BOARD.DEV

DELL Technologies

Building the Leadership Nonprofits Need to Make AI Work

Nonprofit Tech Leadership Insights Brief | May 2026

A co-publication of Board.Dev and Dell Technologies

Data sources: national survey conducted with TechSoup, qualitative interviews, and secondary research from McKinsey, Salesforce, Stanford HAI, Nonprofit Tech for Good, IBM, Bridgespan, Technology Association of Grantmakers, and SHRM.

Executive Summary

Nonprofits are not stuck with AI because they lack interest. **They are stuck because they lack the leadership infrastructure to convert experimentation into durable, governed, and economically rational practice.**

This brief synthesizes findings from a national survey of 180 nonprofits, supplemented by qualitative interviews and secondary research, to understand where nonprofits stand in their adoption of AI and related technologies and what is holding them back from moving forward. While interest and experimentation are widespread, adoption remains uneven and fragile. This report focuses on the leadership, capacity, and governance conditions that determine whether AI experimentation becomes safe, scalable, and mission-aligned.

NOTE

Some of what we found confirmed what practitioners already know. Some of it didn't. The surprises highlight some of the most actionable insights.

KEY FINDINGS

FINDING 01

Adoption is constrained by capacity, not resistance. The top barriers are lack of staff skills (46%), ethical concerns (41%), and lack of funding (28%)—not skepticism. Organizations are not rejecting AI outright. They want to adopt it responsibly and don't yet have the infrastructure or confidence to do that without putting something at risk.

FINDING 02

The sector is stuck in an experimentation plateau. 87% of organizations surveyed are either not using AI organizationally or stuck in early experimentation—only 8% have moved into pilots or scale. Informal use is widespread, but it rarely translates into shared workflows, sanctioned tools, or measurable outcomes.

FINDING 03

Governance barriers start with the organization. 58% of organizations have no AI principles at all. Not incomplete ones, none. Closing that gap requires someone to make the first move—and the data shows it is almost always leadership, not the board. The organizations making progress aren't waiting for board pressure to initiate the conversation. Leadership is building the case and creating the conditions for governance before bringing it to the board.

FINDING 04

Boards determine whether adoption becomes safe, scalable, or stalled. Boards are rarely the spark for experimentation. However, their involvement matters to durable, sustainable adoption. Among organizations where staff and board co-developed AI principles, 60% have reached pilots or scaling, compared to just 5% of those with no principles. **That's a 12× difference.**

12×

DIFFERENCE BETWEEN AI-ENGAGED BOARDS COMPARED TO THE REST

FINDING 05

Size shifts where organizations get stuck, not why. Smaller nonprofits can't see themselves in AI yet; larger ones can see it but can't move. The result is the same: widespread experimentation without institutionalization.

FINDING 06

Board composition predicts adoption outcomes—but presence isn't the same as fluency. Only 24% of boards include members with technology or data backgrounds. Among organizations with no tech-fluent board members, 62% are not using AI organizationally. Among those where a majority are tech-fluent, that number drops to 31%—the presence of even a few credible voices on technology questions appears to cut inaction in half. But the goal is not to recruit a "tech person" and check a box. Organizations with a technology professional on their board who cannot evaluate AI strategy have the appearance of governance without the substance. The organizations moving forward are those where board members, whatever their backgrounds, can engage with AI as a leadership and accountability question, not a technology briefing.

FINDING 07

Cost is the through-line. The most common mistake we see is treating AI as a line item rather than a capacity generator. Organizations are already paying the cost of not adopting AI in staff hours, duplicated effort, and slow reporting. A mid-sized nonprofit delaying adoption is losing the equivalent of 1.2-1.6 full-time employees in unrealized annual capacity. But time is only part of what's at stake. Organizations without AI principles are accumulating liability quietly. Those without digital fluency are becoming less competitive for grants and talent. And every year of inaction widens the gap with peer organizations that are already building the institutional knowledge to make AI work. The cost of inaction is larger—and broader—than most organizations have calculated.

2,400-3,200

HOURS OF UNREALIZED CAPACITY ANNUALLY

The cost of inaction for an average nonprofit is equivalent to 1.2-1.6 full-time employees, or $72,000-$96,000 per year in value your existing staff could be generating.

The opportunity—for partners, funders, and leadership-focused interventions—is to help nonprofits **cross the inflection point from experimentation to institutionalized, governed practice.** This brief provides data, frameworks, and case studies to guide that work.

A NOTE ON TERMINOLOGY

Throughout this report, "AI" refers to technology that learns from data to automate tasks, make predictions, or generate content. It works through three elements: **data** (what organizations already collect), **models** (systems that recognize patterns), and **infrastructure** (the hardware and tools that run them). Most organizations already have all three in some form. The barrier is rarely the technology—organizations often underestimate what they already have and overestimate what they need to start.

Foreword from Dell Technologies

At Dell Technologies, we've long believed that true impact happens when technology, skills, and innovation come together to drive change. Across the nonprofit sector, we see organizations motivated to explore AI, yet constrained by the same barriers: limited capacity, unclear governance, and no shared language for what AI is or what it should enable in their specific context.

This report makes that reality plain: nonprofits are not resisting innovation; they're navigating it without the leadership infrastructure needed to adopt AI safely, sustainably, and in service of mission.

The case studies in this brief show what's possible when leadership pairs curiosity with structure. Whether it's Hopeworks improving employment outcomes or organizations reclaiming staff hours through smarter workflows, the early wins demonstrate that AI's value is measured in real-world results, not in how advanced the technology appears. What moves organizations forward isn't technical sophistication; it's clear problem definition, responsible governance, and a willingness to start small and learn quickly.

Dell Technologies is proud to partner with Board.Dev on this work because it centers what matters most: giving nonprofit leaders and boards the confidence, guardrails, and practical steps to adopt AI responsibly and sustainably. Our hope is that this brief empowers organizations to move from experimentation to intentional, mission-aligned use of AI, grounded in clarity, not complexity.

- For Dell Technologies' framework on AI infrastructure, governance, and workforce readiness at scale, see Dell's AI Blueprint: delltechnologies.com.
- For more on how Dell Technologies gives back to communities, visit dell.com/en-us/lp/dt/community-impact.

Cara Borawski
Head of North America Giving and Impact

Methodology

This brief draws on three data sources:

National survey.

180 nonprofit leaders responded to a structured survey distributed through Board.Dev and TechSoup networks between August 2025 and January 2026. Respondents span 44 states, with representation across human services (23%), education (16%), health (11%), youth development (7%), arts (5%), workforce (4%), and other issue areas. Budget sizes range from under $500K (11%) to over $20M (4%), with the largest segment at $1-5M (44%).

