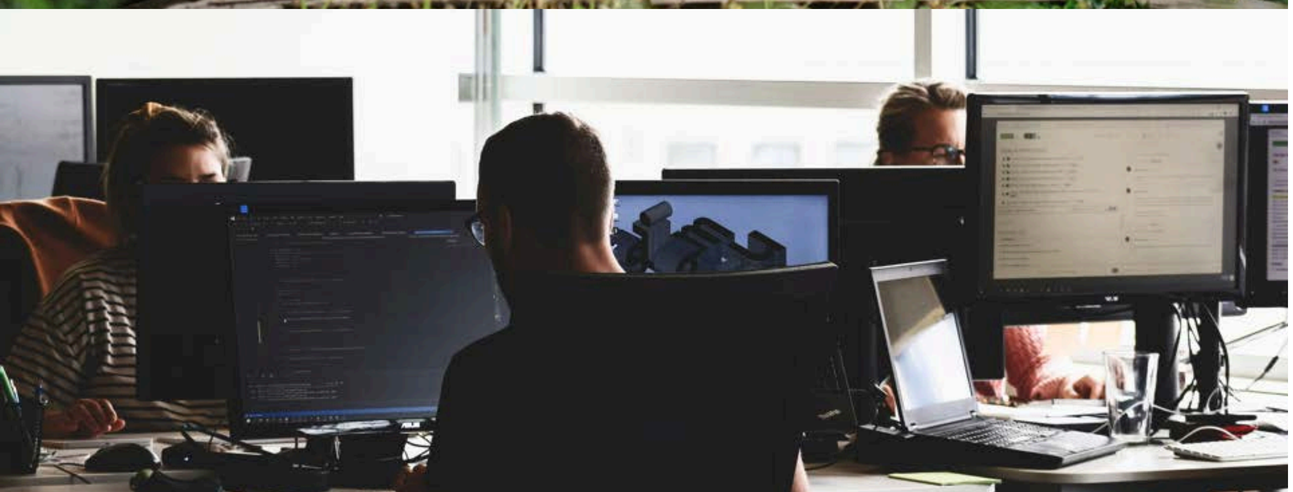

NATION SWELL

In collaboration with

IBM

Responsible Use of AI *for Social Impact*

The nonprofit sector stands at a critical juncture.

Recent research shows that 82% of nonprofits are already using some form of AI, yet 76% lack an AI policy. This gap between adoption and governance creates both risk and missed opportunity.

At IBM, we've spent years working to close this gap.

Our Responsible Technology Board guides our approach to ethical innovation. We've built tools like watsonx.governance to help organizations create transparent, explainable AI workflows, and Granite Guardian to detect and reduce risks in AI use.

We have seen firsthand how these systems can strengthen trust and expand impact when applied thoughtfully.

This playbook builds on that experience and insights from leaders across the field. It aims to bridge the divide between AI's potential and its practical application, ensuring every nonprofit can use AI ethically and effectively to advance its mission.

Progress will require collaboration at every level: nonprofits strengthening AI literacy, funders investing in shared infrastructure, technologists providing trustworthy tools, and academic partners preparing the next generation of practitioners.

AI is reshaping how the social sector delivers on its purpose. The real question is not whether to adopt it, but how to do so responsibly, guided by values, transparency, and the communities we serve.

How to *Use this Guide*

This playbook is designed primarily for nonprofit professionals who need practical, actionable guidance on implementing AI responsibly in their organizations.

Whether you're an executive director exploring AI strategy, a program manager considering AI tools, or a frontline staff member trying to understand AI's impact on your work, this guide provides a pathway forward.

Funders will find valuable guidance in Section 5, which outlines strategies for supporting AI capacity across in your grantees and the civil sector at large.

Find what you need quickly:

New to AI?

Looking to develop your AI skills?

Ready to implement AI in your organization?

Need guidance on responsibly governing AI in your organization?

Considering a university partnership?

Funder wanting to support grantee AI capacity?

Looking for inspiration and real examples?

Key Players in Supporting AI in the Social Impact Sector

Key Elements of Responsible & Ethical AI for Nonprofits

Building responsible AI capacity requires three interconnected elements — literacy, equitable design, and responsible AI development and deployment.

Element 1 – Foundational AI Literacy

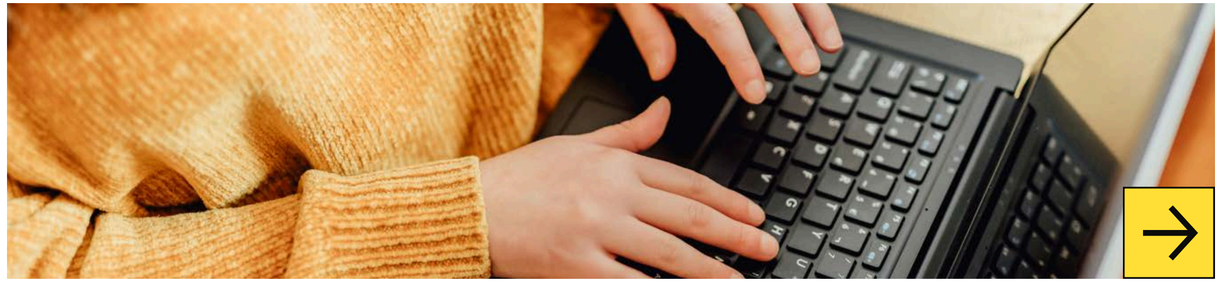

Element 2 – Equitable, Human-Centered Design of AI Systems

Element 3 – Responsible AI Development, Deployment & Governance

Element 1

Foundational AI Literacy

What is AI Literacy

AI literacy is the ability to understand, use, and critically evaluate AI technologies in a responsible and effective manner.

It encompasses:

- Knowledge of what AI is and how it works
- The skills to use AI tools for various tasks
- The capacity to discern the risks and ethical implications associated with its use

Why AI Literacy Matters

Without foundational literacy, organizations cannot distinguish between AI hype and genuine opportunity, assess risks to their organizations and constituents, make informed decisions about adoption, and ensure AI work products are of high quality and a value add to their organization's missions.

