BOARD.DEV

SAN FRANCISCO · INSTALLMENT ONE

Building the Board We Need Now

Thirty nonprofit leaders spent an afternoon naming what a board is for now. They left us with two questions: why have we never built the board we keep describing, and what can my board do that AI can't?

Board.Dev, with **Okta** and **AWS**
Blueprint partner: **Bridgespan**

Only 3% of nonprofits wrote AI principles with staff and board together. Those organizations are 12 times more likely to have AI in pilots or at scale. That's from our national research with Dell Technologies and TechSoup, where we found 87% of organizations stuck in early experimentation or not started at all, and 58% with no AI principles at all—the same stalling McKinsey found across nearly 2,000 companies worldwide, where two-thirds have not begun scaling AI either.

The sector has spent two years treating nonprofit AI adoption as a tools-and-training problem. Our data and partnerships point somewhere else: to leadership, and specifically to boards.

That matters more than usual right now. A third of organizations that serve communities had government funding disrupted in the first half of 2025. Dollars raised grew 5% in 2025 while the number of donors fell again. Fewer people are giving more, which puts more weight on fewer relationships. Demand keeps rising.

Almost every organization in this sector faces a decision about what it needs to become, and yet almost no one is empowering or enabling the board leaders who authorize and resource those decisions.

Why we started with boards

Staff can start AI. Only a board can authorize an organization to become something new, resource the change, and stand behind it.

On June 8, we gathered about 30 nonprofit CEOs, funders, board chairs, and technologists for an afternoon at Okta to explore one question: what is a nonprofit board actually for right now? Leaders from organizations like JVS, TechSoup, NTEN, Fast Forward, Tech Matters, Project Evident, Free Law Project, Draper Richards Kaplan, and Tipping Point worked through the question together.

This was the first of three working sessions we're using to produce a field blueprint this fall, in partnership with Bridgespan. The blueprint is a shared answer the sector can build on: what a board is for now, what it does, and how a board gets from where it is today to there. We're writing it with Bridgespan because they're working on the same problem from the vantage point of the nonprofit executive team, with research that asks what the nonprofit operating model has to become. Ours asks what the board has to become to authorize it. Neither answer is much use without the other. The guidance boards run on today was built for a slower environment, and organizations are now redesigning themselves without knowing what the board that approves the redesign needs to be.

We started with the organization rather than the board. Each group took one archetype—a legacy direct-service organization deciding which programs to retire, a younger organization building AI into its model from the start, an intermediary rethinking what it offers the field, and a funder rewriting what it asks of grantees—and described what it needs to be by 2030. Then we worked backward to the board that gets them there. The instruction was specific: name what a board needs to do to create, drive, approve, or support the change, in actions rather than attributes.

Where these organizations are going

Legacy direct-service organizations described a future with less government money in it. They want to become financially resilient and independent, positioned as a partner to government rather than a dependent of it, and able to demonstrate value in whatever terms a new funding model demands. Getting there means owning their own data under revocable consent, implementing AI that helps people reach services faster, and using AI on low-joy, high-friction staff toil so staff can spend more time with people.

WHAT'S AT STAKE IF IT DOESN'T HAPPEN

A two-tiered sector split between the tech-capable and everyone else, fewer options for the people who need them, and for individual organizations, existence.

New and AI-native direct-service organizations wanted stability alongside funding agility, and spent most of their time discussing where AI belongs, including where human relationship needs to be the service rather than the delivery mechanism. Their answer was discernment: use AI to drive efficiency without replacing the human contact the work depends on.

WHAT'S AT STAKE IF THEY GET IT WRONG

If they build new models without that discernment, power stays with the vendors they're building on, and the model serves the vendors' priorities rather than addressing the market failures these organizations were started to correct.

This group also named something about boards no other group did. For an organization this young, the board's outward work matters more than its internal oversight, because they need to build credibility and have almost no track record to oversee.

