Deloitte.

AI for NPOs Toolkit

A practical guide for non-profit organisations to understand, evaluate, and begin using artificial intelligence developed by Deloitte Belgium

2026

Part I — Introduction

1.1. Why does this toolkit exist?

Artificial intelligence is no longer a technology reserved for large corporations or well-funded research institutions. It is increasingly accessible, often free, and practical enough to help organizations of any size work more effectively. Yet for many non-profit organizations, the topic still feels abstract, intimidating, or simply too distant from the daily pressures of delivering programs and securing funding.

This toolkit exists to change that. Because we recognize that technology is reshaping opportunity; but not evenly. The path forward is becoming more complex and uncertain for many organizations and communities. That's why, as part of Deloitte's global Impact Strategy, we're committed to creating pathways to progress, and the AI for Good Toolkit is a part of that mission. Over the past months, we've worked closely with our societal impact partners, organizations like Vlajo (Vlaamse Jonge Ondernemingen) and Give a Day, to understand what's really needed. We've identified real challenges: data quality concerns, staff readiness, integration complexity, and ethical questions. This toolkit reflects those learnings, offering not just opportunities, but honest guidance on navigating the obstacles.

Why does urgency matter? AI capabilities are advancing quickly, but the evidence should be interpreted carefully. Recent research suggests that frontier AI systems are becoming able to complete increasingly longer software-oriented tasks, with measured task horizons roughly doubling every seven months. This does not mean AI can reliably perform every type of human work autonomously. It does mean that organizations will benefit from building practical experience now, while the stakes and costs of experimentation are still manageable.

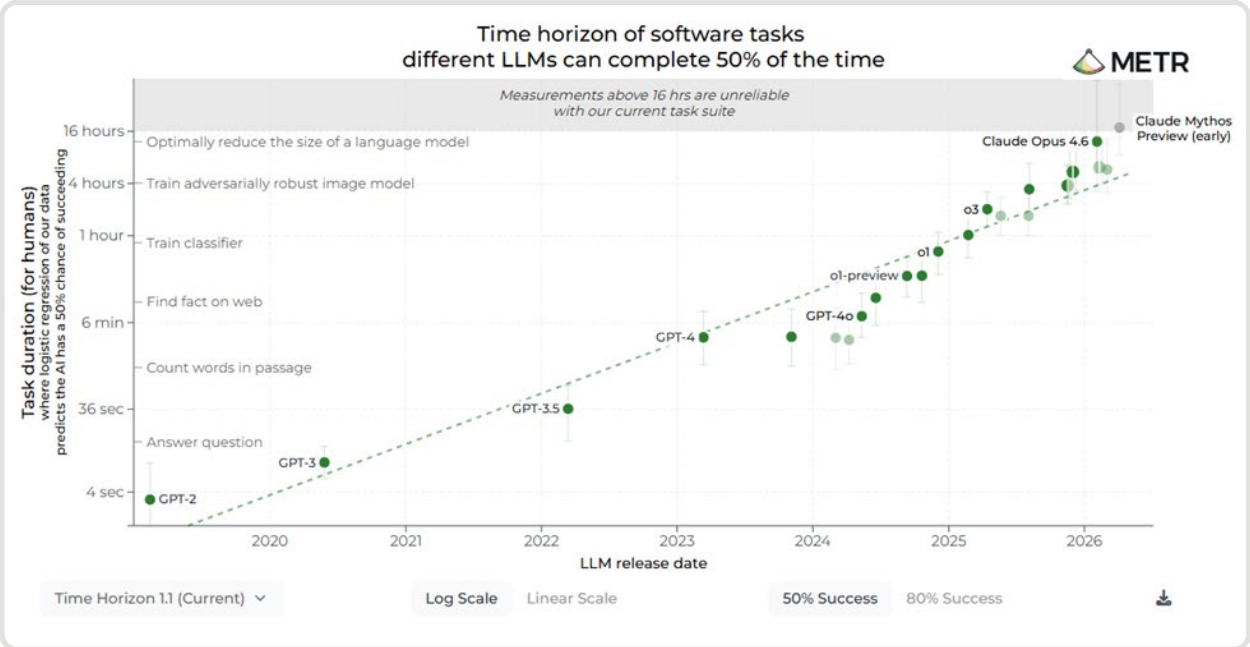

Source: Measuring AI Ability to Complete Long Tasks, arXiv:2503.14499, 21/05/2026

What you will find in this toolkit

Part 2	Part 2 explains what AI is, how it works, and what the different types mean in plain terms.
Part 3	Part 3 presents 11 concrete use cases drawn from organizations like yours, with tools, examples, and effort estimates for each.
Part 4	Part 4 provides three practical tools: an AI readiness checklist, an ideation canvas for identifying your own use cases, and our AI Ideation assistant: Noor
Part 5	Part 5 is a list of frequently asked questions we hear in our conversations with NPOs.
Part 6	In part 6, you can apply for an AI Ideation session, which then might be selected as a pro bono project by Deloitte!

You do not need a technological background to use this toolkit: it is designed for non-technical readers. The language is kept plain throughout. The examples are illustrative and intended to show realistic starting points, not to imply that every organization should adopt every use case. The tools referenced should be assessed against your own privacy, procurement, accessibility, and governance requirements.

This toolkit offers a starting point for organizations that want to understand what AI can realistically offer them, make informed decisions about where to begin, and avoid common pitfalls along the way.

Whether you are completely new to the topic or have already started experimenting with a few tools, this toolkit is designed to meet you where you are and help you take a considered next step.

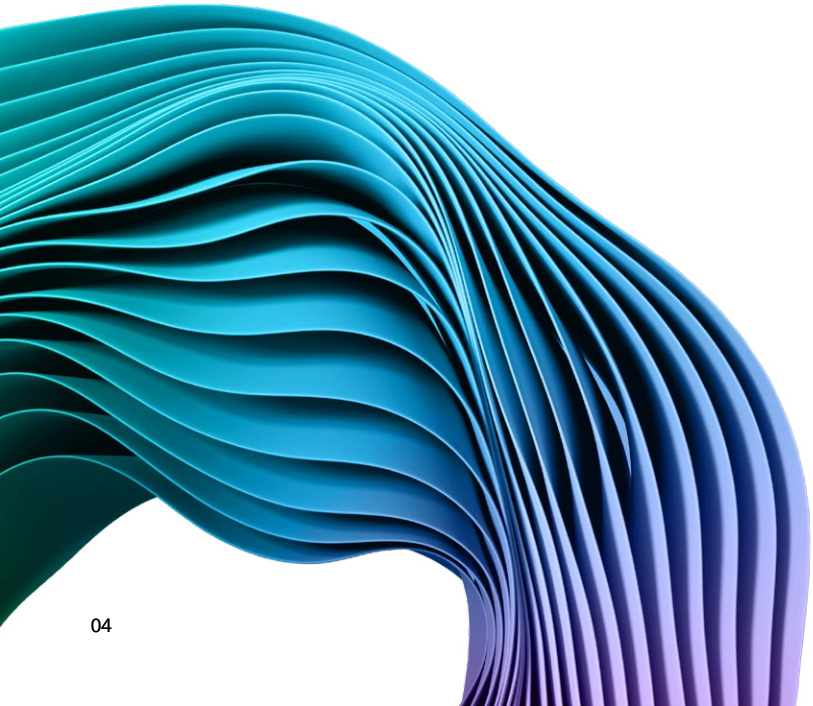

Part II
Understanding Artificial Intelligence

2.1. What is Artificial Intelligence?

Artificial intelligence refers to computer systems that can perform tasks which would normally require human intelligence. These include understanding language, recognizing images, making predictions, solving problems, and generating written or visual content.

