all meeting participants are informed in advance before Granola is activated – given the sensitivity of conversations involving young people and employment circumstances. They apply a Green/Amber/Red data classification: beneficiary names and case notes are Red and never enter free tools; internal documents are Amber; published content is Green. They are moving from informal norms to a documented policy.

[E] One page maximum. Six sections: Purpose, Approved Tools, Data Classification, Verification Requirements, Trustee Visibility and Sign-Off, Review Cycle. Plain English, no jargon. Suitable for sharing with trustees.

WHAT TO DO WITH THE OUTPUT

Review, edit, and adapt. The AI drafts the structure, you supply the accuracy. Run it through GUARD before sharing with your board.

TAKEAWAY Most charities are at Stage 1 or 2. Moving to Stage 3 takes a half-day workshop and one trustee agenda item, not a consultant.

Working effectively with data using CRAFT

From messy spreadsheet to usable insight. Building a simple CRM
with synthetic donor data.


CRAFT Recap


Collect & Refine


Analyse & Visualise


Lightweight CRM


THE PIPELINE

C  **Collect** Bring your data together in one place.
Export your spreadsheet donor list, beneficiary contact log, and last year's grant report into one document.

R  **Refine** Clean, standardise, and de-duplicate.
Prompt: "Identify duplicate rows, standardise date formats, and flag blank required fields."

A  **Analyse** Ask questions of your data.
"Which donor segments gave the most in the last 12 months? What's the average gift by acquisition channel?"

F  **Fact-check** Verify AI-generated numbers against source data.
Never publish AI-calculated figures without checking one row manually.

T  **Transform** Convert output into the format you need.
Produce a board-ready summary table, a funder impact report, or a segmented email list.

THE MOTIVATING PROBLEM

45%

Say they need support with AI for data analysis

13%

Currently using AI for numerical analysis

CRAFT CLOSSES THIS GAP

It gives you a repeatable pipeline that works whether you have 100 records or 10,000. Each step is a plain-English instruction to your AI assistant.



GROUND TRUTH

Every dataset you clean and verify becomes a ground truth asset. You use it in future board papers, funder reports, and grant applications without re-verifying from scratch. This is organisational memory.

TAKEAWAY CRAFT is not a one-off exercise. It's a repeatable pipeline. Run it once and you have a template for every future data task.

BEFORE: RAW DONOR DATA

COLLECT DONOR DATA: TYPICAL RAW STATE BEFORE CRAFT

Donor ID	Gift Date	Amount	Channel	Notes
001	13/04/26	£120	online	regular
002	April 13	MISSING	post	called twice, left voicemail
003	2026-04-13	£50	event	first-time gift
002	14/04/26	£75	post	DUPLICATE
004	15 Apr 2026	£500	major	board connection
005	16/04/26	£30	MISSING	lapsed reactivation
006	17 Apr 2026	£150	event	MISSING
007	18/04/26	£200	direct	regular
008	19/04/26	£25	post	one-off
009	20 Apr 2026	£80	online	monthly

Issues: mixed date formats, missing gift amounts, missing channel, missing note, duplicate row (highlighted), free text in notes

LIVE DEMO MOMENT 2

I will switch to Claude to show how to build a script to clean up this data.

AFTER: BOARD-READY SEGMENT ANALYSIS

COLLECT: [paste your raw donor CSV data here]

REFINE: Standardise all dates to DD/MM/YYYY. Remove duplicate rows (keep first occurrence of each Donor ID). Exclude rows with a missing Amount.

ANALYSE (apply in order):

1. Major donor: Amount \geq £500
2. Regular (monthly): Notes contain "regular" or "monthly"
3. Lapsed (12m+): Notes contain "lapsed"
4. One-off (recent): all remaining rows

FACT-CHECK your totals against these figures before proceeding:

Lapsed (12m+): £30. Regular (monthly): £120 + £200 + £80 = £400.

One-off (recent): £50 + £150 + £25 = £225. Major donor (>£500): £500.

Total: £30 + £400 + £225 + £500 = £1,155.

TRANSFORM: Produce a table – Segment | Count | Total Giving | Avg Gift. Monetary values prefixed £, averages rounded. Write 200 word Board report.