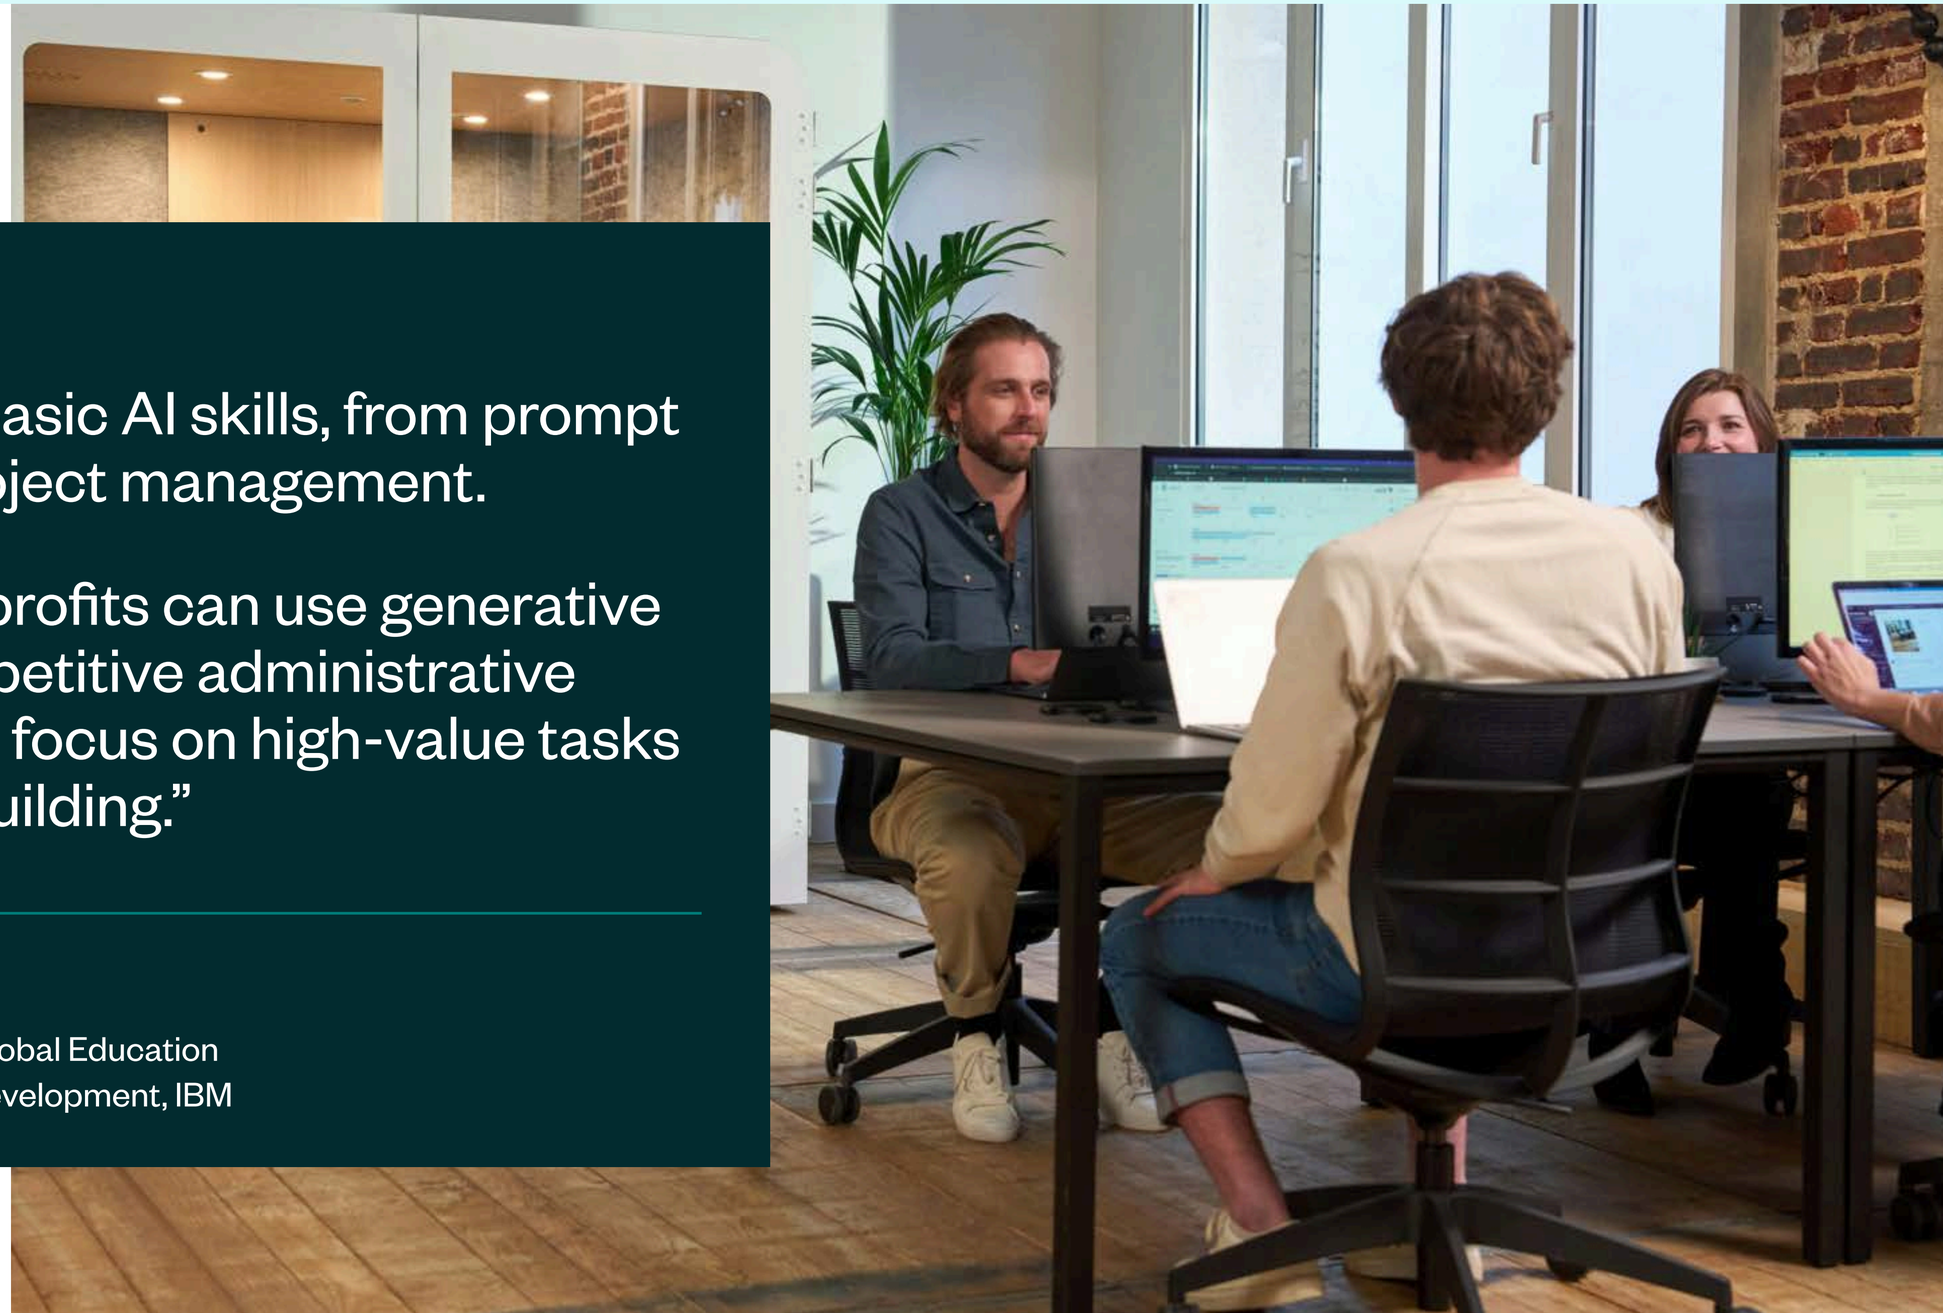

"Everyone needs basic AI skills, from prompt engineering to project management.

Even smaller nonprofits can use generative AI to automate repetitive administrative tasks, so staff can focus on high-value tasks like relationship building."

