Grant Thornton

Charities and AI: Governance needs to catch up with use

2026



Introduction

Charity audit and risk committee chairs tell us the debate has moved on from whether to use artificial intelligence (AI) to how to use it well. The harder questions now are who is checking the answers, and whether boards know how AI is really being used.

AI has arrived in the charity sector. Not through a grand strategy, but through software updates and the laptops of staff who often got there before their boards did.

On 15 September 2026, we hosted a roundtable of charity audit and risk committee chairs to discuss what that means. Participants joined from organisations of very different shapes and sizes: learned and professional bodies, health and medical research charities, and charities focusing on conservation and environmental missions.

We discussed how AI is being used, the risks it brings and where the real opportunities lie. As always, to encourage an open discussion, we worked under the Chatham House Rule, and therefore have not attributed comments to individuals or their organisations.

Key takeaways from the roundtable

- The question is no longer whether to use AI, but how to use it well, with governance and controls keeping pace
- Cyber risk is accelerating, and what counts as good enough today may not be in 12 months. AI helps attackers. It makes phishing emails more convincing, speeds up the search for weaknesses and lowers the skill needed to launch an attack
- Human judgement remains the control that matters most
- Chairs want practical, proportionate tools to help them ask the right questions

From ‘should we?’ to ‘how should we?’

Most chairs described a clear change in tone over the past two years, with a sharper focus on getting AI governance and controls right.

The uses are broad. Charities are using AI to analyse large datasets, which is particularly valuable in scientific settings, as well as to generate content, summarise documents and bring together knowledge from different sources. One charity described using AI to produce meeting minutes, exploring chatbots for its helpline and mining information from research funding applications. It is also weighing whether its years of data could create a commercial opportunity and help diversify its income in a way that complies with UK data protection law.

Working with a university team, one conservation charity in the room has developed AI that monitors wildlife in the field in real time. Instead of researchers spending days reviewing footage after the event, the system alerts them within about 30 seconds when a target species walks past a camera.

One other charity uses Microsoft Copilot to speed up subject access requests, which used to be slow and labour-intensive. The aim, the chair stressed, isn’t to cut jobs but to make existing teams more efficient; freeing up time to focus on other things.

Several chairs were clear that AI is a tool, not a substitute for expertise. One charity, which made AI the focus of its trustee strategy day this year, compared it to having a new graduate on the team: useful, but someone qualified needs to check the work.

When everyone has AI

The consensus in the room was that AI doesn’t only help charities. It helps everyone who deals with them, and not always in welcome ways.

One chair explained that AI is making it easier for the public and stakeholders to raise complaints and submit requests for information, which creates a new burden for staff. Fake research has become a problem across academia, and at one organisation the team handling fraudulent submissions has grown significantly. There was a consensus that time saved in one place is being spent in another.

Fundraising is feeling similar pressure. When anyone can produce a funding application in seconds, funders are deluged and charities face more competition for attention.



The risks on the audit committee's agenda

Cyber security is moving fastest

Cyber was the risk raised most often. One chair warned that what's good enough now probably won't be in 12 months. Budgets are rising as a result. One institution has judged its cyber budget too low and is reviewing it regularly. It is also moving cyber and digital services to an external provider to strengthen security. But every pound spent on IT is a pound not spent on the charitable mission, and as one chair observed, you'll probably never feel you've spent enough.

Another chair asked a question many boards will recognise. If AI-enabled attacks start defeating multi-factor authentication tomorrow, how do audit committees give their boards assurance that they've done everything they reasonably can? Alex Hunt, our Data & AI Services Director, pointed out that AI also helps defenders: it can find an organisation's own vulnerabilities and help prioritise which gaps to close first.

Data, regulation and reputation

Chairs highlighted data privacy, data security, regulatory compliance and the protection of intellectual property. The shared worry was loss of control. Once personal, beneficiary or funder information goes into an AI tool, can the charity say where it goes, who can see it, whether it trains the model and how long it is kept? Reputation was a recurring concern. If confidential material reaches the public domain through AI, the damage can be fast and hard to reverse. Trust takes years to build and seconds to lose.

From our experience, practical starting points are to classify data and set which tools can be used for which types of data, to prefer approved tools with contractual data protections over free versions, to run data protection impact assessments for new uses involving personal data, and to check supplier terms on training and retention.

Trusting the output

Model accuracy and over-reliance came up again and again. One chair described using AI to analyse a large dataset and getting an answer that was out by an order of magnitude, caused by a misplaced bracket in a spreadsheet formula. They only spotted it because they knew roughly what the answer should be. Another had seen AI-generated recommendations that sounded sensible but weren't based on the available information.