What makes modern AI different from older software is how it is built and used. Traditional software follows rules that a programmer writes in advance. If the invoice total is over 1,000, flag it for review. Many AI systems, by contrast, are trained on examples and learn statistical patterns from data. This training happens before you use the tool. When you prompt an AI tool, it is usually applying patterns it has already learned, not automatically learning from your individual prompt unless the provider explicitly says so.

This shift from rule-following to pattern-learning is why AI can now do things that were previously considered too complex or too unpredictable for computers: translate languages naturally, answer questions about photographs, summarize long documents, or predict which donors are most likely to give.

The different types of AI

AI is not a single concept. It is a broad field that contains several distinct approaches, each suited to different kinds of problems. The diagram below shows how these types relate to one another. Think of them as nested categories, each building on the one that contains it.

Machine learning is the most widely used form of AI today. Rather than following a fixed set of rules, a machine learning system is trained on examples and learns to make predictions or decisions based on patterns it finds in that data. An email spam filter is a classic example: it learns what spam looks like from thousands of flagged examples and applies that knowledge to new messages.

Deep learning is a subset of machine learning that uses structures loosely inspired by the human brain, called neural networks. These networks can handle much more complex patterns than earlier approaches, which is why they are behind most modern image recognition, speech processing, and language tools.

Generative AI refers to systems that can create new content: text, images, audio, and video. These systems have learned patterns from large collections of existing content and can produce new material in response to a prompt. Their output can be useful, but it still needs human review, especially when facts, tone, rights, or sensitive audiences matter.

Agentic AI is a newer development where AI systems are given a goal and can work through a sequence of steps to achieve it. Depending on the design, an agent may use tools, search information, analyze documents, write code, or ask for human approval before continuing. This can make AI more useful for multi-step work, but it also increases the need for clear boundaries, logging, permissions, and human oversight.

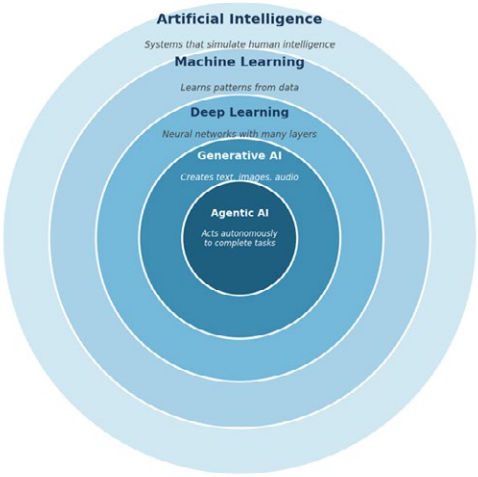

The diagram below explains the difference between a standard AI model and an AI Agent. Comparing these concepts side by side, we see the following differences:

- **Architecture:** A standard AI model (e.g., an LLM) generates outputs based on a single prompt-response interaction. An AI agent builds on top of such models and can plan, reason over multiple steps, call external tools, and manage task execution toward a defined goal.
- **Use case scope:** Standard AI models are typically used for isolated tasks (e.g., summarization, Q&A). AI agents are designed to handle multi-step tasks and can automate broader workflows or business processes.
- **Memory:** AI models are generally limited to context within a session (short-term memory). AI agents can incorporate both short-term memory (context) and long-term memory (stored information across sessions), depending on their design.
- **Tool integration:** AI models are not inherently connected to tools. AI agents are explicitly designed to use and orchestrate external tools and APIs as part of task execution.
- **Data integration:** AI models rely partly on pre-trained knowledge, which can become outdated. AI agents can be designed to access current information through approved sources, databases, or tools, but this depends on the implementation and should not be assumed by default.
- **Accuracy:** AI models generate outputs probabilistically and do not inherently guarantee correctness. AI agents can improve reliability by checking outputs against tools, databases, or defined rules, but they can still fail. Verification remains essential for important decisions.

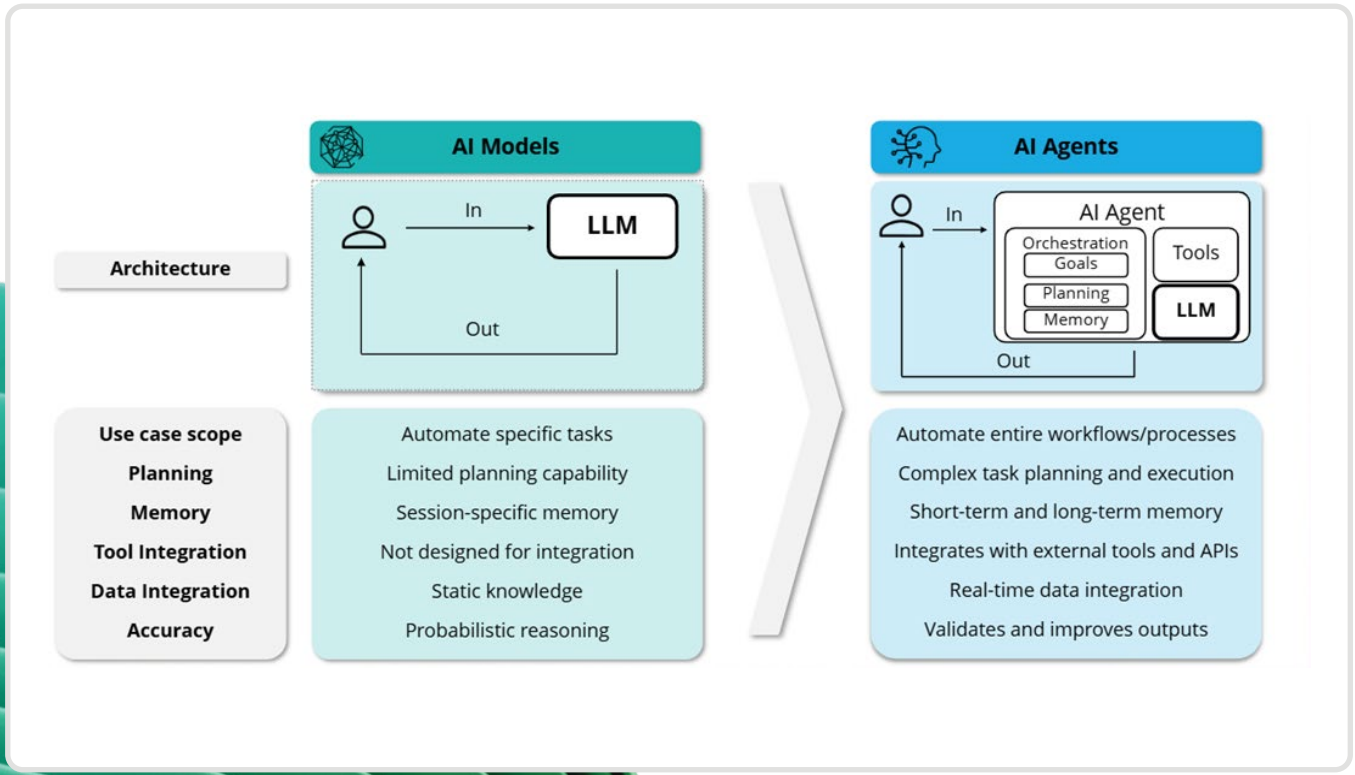

2.2. How AI Tools Work

Understanding what happens when you use an AI tool helps you get better results from it, spot when something has gone wrong, and make informed decisions about where it is appropriate to use one.

