BC

IIIBC RESEARCH · NON-PROFIT AI ADOPTION

Non-profit AI adoption: Money and skills slow it down, missing safeguards make it fail

For executive directors and boards deciding how their organization should use AI. We reviewed 147 graded sources and 32 cases to find what holds non-profits back and what separates success from failure.

Published: November 27, 2026 **Evidence:** 147 sources · 32 cases

Window: January 2023 to November 2026

4%
of non-profit staff and executives call AI fully integrated

53%
report staff using AI tools outside approved channels

38%
of executives have written guidance on safe AI use

90%
of foundations give grantees no AI implementation support

Sources: NTEN and The Bridgespan Group, *2026 State of Nonprofit AI*, September 2026 (917 staff and executives; executive items, 404). Center for Effective Philanthropy, *AI With Purpose*, September 2025 (209 foundation leaders). [C-001–C-003, C-006–C-007, C-031, C-033]

Almost every non-profit now uses AI, but very few have built it into how they work. Non-profits blame money, skills, and privacy fears, and the surveys back them up. The cases show something different: what decides whether AI goes well is whether people check its output, test it before scaling it, and tell the people it affects. Nine of the 10 measured successes kept a person reviewing the output and tested the tool first. Eleven of the 13 failures skipped telling or asking the people affected. Leaders should write those three safeguards down before they buy more tools. [I-008, I-010, I-013–I-014, I-023]

01 · ADOPTION

Use is common, but almost no one has built AI into how they work

Every grade-A survey that measures full integration puts it between 4 and 9 percent (Exhibit 1). In the NTEN and Bridgespan survey of 917 non-profit staff and executives, fielded from late April to mid-June 2026, 4 percent call AI fully integrated, 33 percent run organized pilots, and 37 percent use AI only informally or individually. The UK Charity Digital Skills Report 2026 found 4 percent of 807 charities using AI strategically. [C-001, C-003–C-005, C-023–C-024, C-038, C-040]

Use runs ahead of rules. In the same survey, 53 percent of respondents report shadow AI use, meaning staff using AI tools outside approved channels. Only 38 percent of executives say written guidance on safe use is in place, and only 24 percent say their data is organized enough to support AI. We read this as an operating gap, not a technology gap: staff already have the tools, but no one has decided how the organization should use them. [C-006–C-008]

EXHIBIT 1

Full integration is rare in every grade-A survey

Survey	Population	Sample	Measure	Result
NTEN and Bridgespan, 2026	US-led non-profit staff and executives	917	AI fully integrated	4%
Charity Digital Skills Report, 2026	UK charities	807	Strategic AI use	4%
Humanitarian Leadership Academy, 2025	Humanitarian workers, 144 countries	2,539	AI widely integrated	8%
Humanitarian Leadership Academy, January 2026	Humanitarian practitioners	1,729	AI widely integrated	9%

Sources: NTEN and Bridgespan, September 2026; Charity Digital Skills Report, July 2026; Humanitarian Leadership Academy and Data Friendly Space, August 2025 and March 2026. All self-reported; each survey defines "integrated" differently. [C-001, C-003, C-023–C-024, C-037–C-040]

02 · BARRIERS

Privacy fears, skills, and money slow adoption, and small organizations fall furthest behind

Privacy and security top the list of major barriers. Among the 404 executives in the NTEN and Bridgespan survey, 55 percent call data, privacy, or security concerns a major barrier, ahead of every other factor (Exhibit 2). The Center for Effective Philanthropy (CEP) found the same concern at 67 percent of 451 US non-profits in April and May 2025. [C-002, C-009, C-028–C-029]

Skills come next. In CEP's survey, 62 percent of non-profits say none or only a few of their staff understand AI well. In the UK, the share of charities naming limited skills as a barrier rose from 43 percent in 2025 to 56 percent in 2026. Money is the most widely felt constraint: 79 percent of NTEN and Bridgespan executives call cost a major or partial barrier, 57 percent have no AI budget, and 41 percent spent under $5,000 on AI in 2025. [C-015, C-018–C-019, C-025, C-030]