Qualitative interviews.

10+ structured interviews with nonprofit executives, including case study subjects featured in this report, conducted between November 2025 and February 2026. Interviews explored leadership decision-making, board engagement, AI adoption experiences, and barriers to progress.

Secondary research.

Findings are benchmarked against the following sources.

SECONDARY SOURCES

- **McKinsey & Company.** "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value." *May 2024.*
- **Salesforce.** "State of IT, 4th Edition–AI and App Development." *2024.*
- **Stanford Institute for Human-Centered AI.** "AI Index Report 2025." *April 2025.*
- **Nonprofit Tech for Good.** "2023 Nonprofit Tech for Good Report." *2023.*
- **IBM.** "Global AI Adoption Index 2024." *January 2024.*
- **IBM.** "Cost of a Data Breach Report 2025." *July 2025.*
- **Technology Association of Grantmakers.** "2024 State of Philanthropy Tech Survey." *October 2024.*
- **Bridgespan Group.** "Closing the Nonprofit Funding Gap in the Age of AI." *2025.*
- **Society for Human Resource Management.** "Employee Turnover and Replacement Costs." *shrm.org. 2024.*

Survey responses were analyzed for frequency distributions, cross-tabulations by organization size and governance posture, and correlation patterns between board engagement and adoption outcomes. All percentages are rounded.

Cost-of-inaction modeling draws on function-level AI efficiency findings from McKinsey's *"The Economic Potential of Generative AI: The Next Productivity Frontier"* (June 2023), applied to a typical nonprofit staffing profile. The resulting estimates represent Board.Dev's modeled projections. Because McKinsey's benchmarks are drawn from corporate contexts, the model applies conservative estimates to account for sector differences in staffing, workflow complexity, and AI tool adoption.

What We Found

The story of nonprofit AI adoption in 2026 is not a story of resistance to new technologies, but a story of capacity.

FINDING 01

Adoption is constrained by capacity, not resistance

What we hear consistently from nonprofit leaders is this: we want to do this, we just don't have anyone who can own it. The data confirms it. The top barriers aren't related to skepticism—they're about capacity.

A small but notable minority told a different story: for a handful of respondents—primarily in arts and education—saying no to AI is currently an active expression of mission. These respondents cited data center energy consumption or fundamental misalignment with organizational values; they represent a different and legitimate conversation about what responsible adoption actually means. For the vast majority of the 41% who cited ethical concerns as a barrier, however, the issue is not opposition to AI but uncertainty about how to adopt it responsibly—a desire not to opt out, but to get it right.

BARRIER	% OF RESPONDENTS
Lack of staff skills	46%
Ethical concerns	41%
Lack of funding	28%
Do not have time to prioritize	26%
Do not see a use case	23%
Unclear governance	22%
Regulatory concerns	16%
Data quality	14%
Lack of board understanding or buy-in	13%

FINDING 01

Adoption is constrained by capacity, not resistance

CONTINUED

These are not the answers of organizations that reject AI. They are the answers of organizations that are stretched thin. The top three barriers—skills, ethics, and funding—tell a consistent story: nonprofit leaders want to get started, but they need confidence that they can do it responsibly, and that someone will help them if it goes wrong.

Ethical concerns rank second—higher than funding, and far higher than regulatory worries. This is significant. It suggests that nonprofits are not simply worried about cost; they are worried about getting it right. In a sector that serves vulnerable populations, handles sensitive data, and operates under public trust, that caution is appropriate. But it also means that without clear guidance and governance, caution can become paralysis.

How does this compare to the broader market? Global AI adoption benchmarks from McKinsey and Salesforce show 88-89% of organizations reporting some AI use. In our survey, roughly half of nonprofits have started organizational AI use—a gap of nearly 40 percentage points. But the reason for that gap is not attitude. It is capacity: the skills, funding, and leadership needed to move from individual experimentation to organizational practice.

A note on infrastructure: much of the public conversation about AI adoption focuses on compute, storage, and technical readiness. Those issues matter. But for the organizations in this survey, infrastructure isn't the binding constraint. The barrier sitting above unclear governance, above data quality, and above lack of board buy-in is skills: 46% of respondents cite it first. The technical foundation can be built once an organization knows what it's building toward. Leadership clarity comes first.

> **We're not resistant to AI. We're resistant to doing it badly. And right now, nobody is helping us understand the difference.**
>
> **— Nonprofit CEO, survey respondent**

FINDING 02

The sector is stuck in an experimentation plateau

Nearly every nonprofit we spoke with is experimenting with AI in some form. Almost none of them have moved beyond pilots. The picture that emerges from the data isn't one of reluctance; it's one of organizations stuck at the starting line.

STAGE OF ADOPTION	% OF RESPONDENTS
Not using AI at an organizational level	46%
Early experimentation	41%
Multiple pilots	6%
Scaling across multiple functions	2%

Nearly 87% of respondents are either not started or stuck in early experimentation. Only 8% have moved into pilots or scale. This is not a curiosity gap; it's an execution gap. Individual staff members are using ChatGPT to draft emails and summarize meetings. But that informal use rarely translates into shared workflows, sanctioned tools, or measurable organizational outcomes.

The experimentation plateau matters because it is where organizations spend time and energy without capturing returns. Staff experiment on their own; leadership does not know what is being used, how, or whether it is safe; successes are not replicated, and organizations don't learn from their peers so they can move faster. And because there are no shared goals or metrics, nobody can demonstrate ROI, which in turn makes it harder to justify the investment needed to move forward.

The experimentation plateau: 87% of nonprofits surveyed are either not using AI organizationally or stuck in early experimentation. Only 8% have reached pilots or scale.

One interview subject told us that organizations are "curious and open," but without dedicated support, their experimentation stays individual and informal. As Zenae Campbell, CEO of Boys & Girls Clubs of the Austin Area, described, the challenge is not willingness. It is that no one is leading the work in a directed, contextualized way.

FINDING 03

The governance gap starts with leadership ownership

If the first half of the story is about capacity, the second half is about who decides what to do with that capacity—and whether anyone is clearly responsible for the answer.

Here is what the governance gap actually looks like in practice: staff are using AI tools right now, today, without any organizational framework for how to do it safely. 58% of the organizations in this survey have no AI principles at all. Not incomplete principles—none.