TRANSFORMED DONOR SUMMARY: Q1 2026

Segment	Count	Total Giving	Avg Gift
Lapsed (12m+)	1	£30	£30
Regular (monthly)	3	£400	£133
One-off (recent)	3	£225	£75
Major donor (£500+)	1	£500	£500

8 usable gifts, £1,155 total, average gift £144. One duplicate removed; one amount missing.

TAKEAWAY Eight minutes from messy spreadsheet to board-ready segment analysis. That's what CRAFT gives you.

THE PROBLEM & WHAT A FUNDRAISING CRM NEEDS TO DO

30%

Running on poor or non-existent systems

500+

Active records: when to graduate to proper CRM

CDSR 2025, P.12

"30% say that their systems and databases (covering everything from internal management tools to client databases and CRMs) are poor or non-existent."

If we extrapolate then potentially 30% of organisations are managing fundraiser data in spreadsheets. For them a fundraising CRM is not a nice-to-have. It is a governance gap. Here is what a fundraising CRM needs to support:

- ✓ Contact unique ID
- ✓ Contact profile
- ✓ Relationship history
- ✓ Pledges / Contributions
- ✓ Giving channel / campaign type
- ✓ Contact Engagement Plan

METAPROMPTING: TRY THIS NOW

I want you to help me build a basic CRM system for a UK charity. It should be populated with synthetic data for around 1000 fundraiser contacts. Research and find out what kind of information such a resource would expect to contain. Build an xlsx spreadsheet with that content in. The spreadsheet must include a Status History sheet that records every donor status transition (Prospect → Active Donor → Lapsed etc.) for each contact, so that conversion rate analysis is possible. Please build out a more detailed prompt I can use to seed the process.

WHAT TO DO WITH THE OUTPUT

Use the generated prompt to seed your spreadsheet. Run it through CRAFT to clean and analyse. Then build a visualisation over the spreadsheet.

● LIVE DEMO MOMENT 3

I will walk through the process of creating a synthetic donor dataset from a prompt, then building a simple visualisation over it.

BrightHope CRM is the result. It's a proof-of-concept fundraising CRM aimed at small to medium charities (up to ~20 staff), built to show what is possible with AI-assisted development.

brighthope-crm.lovable.app

TAKEAWAY *A well-designed AI-assisted spreadsheet beats a poorly-used CRM every time. Start simple. Graduate when you're ready.*

SIX APPROACHES: PICK BASED ON GOVERNANCE AND LONGEVITY.

1 Spreadsheet + Low-Code ([Airtable](#) / [Baserow](#) / [Glide](#) / [Retool](#)) **Phase 1**
Fastest route. Low engineering overhead. Best for Phase 1 discovery and proving data model assumptions. Glide is strong for mobile-first operational apps.

2 Spreadsheet + Custom Web App ([Lovable](#) / [Replit](#)) **Phase 1-2** **Live demo**
Full control, proper relational integrity, best governance. Appropriate for small, innovative charities. Beware development and maintenance burden.

3 Internal Tool Platforms ([Retool](#) / [Appsmith](#) / [ToolJet](#)) **Phase 1-2**
Reasonable middle ground. Excellent for operational consoles and fundraising ops teams. Often highest ROI for constrained budgets.

4 Microsoft ([PowerBI](#), [Power Apps](#), [Dataverse](#), [Dynamics 365](#)) **Phase 1-3** **Default**
Best for charities already on M365. Strong data governance around existing tools, but beware vendor lock-in, hidden IT overhead, and AI that is not leading edge.

5 Open Source CRM ([CiviCRM](#) / [ERPNext](#) / [Odoo](#)) **Phase 2-3**
Low licensing cost. Significant implementation effort required.

6 Direct SaaS CRM ([Beacon](#) / [Salesforce NPSP](#) / [Raiser's Edge NXT](#)) **Phase 2-3**
Fastest to operational maturity. Sector-specific workflows. Buying is economically rational unless fundraising model is highly differentiated.

RECOMMENDED STRATEGY BY STAGE

Phase 1 Get Organised

You're managing supporter relationships in spreadsheets. The priority is agreeing what data you need, who owns it, and how it stays clean. The constraint here is process, not platform.

Phase 2 Get Operational

You have consistent data and need a shared system with proper access controls and an audit trail. This is where a dedicated CRM earns its keep.

Phase 3 Get Strategic

Your CRM runs reliably. Now you can use AI-assisted communications, donor segmentation, and analytics to drive fundraising strategy.