Lydia Logan
Vice President, Global Education
and Workforce Development, IBM

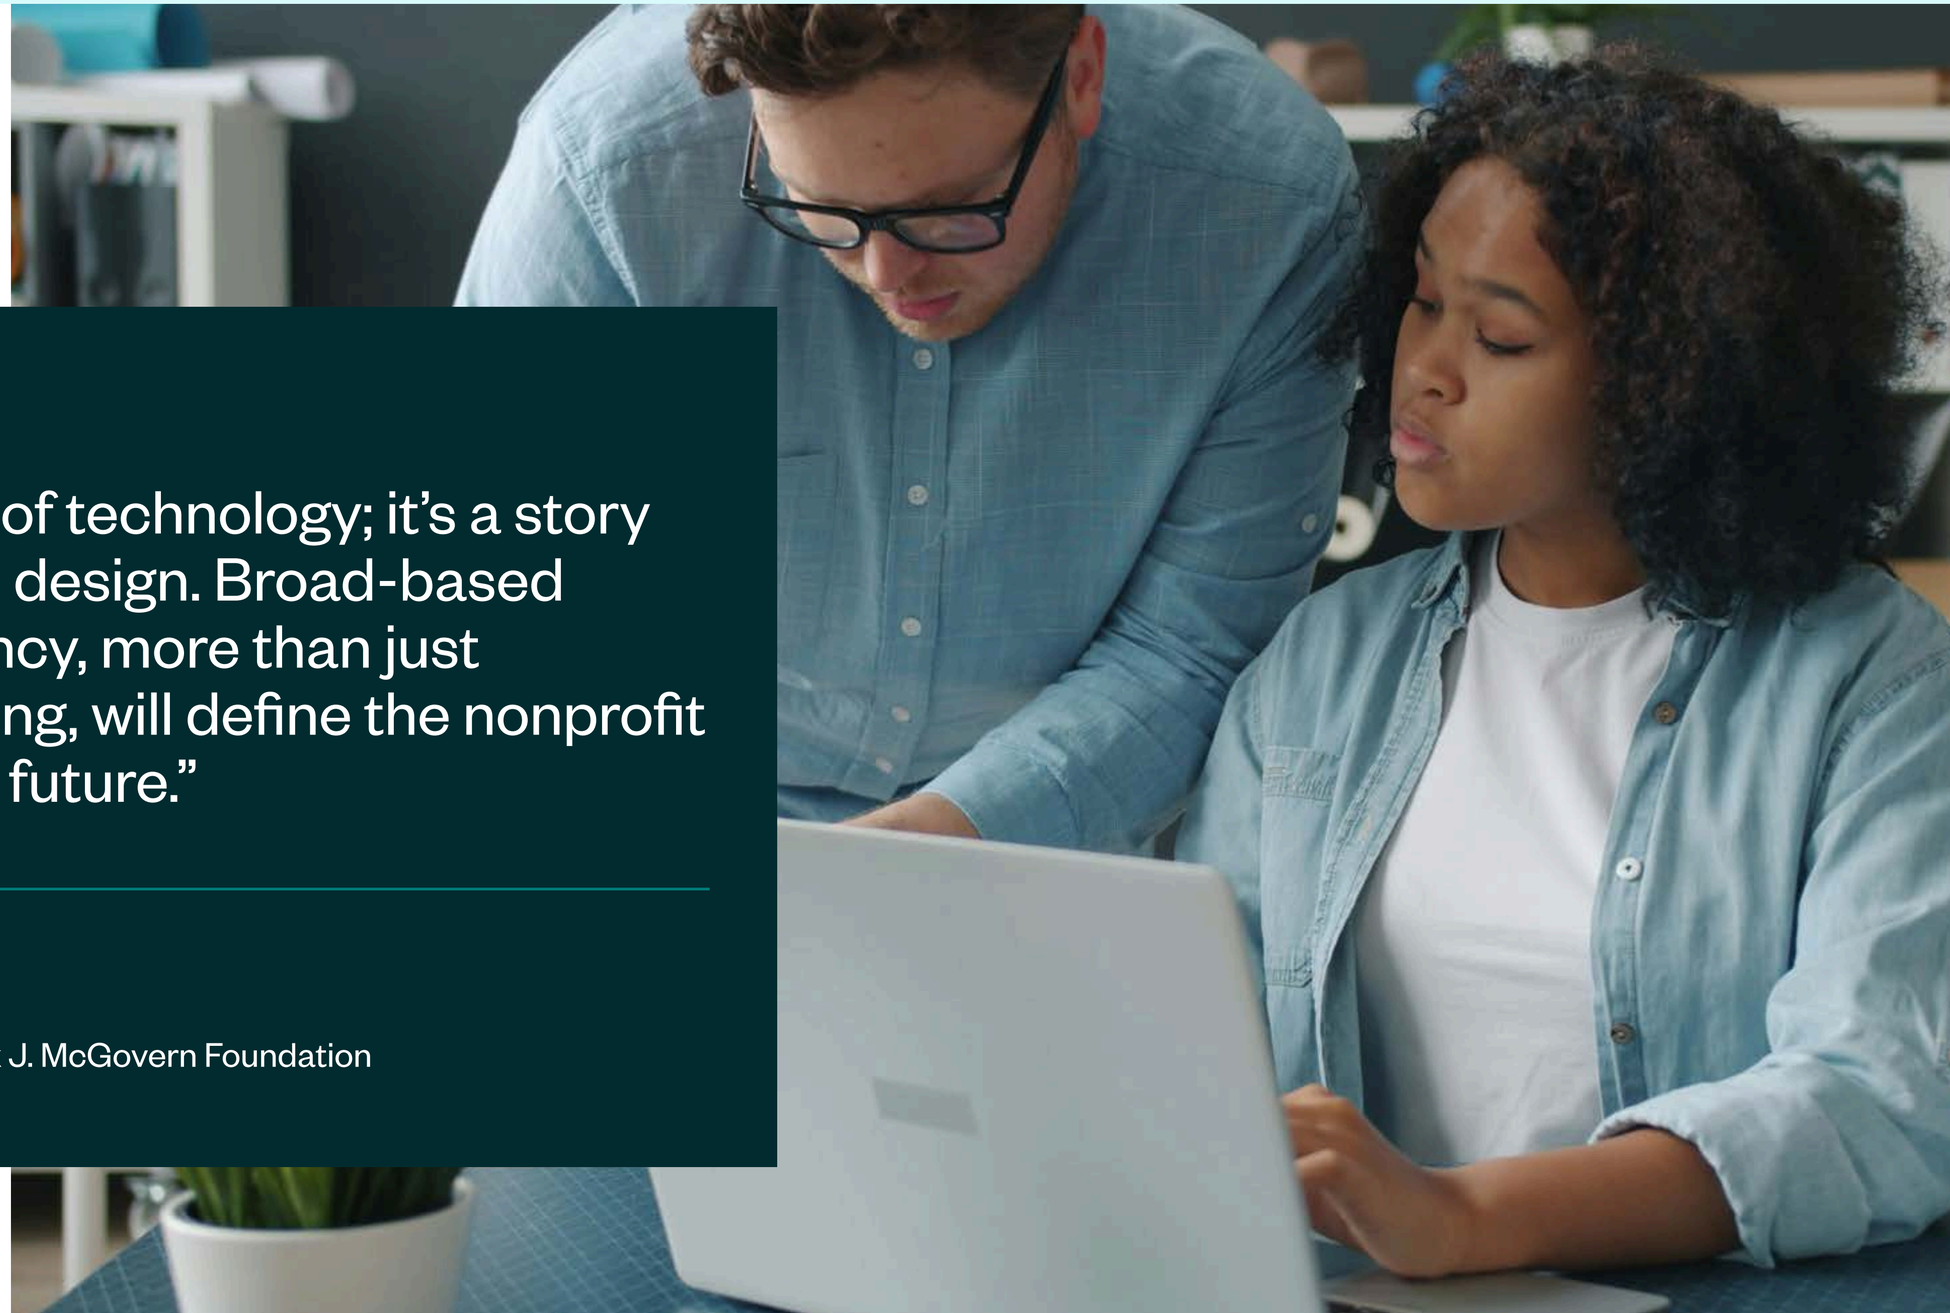

"This isn't a story of technology; it's a story of organizational design. Broad-based literacy and fluency, more than just technical upskilling, will define the nonprofit workforce of the future."

Vilas Dhar
President, Patrick J. McGovern Foundation

AI Competencies Across Your Organization

The AI Alliance's Guide to Essential Competencies for AI, developed in partnership with IBM, identifies five critical AI competency areas that organizations need to develop across their workforce:

Identifying Data Limitations

Ethical AI usage that protects security, mitigates negative impacts, and prioritizes mindful interaction

Data Analysis

Selecting, using, and managing data with appropriate methods to meet goals

Machine Learning

Understanding concepts and applications core to AI programming and GenAI tools

AI Logic

Working with symbolic/ logic-based AI and making strategic decisions about AI approaches

Responsible Use

Ethical AI usage that protects security, mitigates negative impacts, and prioritizes mindful interaction

The Four Competency Levels

Each competency area progresses through four levels.