The interaction between you and an AI tool follows a straightforward cycle, regardless of which tool you are using or what task you are performing:

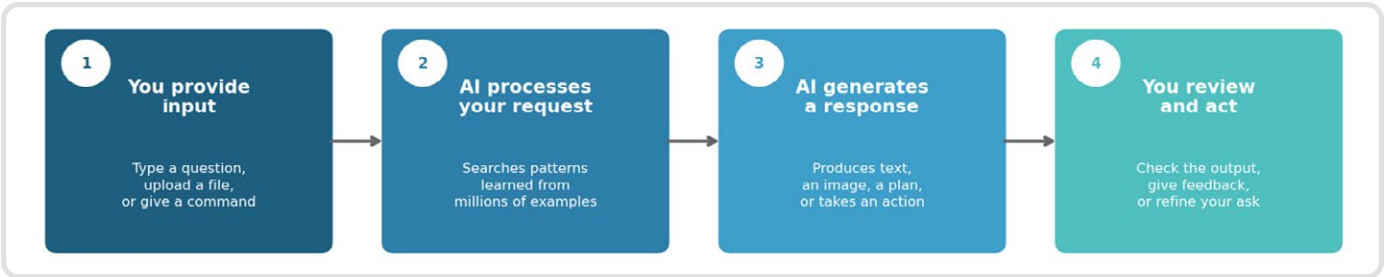

Prompts and prompt engineering: how you communicate with AI

The instructions you give to an AI tool are called a prompt. The quality of your prompt has a direct effect on the quality of the response. A vague prompt produces a generic answer. A specific, well-structured prompt produces a more useful one.
Below components make a good prompt:

1. **Role**
   What role, experience or expertise do you want the AI to have?
2. **Context**
   What information and background can you give?
3. **Task**
   What do you want the AI to do?
4. **Output format**
   How do you want the results?
5. **Tone**
   What is your desired tone of voice?
6. **Example (optional)**
   Do you have an example of the output you want?
7. **Improve your prompt**
   Use an LLM as a prompt coach to improve your prompts.

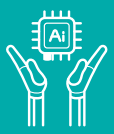

Prompting: a practical principle

Think of writing a prompt as briefing a capable but uninformed colleague. The more context you give them about your organization, your audience, and exactly what you need, the more useful their output will be. Tell the AI who you are, what you want, and what constraints apply. For example:

For example:

"[Role] You are working in external communications for an NPO. [Context] The NPO is focused on urban food poverty in Brussels. [Task] Write a 200-word appeal for our end-of-year donor email. [Output Format] Ensure the email is a coherent text without bullets. [Tone] The tone should be warm but not sentimental. Focus on the practical impact of a 50-euro donation."

After the output to your initial prompt, you might identify room for improvement in the output by changing specific elements in your prompt. This process is called prompt engineering: Designing, refining, and testing prompts to optimize AI model performance. Making this standard practice when you're interacting with AI models is a helpful tool to improve the accuracy of your desired input.

2.3. Understanding the Risks of AI

Like any tool, AI comes with both genuine advantages and real risks. Neither should be exaggerated. The organizations that get the most value from AI are those that understand both sides clearly and apply the technology with intention. The risks listed below serve as examples and are not limited to the key risks associated with the use of AI.

Understanding and mitigating the key risks

AI systems can produce content that sounds authoritative and well-formed but is factually wrong, incomplete, or misleading. This happens because many AI tools generate likely responses based on patterns, not because they always check facts against reliable sources. Treat AI output as a draft or hypothesis unless it has been verified.

How to manage hallucination

- Make the people in your organization aware of risks like hallucination, and make sure they understand the responsibilities tied to publishing AI-generated content.
- When facts, figures, or citations matter, verify them independently.
- If accuracy is critical, prefer tools that show you their sources.

Bias in AI outputs reflects bias in training data. If a model is trained predominantly on content from certain languages, cultures, or demographics, its outputs will reflect those patterns.

How to manage bias

- Review AI outputs with the communities you serve in mind. Ask: whose perspective is missing here?
- For high-stakes decisions affecting beneficiaries, do not rely on AI outputs alone.
- Acknowledge bias as a topic in your AI literacy training for staff.

Data privacy is a critical concern for NPOs that hold sensitive information about beneficiaries, donors, staff, volunteers, or vulnerable populations. Many AI tools process data on external servers. Unless your organization has approved the tool, reviewed its contractual terms, and confirmed how data is handled, assume that sensitive personal data should not be pasted into it.

How to manage data privacy

- Set a clear rule: personally identifiable information about beneficiaries, donors, or staff never goes into general-purpose AI tools.
- Read the data handling policies of any tool before introducing it to your workflow.
- Consider paid (often called enterprise) plans that include data privacy guarantees if your work is sensitive.
- Always keep in mind GDPR laws.

Additional risks NPOs should consider include overreliance, copyright and intellectual-property issues, security threats such as prompt injection or unsafe file handling, reputational damage from inaccurate or tone-deaf content, and digital exclusion when tools perform less well for certain languages, literacy levels, or accessibility needs.

For any use case that could affect beneficiaries, eligibility, funding, employment, complaints, or access to services, AI should support human decision-making rather than replace it. Keep a person accountable, document the reasoning, and provide a clear route for challenge or correction.

Trustworthy AI: Beyond compliance

Deloitte's Trustworthy AI™ Framework A identifies six dimensions that every organisation should consider when deploying AI-powered systems or processes.

- **Fair and impartial:** AI applications include internal and external checks to help ensure equitable application across all participants
- **Robust and reliable**: AI systems can learn from humans and other systems and produce consistent and reliable outputs
- **Private:** Consumer privacy is respected, and customer data is not used beyond its intended and stated use; consumers are able to opt in/out of sharing their data
- **Safe and secure:** AI systems can be protected from risks (including Cyber) that may cause physical and/or digital harm
- **Responsible:** The technology is created and operated in a socially responsible manner.
- **Accountable:** Policies are in place to determine who is held responsible for the output of AI system decisions
- **Transparent and explainable:** All participants are able to understand how their data is being used and how AI systems make decisions; algorithms, attributes, and correlations are open to inspection

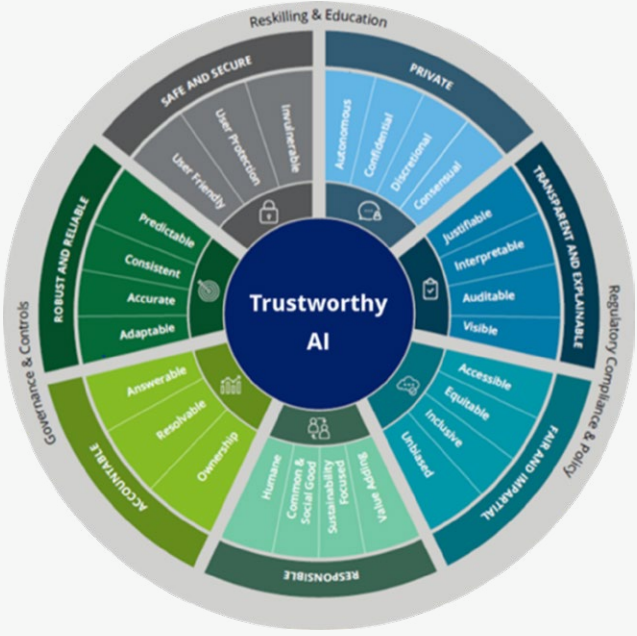

2.4. AI Governance and The EU AI Act

The EU AI Act

As part of the EU Digital Strategy, the European Commission has introduced legislation to manage and support digital technologies. The AI Act is a horizontal legal framework for AI in the EU. It regulates AI systems according to risk and imposes different obligations on providers, deployers, importers, distributors, and certain users depending on the role and use case.