Size widens every gap. Written AI guidance is in place at 27 percent of organizations with budgets under $10 million, against 56 percent of larger ones. In Imagine Canada's 2025 survey, 5 percent of organizations with income under $500,000 have an AI policy, against 40 percent of those over $25 million. Its regression analysis of 963 organizations adds a useful distinction: staff time and knowledge predict whether an organization starts using AI, while funding predicts how far use spreads. Small organizations need rules and skills first; money matters more later. [C-020, C-034, C-036]

EXHIBIT 2

Privacy is the barrier executives most often call major

Barrier	Major barrier	Somewhat a barrier
Data, privacy, or security concerns	55%	36%
Staff capacity or time to learn	34%	51%
Trust or equity with communities served	32%	n/a
Limited understanding of use cases	30%	49%
Concern about biased outcomes	29%	n/a
Cost or funding constraints	27%	52%
Low confidence in output quality	20%	57%
Lack of leadership alignment	16%	n/a

Source: NTEN and The Bridgespan Group, September 2026, 404 non-profit executives, fielded late April to mid-June 2026. "n/a" marks values we didn't extract. [C-009–C-014, C-016–C-017]

AI policy rates depend on who you ask

Policy rates range from 10 percent to 38 percent across grade-A surveys (Exhibit 3). Question wording, survey year, organization size, and sample explain most of the spread, so we don't average them. Every survey agrees on the direction: most non-profits use AI without written rules. [C-007, C-026, C-035, C-041, C-043, C-045]

EXHIBIT 3

One in 10 to four in 10 non-profits have AI rules, depending on the survey

Survey	Population	Sample	Date	Result
Imagine Canada	Canadian charities and non-profits	963	Summer 2025	10% formal policy
Association Societies Alliance	US association professionals	305	November 2024	13% AI policy
Infoxchange	Australian and New Zealand not-for-profits	824	November 2025	14% policy or guideline
Charity Digital Skills Report	UK charities	807	2026	18% formal policy
Humanitarian Leadership Academy	Humanitarian practitioners (individuals, self-selected)	1,729	January 2026	23% report a formal policy (local 13%, international 31%)
NTEN and Bridgespan	Non-profit executives	404	2026	38% written guidance in place

Sources linked in the table. Each survey asks a different question, so compare directions, not exact rates. [C-002, C-007, C-023, C-026, C-034–C-035, C-039, C-041–C-045]

03 · FUNDING

Funders are largely absent

Funders are the least engaged party. In CEP's 2025 survey, 90 percent of 209 foundation leaders say they give no AI implementation support to any grantee, and 83 percent of non-profits say none of their funders have talked with them about AI. In the 2026 NTEN and Bridgespan survey, 63 percent of executives received no AI-specific funding in any form. The share of UK charities receiving digital funding from trusts and foundations fell from 30 percent to 17 percent in a year. [C-021, C-027, C-031–C-033]

Large new philanthropic money for AI mostly goes elsewhere. The Humanity AI coalition of 10 foundations committed $500 million over five years in October 2025, mainly to public-interest AI and governance, not to non-profits' own adoption. Where restricted AI grants do reach non-profits, executives report they go mostly to infrastructure (54 percent), then training (24 percent) and governance (10 percent). [C-022, C-046]

We think this is the widest open gap in the sector: non-profits need a plan a funder can back, and funders lack the confidence to judge one. In Project Evident's March 2025 study of 38 AI-active foundations, as reported by the Chronicle of Philanthropy, nearly two-thirds of program officers said they lack confidence assessing AI proposals. An organization that brings a costed plan with risk controls will stand out. [C-047, C-075]

04 · SUCCESSES

Successes share two habits: A person reviews the output, and the tool is tested before it scales

Nine of the 10 measured successes kept a person between the AI and the final output. Nine tested the tool against a baseline or a control group before scaling it (Exhibit 4). The cases rarely credit leadership sponsorship or funding. Safeguards and measurement did the work. [I-008, I-013–I-014]