Field-leading intermediaries described themselves as always in beta. They expect increasing pressure to make sense of the ecosystem for everyone downstream, and they want to hold a strategic direction rather than execute a strategic plan, since those plans are out of date almost as soon as they're published. They called intermediaries the trusted cartilage between organizations and communities, and one group pushed that even further, floating a union-like model with regulatory or bargaining power for the field. They didn't fully define the idea, but we'll take it up in future discussions: intermediaries can feel the gap between influence and authority, and think they are positioned to close it.

WHAT'S AT STAKE IF THEY GET IT WRONG

Moving too fast leaves people behind, grassroots organizations get more fragile, and the field loses a consolidated read on what communities actually need.

Funders and grantmakers named the shift they need most directly, from legacy funder to tech-forward funder. Tech-forward means AI and data fluency, a stated risk threshold, an AI rubric, portfolio assessment, and an honest look at whether they hold data too sensitive to keep.

WHAT'S AT STAKE IF THEY GET IT WRONG

Their mistakes don't stay in-house. A funder that moves too slowly leaves grantees to figure out AI alone; one that moves too fast pushes organizations toward AI before they can absorb it; and one that doesn't address infrastructure or tech foundations loses control of sensitive grantee data before it figures out how to secure it.

Asked what needs to stay true, nearly every group listed the same answer: mission. How the work gets done changes, but the reason for doing it doesn't.

In San Francisco, we worked with these four archetypes, but we'll take on two more in New York and Chicago: networked and franchised organizations delivering through local affiliates, and advocacy and systems-change organizations. Each of these organization types will have unique challenges and opportunities in the AI era.

Group after group arrived at the same two questions, unprompted. They point to where boards need to go.

Question 1: Why have we never built the board the field has been describing for 30 years?

Asked to describe the board this moment demands, these experienced leaders wrote:

> Diverse and representative
>
> Forward-thinking, agile
>
> Wildly engaged
>
> An exemplar
>
> Trusted advisors

Nothing on that list is wrong; every answer describes a board any of us would want. But these are also the answers the field has been giving for 30 years, and most boards still don't deliver them. This is our core finding: when a room with this many strong leaders gives the same answers the field has been calling for since the nineties, the problem isn't naming what a board should be. We knew that already.

Question 2: What can my board do that AI can't?

The second question produced a totally different list:

> Values-aligned judgment
>
> Support for a CEO under pressure
>
> Relationships and networks the organization can't reach by itself
>
> Sense-making
>
> Emotional stability
>
> Being the mirror that tells leadership the truth
>
> Championing the organization in rooms staff can't enter

Not one of those answers is about information or expertise, which is what good boards have been built on for decades. You recruit a lawyer, an accountant, you get the knowledge in the room. AI makes that knowledge cheap and fast. What the room wanted instead is the part we don't yet know how to recruit for.

Two moves for stronger boards

The room also produced two calls to action, one for each question. Start by answering what a board can do that AI can't, because it's the harder topic of the two. Then move to unfinished business—why we haven't built the board we keep describing? These evergreen attributes are jobs that used to be ideal but optional. They aren't anymore, and with new tools and new people at the table, they're finally possible.

The first move

What boards do that AI can't is outward, relational, human work, and most boards barely do it. Making an introduction to open a door staff can't open on their own; sitting with a CEO through a tough decision that isn't in any playbook; reading a chaotic funding and political environment and telling the organization what it means; being the mirror that tells leadership the truth no one else will; putting the board's own credibility behind the organization in rooms it can't get into by itself. Underneath all of it sits the move only a board can make: authorizing—and prompting—reinvention, and giving a CEO the mandate, the cover, and the resources to change what the organization is. As AI absorbs more of the inside work, this becomes the majority of the job rather than the part you get to after the finance report.

The second move

This is finishing the unfinished business. None of those evergreen words the group listed has ever had a definition a board could fail, which is most of why 30 years of unfinished work stayed unfinished—they weren't actionable. Giving them definition is what boards need to do now. Agile means a named group can decide key issues between meetings, on a fixed clock, with delegated authority and written consent. Engaged means a set of mandates a board either meets or doesn't.