THE MAJOR STRATEGIC MISTAKE

Charities rarely have a technology problem. They have a process and data quality problem. Buying an enterprise CRM before fixing data entry and workflows creates an expensive system nobody trusts. Fix the process first. The platform follows.

THE KEY TRADE-OFF

Phase 1 favours speed and low cost. Phase 2 adds reliability and governance. Phase 3 delivers intelligence and growth. Don't jump to Phase 3 before Phase 1 is solid.

TAKEAWAY Start with Phase 1. Most charities over-engineer too early. Prove the data model before building the platform.

Board & Trustee Governance

What trustees need to know, what to ask, and how to bring AI onto the board agenda.

[Trustee Oversight](#)

[The Four Questions](#)

[GUARD in Governance](#)

[Board Paper Template](#)

THE FOUR QUESTIONS TO KEEP ASKING

Trustees do not need to be AI experts. They need a simple rhythm for asking the right questions and knowing what good answers look like.

1 1. Which AI tools are staff using?
Good: Approved tools list, owner, purpose, data boundary, and review date.
Poor: I think most people use ChatGPT but I'm not sure.

2 2. What data is going into those tools?
Good: Green/Amber/Red classification. No named beneficiary data in free tools.
Poor: We use AI for drafting. I don't think anything sensitive goes in.

3 3. Who is accountable for AI-assisted work?
Good: Named policy owner, human sign-off for high-risk outputs, incidents logged.
Poor: I assume staff check it before sending.

4 4. What has changed since the last review?
Good: Six-month review date, current tools checked, new risks escalated.
Poor: We don't have a written policy yet.

- 1. Does our data improve the tool without benefiting us?
- 2. Are we adopting it for impact or impressiveness?
- 3. Can we unwind decisions it influences?

GOVERNANCE MECHANISMS THAT WORK

Quarterly board update

A short standing item: tools in use, incidents or near misses, new risks, training completed, and one example of AI improving delivery.

AI register

A lightweight log of approved tools, owners, purpose, licence tier, data classification, DPA status, and next review date.

Exception and incident route

Clear escalation when someone wants to use a new tool, handle Red data, publish AI-assisted output, or report a mistake.

Six-month policy review

Keep the policy short, but review it twice a year. Check tools, data flows, roles, assurance evidence, and trustee visibility.

KEEP IT PROPORTIONATE

For a small charity, good governance is not a committee-heavy process. It is a visible owner, a short register, a review rhythm, and clear escalation when the risk changes.

TAKEAWAY Board AI governance works best as a rhythm: four questions, good evidence, a register, quarterly visibility, and clear escalation when risk changes.

Ethics & Responsible AI

Bias, environmental cost, and where AI must not replace human relationships.

The Responsible AI Spectrum

Bias in Practice

Environmental Cost

Human Connection

THREE POSITIONS

Minimalist

"We use AI only where strictly necessary, with maximum human oversight at every step."

Risk profile: Low adoption risk, high opportunity cost. Organisations here may fall behind peers in efficiency and fundraising capability.

Responsible ← You should be here

"We use AI proactively where it adds genuine value, within clear governance boundaries, with human sign-off on consequential decisions."

Risk profile: Balanced. This is where the sector should be heading.

Maximalist

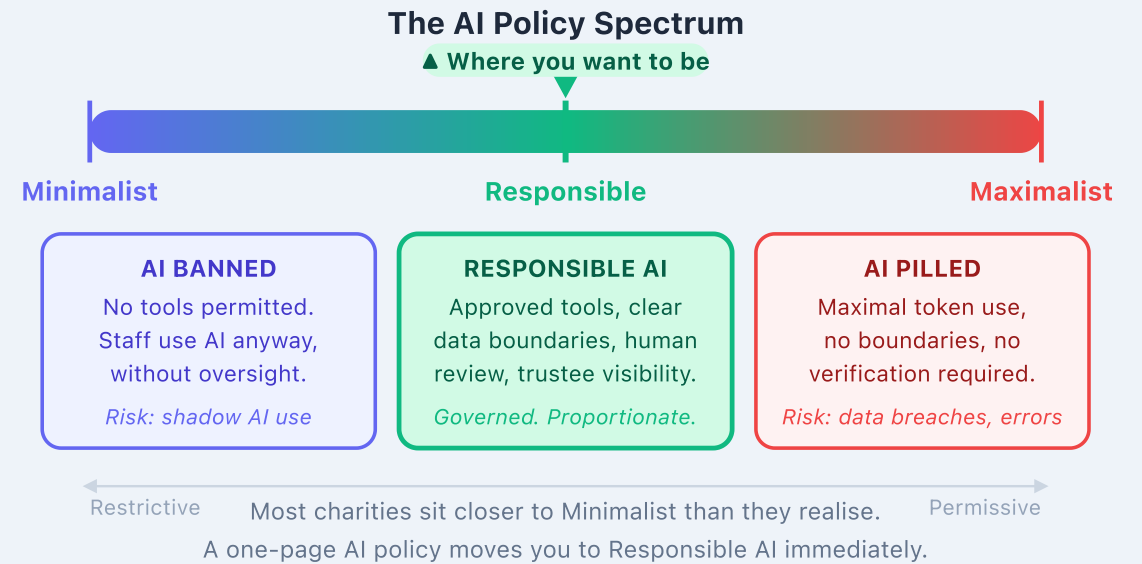
"We automate whatever we can as fast as possible to drive efficiency."