The clearest examples come from think tanks, associations, and convening organizations:

> **International Conference on Learning Representations (ICLR), a convening non-profit.** It ran a randomized trial across about 45,000 peer reviews in its 2025 cycle. About one in four reviewers who got AI feedback revised their reviews, and blinded raters judged quality improved in 89 percent of cases. Acceptance rates didn't change. The AI only advised; reviewers decided (ICLR, April 2025). [C-052–C-054]
>
> **Pew Research Center, a Washington, DC think tank.** It checked AI coding of open-ended survey answers against 1,200 answers coded by people before using it, and reached 89 percent agreement. It publishes an AI-use policy, bars respondents' personal data from public AI tools, and discloses AI use in its methods (Pew Research Center). [C-055, C-070]

Association for the Advancement of Artificial Intelligence (AAAI), an association. It added one labeled AI review for each of 22,977 submitted papers, at under $1 per paper, without replacing any human reviewer. Of 5,834 survey responses, 53.9 percent found the AI reviews useful, against 20.2 percent who didn't (AAAI). [C-056–C-057, C-071–C-072]

Citizens Advice, a UK advice service. Its SORT assistant drafts answers to advisers' questions, and a supervisor approves each one. Supervisor time per query fell from 10 minutes to four (Computing, 2025). [C-058]

MIT Solve, a convening and fellowship program. AI screening cut first-round review of 2,901 applications from about 24 working days to about 10. A human audit agreed with 70 of 80 AI decisions (MIT Solve). [C-059–C-060, C-073]

Most success evidence is self-reported. Only ICLR's trial and an outside survey of Digital Green's farmer chatbot rest on independent measurement, so treat outcome sizes as indicative. The full list is in the case library.

EXHIBIT 4

Successes rely on human review and testing; failures skip disclosure and consent

Condition	Credited in successes (of 10)	Missing or a cause in failures (of 13)
Human review and safeguards	9	7
Scoped test with metrics before scaling	9	3
Disclosure, consent, and consultation	4	11
Data rules and AI policy	4	6
Leadership and mission fit	1	4
Skills and training	3	1
Funding	2	1
Technology infrastructure	1	1

Source: IIIBC case library, November 2026. Counts are cases, not survey respondents; one case can carry several conditions. Failures are the 13 non-profit pitfalls; the four adjacent cases are excluded. [I-008, I-010–I-011, I-013–I-028]

05 · FAILURES

Failures share one blind spot: The people affected weren't told or asked

In 11 of the 13 non-profit failures, the people affected weren't told about the AI use, didn't consent to it, or weren't consulted before it. Surveys rarely name this as a barrier, but it appears in more failures than any other cause. The pattern runs across every function: [I-010, I-023]

Programs. The National Eating Disorders Association took its "Tessa" chatbot offline on May 30, 2023, after users reported it advised counting calories and restricting their diet. Koko sent about 4,000 people replies partly written by AI, marked only with a short note, and drew criticism for testing on vulnerable people without informed consent. [C-066]

Data. Crisis Text Line announced on January 31, 2022, that it had ended data sharing with its for-profit spinoff Loris.ai, after public criticism that people in crisis couldn't meaningfully consent. It predates our window, but it's still the precedent most often cited.

Fundraising and campaigns. Charity Right's unlabeled AI images of children led viewers to suspect fraud, and by its own figures its team photos drew more clicks per impression than the AI images. [C-067] Amnesty International and Plan International Netherlands withdrew or revised AI imagery after criticism.

Community and governance. NaNoWriMo's September 2024 AI statement, issued without consulting its community, led to resignations and lost support before the organization closed in 2025. Seattle Worldcon 2025 vetted proposed panelists with ChatGPT without telling them, then re-vetted everyone by hand after resignations.

Publications. The Chicago Sun-Times, owned by the non-profit Chicago Public Media, printed an AI-generated summer reading list of books that don't exist on May 18, 2025. In 2026, consulting firms withdrew published reports with AI-invented citations. For think tanks and associations whose product is credible research, this is the most transferable risk.