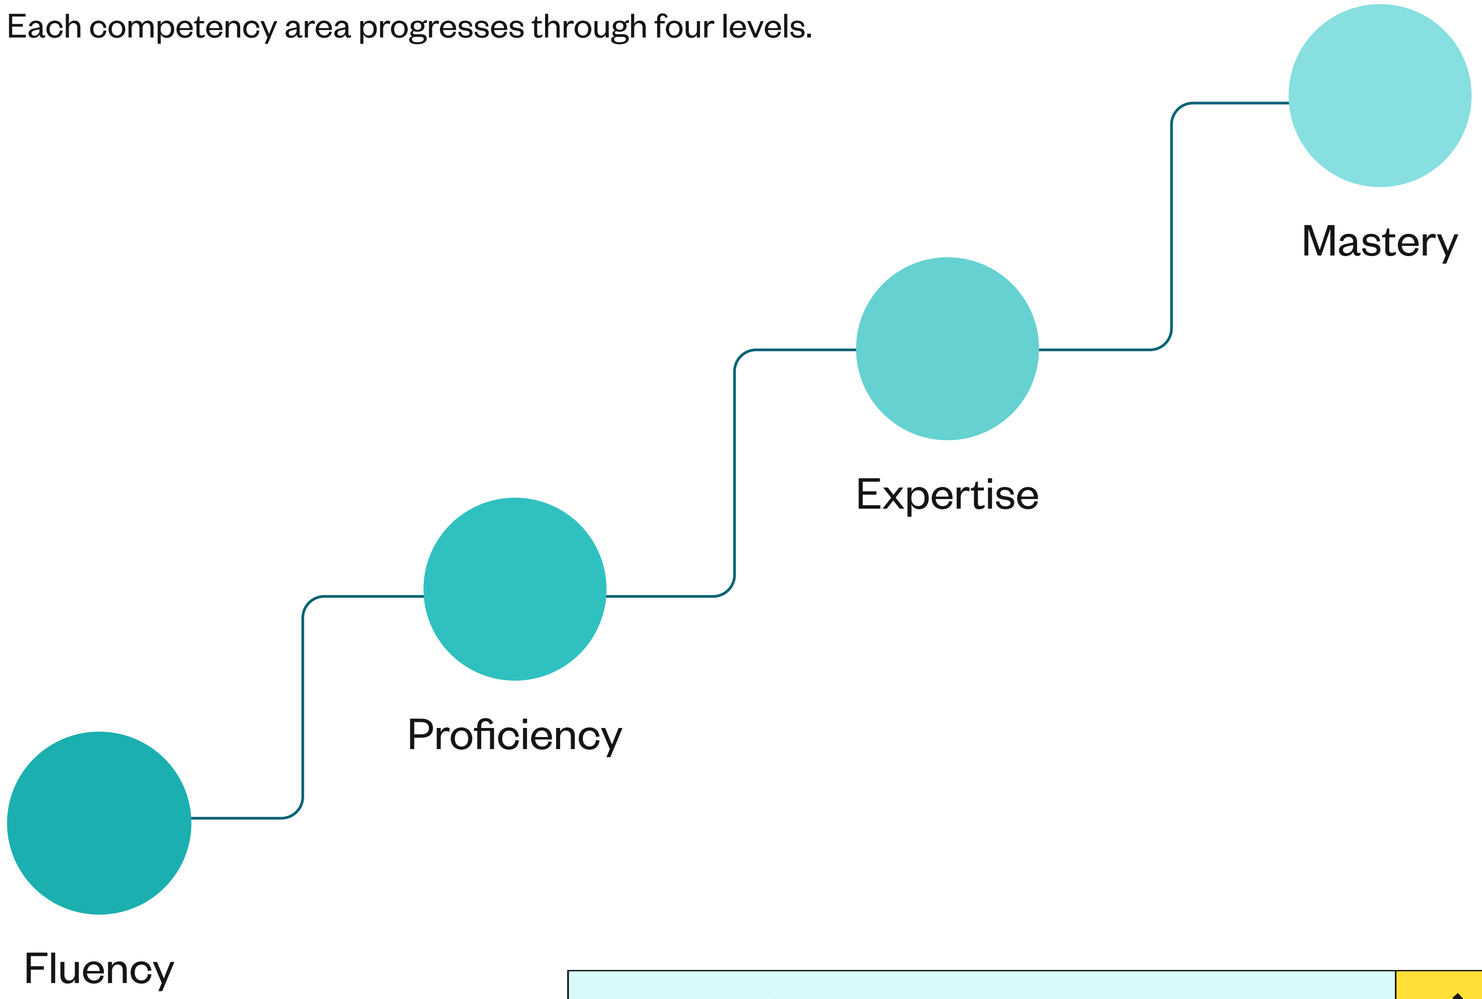

Level 1

Fluency

Basic understanding and ability to use AI tools

Level 2

Proficiency

Critical use of AI to solve problems

Level 3

Expertise

Working within AI domains, analyzing, and adapting

Level 4

Mastery

Leading innovation and teaching others

Take the IBM SkillsBuild AI Level Up assessment to determine your AI readiness and get access to curated learning plans

Key Role Profiles and Development Paths

Novice User

Who

General employees, students, frontline workers

Focus Areas

Fluency in Responsible Use and Identifying Data Limitations

Essential Skills +

AI User on the Job

Who

Professionals using AI tools in daily work

Target Competencies

Proficiency in Responsible Use and Identifying Data Limitations; Fluency in Data Analysis

Essential Skills +

AI-Informed Decision Maker

Who

Organizational leaders, policymakers, managers

Target Competencies

Expertise in Responsible Use and Identifying Data Limitations; Proficiency in Data Analysis, Machine Learning, and AI Logic

Essential Skills +

Learn more

AI Alliance: Guide to Essential Competencies for AI

AI Literacy: Closing the Skills Gap

Element 2

Equitable, Human-Centered Design of AI Systems

Why It's Important

IBM's approach to human-centered design puts people first, not technology. The principle "good design is good business" from IBM Chairman Thomas J. Watson Jr. remains central to IBM's philosophy today. This approach requires an unrelenting focus on people that transforms not just design teams but entire organizations.

IBM's Design Thinking Framework

"Don't start with technology — start with the user, their problem, and how you will address it in the marketplace."

Michael Jacobs
Head of Social Innovation, IBM

Three Core Principles

01

Focus on User Outcomes

Drive organizational outcomes by helping users achieve their goals. “We’re not measured by the features and functions we ship. We’re measured by how well we fulfill our users’ needs,” -IBM Enterprise Design Thinking Framework.

Every AI initiative must start with clear user outcomes. What will beneficiaries, staff, or donors be able to achieve that they couldn’t before?

02

Restless Reinvention

Stay agile by treating everything as a prototype - even deployed solutions. This mindset prevents nonprofits from getting locked into inflexible systems.

03

Multidisciplinary Teams with Wide-Ranging Expertise

Move faster by empowering multidisciplinary teams with backgrounds across fields including technology, compliance, and community needs to act. "If you want a breakthrough idea, you're more likely to get it with a diverse team."- IBM Enterprise Design Thinking Framework

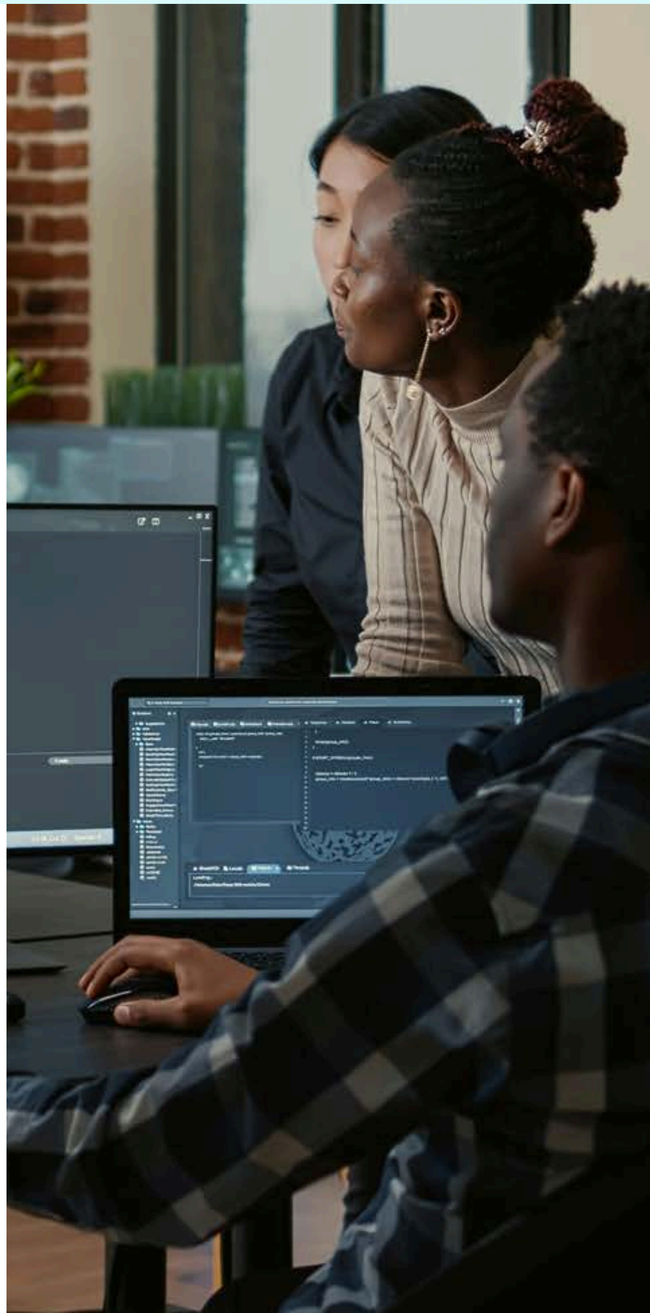

The Iterative Process: Observe, Reflect, Make

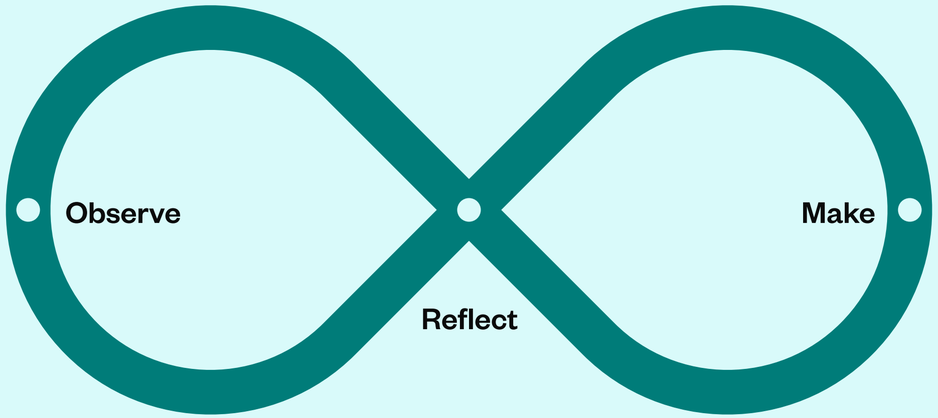

Observe

Immerse yourself in the real world

Get out of the building and meet users where they are. “Understanding can’t be delegated” - the team needs direct user contact.

Key Actions +

Reflect

Come together and look within

Synthesize what you’ve learned and align as a team. “An insight isn’t restating an observation — it’s a leap in clarity, reframing your point of view.”

Key Actions +

Make

Give concrete form to abstract ideas

“The earlier you make, the faster you learn.” Don’t wait for perfect understanding - build to learn.

Key Actions +

Source: IBM Enterprise Design Thinking Framework

"Respectful engagement is key — field interviewers are trained, often hired from within communities, to avoid cultural missteps."