Diverse is one value that doesn't need redefining. The field has known what it means for 30 years, and the failure was never comprehension; it was recruiting. Boards fill seats from the networks their members already have, which reproduces the board that already exists. That's a delivery problem, and two of the ten mandates are aimed at it, including the one about board terms below.

What is new is an expanded skill set that sits beside diversity rather than inside it. A board now needs enough fluency to ask a real question about a technology decision. That's an addition to the skills list—new expertise. And it isn't a job for one person in the corner or for the one person with IT in their title, since a board with a single AI expert has outsourced the judgment it exists to exercise. One group wrote "minor in AI," which comes close: a competence sitting alongside whatever each member actually brings. Finance, legal, HR, operations, marketing, and program all contain people who have already made these decisions somewhere else.

A good share of this second move costs nothing. Meeting more than four times a year, delegating authority so a named group can decide between meetings, writing a real job description, putting a sunset date on standing commitments—no board needs money or permission for any of it. The challenge is cultural rather than mechanical. These are habits nobody owns and nobody needs, and yet nobody feels entitled to break them.

Neither move is sufficient on its own. A board that only adds new areas of focus is skipping 30 years of unfinished work. A board that only gets better at evergreen attributes will function smoothly inside a job that keeps getting smaller and less relevant.

Mandates: what this looks like in practice

So as a CEO or a board member, what do you do to evolve? The room named qualities. We turned them into behaviors. From the session and our work since, we drafted ten mandates, each built the same way: a board needs to ___ so the organization can ___.

We used four rules to decide what qualifies:

1. **It's a verb.** A board does it, rather than has it.
2. **The second half names something the organization gains.** If you finish that clause with something about the board, not the nonprofit, it isn't a mandate. The only reason to ask what a board should be is to move what the nonprofit needs to be.
3. **Someone outside could tell whether you did it this year.** Measurable—belief doesn't count, and values aren't enough.
4. **A board is the only body positioned to do it.** Staff can't, or staff shouldn't have to, and neither can AI.

And one check: if a CEO would experience a mandate as inspection, it's not right. That test eliminated two mandates from our first draft.

The full set of mandates will appear in the final blueprint. Three examples here:

1. **A board needs to take a position on how much the organization should change,** and tell the CEO, so the CEO isn't guessing whether the board will back a bet or punish it. Almost every AI decision inside an organization is a small version of this question, and most of them stall in the gap where nobody can tell what the board would say.

2. **A board needs to make real decisions between meetings,** with a named group and a fixed clock, so a fast-moving organization isn't waiting on the calendar. Michael Lissner of Free Law Project runs his board on Slack. Most boards could do it by Tuesday and don't.

3. **A board needs to require every member to make the case for their next term,** so the organization isn't led by a board built for problems and opportunities it no longer has. If nobody has failed to be renominated in five years, the board isn't doing this. You can't authorize reinvention of the organization you serve if you're afraid of reinventing yourselves. For a board that carries significant giving obligations, this is the hardest thing on the list, and it's still the one that decides whether the rest is possible.

One finding from testing the full set: some mandates mean the same thing everywhere, and others mean opposite things depending on the organization.

Setting risk appetite in writing **constrains an AI-native startup** that's already moving fast,

but **frees a legacy direct-service organization** where staff are used to everything needing permission.

This split is a finding rather than a flaw, and we'll figure out where it matters in the remaining working sessions.

You'll notice nothing here says: boards need to raise money. We didn't skip fundraising because it stopped mattering, or because most boards do it well. It's absent because give-and-get was never the part boards misunderstood, and because the outward work above is where a board's fundraising value now sits: the introduction, the vouching, the credibility spent in a room the CEO can't get into.

A board that does that work well is fundraising, whether or not anyone calls it that.

Boards built for questions nobody has faced before

After the working session, about 80 people gathered at the AWS Builder Loft. Two organizations in the room that evening had already made some of these calls.

Buddy Shah of Anthropic's Long-Term Benefit Trust with Charlotte Wylie of Okta.

Hilary Noon of AWS with Michael Lissner of Free Law Project.