Public opinion explains why disclosure matters. In Independent Sector's July 2026 trust report, based on 3,000 US adults surveyed April 3–20, 2026, 37 percent would trust a non-profit they support less if it used gen AI widely, and 22 percent would trust it more. In a YouGov survey of 1,291 US adults fielded August 20–21, 2026, 55 percent called fictional AI stories or images of the people a charity helps unacceptable. [C-048–C-051]

Missing human review drives severity. Both high-severity non-profit cases, NEDA's chatbot and a pending lawsuit alleging that Sharp HealthCare recorded patients with an AI scribe without consent, involved vulnerable people and no human check or consent.

WHAT TO DO

What leaders should do now

Each action answers one finding above. None needs new technology.

01 **Publish a short AI-use rule before you buy more tools.** Name what data may go into AI tools, which uses need human review, and who owns AI decisions. Your staff likely already use the tools: 53 percent of survey respondents report shadow use. [C-006]

02 **Tell and ask the people affected.** Disclose AI use in content, campaigns, and services, and consult members, beneficiaries, and donors before AI touches them. This was the most common failure in the cases.

03 **Keep a person between the AI and anything that leaves the building.** Review every published citation, every message to a vulnerable person, and every decision about an applicant.

04 **Test one use against a baseline before you scale it.** Set a metric, such as time saved, error rate, or agreement with experts, and run a scoped pilot. Nine of the 10 successes did this. [I-008, I-014]

05 **Take a fundable plan to your funders.** With 83 percent of non-profits never asked about AI by a funder, organizations that bring a costed plan with risk controls will stand out. [C-032]

CASE LIBRARY

Thirty-two cases, each traced to a public source

A success needs a named organization, AI in regular use for at least six months or across a team, a measured result against a baseline, and grade-A or grade-B evidence. Cases that meet the first two tests on weaker evidence are listed as promising. A pitfall is a case where AI use led to harm, a public withdrawal, legal or regulatory action, a data incident, documented reputational damage, or an abandoned project. [I-007–I-011]

EXHIBIT 5

Ten successes: Each kept people in control and measured the result

Organization	What the AI did	Reported result	Grade
Citizens Advice (UK)	Drafts answers to advisers' questions; a supervisor approves each	Supervisor time per query fell from 10 to 4 minutes [C-058]	B
Pew Research Center (US)	Codes open-ended survey answers	89% agreement with 1,200 human-coded answers [C-055, C-070]	A
AAAI (US)	Adds one labeled AI review per submitted paper	22,977 papers at under $1 each; 53.9% of respondents found it useful [C-056, C-071–C-072]	A
ICLR (US)	Gives reviewers optional feedback on their reviews	Randomized trial: quality improved in 89% of rated cases; acceptance unchanged [C-054]	A
MIT Solve (US)	Screens challenge applications in the first round	About 24 working days cut to about 10; humans agreed with 70 of 80 decisions [C-060, C-073]	B
TED (US)	Dubs talks into other languages; volunteer translators review them	Views on dubbed talks up 115% [C-061]	B
Tech for Campaigns (US)	Helps staff draft fundraising emails	Three to four times more dollars raised per work hour [C-074]	B
Médecins Sans Frontières	Antibiogo app reads antibiotic resistance tests	88.8% agreement with a blinded expert across 378 samples [C-062]	A
Jacaranda Health (Kenya)	Triages mothers' text messages for risk	About 17% more accurate than its earlier model; help-desk workload down about 12% [C-063]	B
Digital Green (Kenya)	Farmer.Chat answers farmers' questions	61% of 450 surveyed users showed meaningful use, against a 27% benchmark [C-064]	B