Purpose	Approach	Scope
• Promote the **trustworthy and human-centric** application of AI – while supporting innovation within the EU. • **Protect** fundamental rights, health, safety, democracy and the rule of law, and the environment from potential harmful effects. • Promote the **development** of AI systems in a transparent and compliant manner.	• **Risk-based approach** that classifies AI systems by **use case and risk level**; the category determines the obligations that apply. • **General-purpose AI models** (GPAI) can have **additional obligations**, especially where they pose systemic risk.	• **Any AI system affecting a person located in the EU**, regardless of where the system was developed or where the system is hosted, operated or monitored. • The Act involves **many stakeholders**. It mainly applies to providers (developers), deployers (professional users), importers, and distributors of AI systems.

The AI Act entered into force on 1 August 2024 and applies gradually. Prohibited AI practices and AI literacy obligations started to apply on 2 February 2025. Governance rules and obligations for general-purpose AI models started to apply on 2 August 2025. The Act is generally applicable from 2 August 2026, with specific transition periods for certain high-risk AI systems. Obligations for high-risk AI systems apply from 2 December 2027, while high-risk AI systems embedded in regulated products benefit from an extended transition period until 2 August 2028.

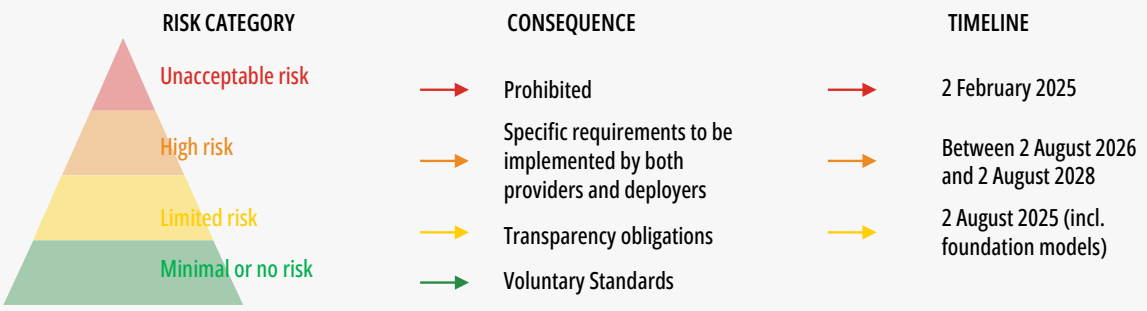

AI Governance

Clear AI governance is not just a nice-to-have, it is a necessity for ensuring your organization's compliance with the EU AI Act. The recommendations below are intended as a starting point for responsible AI use and do not constitute legal advice. They are not exhaustive; the steps required to establish robust AI governance will vary depending on your organization's size, sector, and how deeply AI is embedded in your operations.

1. **Audit your current AI use.** Take stock of how AI is being used across your organization: by whom, for what purpose, and in what context. Map each use case against the EU AI Act's risk classification to understand where your obligations lie.
2. **Avoid high-risk AI applications** unless your organization can meet the required governance, documentation, human oversight, monitoring, and accountability obligations. For most NPOs, early AI adoption should focus on low-risk productivity and support use cases before moving into decisions that affect people's rights, access to services, or eligibility.
3. **Eliminate Shadow AI**. Actively discourage and monitor for the unsanctioned use of AI tools outside of approved channels. Shadow AI (employees using personal or unapproved AI tools for work tasks) creates significant compliance, security, and data privacy risks that are difficult to audit or control.
4. **Invest in AI literacy.** Run regular training and awareness activities so that all members of your organization understand what AI can and cannot do, how to use it responsibly, and what the consequences of misuse can be.
5. **Establish clear internal policies**. Develop and communicate an AI use policy that defines which tools are approved, what data may be used as input, who is accountable for AI-assisted decisions, and how incidents or concerns should be reported.
6. **Assign governance ownership**. Designate a responsible person or working group, such as an AI Officer or a cross-functional AI governance committee, to oversee compliance, keep up with regulatory developments, and update internal policies as the landscape evolves.
7. **Document and review continuously**. AI governance is not a one-time exercise. Keep records of your AI use cases, risk assessments, and policy decisions. Schedule regular reviews to reassess risk classifications and ensure your practices remain aligned with the EU AI Act as it comes into full effect.

Part III
AI for NPOs

3.1. AI Use Cases for NPOs

The following pages present a selection of practical AI use cases for NPOs. Each card outlines what the application could do, the potential benefits, a safe first experiment or a real-world example, and the type of tool category to explore. The aim is not to promote specific vendors, but to help teams identify where AI may add value and where additional safeguards may be needed.

Use cases 1 to 4 present real-world examples shared by Give a Day, Odoo, and Get in Touch, who collaborated with us on this initiative. Use cases 5 to 11 illustrate common AI applications for NPOs and include a suggested first experiment to help organizations begin exploring AI in a safe and practical way.

Each use case includes three metrics to support prioritization. Time gain and efficiency gain indicate the potential value when a use case is implemented effectively. Implementation complexity reflects not only technical effort, but also factors such as data sensitivity, governance requirements, integration effort, and the level of human oversight required.

Please note that this list is illustrative rather than exhaustive. The most valuable AI use cases are typically those linked to a concrete bottleneck or organizational need. Before adopting any AI tool, organizations should assess data sensitivity, accessibility, review requirements, user impact, and whether the use case could meaningfully affect beneficiaries, donors, staff, or volunteers.

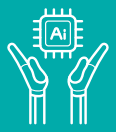

Reading the use case cards

Time Gain / Efficiency Gain:

High	significant potential time or quality impact
Medium	moderate potential impact

Implementation Complexity:

Low	can be tested by staff using approved tools and non-sensitive data
Medium	needs configuration, process change, or additional review
High	requires technical support, integration, sensitive-data controls, or a formal implementation project

1. Volunteer Matching & Management Automation

Management & Operations | Community development, environmental, social service, and event-based NPOs

Time Gain	Efficiency Gain	Implementation Complexity
Medium	Medium	Medium

What it does

AI can help match volunteer skills, availability, location, and preferences to open roles, and can support scheduling, reminders, and follow-up.

Key benefit

Reduces coordination effort and can improve volunteer fit and retention. The matching logic should be transparent and easy for coordinators to override.

Real-world example

AI for volunteer coordination to match thousands of volunteers to building projects, allowing coordinators to focus on relationships rather than spreadsheets.

Tools to explore: Give a Day offers a 'Tinder for volunteers' where availabilities and interests are matched to volunteer vacancies

2. Task assignment in Odoo

Project management | All NPOs

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Medium

What it does

Odoo AI automatically analyzes incoming project tasks, donor requests, or volunteer tickets and matches them against team profiles. By understanding natural language descriptions—such as "User X is an expert in grant compliance" or "User Y handles community outreach"—the AI instantly assigns the task to the most qualified team member without requiring manual triage.

Key benefit

Eliminates administrative bottlenecks and maximizes lean resources. NGOs often operate with limited staff or a shifting network of volunteers. This automated routing ensures that time-sensitive requests (like donor inquiries or grant deadlines) are never lost in an unassigned queue, allowing your team to focus on mission-driven work rather than manual project management.