Mariana Paál Martinato
Former Executive Director, Instituto Yarandu

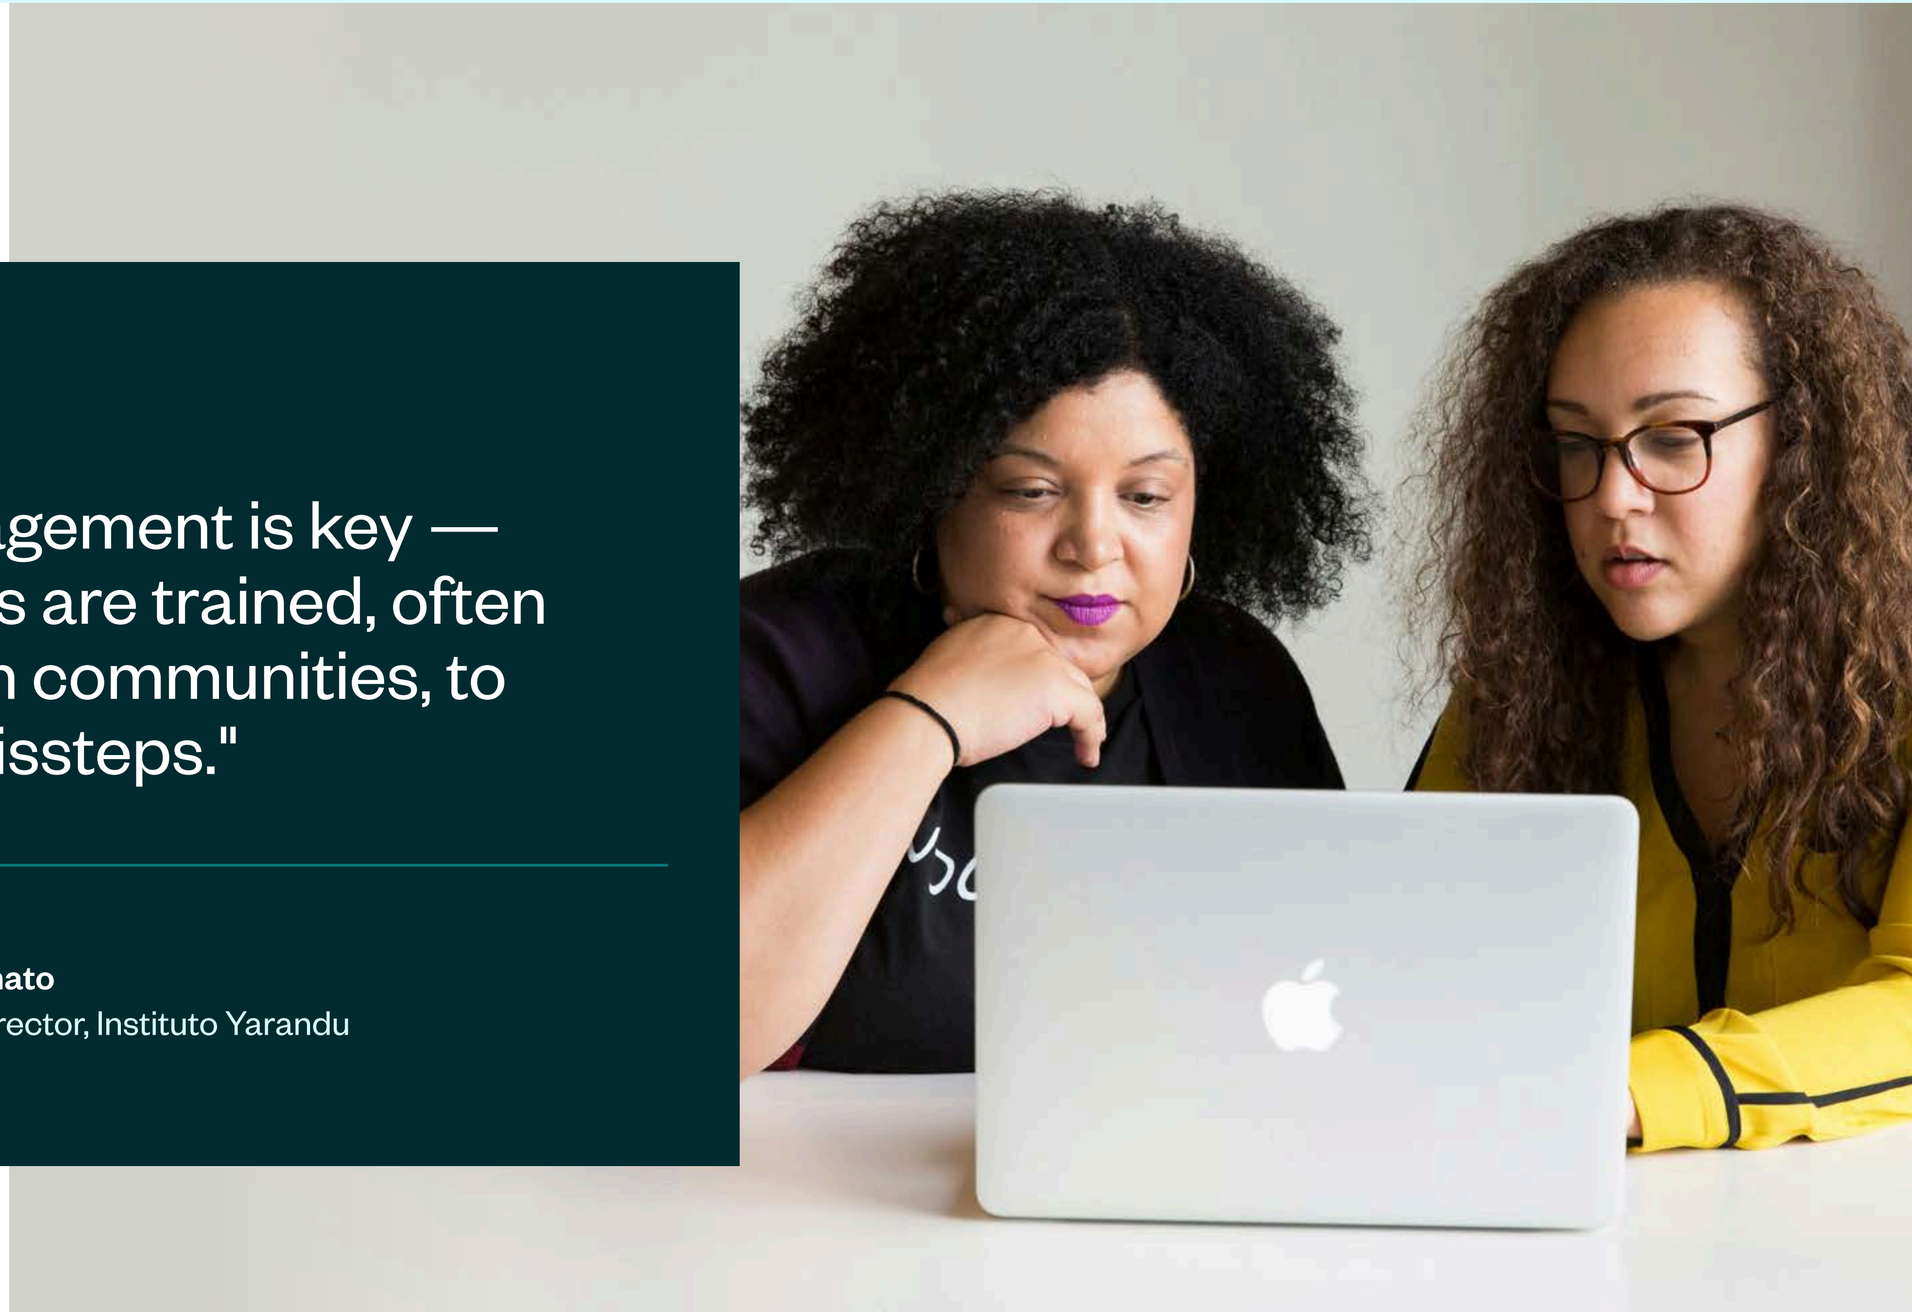

Learn more

IBM Enterprise Design Thinking Framework

IBM SkillsBuild Enterprise Design Thinking Practitioner

IBM Human-Centered AI

Element 3

Responsible AI Development, Deployment & Governance

Why Responsible AI Matters

AI governance encompasses the policies, processes, and practices that ensure AI systems are developed and used ethically, safely, and in alignment with organizational values.

It creates accountability structures and decision-making frameworks that guide how AI is integrated into operations.

Nonprofits face unique challenges with AI adoption. A biased algorithm denying services or a privacy breach exposing undocumented individuals could cause unknown harms. Deploying and using AI responsibly requires a sociotechnical lens, aligning the organizational culture with people, processes, and tools to manage the technology.

Below are examples of IBM Principles of Trust and Transparency. Many companies have developed their own Principles that help their employees, clients, and society have their own ‘north star’ to work toward.

If you’re getting started on building your own Principles, IBM can help

Principles of Trust and Transparency

Principle 1

The purpose of AI is to augment human intelligence

AI should augment, not replace human intelligence.