Buddy Shah, who chairs Anthropic's Long-Term Benefit Trust and runs the Clinton Health Access Initiative, talked with Charlotte Wylie of Okta about a board built for a question almost no board has faced: how do you hold a company accountable to its mission when the technology it builds becomes powerful enough to shape outcomes far beyond the company itself? The Trust is an independent body with no financial stake that appoints a majority of Anthropic's directors. Buddy walked through how that model is evolving under public scrutiny, concentrated resources, and enormous social need.

Hilary Noon of AWS came at the same question from the other end, in conversation with Michael Lissner, co-founder of Free Law Project. Free Law Project has collected American case law since 2009, with an archive running back to 1650, and now builds technology for courts, institutions with little history of doing technology well and rising caseloads driven by problems that should have been solved somewhere else.

> We could solve a lot of this with social workers, and we don't, so people end up suing each other.
>
> MICHAEL LISSNER, FREE LAW PROJECT

Funders told the team at the start that the work was impossible.

Lissner described two board decisions most organizations never make. The first was dropping the requirement that board members give. Free Law Project recruited people it wanted to work with who brought the skills the mission needed, and treated giving as welcome rather than expected. He called it a risk, and said he had the privilege of being able to glide for a while as they figured it out. It produced a board with the range the work demanded, including a co-founder who is a lawyer, a technologist, and works on AI in law.

The second is the Slack decision from the mandates above: his board lives on Slack. Most members are in the channel, and he talks to them when an opportunity appears instead of waiting for the quarterly meeting. His reason names one of the core board problems precisely—change arrives from technology, government, and culture at the same time, and the agility makes the difference.

For leaders who inherited a board rather than building one, his advice was practical. Write down the skills the organization needs. Work your networks. Where you can't recruit, upskill: ask who around the table already pays attention to AI, and give them the role. That's where most boards will start, even if it isn't where they should end up.

The questions we didn't answer

The conversation opened a few issues the room didn't answer. We're including them here unresolved, because the disagreements are material—the decision you make on each changes what a board should do.

- Should a board set strategic direction or approve a strategic plan?
- Can you effectively champion an organization externally without being heavily involved on the inside?
- Where does oversight end and operating begin, particularly on AI?
- Can the same board both raise substantial money and do judgment-heavy work, or are we asking one group of people to be two different kinds of useful?
- Legacy funders often want stewardship and tech-forward funders want speed. What does an organization need from a board when it takes money from both?

One question for your board

If you're a CEO reading this at the end of a long week, the honest version is this. You already spend real time on your board. You hold meetings, build a pre-read, do the prep, and none of that is optional. What's variable is whether you get that effort back in air cover for a hard decision, an opened door that was closed to you, or someone telling you the truth before a funder or your community does. Nothing we're naming here asks for more of your time, or your board's. It asks for different time.

Before your next meeting, ask the people around your table: what is one thing this board should do in 2026 that we're not doing, and what is one norm we could change at the next meeting? If the answers come back evergreen—agility, championing, engagement—that isn't a failed discussion; it's where you can start. Ask what you can do to make those aspirations real this year.

What's next

Over the summer, we'll test the two recommended moves and the ten mandates through a survey, CEO interviews, and two more convenings: New York on September 29, headlined by Shamina Singh, Zia Khan, Nabiha Syed, and Anil Dash, and Chicago in October during SOCAP week. In each session, we'll test the mandates against a new group, and then publish what we learn. In Q4, we'll release the full blueprint with Bridgespan.

> Help shape the outcome. We're interviewing CEOs and board chairs at organizations in the middle of real decisions about what to become, and the survey opens in September. Put 30 minutes on my calendar, or write to info@board.dev and we'll send you the survey when it goes live.

ABOUT BOARD.DEV

Board.Dev places tech-fluent leaders on nonprofit boards and builds AI leadership capacity at the board level, with corporate partners including Okta, Dell, Morgan Stanley, and CDW.

A NOTE ON SUPPORT

Okta hosted the working session and AWS hosted the reception. Free Law Project is an AWS Imagine Grant recipient. Both companies put real resources behind the leadership this sector needs, and build infrastructure the field runs on.