EXHIBIT 6

Five promising cases still lack measured results

Organization	What the AI does	What's missing	Grade
Full Fact (UK)	Monitors claims for fact checkers; tools used by 29 organizations in 14 countries [C-065]	Scale is reported, not a measured effect	C
International Rescue Committee	Signpost AI agents answer information requests	Productivity gains reported without a baseline	C
Lone Star Legal Aid (US)	Three in-house chatbots for staff and the public	No outcome published	C
UPR Info (Switzerland)	Extracts human rights recommendations from UN documents	Before and after workloads differ, so not a like-for-like comparison	C
American Society of Clinical Oncology	Guidelines assistant for members	No usage or outcome figures published	C

EXHIBIT 7

Thirteen non-profit pitfalls: Most began with people who weren't told or asked

Organization	What happened	Severity	Grade
National Eating Disorders Association	Helpline chatbot withdrawn in May 2023 after it gave dieting advice	High	B
Crisis Text Line	Ended sharing crisis-conversation data with an AI spinoff after criticism (2022 precedent)	Medium	B
Koko	Sent about 4,000 people AI-assisted support replies without clear consent [C-066]	Medium	B
Amnesty International	Removed AI-generated images of real protests after criticism	Low	B
Plan International Netherlands	Criticized for AI images of girls in a child-marriage campaign; since advises against AI images of individual children	Low	C
Charity Right (UK)	Unlabeled AI images of children led viewers to suspect fraud; team photos performed better [C-067]	Low	B
Chicago Public Media	Its Sun-Times printed an AI-generated reading list of books that don't exist	Medium	B
Wikimedia Foundation	Paused AI article summaries after volunteer editors objected	Low	B
NaNoWriMo	AI statement issued without consulting members led to board resignations before closure	Medium	B
Seattle Worldcon 2025	Vetted proposed panelists with ChatGPT without telling them; re-vetted by hand	Medium	B
Sharp HealthCare	A pending lawsuit alleges an AI scribe recorded patients without consent; the claims are unproven	High (alleged)	B
Pratham (India)	India's human rights commission issued notices over children's data in an AI education project; no breach reported	Medium	B
Australian university academics	A Senate inquiry submission included AI-invented allegations against two firms; the authors apologized	Medium	B

EXHIBIT 8

Four adjacent cases show the same failures outside the sector

Organization	What happened	Lesson for non-profits	Grade
KPMG and EY	Withdrew published reports with fabricated citations in 2026	Check every citation before research goes out	B
Michigan Unemployment Insurance Agency	An automated system issued fraud findings without human review; a state review of a sample found 93% wrong [C-069]	Never let a system decide about a person alone	B
State Bar of California	Used AI to draft licensing-exam questions without telling the state Supreme Court	Associations that certify members must disclose AI use	B
South Coast air district comment campaign	Public comments were sent in the names of people who say they didn't send them (allegations)	Advocacy tools that write in members' names need their consent	B

Grades: A, primary data with a stated method or an official record; B, reputable reporting of primary data, or an organization's own figures without a stated method; C, vendor or sponsored material, or a named case with no numbers. Severity follows the effect on people, not on the organization.

METHOD AND LIMITS

How we did this research, and what it can't tell you

We ran 202 logged web searches in November 2026, split evenly between success and failure queries, and kept 147 sources: 31 grade A, 73 grade B, 39 grade C, and four grade D. Survey statistics cover January 2025 onward; cases cover January 2023 onward. We coded each case against eight conditions, from leadership to technology, and marked each as an enabler or a barrier. A second reviewer then checked every headline figure in this report against its source and a random 20 percent of the source register. All headline figures matched; three of 29 sampled rows needed a corrected title or wording, and none changed a finding. [I-001–I-006, I-029–I-032]

Every figure in this report carries a marker such as [C-006]. It links to a row in the evidence register that names the source, quotes the passage word for word from a saved copy, and gives that copy's SHA-256 fingerprint. Counts about the research itself, such as "9 of 10 successes", are recounted by script from the case files.

The evidence has five limits, and we'd rather state them than overclaim:

Self-report. Most survey figures are perceptions from samples that weren't drawn at random. Most success results come from the organizations themselves.