In Practice:

- ✓ AI is used to automate routine tasks, freeing up staff to focus on high-impact work
- ✓ AI provides insights and recommendations, but humans make the final decisions
- ✓ AI is used to support human decision-making, not replace it
- ✗ AI is used to make decisions that affect people's lives without human oversight
- ✗ AI is introduced without training or support for staff to work effectively with the technology

Quick Test +

Principle 2

Data and insights belong to their creator

Clients' data is their data, and their insights are their insights.

In Practice:

- ✓ Nonprofits own and control their data, not vendors or third parties
- ✓ Donors and constituents have control over their personal data and can opt-out of sharing
- ✓ Transparency around how data is collected, used, and shared
- ✓ Data is used only for the purposes intended by the nonprofit and its stakeholders
- ✗ Sell or share data without explicit consent
- ✗ Vendors claim ownership of nonprofit data

Quick Test +

Principle 3

Technology must be transparent and explainable

Clarify who trains the AI systems, what data was used in training and, most importantly, what went into the recommendations of the algorithms.

In Practice:

- ✓ "This was written with AI assistance"
- ✓ "AI recommended you for this program because..."
- ✓ Public list of AI tools you use
- ✗ Hidden AI decisions
- ✗ Unexplainable recommendations

Quick Test +

Learn more

IBM AI Attribution Tool

Augmenting Human Intelligence with AI Best Practices

IBM's Principles for Trust and Transparency

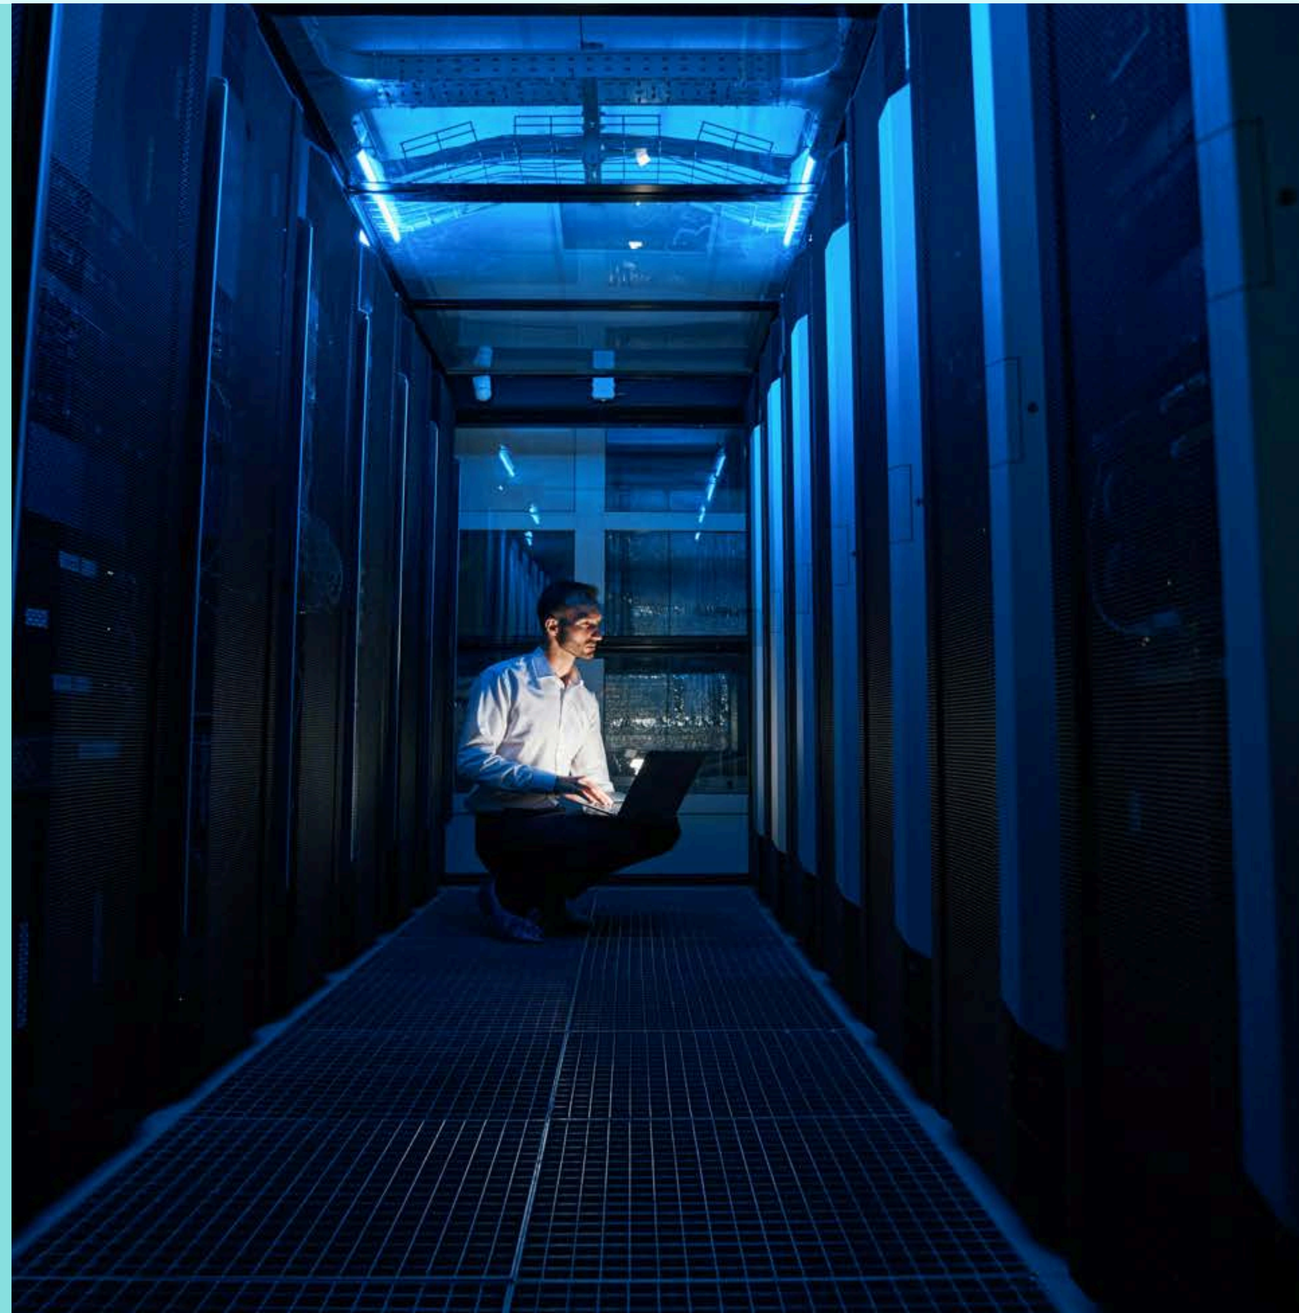

Five Pillars of Trusted AI

Free Responsible AI Tools for Technical Practitioners

01

Explainability

Can you explain every decision to a skeptical donor?

Take the test

Red flags

02

Fairness

Does this outcome prevent harmful, unjustified adverse effects on individuals or groups?

Take the test

Red flags

03

Robustness

Can this system withstand unexpected challenges and changes?

Take the test

Red flags

04

Transparency

Are you open and clear about how you use AI?

Take the test

Red flags

05

Privacy

Would you trust this system with your most sensitive information?