AI PRINCIPLES STATUS	% OF RESPONDENTS
No AI principles	58%
Working on creating principles	21%
Staff-developed, not yet board approved	8%
Staff-developed, board approved	8%
Co-developed by staff and board	3%

What the data shows is that governance doesn't initiate itself. Before a board can approve AI principles, set guardrails, or authorize investment, leadership has to bring the conversation to the table. The organizations that have made progress on governance are not waiting for board pressure to initiate it; they are building the case, defining the questions, and creating the conditions for board engagement to be productive.

Rubber Stamps Aren't Enough

There's a further wrinkle in the governance data: several organizations that report board-approved AI principles also report that no board members are tech-fluent. A signature is not the same as understanding. If boards are approving AI policies they cannot evaluate, that creates a false sense of security that may be more dangerous than no policy at all. Getting a policy signed is not the finish line. If the board doesn't understand what they just approved, the governance is an illusion.

FINDING 03

The governance gap starts with leadership ownership

CONTINUED

Board involvement in AI conversations reflects this pattern. Across our sample, 25% of boards are not involved at all in technology or AI conversations, while 65% are involved only occasionally. Just 8% have AI as a standing agenda item, and only 2% have a dedicated committee.

When asked directly, 71% of organizations said their board had not discussed AI-related risks—privacy, bias, safety, reputation—in the past 12 months. And that matters, because risk conversations are one of the clearest triggers for boards to shift from passive oversight to active governance. When boards are not even aware of the risk landscape, they cannot provide the oversight that makes AI adoption safe and accountable.

FINDING 04

Boards determine whether adoption becomes safe, scalable, or stalled

79% of organizations say board engagement has had no noticeable effect on their AI adoption. That sounds like boards don't matter. The cross-tabulations say the opposite. Board engagement determines whether AI risk—including cybersecurity exposure, data privacy, and liability—gets actively managed or quietly ignored. The organizations in this survey that lack AI principles aren't just missing a governance document; they're missing the mechanism that triggers security conversations at all.

BOARD INVOLVEMENT LEVEL	% THAT HAVE DISCUSSED AI RISKS
Not at all	9%
Occasionally	32%
Regularly (standing agenda item)	62%
Very actively (dedicated committee)	100%

The pattern is clear: boards that are more engaged are dramatically more likely to have addressed the governance questions that unlock safe scaling. And when we look at the relationship between AI principles and actual adoption, the effect is even more striking.

FINDING 04

Boards determine whether adoption becomes safe, scalable, or stalled

CONTINUED

AI PRINCIPLES STATUS	NOT USING AI	EARLY EXPERIMENTATION	PILOTS OR SCALING
No principles	**61%**	**34%**	**5%**
Working on principles	**42%**	**53%**	**5%**
Staff-developed (not board approved)	**33%**	**47%**	**20%**
Staff-developed (board approved)	**13%**	**80%**	**7%**
Co-developed by staff and board	**20%**	**20%**	**60%**

12×

Among organizations where staff and board co-developed AI principles, **60% have reached pilots or scaling.** Among those with no principles, just 5%. **That is a 12× difference.**

Boards are not the engine of experimentation. But they are the gate between informal use and institutional scale. The data suggests a permission catch-22: organizations need principles to deploy AI safely, but they struggle to get principles in place until the board is fluent enough to engage. Breaking that cycle—by giving leadership teams the language and frameworks to bring AI to the board as a governance issue, not a technology request—is where the highest leverage lies.

> *I presented AI to the board as a question about organizational health and accountability, not a question about tools. That reframe changed the whole conversation.*
>
> **— Elaine Lofton, CIO (ret.), African American Youth Harvest Foundation**

FINDING 04

Boards determine whether adoption becomes safe, scalable, or stalled

CONTINUED

When Board Engagement Clogs Progress

11 organizations in this survey said board engagement has actively slowed their AI progress. When organizations are running pilots or in early experimentation, staff is often moving faster than the board's fluency can keep up with, and when that low-fluency board suddenly 'wakes up' to AI risk, the default is often to hit the brakes. Uneducated board engagement can be more paralyzing than board disinterest. This is precisely why leadership has to build the case before asking for permission–and why boards must embrace a learning agenda for AI and the implications for their specific organizational context.

FINDING 05

Size shifts where organizations get stuck, not why

Organization size clearly affects where nonprofits are in their AI journey—but the overall pattern of experimentation without institutionalization holds across the board.

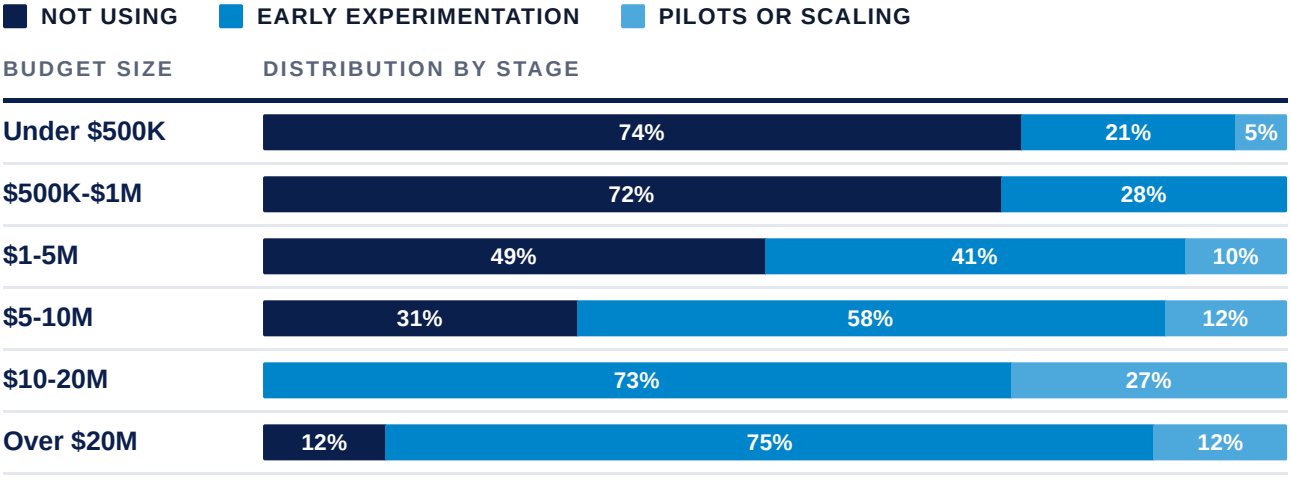

For organizations under $1M, the barrier can be imagination or inspiration: a third cite 'no use case,' a response that virtually disappears among larger organizations. Mid-sized organizations ($1-5M) have figured out how AI is relevant to their work but hit an execution wall; staff skills and time peak as barriers here. Larger organizations ($5M+) come closest to global benchmarks on experimentation, but face the biggest challenges converting that activity into governed, institutionalized practice.