Publication bias. Failures are under-reported, and no near miss met our definition.

Segment gaps. We found no survey of Washington, DC non-profits or of advocacy organizations, no think-tank or association failure, and one association success. Two of the three convening successes are AI research conferences, which know AI unusually well.

Causation. Exhibit 4 shows patterns across 23 non-profit cases. [I-012] It suggests which conditions matter; it doesn't prove them.

Dropped figures. We dropped three widely repeated statistics we couldn't trace to their stated source, including a claim that 76 percent of non-profits have no AI policy. The source it's attributed to measures AI strategy, not policy.

SOURCES

Principal sources

NTEN and The Bridgespan Group, *2026 State of Nonprofit AI: Adoption and Governance Report* https://content.nten.org/wp-content/uploads/2026/09/2026-State-of-Nonprofit-AI.pdf, September 10, 2026. 917 respondents, including 404 executives; fielded late April to mid-June 2026.

Center for Effective Philanthropy, *AI With Purpose* https://cep.org/wp-content/uploads/2025/09/CEP_AI_Layout_FINAL.pdf, September 2025. 451 non-profits and 215 foundations; fielded April and May 2025.

Charity Digital Skills Report 2026, AI findings https://charitydigitalskills.co.uk/report/detailed-findings/artificial-intelligence/ and digital funding findings https://charitydigitalskills.co.uk/report/detailed-findings/digital-funding/, July 2026. 807 UK charities.

Imagine Canada, *The State of Artificial Intelligence Adoption in Canadian Nonprofits* https://imaginecanada.ca/sites/default/files/the-state-of-ai-adoption-in-canadian-nonprofits-en.pdf, 2025. 963 organizations; fielded June 26 to September 6, 2025.

Humanitarian Leadership Academy and Data Friendly Space, 2025 insights report https://www.humanitarianleadershipacademy.org/wp-content/uploads/2025/08/Full-insights-report-How-are-humanitarians-using-AI-in-2025.pdf (2,539 respondents) and January 2026 pulse survey https://www.humanitarianleadershipacademy.org/wp-content/uploads/2026/03/Final-Humanitarian-AI-January-2026-pulse-survey-research-briefing-note-1.pdf (1,729 respondents).

Association Societies Alliance, *AI Usage in Associations* https://fsae.memberclicks.net/assets/docs/AI-in-Associations-Report.pdf, November 2024. 305 association professionals.

Infoxchange, *Digital Technology in the Not-for-Profit Sector* https://www.infoxchange.org/sites/default/files/infoxchanges_2025_digital_technology_in_the_not-for-profit_sector_report.pdf, November 2025. 824 organizations in Australia and New Zealand.

Independent Sector, *Trust in Nonprofits and Philanthropy Report* https://independentsector.org/wp-content/uploads/2026/07/Trust-in-Nonprofits-and-Philanthropy-Report-July-2026.pdf, July 2026. 3,000 US adults; fielded April 3–20, 2026.

YouGov, "AI use in the charity sector: where people in the US draw the line" https://yougov.com/en-us/articles/55996-ai-use-in-the-charity-sector-where-people-in-the-us-draw-the-line, November 19, 2026. 1,291 US adults; fielded August 20–21, 2026.

Project Evident, *Funding the Future: Grantmaker Strategies in AI Investment* https://projectevident.org/wp-content/uploads/2025/03/EXECUTIVE-SUMMARY-Funding-the-Future_-Grantmaker-Strategies-in-AI-Investment.pdf, March 2025. 38 AI-active foundations.

Case sources are linked from each organization in Exhibits 5 to 8. Every figure's source, quote, and page is in the evidence register evidence.html.

Find your organization's gap in 30 minutes

The AI Strategy Diagnostic is a free 30-minute call built on five questions. You leave with one named gap and one action for this week. It isn't a pitch.

> *One of the five: "Has your leadership defined what AI is not allowed to decide without human review?"*

Book the diagnostic: `cal.com/iii-bc-zwbptm` `hello@iiibc.com`