That raises a wider question about skills. If experienced people catch errors because they know what sensible looks like, how do new entrants build that judgement when AI does the groundwork? One chair worried about losing tacit knowledge: the understanding that sits in people's heads and hasn't been written down. It is easily lost when decisions are shaped by whatever documents are fed into a tool. The message was consistent: use AI to do what people can already do, faster, in an augmented way, and value people for their judgement.

Shadow AI and the AI you didn't know you bought

Even well-controlled organisations face 'shadow AI': staff using tools outside approved channels, often on personal devices. Chairs were sceptical that blocking tools works. Their IT providers say broad blocks can get in the way of normal work, and blocking one tool simply sends people to another.

Suppliers are another gap. Almost every new piece of software now includes AI, and chairs questioned whether organisations really understand what's embedded in what they buy. One chair argued that procurement teams should now take the lead, assessing AI suppliers on transparency, security, intellectual property, data rights, service continuity and ethics, not just cost and functionality.

Cost and climate

Two further risks are easy to overlook. AI uses significant amounts of energy and water, which matters to organisations whose missions address environmental harm. And chairs expect costs to rise as AI providers seek returns. Rising costs are driving interest in small language models, which are cheaper to run than general-purpose tools for focused tasks.

Front of house and back of house

One chair offered a way of thinking about this that the group found useful. The 'back of house' is technical safeguarding: cyber security, model testing and access controls. With the right experts in place, it's comparatively easy to manage. The 'front of house' is where AI meets people: staff, members, volunteers and the communities a charity serves. Training goes some way, but behaviour and culture vary widely across disciplines and countries, and they're much harder to control.

Alex added that back-of-house uses give staff a relatively safe place to build confidence, while front-of-house uses are more likely to change what a charity achieves for its beneficiaries, and carry greater reputational risk. Several chairs said they keep AI in the background for now, with a person reviewing anything before it reaches a member, donor or customer.

On controls, Alex outlined options beyond written policy. Data loss prevention tools act like a sniffer dog, flagging confidential information before it leaves the organisation. New governance tools give visibility of AI agents, which chain automated actions together and whose steps can be hard to trace. Enforcing rules through software tends to work better than relying on documents alone.

Where the opportunities lie

Despite the risks, the mood was energetic around using AI. Nobody suggested standing still.

A trustee who recently joined a small charity with very few systems described the chance to design it almost like a start-up: imagining what it could look like in five years with the tools now available, then working backwards.

Others were clear that AI must serve the mission, not redefine it. For the conservation charity, the purpose is connecting people with nature, and a screen can't replace that. But there's a limit to how many people can engage in person, and reaching people further afield means digital channels. In five years, the chair argued, there won't be digital without AI. The health charity described itself as people-first: it wants AI to help its people deliver services, not replace them. Its challenge is moving from piecemeal use cases to a coherent strategy.

Measurement was another theme. Chairs asked whether productivity gains can be proven, and whether trustees should push for benefits to be targeted and measured, using reliable data. Fundraising is likely to be one of the first areas where AI delivers value, through better use of data, content and targeting. But data quality comes first.

Questions for audit and risk committees

Drawing on the discussion, here are questions audit and risk committees may want to ask:

- Do we know where and how AI is being used across the organisation, including outside approved tools?
- Is our AI policy written down, communicated and understood, and is it backed by technical controls?
- Do we understand the AI built into the software and services we buy?
- How are AI outputs checked before they're relied on or reach people outside the organisation?
- Is our cyber budget keeping pace with the threat, and how would we respond to an AI-enabled attack?
- Are we measuring the benefits, and are they worth the financial, reputational and environmental cost?

Thank you to everyone who took part. If you'd like to talk about how AI is affecting your organisation, please contact Nicholas Burton at nicholas.g.burton@uk.gt.com or anyone in the Not for Profit team.



GRANTTHORNTON.CO.UK

© 2026 Grant Thornton UK Advisory & Tax LLP. All rights reserved.

'Grant Thornton' refers to the brand under which the Grant Thornton member firms provide assurance, tax and advisory services to their clients and/or refers to one or more member firms, as the context requires. Grant Thornton UK Advisory & Tax LLP is a member firm of Grant Thornton International Ltd (GTIL). GTIL and the member firms are not a worldwide partnership. GTIL and each member firm is a separate legal entity. Services are delivered by the member firms. GTIL does not provide services to clients. GTIL and its member firms are not agents of, and do not obligate, one another and are not liable for one another's acts or omissions. This publication has been prepared only as a guide. No responsibility can be accepted by us for loss occasioned to any person acting or refraining from acting as a result of any material in this publication. DTSK-10722

Grant Thornton