Risk profile: High. Without governance, maximalist adoption creates data, bias, and reputational exposure, especially for charities working with vulnerable people.

WHERE CHARITIES SHOULD BE

Responsible AI is not a compromise. It enables genuine efficiency gains while maintaining the accountability and human judgement that beneficiaries and funders expect.

THE AI POLICY SPECTRUM



DESKILLING RISK

Over-reliance on AI can erode critical thinking and deskill routine work if people stop checking outputs or stop practising the underlying judgement themselves.

[Yahoo Tech: AI and critical thinking](#)

TAKEAWAY The goal is Responsible, not Minimalist. Fear-based inaction is not an ethical position. It just transfers risk from AI to underfunding.

BIAS, AGENTIC SYSTEMS & ENVIRONMENTAL IMPACT

BIAS IN AI SYSTEMS

- AI systems reflect data they were trained on, reflecting historical biases
- For charities, this matters most in: grant eligibility assessment, beneficiary screening, communication personalisation

"Commitment to environmental principles by digital suppliers is either very important or important to almost two-thirds (65%) of charities."

— CDSR 2025, p.16

ENVIRONMENTAL COST

Raised unprompted by Session 1 respondents. A GPT-4 query uses ~10x the energy of a Google search; 100 queries a day is roughly a short car journey per week. AI has a water footprint — data centres intensify pressure on water-stressed areas.

[Hannah Ritchie: AI and Electricity \(2025\)](#)

[UC Riverside: AI programs consume large volumes of scarce water](#)

[Sasha Luccionj, Turing Lecture 2026: Making AI Sustainable \(YouTube\)](#)

AGENTIC AI

Non-agentic AI answers once. Agentic AI plans, retries, calls tools, and loops across multiple turns. That makes it much more impactful, both positively and negatively, because the system can act repeatedly before a human checks it.

WHERE AI CAN HELP AND WHERE IT SHOULD NOT BE USED

- Use AI for high-value tasks, not trivial substitutions
- Batch similar tasks into one session rather than running repeated queries
- Using AI to build software like BrightHope CRM is a good example of AI usage
- Once the tool is built, the token cost of building is amortised over the user base

DO NOT AUTOMATE

- ❌ Unsupported mental health or crisis communications to beneficiaries
- ❌ Decisions about eligibility for services that significantly affect individuals
- ❌ Communications simulating a human relationship where the recipient expects one
- ❌ Any process where an error would cause direct harm to a vulnerable person
- ❌ Replacing staff who provide human connection as part of the service model

"So-called artificial intelligences do not undergo experiences, do not possess a body, do not feel joy or pain, do not mature through relationships and do not know from within what love, work, friendship or responsibility mean."

— Pope Leo XIV, [Magnifica Humanitas](#), 2026

THE QUESTION TO ASK

Not just "can we?" but "should we?" and that question can extend beyond trustees to the people the charity serves. A lightweight beneficiary or staff consultation on where AI is and is not appropriate turns governance from a closed-door exercise into something tested with the people it affects.

TAKEAWAY *Responsible AI use asks not just 'can we?' but 'should we?' and that question belongs at trustee level.*

Choosing the Right Tools

Copilot vs Claude, what you're actually paying for, and the honest view on specialist charity tools.