The through-line across all size bands is the same gap the data shows everywhere: activity does not lead to institutionalization. The organizations that break the pattern, regardless of size, tend to share one thing: someone in leadership made a deliberate decision to connect AI activity to a governance structure before the board asked them to.

FINDING 06

Board composition predicts adoption outcomes—but presence isn't the same as fluency

The data points to a direct relationship between who sits on a board and whether an organization moves forward on AI. Across our sample, only 24% of boards include members with technology or data backgrounds. A majority of respondents (54%) say less than half of their board is tech-fluent, while 23% have not evaluated their board's tech fluency. And 7% report that no board members are tech-fluent at all.

The adoption consequences are significant. Among organizations with no tech-fluent board members, 62% are not using AI at an organizational level. Among those where a majority of board members are tech-fluent, that number drops to 31%. The presence of even a few board members who can engage credibly on technology questions appears to cut organizational inaction in half.

> **The IT Guy Fallacy**
>
> Presence doesn't equal fluency. 24% of surveyed organizations have technology or data professionals on their boards—yet many of those same organizations rate their board's AI fluency as zero or "not sure." The pattern reflects an old model: organizations recruited a tech professional to the board to handle IT procurement or website infrastructure, and assume that covers AI readiness. It doesn't. Traditional IT expertise and AI strategy fluency are different skills. Having a tech seat on your board is a starting point, not a finish line.

The implication for organizational leaders is precise: the goal is not to recruit a "tech person" to the board and check a box. It is to build a board that can engage with technology as a strategic and governance question—one where members understand enough to ask the right questions, push back on the right assumptions, and recognize when AI decisions require board-level attention.

Dan Rhoton of Hopeworks describes this in practice. His board doesn't engage with AI as a technology topic. They engage with a dashboard of outcomes—using blue lines and red lines to track job placements, starting wages, and 90-day retention—and AI is embedded in the strategy that moves those lines. What his board provides is not technical expertise. It is enough tech fluency to pair accountability to results.

FINDING 06

Board composition predicts adoption outcomes—but presence isn't the same as fluency

CONTINUED

> **Our experience is that tech fluency is pretty easy to acquire, especially as a nonprofit. I can call Dell. I can call Deloitte.**
>
> — Dan Rhoton, Hopeworks

Zenae Campbell of Boys & Girls Clubs of the Austin Area found a different entry point. A prospective donor revealed before a meeting that they had used an AI tool to research the organization—and that the output had already begun shaping their giving decision. "There was a lot of shock. True shock," Campbell recalls. That story became the moment that moved her board from passive awareness to active engagement. Not a technology briefing, but a story about what was already happening in their own donor relationships.

Elaine Lofton of African American Youth Harvest Foundation brought her own CIO background into the room, but arrived at the same conclusion: the conversation that moved her board was not about tools. "I presented it to the board as a question about organizational health and accountability, not a question about tools," she says. "That reframe changed the whole conversation."

The organizations moving forward on AI are not necessarily the ones with the most technically sophisticated boards. They are the ones where board members—whatever their backgrounds—have the language, the framing, and the context to engage with AI as a leadership question.

FINDING 07

Cost is the through-line

The most persistent negative perception we encounter in this work is treating AI as a line item rather than a capacity generator. Organizations are already paying the cost of not adopting—in staff hours, duplicated effort, and slow reporting. They just aren't calling it that.

But the data points to a catch-22. Nonprofits report a lack of time and staff capacity as their top barriers to using AI effectively. Their hesitance means they are missing out on the very tools that could materially reduce administrative burden, streamline workflows, and reclaim staff hours. In practice, the cost of not adopting AI, including lost time, duplicated effort, and unmanaged risk, is often higher than the cost of responsible adoption.

Modeling the cost of inaction

Applying AI efficiency benchmarks to common nonprofit functions—program delivery, communications, fundraising, finance, and HR—we estimate that a typical organization delaying adoption is losing the equivalent of 1.2-1.6 full-time employees in unrealized annual capacity. That gap is not theoretical. It reflects what peer organizations are already recapturing through responsible AI adoption.

Using function-level AI efficiency findings from McKinsey & Company's *The Economic Potential of Generative AI: The Next Productivity Frontier* (2023), Board.Dev modeled the potential time savings for a typical mid-sized nonprofit ($1-5M budget, 20-40 staff) that has not yet adopted AI organizationally. The model applies conservative and moderate estimates of time savings across common nonprofit functions, and the resulting projections are Board.Dev's own:

FUNCTION	TYPICAL STAFF	CONSERVATIVE HOURS SAVED/YEAR	MODERATE HOURS SAVED/YEAR
Fundraising & development	3-5	400-500	**600-800**
Finance & administration	2-4	300-400	**500-600**
Communications & marketing	2-3	350-450	**500-650**
Programs & operations	8-15	800-1,000	**1,200-1,500**
HR & people management	1-2	150-200	**250-350**
Executive & strategy	1-2	100-150	**150-250**
TOTAL	**17-31**	**2,100-2,700**	**3,200-4,150**

FINDING 07
Cost is the through-line

CONTINUED

We use a conservative-to-moderate estimate of 2,400-3,200 hours lost per year (the midpoint of our conservative scenario through the lower bound of our moderate scenario) to anchor the cost-of-inaction model in relatable, not aspirational, figures. It reflects the gap between what organizations are doing manually and what peer organizations are already accomplishing with responsible AI adoption.

$72,000-$96,000

PER YEAR

The estimated cost of inaction for a mid-sized nonprofit: unrealized capacity from delayed AI adoption, equivalent to **1.2-1.6 full-time employees.**

Two challenges that keep nonprofits stuck

The "no time" challenge

Organizations say they cannot prioritize AI because they are too busy. But the busyness itself is partly a function of the manual processes they are tied to that AI could automate. The tool that would save time is being deprioritized because of a lack of time. Breaking this cycle requires leadership to carve out dedicated capacity for a bounded pilot—even a small one—let it run till it's gotten the kinks out, and then measure the time returned.

The "no money" challenge

Organizations say they cannot afford AI tools. But they are already bearing the cost of not having them—in duplicated effort, slow reporting, manual data entry, and staff burnout. Until leadership frames AI as a cost-recovery investment rather than a new line item, funding will remain the barrier.

The cost scales with organizational size

The time-loss model is not uniform. Smaller organizations lose proportionally more capacity per staff member. Larger organizations lose more in total dollars. Across all sizes, the gap between current practice and AI-enabled peers compounds annually.