Take the test

Red flags

Learn more

IBM Pillars of Trusted AI

IBM AI Governance Consulting

watsonx.governance

AI Ethics

Ethical Considerations for Generative AI

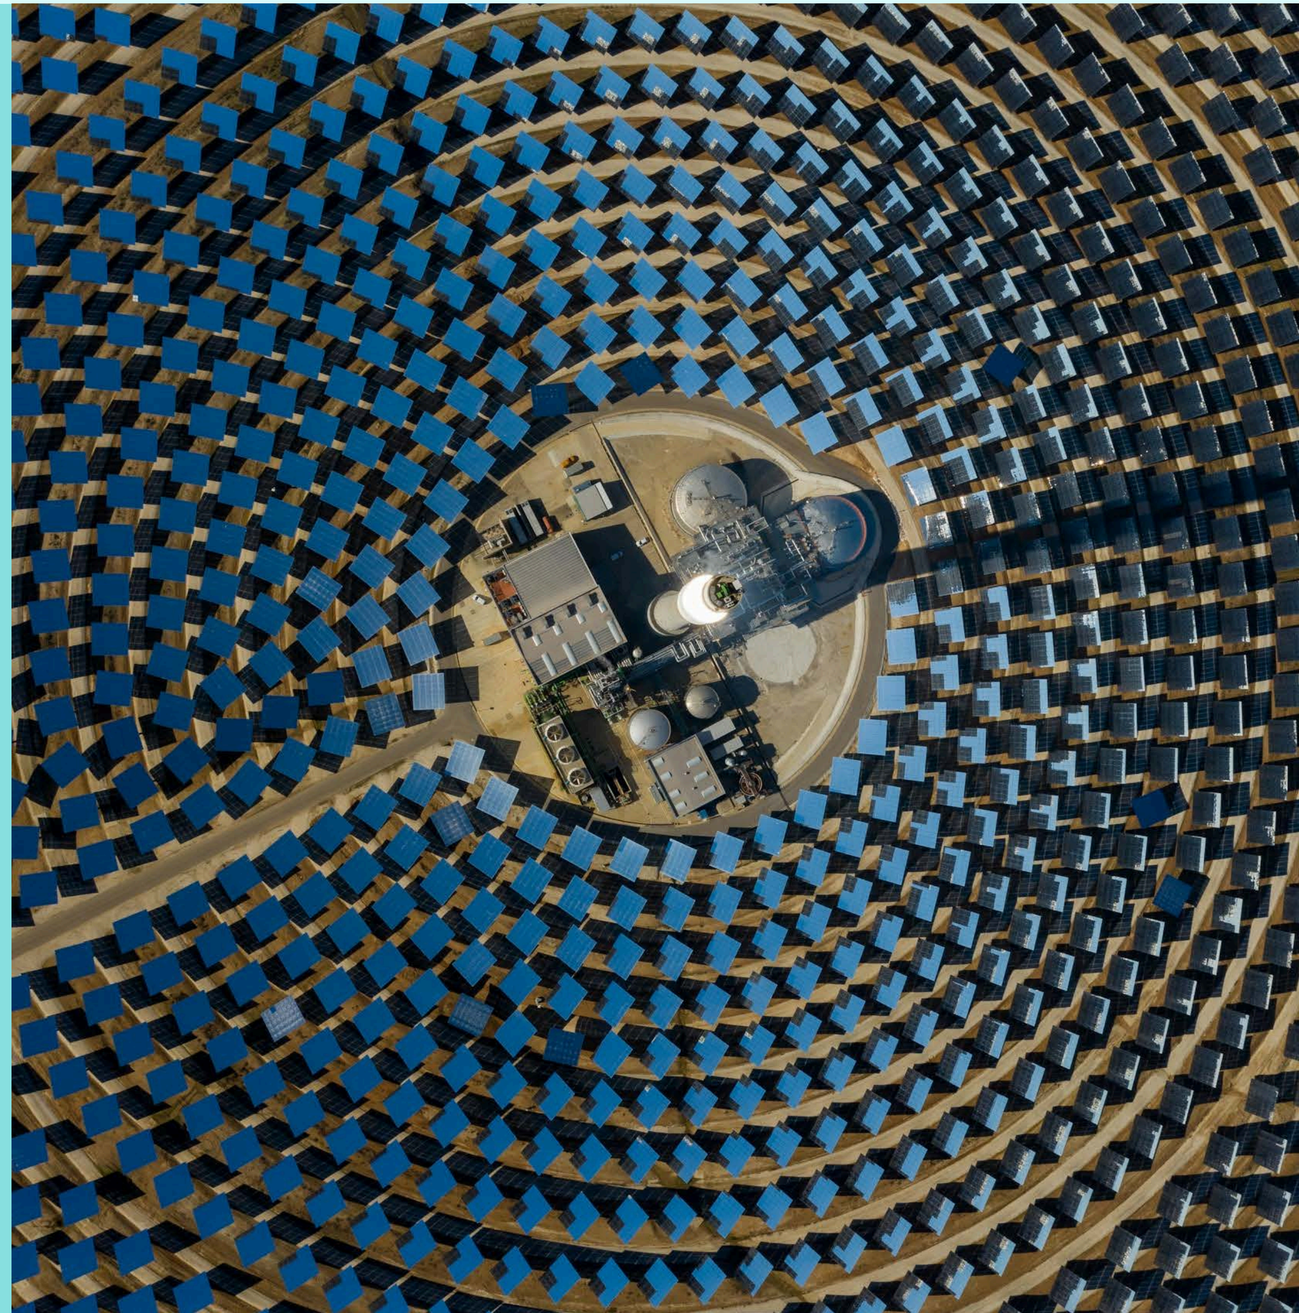

Why Every Organization Needs a Responsible Technology Board

A Responsible Technology Board is a cross-functional group that provides governance, standards, and practical guidance on the development and deployment of AI and emerging technologies. This includes reviewing and providing guidance on AI use cases to enable them to align with your organization's values and protect those you serve.

Responsible Technology Boards consider questions including:

- Is this use of AI fair to the communities we serve?
- Could this technology unintentionally harm our constituents or be repurposed for a secondary use to cause harm?
- Are we being transparent enough in our use of AI with our stakeholders?

Building Your Board: The Multidisciplinary Imperative

A multidisciplinary board is key to success. This may be setting up a new board or repurposing an existing governance committee or board. To make informed decisions about AI and emerging technologies, you need a range of perspectives and expertise, including:

Technologists

who understand the possibilities and limitations of AI

Leaders

who can align AI initiatives with the organization's strategy and mission

Frontline staff

who bring a deep understanding of the needs and realities of the communities you serve

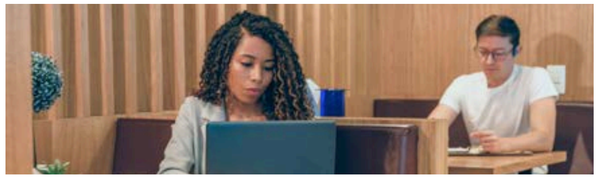

Other stakeholders

who can provide unique insights and expertise

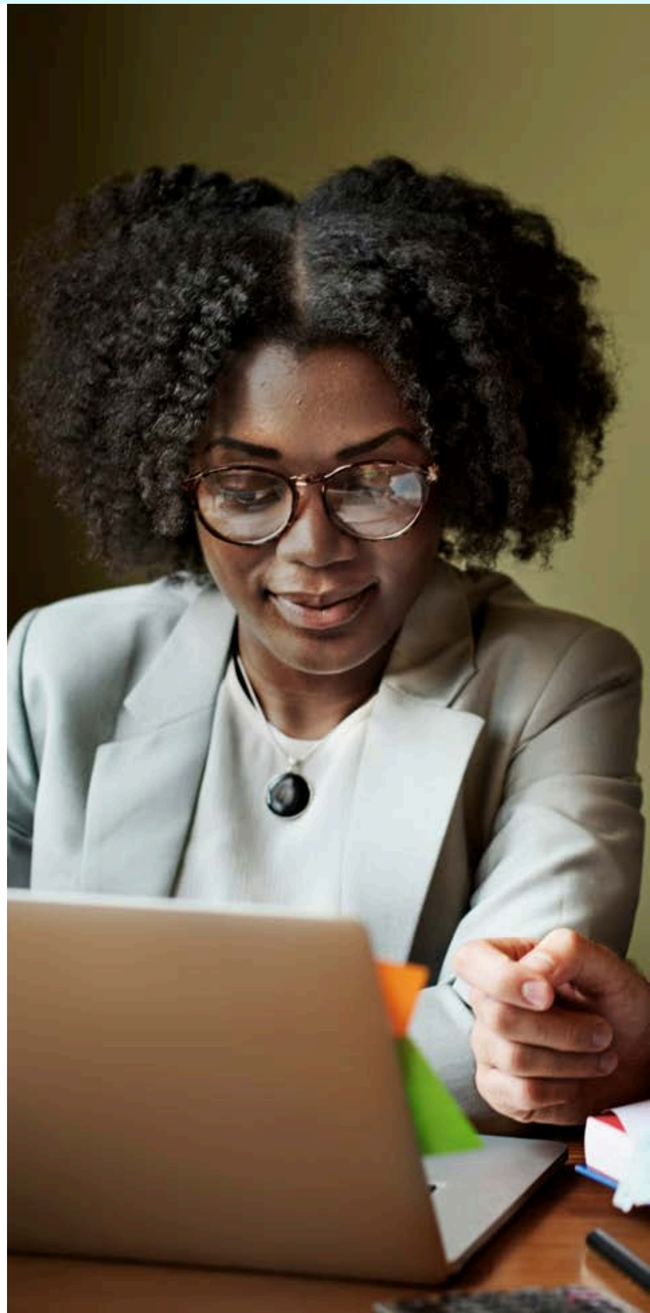

Responsible Technology Board Meeting Essentials

When making decisions about AI and emerging technologies, it's essential to consider multiple perspectives and factors.

Here are some key considerations to keep in mind:

- **Define clear goals and objectives**

  Establish a clear understanding of what you hope to achieve with AI and emerging technologies, and align stakeholders.

- **Assess risks and benefits**

  Consider the potential risks and benefits of AI and emerging technologies, and develop mitigation strategies.

- **Foster open communication**

  Encourage open and transparent communication among stakeholders, and for all voices to be heard.

- **Monitor and evaluate**

  Regularly review and assess the impact of AI and emerging technologies, and make adjustments as needed.

By considering these key factors, nonprofits can make informed decisions about AI and emerging technologies that align with their mission and values, and benefit the communities they serve.