[Copilot vs Claude](#)

[The Tier Decision](#)

[Specialist Tools](#)

[Q&A](#)

Copilot vs Claude: An Honest Comparison

MICROSOFT COPILOT

75% of Session 1 attendees are already using Copilot: this is your anchor tool.

STRENGTHS FOR CHARITIES

- ✔ Deep M365 integration across Word, Excel, Outlook, and Teams
- ✔ Copilot Chat is included at no additional cost for eligible Microsoft 365 subscriptions
- ✔ Excellent for meeting summaries, email drafting, document editing, and admin workflows
- ✔ Enterprise controls include audit logs, Purview compliance, and role-based access

LIMITATIONS

- ⚠ Less flexible for long-document synthesis or novel tasks outside M365
- ⚠ Best value is inside a Microsoft workflow, not as a general-purpose deep-reading assistant
- ⚠ Check your tenant and licensing, because access and retention depend on the Microsoft 365 setup

REFERENCES

- [Microsoft 365 Copilot pricing](#)
- [Anthropic Claude Pro](#)
- [Anthropic Team plan](#)
- [Anthropic Projects](#)
- [Anthropic data handling](#)

CLAUDE (ANTHROPIC)

26% of Session 1 attendees already using Claude: growing rapidly.

STRENGTHS FOR CHARITIES

- ✔ Exceptional for long-document analysis, uploaded files, and project-based work
- ✔ Strong at nuanced writing, grant narrative, impact framing, and trustee reports
- ✔ Pro plans include Projects and Knowledge Bases, which help you work from your own source material
- ✔ Claude Pro conversations are not used to train models by default

LIMITATIONS

- ⚠ No native Microsoft 365 workflow integration like Copilot
- ⚠ Team features such as SSO, role-based permissions, spend controls, and enterprise search sit on Team and Enterprise plans
- ⚠ Enterprise search and collaboration are better suited to team plans than Pro

NOT EITHER/OR

Copilot is your anchor inside Microsoft 365, but it does not exclude other assistants. Many people in the tech world use multiple AI tools heavily, switching by task, document type, and workflow.

THE SUMMARY RULE

Use Copilot for work that lives inside Microsoft 365 and needs tenant-level controls. Use Claude for long documents, uploaded sources, and project-based analysis. They complement each other, you don't have to choose.

TAKEAWAY *The best tool is the one your team will actually use: consistently and safely. Platform loyalty matters less than governance clarity.*

The Tier Decision: What You're Actually Paying For

THE THREE TIERS EXPLAINED

Free Tier

£0/month

- ✓ Good for personal skill-building and low-stakes tasks
- ✗ Prompts may train the model in some tools
- ✗ No DPA, organisational data should not be entered
- ✗ No audit logs, no admin controls

Paid Personal

~£15–16/month per user

- ✓ Your conversations do not train the model
- ✓ Appropriate for organisational content with anonymised data
- ✓ The minimum tier for professional use
- ✗ No team admin controls or audit logs

Team / Enterprise

~£14–20/month per user

- ✓ Data processing agreement in place
- ✓ Admin console, see who's using what
- ✓ Audit logs for compliance
- ✓ Required if any Amber data is being used at scale

THE GOVERNANCE IMPLICATION

THE FUNDAMENTAL RULE

The tier determines your governance ceiling, not the brand name. A free ChatGPT account and a £30/month enterprise account are fundamentally different governance instruments.

THE MICROSOFT ALERT

Microsoft is phasing out some free Copilot tiers in 2026. If your organisation relies on free Copilot access, check your current entitlements. You may already be on a paid M365 plan that includes Copilot at no extra cost.

BUDGET GUIDANCE

For a team of 5 using AI seriously: ~£75–80/month (paid personal) or ~£70–100/month (team tier). Frame against time saving: if one staff member saves 3 hours a week on grant admin, a ~£20/month subscription has a strong ROI within a month.

PRICING BASIS

- [OpenAI pricing](#): ChatGPT Plus \$20/month; ChatGPT Business \$300/user/year billed annually or \$30/user/month billed monthly.
- [Anthropic Claude Pro](#): \$20/month.
- [Anthropic Team plan](#): \$240/member/year or \$25/member/month.
- [Microsoft 365 Copilot](#): Copilot Business £165.60/user/year or £19.32/user/month.

TAKEAWAY Free is not free when your beneficiary data is the cost. £15–20/month per user is the minimum for responsible professional use.