FINDING 07
Cost is the through-line

CONTINUED

ORGANIZATION SIZE (BUDGET)	ANNUAL HOURS LOST	FINANCIAL VALUE OF LOST CAPACITY	FTE EQUIVALENT
Small (<$1M)	600-900 hours	**$18,000-$27,000**	0.3-0.5 FTE
Mid-Size ($1M-$5M)	2,400-3,200 hours	**$72,000-$96,000**	1.2-1.6 FTE
Large ($10M+)	8,000+ hours	**$240,000+**	4+ FTE

Time is only part of what's at stake

The hours-and-salary model captures the most measurable cost of inaction. But it leaves out several compounding risks that organizations rarely price in when deciding whether to prioritize AI adoption.

Risk exposure

Organizations without AI principles or data governance policies are accumulating additional liability quietly. Staff are using AI tools right now, without organizational oversight, and 61% of organizations in our survey report concern that AI will expand their cybersecurity vulnerability, yet most have not acted on that concern at the policy level. That gap has a measurable cost: IBM's 2025 *Cost of a Data Breach Report* found that organizations with high levels of shadow AI use face an additional $670,000 in breach costs on average–and 1 in 5 breaches now involves unsanctioned AI tools. For a nonprofit serving vulnerable populations, exposure isn't just financial. A breach involving client data or an AI misuse incident can cost more than the organization can absorb, but the reputational damage can be more costly still.

Grant competitiveness

A growing number of funders are beginning to ask grantees how they govern AI, not just whether they use it. As the sector matures, digital readiness is moving from a nice-to-have into a signal of organizational health and grant-readiness. (Technology Association of Grantmakers, 2024; Bridgespan, 2025.) As AI fluency becomes a standard expectation, organizations that have not begun to demonstrate it face a quiet eligibility risk—not always rejection, but taking second place in funding relative to peers who can show they are ready to steward tech-enabled programs responsibly.

FINDING 07
Cost is the through-line

CONTINUED

Talent attrition

Staff who are not developing AI fluency in their current roles will seek it elsewhere. This is especially acute for program and data staff, who increasingly view AI literacy as a core professional skill. The cost of replacing a mid-level nonprofit employee is estimated at 50-75% of annual salary for entry to mid-level roles, according to SHRM; it's a turnover risk that is rarely attributed to tech stagnation, but increasingly connected to it. This is not just an organizational problem; it reflects a sector-wide gap in AI skills infrastructure that no single nonprofit can solve alone, and that funders and workforce partners are increasingly being asked to address.

The compounding gap

Every year of inaction is a year of organizational learning that peer organizations are accumulating. Early adopters are building institutional knowledge, refining workflows, and developing board-level oversight capacity. The gap between them and organizations that are still debating whether to start does not stay constant—it widens. The cost of catching up grows with each cycle.

Case Studies

The following organizations are not waiting for perfect conditions. Each has found a way to move from experimentation to action—driven by leadership clarity, practical governance, and a focus on results over tools.

FEATURED ORGANIZATIONS

CASE 01

Hopeworks Camden

When the current economic order doesn't work for you, disruption is the strategy

CASE 02

Boys & Girls Clubs of the Austin Area

From curiosity to capacity: building the case for guided AI adoption

CASE 03

African American Youth Harvest Foundation

A retired CIO brings tech governance to a community nonprofit

CASE 04

Norby: AI for Language Learning

How a social purpose small business used AI infrastructure to scale from 20 to 40+ languages

HOPEWORKS CAMDEN

When the current economic order doesn't work for you, disruption is the strategy

THE CHALLENGE

The mission of Hopeworks is to move chronically unemployed young adults into life-changing jobs. When AI emerged as a potentially transformational technology, Executive Director Dan Rhoton did not see risk; he saw an opportunity to overturn an economic order that has never worked well for the young adults Hopeworks serves. The question was not whether to adopt AI, but how to evaluate if—and where—it worked.

THE APPROACH

Hopeworks started with a specific, measurable problem: mock interviews. Human-led mock interviews were inconsistent, could not be customized for different job descriptions, and required new volunteers for every role. AI could do it better. They partnered with Dell Technologies to deploy an AI-powered mock interview tool that uses digital assistants, ran it with real participants, and tracked the outcomes against their existing metrics.

What happened next was decisive. The data showed **30% higher starting salaries** and higher placement rates for participants who used the AI tool. Armed with that evidence, Hopeworks expanded rapidly: training young adults to use AI to compensate for resume gaps (writing, professional jargon, proposals), launching AI services for other nonprofits and financial institutions, and building on-premise AI tools for sensitive compliance and data work.

Organization Details

SECTOR
Workforce development

LOCATION
Camden, NJ (expanding to Newark, Baltimore)

BUDGET
$5M-10M

SCALE
3 locations; 93% 12-month job retention; avg. starting salary $43K+

AI FOCUS
Mock interviews, workforce training, AI services for hire

THE GOVERNANCE MODEL

Rhoton does not bring “AI initiatives” to his board. He brings a dashboard of outcomes—a set of blue and red lines tracking job placements, starting wages, and retention rates—and the board conversation centers on where the actual performance is diverging from where it should be. AI is embedded in strategy, not bolted on as a separate topic. When a board member who is a former Fortune 500 CEO asks about a dip in placements, the response is grounded in data, not technology jargon.

> *“We have no particular confidence in AI. But what we do have confidence in is our ability to collect, manage, and evaluate data very fast.*
>
> **— Dan Rhoton, Executive Director, Hopeworks**

KEY INSIGHT

Hopeworks illustrates Finding 6 in practice: the organizations absorbing the highest cost of inaction are the ones still treating AI as a discretionary add-on. When an AI director is budgeted alongside an HR hire, and both are presented as infrastructure for growth, the question shifts from “can we afford this” to “what are we losing without it.” That reframe is a leadership decision before it is a budget decision.

WHAT HOPEWORKS LEARNED FROM DELL TECHNOLOGIES

Working with Dell Technologies shifted Hopeworks’s sense of what was possible. As Rhoton describes it: the operational rigor at Dell Technologies—how they think about repeatability and scale—became a model for Hopeworks’s own operations. The exposure to a global technology company’s pace and standards helped Hopeworks reimagine what scale could look like for a nonprofit.

BOYS & GIRLS CLUBS OF THE AUSTIN AREA

From curiosity to capacity: building the case for guided AI adoption

THE CHALLENGE

Under CEO Zenae Campbell's leadership, BGCAA underwent a transformation: from an organization that could not even track its own metrics to one serving 19,000 young people a year with clear data visibility. As AI entered the conversation, Campbell was curious and open, but the organization lacked the dedicated expertise and consultation to move from curiosity to a concrete plan.