Real-world example

An international NGO receives an urgent inbound message via their Odoo portal regarding a complex financial reporting discrepancy for an active European Union grant. Odoo AI parses the message, identifies it as a "high-priority finance and compliance" issue, and instantly assigns it to Sarah (the Finance Lead) rather than Tom (the Marketing Coordinator). Sarah is notified immediately, ensuring the grant compliance deadline is met without a single manager needing to route the email.

Tools to explore: Odoo

3. HR applicant screening in Odoo

Human Resources| All NPOs

Time Gain	Efficiency Gain	Implementation Complexity
Medium	High	Low

What it does

Odoo AI scans incoming job or volunteer applications, parsing resumes and cover letters to evaluate candidates against the specific requirements of the role. It automatically assigns a qualification score (e.g., 85/100) based on skills, experience, or certifications, and dynamically suggests the next best action—such as "Schedule technical interview" or "Send automated polite rejection email."

Key benefit

Drastically reduces time-to-hire and onboarding friction for lean HR teams. NGOs often face massive influxes of applicants during major advocacy campaigns or emergency hiring surges, but lack dedicated recruitment staff. This feature filters out unqualified profiles instantly, ensuring recruiters spend their limited time interviewing the highest-potential advocates and professionals.

Real-world example

A humanitarian NGO launches a rapid response drive for field volunteer coordinators following an environmental crisis, receiving over 400 applications in three days. Instead of a single HR person reading every resume, Odoo AI instantly scores applicants based on critical keywords like "first-aid certification," "crisis management," and "local language fluency." The AI tags the top 40 candidates with a 90+ score and prompts the coordinator with a single click: "Suggested next step: Send bulk invitation to standard phone screening."

Tools to explore: Odoo

4. Getintouch

Communication | All NPOs

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Medium

What it does

Get In Touch makes WhatsApp work for your organisation. Reach your audience where they already are. Forget email and social media. WhatsApp has open rates >94% on average. Claim your username and start conversations with your audience.

Key benefit

You don't want more communication channels. You want the one we all use the most. Integrate your professional WhatsApp account in Google, Facebook, Instagram and even chatGPT, Gemini, ...

Real-world example

Team Belgium (olympics): communication to all stakeholders during Summer and Winter games
De Warmste Week: communication to all fund raisers

Tools to explore: Getintouch

5. Donor Prospect Research & Predictive Scoring

Fundraising | All fundraising NPOs

Time Gain	Efficiency Gain	Implementation Complexity
Medium	Medium	Medium

What it does

AI can help structure public or consented donor information, identify outreach priorities, and suggest segments based on transparent and relevant criteria.

Key benefit

Can help fundraising teams focus limited time on the most relevant outreach, but only when data quality, lawful basis, transparency, and profiling risks are properly addressed.

Safe first experiment

Use AI to summarize publicly available organization-level information or past campaign patterns. Do not upload personal donor data into unapproved tools or use AI scoring as the sole basis for outreach decisions.

Tools to explore: CRM analytics, segmentation, and fundraising workflow tools with appropriate privacy controls

6. AI-Powered Chatbots for Beneficiary & Donor Support

Engagement | Humanitarian, health, education, advocacy NPOs

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Medium

What it does

Conversational AI can answer routine questions, triage requests, and guide users to approved information outside office hours.

Key benefit

Can reduce repetitive enquiries and improve availability, especially during campaigns or events. The safest starting point is a public FAQ bot using approved content only.

Safe first experiment

Deploy an internal or website FAQ assistant for non-sensitive questions such as opening hours, donation methods, volunteer sign-up, or event logistics. Do not allow it to provide crisis, health, legal, or eligibility advice without strong human oversight.

Tools to explore: FAQ chatbot, knowledge-base assistant, and service-desk automation tools

7. Personalized Learning & Adaptive Education Support

Education | Education NPOs, training providers, international development organizations

Time Gain	Efficiency Gain	Implementation Complexity
Medium	High	Medium

What it does

AI can help tailor exercises, explanations, quizzes, and feedback to different learning levels, languages, or support needs.

Key benefit

Supports differentiated learning and helps educators prepare materials faster. It should complement, not replace, qualified educators or safeguarding processes.

Safe first experiment

Ask AI to generate alternative explanations, reading-level variants, or practice questions from approved learning material, then have an educator review before use.

Tools to explore: Adaptive learning, tutoring support, translation, and educational content-generation tools

8. AI-Generated Content & Communications

Creativity & Communications | All NPO types, especially those with small communications teams

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Low

What it does

Generative AI can help draft newsletters, donor emails, social posts, grant narratives, website copy, and campaign materials from human-provided context.

Key benefit

Reduces first-draft time and helps small teams maintain regular communication. Human review remains essential for accuracy, dignity, tone, and factual claims.

Safe first experiment

Create three versions of a donor email from an approved campaign brief, then have the communications owner select, edit, and fact-check the final version.

Tools to explore: Approved text, design, translation, and content-assistance tools

9. Grant Writing Assistance & Funder Matching

Fundraising | All grant-seeking NPOs, especially small and medium organizations

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Medium

What it does

AI can help structure proposal drafts, adapt language to funder requirements, summarize prior project evidence, and compare opportunities against organizational fit.

Key benefit

Can reduce drafting time and improve consistency, but funder claims, budgets, eligibility, and impact evidence must be verified by the responsible team.

Safe first experiment

Use AI to turn an approved project brief into a first proposal outline and checklist of missing evidence. Avoid uploading confidential budgets, personal stories, or unpublished strategies into unapproved tools.

Tools to explore: Proposal drafting, document analysis, opportunity tracking, and CRM-supported fundraising tools

10. Financial Forecasting & Anomaly Detection

Management & Operations | NPOs managing multiple budgets, grants, or programs

Time Gain	Efficiency Gain	Implementation Complexity
Medium	High	Medium

What it does

AI can support budget forecasting, variance explanation, cash-flow analysis, and anomaly detection in financial data.

Key benefit

Helps finance teams spot unusual patterns earlier and spend more time on analysis. Payment execution should remain controlled by existing approval, segregation-of-duties, and audit processes.

Safe first experiment

Use anonymized or test financial exports to generate a variance commentary or list of transactions that require human review.

Tools to explore: Financial planning, anomaly detection, reporting, and analytics tools with audit and access controls

11. Impact Measurement through Feedback & Data Analysis

Engagement | All NPOs

Time Gain	Efficiency Gain	Implementation Complexity
High	High	Low

What it does

AI can help analyze survey responses, open comments, interview notes, and program feedback to identify themes, sentiment, outliers, and evidence of impact.

Key benefit

Helps teams turn qualitative feedback into structured learning more quickly. Results should be checked for bias, missing perspectives, and over-simplification.

Safe first experiment

upload anonymized feedback exports and ask AI to summarize themes, representative quotes, concerns, and follow-up questions for human review.

Tools to explore: Text analysis, survey analytics, reporting, and approved document-analysis tools

Part IV
How to Start Using AI in Your Organisation

4.1. The AI Readiness Checklist

Before deploying any AI tool at scale, it helps to assess your organization's readiness across six areas: strategy, people, use cases, data classification, technology, and governance. The checklist below is designed to be completed by one or two people in a single sitting. It does not require technical expertise.

Do not wait until every box is ticked before starting. Most NPOs begin experimenting before they have everything in place. The point of this checklist is to identify your gaps early so you can address them deliberately, rather than discovering them under pressure.

Strategy and Leadership

- ☐ **Appoint an AI champion**

  Identify one person, even part-time, who is curious about AI and give them protected time each week to experiment and report back. This does not need to be a technical role.