Specialist Tools Worth Being Aware About

SPECIALIST TOOLS BY TASK

These tools are strongest when the task is repeated, bounded, and worth formalising. Use them for grant work, meeting records, beneficiary support, and similar workflows, not for novelty.

Charity Excellence Framework Free
Structured bid writer built on ChatGPT. Best for one-off bids and charities new to grant writing. Review everything carefully.

CharityIQ £49-349/month
Integrated platform for grant writing, compliance checking, and impact reporting. Best for UK charities applying regularly.

FundRobin £15-399/month
Grant matching plus proposal writing. Good entry point if grant discovery is the priority.

Grant Genie Free / £25-35/month
Transparent pricing, genuine free tier, and optional bid-writing add-on. Good for smaller teams exploring AI grant tools.

Beam.org AI Sector-specific
Beneficiary support coordination. Appropriate only within its designed context and service model.

Otter.ai / Granola Low-cost
Meeting transcription and summaries. Useful for governance records; check consent and retention policy.

RESEARCH NOTE

The full research note used to shape this slide is here:

[Perplexity Deep Research](#)

HOW TO DECIDE

Apply the same five questions before you adopt any specialist tool:

- 1 Data**
What data enters this tool, and where does it go?
- 2 DPA**
Is there a Data Processing Agreement available?
- 3 Tier**
Is there a paid tier with appropriate data protections?
- 4 Transparency**
Does the tool explain what it does with your prompts?
- 5 Necessity**
Does this genuinely improve on what you already have?

THE HONEST CAVEAT

No specialist tool beats a general-purpose tool used well. Before subscribing to anything new, ask whether Claude or Copilot with a good prompt achieves the same result.

TAKEAWAY *New tools emerge every week. The evaluation framework matters more than any specific tool recommendation.*

WHAT TO REMEMBER

1 Governance first

A one-page AI policy, approved tools list, and Green/Amber/Red data rules are enough to move from informal use to managed adoption.

2 Practice beats theory

Teams build confidence by using AI on real work, then checking outputs against sources and organisational judgement.

3 Data quality is the multiplier

CRAFT only works when the underlying spreadsheets, CRM records, and source files are clean enough to trust.

4 Use a mix of tools

Copilot, Claude, and specialist products all have different strengths. A diverse toolkit is healthier than over-relying on any one assistant.

SECTOR BASELINE & CLOSING PRINCIPLE

Most charities are still early. If you leave with a policy draft, an approved tools list, and one live workflow to practise, you are already ahead of the sector baseline. AI readiness is not about buying the perfect tool. It is about building a repeatable habit: clear task, safe data, checked output, human accountability.

WHAT TO DO NEXT

TODAY

Give feedback & take the REACH™ self-assessment. See air.netlify.app for more.

BrightHope CRM is a proof-of-concept lightweight CRM for small to medium charities, built live during this session. Try it at brighthouse-crm.lovable.app

Pick one workflow this week.

Choose something real, low-risk, and repeatable: meeting notes, grant drafting, donor segmentation, policy review, or board-paper preparation.

Within 48 hours

Draft or refine your one-page AI policy and approved tools list.

Within 7 days

Run one CRAFT workflow on synthetic or low-risk data and verify the output.

At the next board meeting

Put the four AI governance questions on the agenda.

Within 30 days

Decide which tool tier is appropriate for the data your team actually uses.

Thank You, Feedback and Q&A

CLOSE

GET IN TOUCH

Mal Minhas · Diorama Consulting

CONNECT

dioramaconsulting.co.uk

[LinkedIn](#)

[CxAI Newsletter](#)

[CxAI Substack](#)

REACH™ SELF-ASSESSMENT

Take the free self-assessment at dioramaconsulting.co.uk/ailr — register with **ThirdSector** as the org name.

FEEDBACK

Your feedback shapes the next session. Please scan the QR code on the right to complete the short [feedback form](#).

SESSION 2: GOING DEEPER

A follow-up session is being planned on developing an organisational AI roadmap. Your feedback today will help shape it.

NOW Q&A

HOW Q&A WORKS

Post questions in the Zoom chat or raise your hand.

RECORDING NOTE

Today's session has been recorded. The presentation and resources will be shared with all attendees.

SCAN TO GIVE FEEDBACK



[feedback form](#)

Thank you.

Now go and try one thing today!

TAKEAWAY *Thank you. One thing this week. Connect on LinkedIn. Complete the feedback form.*