THE APPROACH

Rather than treating AI as a standalone initiative, BGCAA framed it as part of organizational health. Campbell identified that the most valuable first step was not buying tools but getting consultation: understanding where the organization stood, what it already had, and what was actually needed. Through the Dell Technologies AI Day of Learning for Nonprofits program and Board.Dev's cohort, Campbell began building the vocabulary and confidence to move forward.

The first concrete application was an AI fundraiser agent, part of a 25-organization pilot with Blackbaud. The agent manages donor outreach for a segment of small-to-mid-level donors that BGCAA's development team previously had no bandwidth to contact. The team gave the agent a persona, trained it on the organization's voice, and set up a contact cadence—all while keeping the risk profile low, since these were donors the organization was not previously reaching at all.

Organization Details

SECTOR
Youth development

LOCATION
Austin, TX

BUDGET
$5M-10M

SCALE
19,000 young people served annually; 2,000+ daily

AI FOCUS
Organizational operations, fundraising, workforce planning

THE GOVERNANCE MODEL

Campbell's board became engaged through storytelling, not a formal AI strategy presentation. She shared a story from a donor visit where the prospective donor revealed they had used an AI tool to research BGCAA before the meeting, and the AI's output had shaped their giving decision. That story created a boardroom "aha": if donors are using AI to evaluate nonprofits, then what the organization publishes about itself—its website, its data, its public reporting—is now being read by machines that shape funding decisions before any human conversation happens.

> *"I'm okay with somebody coming in to lead and pull. That is okay for us right now in this space, because this is not our strength.*
>
> **— Zenae Campbell, CEO, Boys & Girls Clubs of the Austin Area**

KEY INSIGHT

BGCAA puts a name to Finding 1: the barrier is not resistance, it is directed capacity. Campbell is not waiting to be convinced that AI matters. She is waiting for the right expert guidance, the right context, the right starting point for her specific organization. That is a very different problem than skepticism—and it has a different solution. Federated models like Boys & Girls Clubs also present a structural opportunity: when one chapter figures something out, every chapter benefits.

AFRICAN AMERICAN YOUTH HARVEST FOUNDATION

A retired CIO brings tech governance to a community nonprofit

THE CHALLENGE

African American Youth Harvest (AAYH) Foundation is a 20-year-old nonprofit serving as a one-stop shop for families in Austin. When CIO (retired) Elaine Lofton—who spent 30 years in state government IT, including running the technology infrastructure for the Texas Department of Public Safety—began to explore AI for the organization, she initially shared the skepticism many leaders feel. Her first impression was that AI was primarily being used by students to write essays for them.

THE APPROACH

Once Lofton dug deeper, her technical background allowed her to see AI's potential for operational decision-making. She began with a concrete test: using ChatGPT to model staff utilization across multiple program sites, incorporating travel time, campus schedules, and student volumes. The result gave her, for the first time, a clear picture of where she had available capacity and where she was stretched thin, without paying for enterprise resource management software.

From there, AAYH expanded into grant research (using AI to match funding opportunities to the organization's mission), program data tracking (measuring client touchpoints across services to prove the "one-stop shop" model with data), and youth programming (including autonomous robotics using lane detection and an AI-powered greenhouse project through their Girls Who Game program, developed in partnership with Dell Technologies and Microsoft).

Organization Details

SECTOR
Youth and family services

LOCATION
Austin, TX

BUDGET
$1M-5M

SCALE
4,000+ families served directly; 14,000-16,000 indirect touchpoints

AI FOCUS
Staff planning, grant research, data reporting, STEM youth programs

THE GOVERNANCE MODEL

Lofton took a risk-tiered approach. For the organization's mental health and trauma recovery work, which involves HIPAA-protected data and clients fleeing domestic violence, she maintains strict controls and personally reviews any data that could be exposed. For program-side operations and fundraising, she is more flexible, allowing experimentation within defined boundaries. She engaged the board through a combination of storytelling and practical demonstrations, with Lofton framing AI as a question of accountability and organizational health, not technology adoption.

> *"At this point in time, we can't afford to be late on AI and we can't be careless at the same time. The future belongs to organizations that pair innovation with governance.*
>
> **— Elaine Lofton, CIO (ret.), African American Youth Harvest Foundation**

KEY INSIGHT

AAYH shows what Finding 3 looks like when it goes right. Lofton did not wait for the board to initiate the AI conversation: she built the case, demonstrated results, and brought her board into a conversation they were already prepared to have. Her advice to other leaders is straightforward: start with your top three to five operational pain points, show what AI can do for one of them, and let the evidence do the governance work.