- ☐ **Connect AI to your mission**

  Write down one or two specific goals you want AI to help you achieve. For example: reduce time spent writing reports by 30%, or respond to donor enquiries faster.

- ☐ **Set a realistic timeline**

  Plan to spend the first three months exploring tools and running small tests, not deploying organisation-wide solutions. Progress over perfection.

People and Skills

- ☐ **Run a short AI literacy session**

  Host a two-hour session for your team covering what AI can and cannot do, where it helps most in your work, and what risks to watch for. Use simple, practical examples.

- ☐ **Identify who will use AI tools**

  List the roles in your organisation that would benefit most. Start with the people doing repetitive tasks: grant writing, donor communications, data reporting.

- ☐ **Create a safe space to experiment**

  Let staff try tools without fear of making mistakes. Errors in a controlled setting are how you learn what works before rolling anything out broadly.

Use Cases and Prioritisation

- [ ] **Identify your top three use cases**

  Use the AI Ideation Canvas in section 4.2 to map out your problems and potential AI solutions. Pick the three that offer the highest impact with the lowest effort.

- [ ] **Start with low-complexity, high-value tasks**

  Focus first on use cases that require no technical integration, such as content drafting, meeting summaries, or email responses. These deliver quick wins and build confidence.

- [ ] **Define what success looks like**

  For each use case, write down a measurable outcome. For example: grant applications drafted in half the time, or donor newsletters produced weekly instead of monthly.

Data and Technology

- [ ] **Audit what data you have**

  List the information your organisation holds: donor records, programme reports, beneficiary data, survey results. Understanding your data helps you identify which AI applications are realistic.

- [ ] **Check what tools you already have access to**

  Many tools your organisation already uses, such as Microsoft 365 or Google Workspace, now include built-in AI features. Check what is already available before buying anything new.

- [ ] **Start with free or low-cost tools**

  Before committing to paid subscriptions, test with free versions. Claude.ai, ChatGPT, Canva AI, and Notion AI all have free tiers that are sufficient for most NPO use cases.

Governance and Ethics

- [ ] **Draft a basic AI use policy**

  Write a one-page document outlining which AI tools staff may use, what data must not be shared with external tools, and how outputs must be reviewed before use. (see Annex 1: AI Use Policy example)

- [ ] **Establish a data privacy baseline**

  Identify which data in your organisation is sensitive (beneficiary data, donor personal information). Set a clear rule: sensitive data does not go into general-purpose AI tools.

- [ ] **Plan for human oversight**

  Every AI output in your organisation should be reviewed by a person before it is acted on or published. AI is a first draft, not a final answer.

Budget and Sustainability

- [ ] **Estimate a small AI experimentation budget**

  You do not need a large budget to start. A realistic monthly budget for an NPO should cover a few tool subscriptions for multiple people to explore these tools.

- [ ] **Look for grant funding for AI capacity**

  A growing number of foundations, including Google.org, Salesforce.org, and various national technology funds, offer grants specifically for NPO digital and AI capacity building.

- [ ] **Plan a review at six months**

  Set a date six months from now to assess what is working, what is not, and where you want to invest next. Document your learnings as you go.

4.2. The AI Ideation Canvas

The AI Ideation Canvas, developed by Give a Day, is a structured tool for identifying and developing AI use cases from within your own organization. It walks you through a thinking process that starts with a real problem you face and ends with a concrete experiment you could run within two weeks.

Use it in a workshop with two to five colleagues or complete it individually when you encounter a bottleneck in your work. You do not need to know anything about AI to fill it in. The goal is to clarify your thinking, not to produce a technical specification.

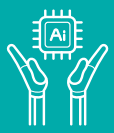

Tips for getting the most from the canvas

- **Be specific in box 2.** "Writing is time-consuming" is too vague. "It takes three staff days per month to compile the quarterly donor report" is specific and actionable.
- **Do not skip box 5.** Constraints are often what determine whether an idea is realistic. Naming them early prevents wasted effort later.
- **Keep box 6 small.** The best first experiments can be run by one person in an afternoon with a free tool. Resist the temptation to plan a full implementation before you have validated the idea.
- **Run more than one.** The canvas works best as a repeatable practice. Complete it for three to five different problems to build a pipeline of ideas you can prioritise over time.

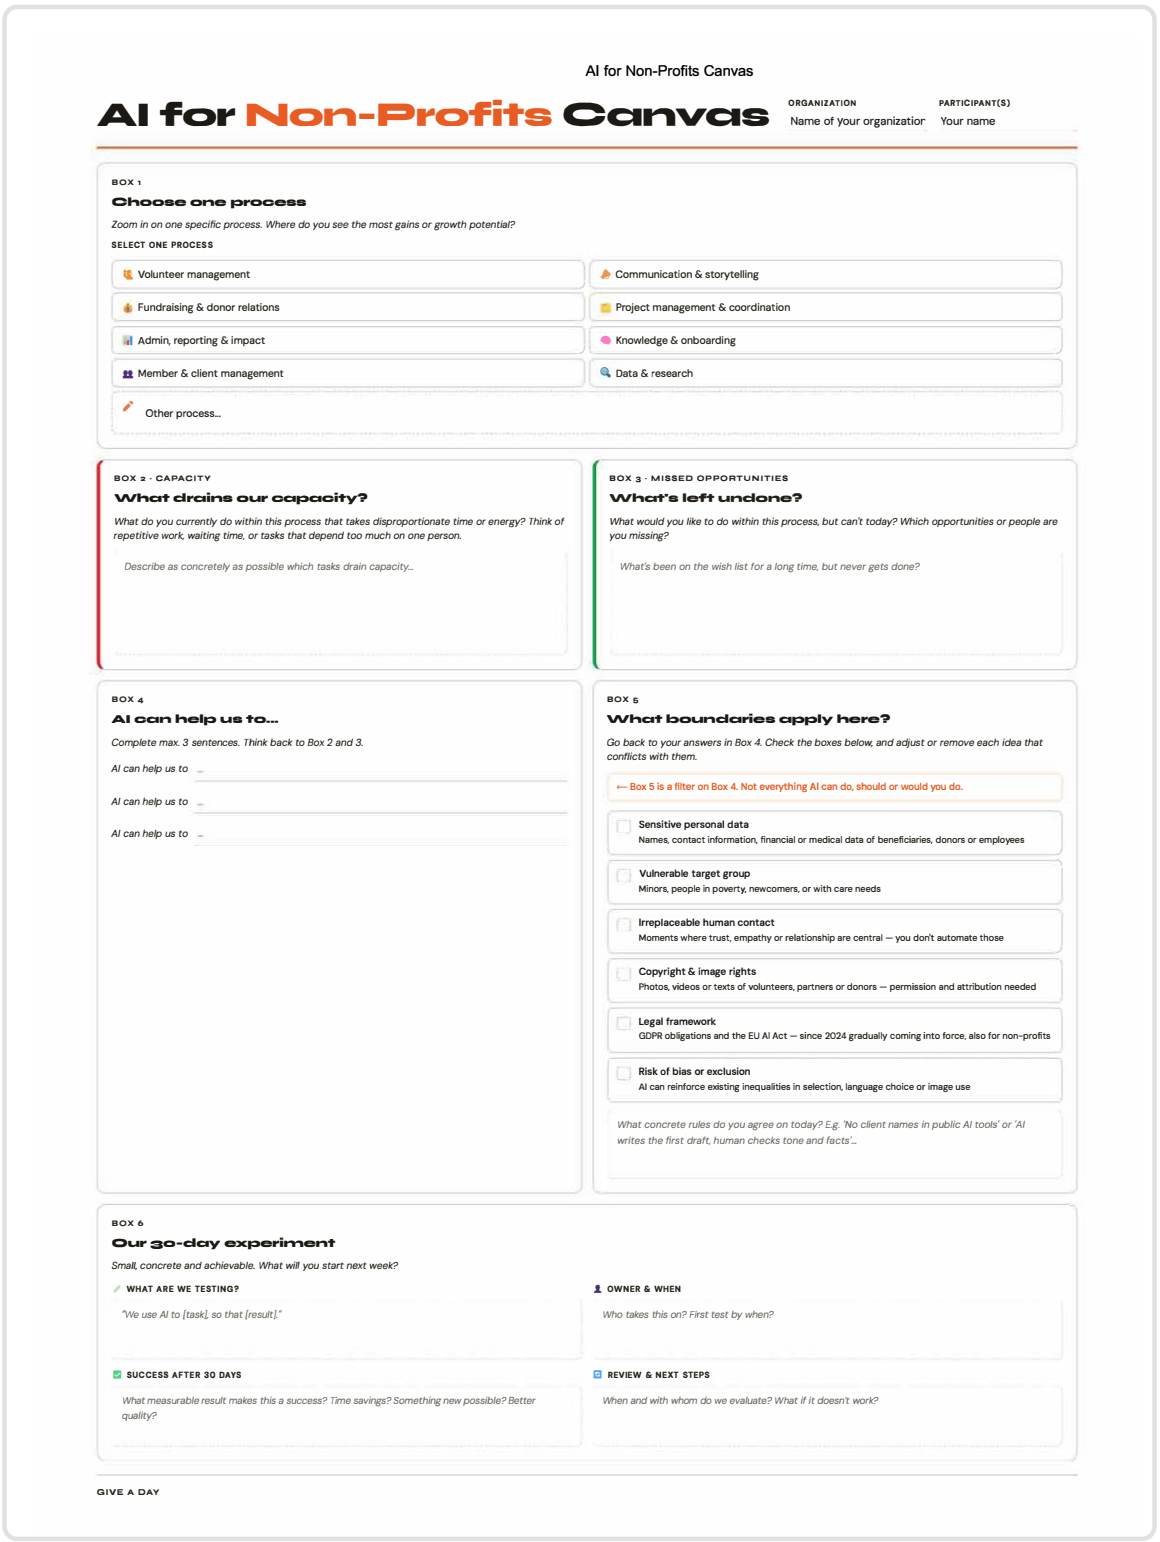

AI for Non-Profits Canvas

AI for Non-Profits Canvas

ORGANIZATION: Name of your organization

PARTICIPANT(S): Your name

BOX 1

Choose one process

Zoom in on one specific process. Where do you see the most gains or growth potential?

SELECT ONE PROCESS

- Volunteer management
- Communication & storytelling
- Fundraising & donor relations
- Project management & coordination
- Admin, reporting & impact
- Knowledge & onboarding
- Member & client management
- Data & research
- Other process...

BOX 2 · CAPACITY

What drains our capacity?

What do you currently do within this process that takes disproportionate time or energy? Think of repetitive work, waiting time, or tasks that depend too much on one person.

Describe as concretely as possible which tasks drain capacity...

BOX 3 · MISSED OPPORTUNITIES

What's left undone?

What would you like to do within this process, but can't today? Which opportunities or people are you missing?

What's been on the wish list for a long time, but never gets done?

BOX 4

AI can help us to...

Complete max. 3 sentences. Think back to Box 2 and 3.

AI can help us to ...

AI can help us to ...

AI can help us to ...

BOX 5

What boundaries apply here?

Go back to your answers in Box 4. Check the boxes below, and adjust or remove each idea that conflicts with them.

← Box 5 is a filter on Box 4. Not everything AI can do, should or would you do.

- ☐ **Sensitive personal data** — Names, contact information, financial or medical data of beneficiaries, donors or employees
- ☐ **Vulnerable target group** — Minors, people in poverty, newcomers, or with care needs
- ☐ **Irreplaceable human contact** — Moments where trust, empathy or relationship are central — you don't automate those
- ☐ **Copyright & image rights** — Photos, videos or texts of volunteers, partners or donors — permission and attribution needed
- ☐ **Legal framework** — GDPR obligations and the EU AI Act — since 2024 gradually coming into force, also for non-profits
- ☐ **Risk of bias or exclusion** — AI can reinforce existing inequalities in selection, language choice or image use

What concrete rules do you agree on today? E.g. 'No client names in public AI tools' or 'AI writes the first draft, human checks tone and facts'...

BOX 6

Our 30-day experiment

Small, concrete and achievable. What will you start next week?

WHAT ARE WE TESTING?

"We use AI to [task], so that [result]."

OWNER & WHEN

Who takes this on? First test by when?

SUCCESS AFTER 30 DAYS

What measurable result makes this a success? Time savings? Something new possible? Better quality?

REVIEW & NEXT STEPS

When and with whom do we evaluate? What if it doesn't work?

GIVE A DAY

4.3. Meet Noor Impact: your AI ideation assistant

At Deloitte, we believe making an impact matters most. That's why we have developed Noor Impact. A prompt you can copy and paste into your favourite AI assistant to help you identify relevant use cases based on your daily tasks and problems. Want to try for yourself? Simply head over to www.deloitte.com/be/ai-for-good, download the system prompt, copy and paste it into your favorite AI assistant, and you can start using it!

Part V
FAQ

These are the questions we hear most often when working with NPOs on AI adoption. The answers are kept deliberately concise. Where a topic is covered in more depth elsewhere in this toolkit, a reference is provided.

Is my data safe when chatting with a public AI tool?

It depends on the tool, the plan, and the contract terms. As a practical rule, never paste personally identifiable or sensitive information about beneficiaries, donors, staff, or volunteers into a public or unapproved AI tool. Use public tools only for non-sensitive work such as drafting public communications, summarizing published materials, or brainstorming. For sensitive use cases, use an approved organizational tool with appropriate contractual safeguards, access controls, retention settings, and a Data Processing Agreement where required.

How can I be sure companies do with my data what they promise?

You cannot fully verify every technical claim yourself, but you can reduce risk through procurement discipline. Ask for a Data Processing Agreement, data-retention terms, subprocessors, training-use policy, hosting location, security certifications, incident-notification commitments, and audit or assurance evidence where available. Keep written answers on file. For high-risk or sensitive use cases, involve legal, privacy, security, or governance support before adoption.

Do we need an on-premise AI solution if we work with personal data?

Not necessarily. The first question is not whether AI is on-premise or cloud-based, but what data is involved, what the tool does with it, and what controls are in place. On-premise AI can provide more technical control, but it is expensive and requires specialist maintenance. For many NPOs, an approved enterprise cloud solution with a Data Processing Agreement, strong access controls, no training on customer data, and anonymization or pseudonymization before upload will be more realistic. On-premise solutions are mainly relevant for organizations with very sensitive data, high volumes, and the capacity to operate the infrastructure securely.

How can we build an AI strategy when the technology is evolving so fast?

Build around your problems, not around the tools. The specific tools available today will continue to change. Your operational challenges will not. The time it takes to write grant reports, the difficulty of responding to beneficiary enquiries at scale, the burden of compiling financial data manually: these are stable problems. A strategy grounded in the problems you want to solve remains relevant even as the tools you use to address them evolve. In practice, this means identifying your three to five priority use cases, running small experiments with whatever tools are available now, and reviewing your approach every six months rather than attempting to create a multi-year technology plan. Treat AI strategy as a living document, not a one-time exercise.