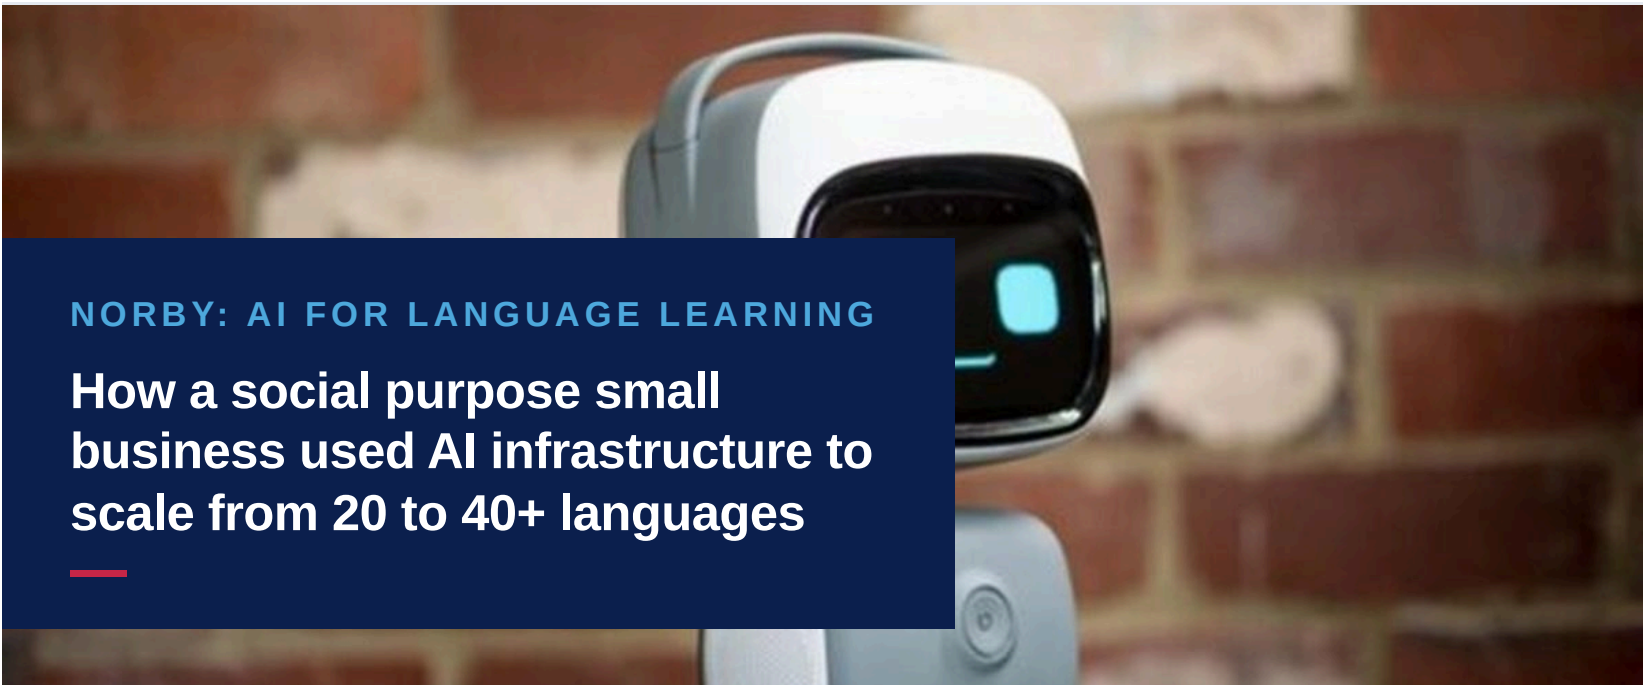

NORBY: AI FOR LANGUAGE LEARNING

How a social purpose small business used AI infrastructure to scale from 20 to 40+ languages

THE CHALLENGE

Norby was founded to solve a problem its creator, Adrian Mullan, experienced firsthand: the ineffectiveness of rote language learning. He wanted to build an interactive, AI-powered robot that could hold natural conversations tailored to each user, whether an eight-year-old learning English or an 80-year-old stroke patient relearning speech. The challenge was immense: multiple large language models, speech recognition, personalization, hardware design, and data privacy all had to come together.

THE APPROACH

Norby used Dell Technologies Precision workstations with NVIDIA RTX GPUs to train and fine-tune models locally, which provided both the speed for rapid iteration and the privacy required for sensitive applications. Tasks that previously took overnight now complete within an hour, enabling the team to iterate quickly across the software learning function, physical robot design, and marketing production.

THE RESULT

Norby has expanded from supporting 20 languages to over 40, with more being added regularly. The company is preparing pilot programs in language learning centers and elderly care facilities, and is exploring cloud-based LLM access for edge deployment in healthcare settings where data compliance is paramount.

Organization Details

SECTOR
Education technology (for-profit SMB)

LOCATION
Melbourne, Australia

AI FOCUS
Conversational language learning robot powered by LLMs

TECHNOLOGY
Dell Technologies Precision AI-Ready workstations, NVIDIA RTX GPUs

THE RESULT

Norby has expanded from supporting 20 languages to over 40, with more being added regularly. The company is preparing pilot programs in language learning centers and elderly care facilities, and is exploring cloud-based LLM access for edge deployment in healthcare settings where data compliance is paramount.

WHY THIS MATTERS FOR NONPROFITS

Norby's most transferable lesson isn't about language learning—it's about data control. Training models locally rather than sending data to the cloud gave Norby both speed and the privacy protection required for sensitive work. Nonprofit organizations handling client mental health records, immigration status, or domestic violence disclosures face the same tradeoff. On-premise AI infrastructure isn't only a cost or complexity question; for organizations serving vulnerable populations, it may be the governance decision that makes responsible adoption possible at all.

What Leaders Can Do

What separates the organizations in this report that are moving forward from the ones that are still stuck isn't budget or technical sophistication. It's whether someone in leadership decided to connect AI activity to a governance structure before anyone asked them to. The two tools below are designed to help leadership teams do exactly that.

Board AI Readiness Scorecard

This scorecard is designed to help boards assess where they stand on the governance and leadership conditions that this research identifies as most critical. It is not a technology assessment; it is a leadership assessment.

	NOT STARTED	EMERGING	ESTABLISHED
01 **Leadership ownership: someone is clearly responsible for AI direction**	No one owns this	ED/CEO owns it informally	Designated leader with allocated time
02 **AI principles or guardrails exist (and are genuinely understood by the board, not rubber-stamped)**	No principles	Draft in progress	Board-approved principles in use
03 **Board has engaged with AI risk in a structured way**	No discussion	AI risks presented once in a dedicated agenda item	Recurring, documented risk conversation with named risk areas (privacy, bias, cybersecurity, liability)
04 **Board includes tech-fluent members** *Note: traditional IT procurement experience doesn't automatically equal AI risk fluency.*	No tech-fluent members	1-2 members	Multiple, including committee role
05 **AI is connected to strategy, not siloed**	AI discussed as “tech topic”	AI connects to specific goals with named owners and metrics	Embedded in strategic plan

Board AI Readiness Scorecard

	NOT STARTED	EMERGING	ESTABLISHED
06 **Organization tracks AI outcomes**	No measurement	Anecdotal tracking	Defined metrics and reporting
07 **Staff have sanctioned tools and guidance**	Individual, unsanctioned use	Some approved tools	Shared tools with usage policies
08 **Budget includes technology and innovation investment, including R&D**	No budget for AI	One-time allocation	Recurring investment in plan
09 **Risk tolerance is articulated**	No risk discussion	Informal comfort level	Board has agreed on what AI decisions require board-level approval vs. staff discretion
10 **Peer learning or external support**	No external input	Occasional events	Active cohort or advisor

Organizations scoring mostly **Not Started** in dimensions 1-3 should prioritize governance before adding new tools. Mostly **Emerging** suggests the organization is ready for a first bounded pilot. Mostly **Established** suggests an organization is ready to build institutional practice and scale what's working.

TAKE IT FURTHER

Get a personalized score and recommended next steps

Complete the online assessment based on this scorecard in 10 minutes and receive recommendations for your organization. Or download the worksheet version to walk through it with your board.

board.dev/scorecard

A Framework for Responsible Adoption

Based on the case studies and research in this report, organizations that successfully move from experimentation to governed practice tend to follow three principles:

1 Rigorous problem definition.

Start with the pain point, not the tool. Hopeworks did not adopt AI because it was new; they adopted it because mock interviews were inconsistent and they could measure the difference. AAYH started with staff utilization modeling because that was a concrete operational bottleneck. Every successful case begins with a specific problem that AI can address measurably.

FIRST ACTION Write down your organization's three most time-consuming, repeatable tasks. That list is your AI roadmap.

2 Data readiness as the ethical check.

Before deploying AI, ask: do we have the data governance to support it? AAYH's risk-tiered approach—strict controls for HIPAA data, more flexibility for program operations—is a feasible model. Organizations that skip this step end up with the governance debt that our survey shows: widespread experimentation without guardrails.

FIRST ACTION Categorize your data by sensitivity level before your next AI pilot. It takes one meeting and prevents most of the problems.

3 Partnerships over solo builds.

No nonprofit needs to figure this out alone. Hopeworks learned operational discipline from Dell Technologies. BGCAA is learning through a national pilot with Blackbaud and a multi-organization cohort. AAYH leveraged the Dell Technologies certification program. The most effective approach is what Zenae Campbell described: guided, contextualized consultation alongside peers who are working through the same questions. This includes being thoughtful about tool selection: nonprofits don't need to choose between sophistication and simplicity, but they should evaluate tools against their data governance posture, vendor transparency, and mission alignment, not just price and ease of use.

FIRST ACTION Identify one peer organization in your sector that is further along on AI adoption and ask them what they wish they had known at the start.

Recommendations for Funders

The organizations in your portfolio are not failing to adopt AI because they lack access to tools. Most of them have tools. What they lack is the leadership capacity, governance infrastructure, and guided support to make those tools stick.

Fund capacity, not just tools.

The top barriers in this survey are people problems, not technology problems: staff skills, time, and the inability to absorb another initiative. Funders who want to accelerate AI adoption should invest in dedicated consultation, assessment, and implementation support, not just software licenses. As one case study subject put it: *"When nonprofits are funded, don't always look at it from a programmatic standpoint. Look at what all the things needed to run that program."*

CONCRETE ASK Add a line in your next technology grant for implementation support—covering a part-time lead, a cohort facilitator, or dedicated consultation hours.

Invest in board-level leadership development.

The 12× difference in adoption outcomes between organizations with and without board-engaged AI principles is one of the clearest findings in this research. Funders can support board education, cohort-based learning for board chairs and executive directors, and tools like the readiness scorecard to structure board-level conversations.

CONCRETE ASK Require grantees to complete a board AI readiness assessment as part of their tech grant reporting. It costs nothing and surfaces the leadership gaps that predict whether the investment holds.

Support governance infrastructure alongside technology.

Add one question to your next grantee survey: "Does your organization have a written policy governing staff use of AI tools?" The answers will tell you more than any technology assessment. 58% of nonprofits have no AI principles—not incomplete ones, none. Principles are not a nice-to-have. They are the single biggest predictor of whether organizations stay stuck or move forward. Funders can accelerate progress by requiring or incentivizing basic AI governance as part of technology grants, and by funding practical, adaptable frameworks that organizations can implement without specialized legal or technical expertise.

CONCRETE ASK At your next grantee convening, set aside 30 minutes for organizations to share their AI governance status with each other. Peer accountability accelerates what individual grant requirements can't.

Enable peer cohorts and shared learning.

The case studies in this report consistently show that peer learning and external support accelerate adoption. Funders can sponsor cohort-based programs where nonprofit leaders learn from each other's experiences, share governance templates, and build collective confidence. This is particularly powerful for federated organizations and sector-specific networks.

CONCRETE ASK If you fund three or more nonprofits working in the same issue area, consider convening their executive directors and board chairs together for a half-day session on AI leadership. Peer accountability at the leadership level accelerates what individual grants can't.

Reframe AI as a cost-recovery strategy.

As long as AI is treated as a discretionary expense, it will be deprioritized by organizations that are already running lean. Funders can help shift this framing by requiring or supporting cost-of-inaction analyses alongside grant proposals, highlighting ROI case studies, and positioning AI investment as a pathway to sustainability rather than a new spending category.

CONCRETE ASK Ask grantees to estimate their cost of inaction before their next grant renewal. The exercise alone tends to reframe the conversation from 'can we afford this' to 'what are we losing without it.'

What Comes Next

Every organization in this report started somewhere small: a single pain point, one honest conversation with a leadership team, a pilot team that was allowed to fail and try again. The inflection point—from experimenting to owning outcomes—doesn't require a big budget or a technology strategy.

> **It requires a decision to act before conditions are perfect.**

START THIS WEEK

The Tech-Fluent Board

Before your next board meeting, work through The Tech-Fluent Board, a playbook from Board.Dev. It walks board chairs and CEOs through the four areas every nonprofit board needs to own on technology: strategy and vision, risk and ethics, resources and talent, and community and impact. You do not need a consultant or a tech background to start. You need thirty minutes and a willingness to ask the right questions.

board.dev/playbook

TAKE THE ASSESSMENT

Board AI Readiness Assessment

Board.Dev offers an online assessment based on the readiness scorecard in this report. Complete it in 10 minutes and receive a personalized score with recommendations for your organization. Share it with your board chair before your next board meeting.

board.dev/scorecard

JOIN THE CONVERSATION

Webinar Series

Board.Dev and Dell Technologies are hosting a webinar series to explore the findings in this report with nonprofit leaders, funders, and partners. Sessions will cover practical steps for moving from experimentation to governed practice, featuring case study leaders and Dell Technologies experts.

board.dev/webinars

BRING IT TO THE BOARD

Share This Report With Your Board

If one thing is clear from this research, it is that the conversation about AI needs to happen at the board level—and it needs to happen as a leadership conversation, not a technology briefing. Forward this report to your board chair and executive committee with a simple ask: where do we stand on the readiness scorecard, and what is our next step?

Forward this brief

Stay Connected

NEWSLETTER	Subscribe to Board.Dev's **SIGNALS** newsletter for weekly briefings translating technology developments into actionable decisions for nonprofit boards. **Subscribe**
BOARD.DEV	Visit **board.dev** to learn more about AI leadership cohorts, board placements, and the annual Nonprofit Tech Governance Congress.
DELL TECHNOLOGIES	For Dell Technologies' framework on AI infrastructure, governance, and workforce readiness at scale, see the Dell Technologies AI Blueprint: **delltechnologies.com**

Board.Dev strengthens nonprofits by equipping boards with tech-fluent leadership, and helps companies grow disruption-ready leaders through real-world board experience. Learn more at board.dev.

CONTACT **Alethea Hannemann** · Co-Founder & CEO · alethea@board.dev