Do we need a technical person on our team to start using AI?

No technical person is needed for the first wave of low-risk AI use cases. Drafting, summarizing, translating, brainstorming, and preparing first versions of communications can usually be tested through simple interfaces, provided staff follow the AI use policy and avoid sensitive data. Technical expertise becomes important when you want to integrate AI with internal systems, automate workflows, connect to databases, handle sensitive information, or deploy a production solution.

Can AI replace our staff?

AI should not replace human accountability, care relationships, or judgement-heavy work in an NPO. It can, however, automate parts of roles and change how work is organized. The best framing is augmentation: use AI to reduce administrative burden, speed up first drafts, structure information, and free staff for relationship-building, field work, advocacy, and decisions that require context and responsibility.

How much does it cost to start using AI?

It can cost very little to start, but costs can rise once you need enterprise plans, governance, training, integrations, or support. Start by checking AI features already included in existing office, CRM, design, or collaboration tools. Then run small experiments with approved tools and non-sensitive data before buying new subscriptions. Budget should include not only licences, but also staff time, training, review, and governance.

Part VI
Apply for a Deloitte AI Ideation session

This toolkit is designed to inspire and guide you. But sometimes, the best way to move forward is with expert support tailored to your specific context.

That's why you can apply for an AI Ideation session for non-profit organizations. Each year, we will select a number of non-profits to participate in a facilitated half-day workshop.

What can you expect?

- **Map your challenges** — Identify where AI can genuinely add value in your operations
- **Explore realistic solutions** — Discover concrete use cases that fit your resources and capacity
- **Build your roadmap** — Create a clear, actionable plan to get started
- **Answer your questions** — Address concerns about data privacy, costs, implementation, and more

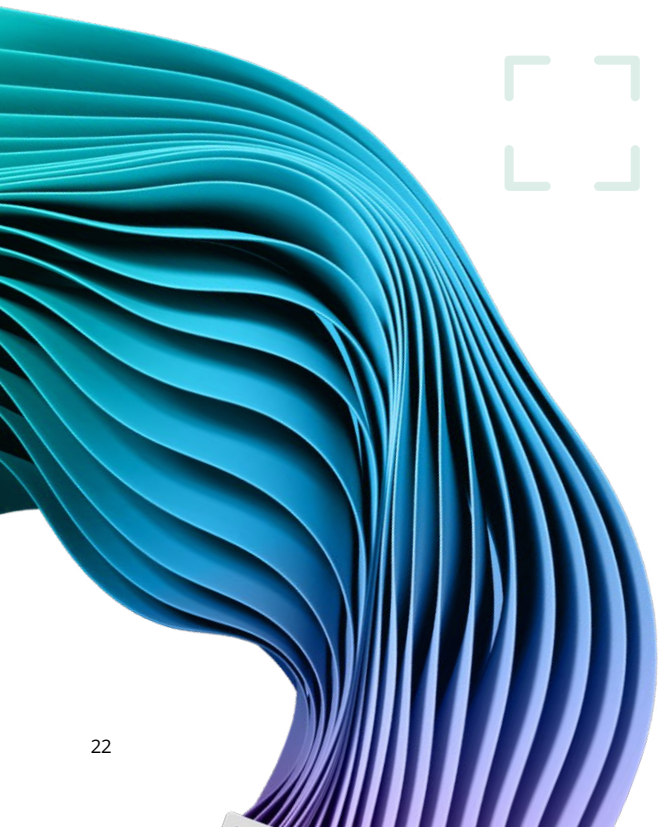

What we need to know

To help us prepare the most relevant session for your organization, please provide the following information:

About Your Organization

- Organization name and mission
- Number of staff (approximate)
- Primary focus area (education, health, community development, advocacy, etc.)

Your AI Readiness

- How familiar is your team with AI? (No experience / Some awareness / Already experimenting)
- Have you identified specific challenges you'd like to address? (If yes, please describe briefly)
- What's your biggest concern about AI? (Cost / Data privacy / Technical complexity / Staff resistance / Other)

Your Capacity

- How many staff members would attend the session?
- Do you have budget allocated for AI tools? (No / Small budget / Flexible budget)
- What's your timeline? (Urgent / Within 3 months / Flexible)

How to apply

Email us at bepurpose@deloitte.be with the information above. Subject line: "AI Ideation Session Application — [Your Organization Name]"

What happens next?

- Our team will review your application and will contact you to schedule a session. Please note that we commit to organizing 3 to 5 ideation sessions per year, so places are limited.

Annex 1: AI Use Policy example

The AI tools referenced in this annex are provided **as examples only** and do not constitute an exhaustive or definitive list of approved solutions. The inclusion of a tool in this annex does not guarantee its suitability for your specific use case, and the exclusion of a tool does not prohibit its use if it meets the approval criteria outlined in the main policy.

1. Approved AI Tools
Staff may use the following AI tools for work purposes only where the specific plan, data settings, and use case have been approved.

Tool	Use Case	Approval Level
ChatGPT / Claude	Drafting, brainstorming, summarizing, research	**Approved**
Microsoft copilot	Document editing, email drafting, presentations	**Approved**
Canva AI / DALL-E	Design and image generation (non-sensitive only)	**Restricted**
ChatGPT Advanced Data Analysis	Data processing and analysis	**Restricted**

New tools: Request approval before use. The responsible manager, IT, privacy, or compliance owner should assess data security, contractual terms, retention, training use, and suitability for the intended use case.

2. Prohibited Data — Never Share with AI Tools

CRITICAL — DO NOT PASTE:
- **Beneficiary information:** Names, contact details, addresses, health data, case histories
- **Donor data: Names, email addresses, donation amounts, financial information**
- **Staff information:** Personal details, salary information, performance reviews
- **Strategic documents:** Unpublished strategies, internal budgets, confidential reports
- **Passwords, API keys, or authentication credentials**

SAFE TO SHARE (after anonymization):
- Beneficiary stories (remove all identifying information)
- Aggregate statistics and anonymized data
- Published reports and public information
- General organizational processes and workflows

3. Mandatory Review Process
Before sharing AI outputs, follow this process based on output type:

- **High-Stakes Outputs (External Communications)**
  Examples: Grant proposals, public statements, social media, website content
  1. AI generates draft
  2. Subject matter expert reviews for accuracy and tone
  3. Manager approves before external sharing
  4. Document review — note that AI was used in creation

- **Medium-Stakes Outputs (Internal Communications**
  Examples: Internal memos, meeting summaries, team presentations, project plans
  1. AI generates draft
  2. Team lead reviews for accuracy
  3. Ready to use (no manager approval needed)

- **Low-Stakes Outputs (Personal Productivity)**
  Examples: Email drafts, brainstorming notes, personal research, formatting
  1. AI generates content
  2. You review for relevance
  3. Ready to use (no additional approval needed)

Documentation: For high-stakes outputs, note in the document: "This content was created with AI assistance and reviewed by [Name, Role]."

4. Consequences & Support

Violations: Sharing prohibited data with AI tools may result in:
- Immediate suspension of AI tool access
- Mandatory retraining
- Escalation to management and compliance team
- Potential disciplinary action

Support & Questions:
- Tool access issues? Contact IT: [IT email]
- Unsure if data is safe to share? Ask your manager or compliance team
- Need training? Attend monthly AI literacy sessions

Policy Review: This policy will be reviewed and updated every 6 months based on emerging risks and new tools.

Deloitte.