Before	**Preparation** • Knowledge of what AI is and how it works • The skills to use AI tools for various tasks • The capacity to discern the risks and ethical implications associated with its use
During	**Decision-Making** • Discussion, not presentation • All perspectives heard • Formal vote taken • No majority = no approval
After	**Potential Outcomes** • Approve with monitoring • Defer for information • Require additional safeguards • Reject as misaligned

"You want different perspectives and opinions and input.

On the IBM AI Ethics Board, you have people from sales, from consulting, from corporate responsibility, from HR, from government relations, and I think that's what makes it robust and where you get the best decisions."

The Lesson

Justina Nixon-Saintil
Vice President, Corporate Social Responsibility, and Chief Impact Officer, IBM

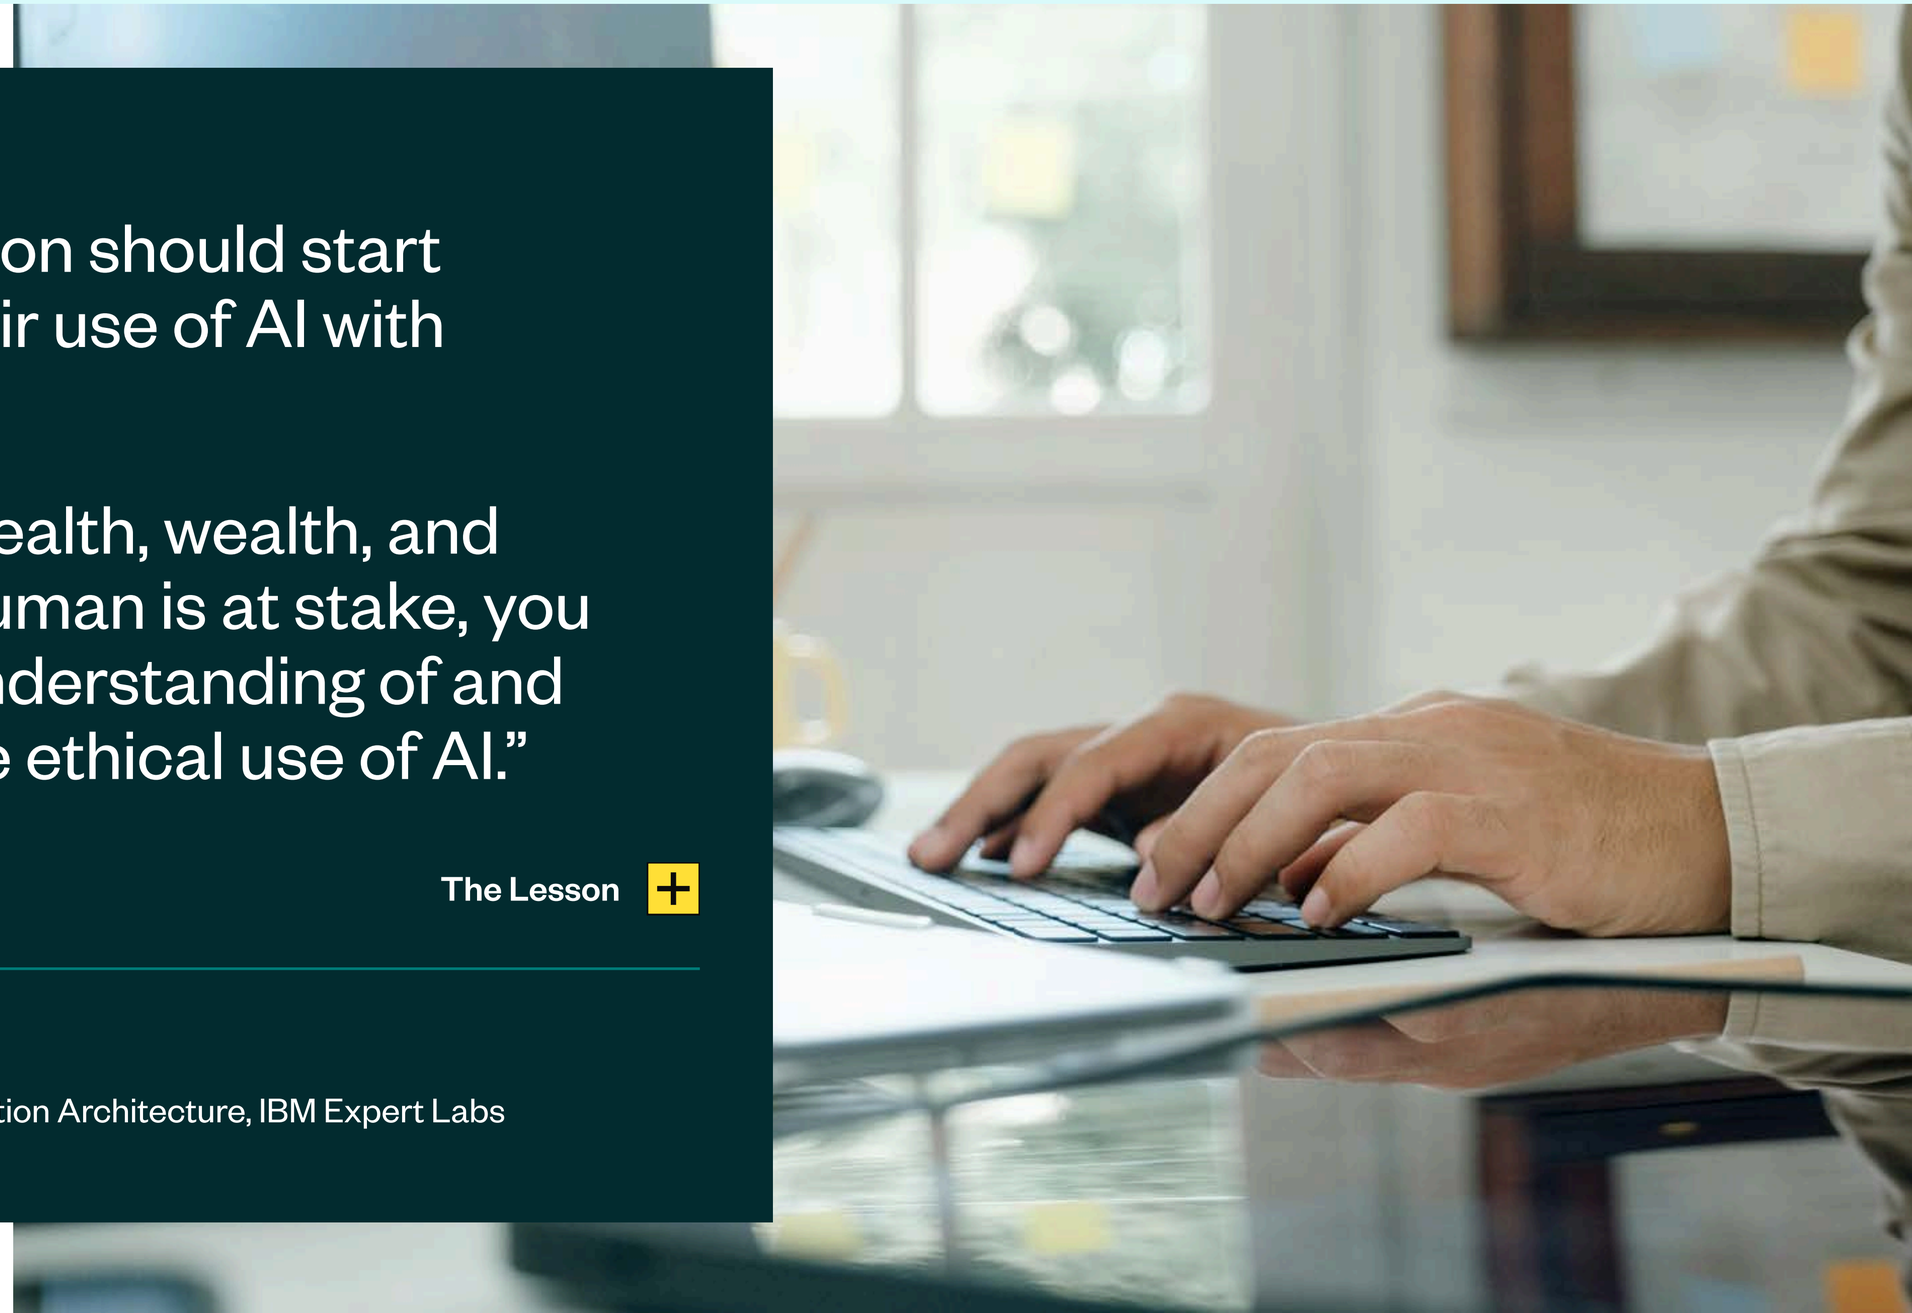

"Every organization should start with aligning their use of AI with their values.

As long as the health, wealth, and wellbeing of a human is at stake, you must have an understanding of and advocate for the ethical use of AI."

The Lesson

John Thomas
Vice President, Solution Architecture, IBM Expert Labs

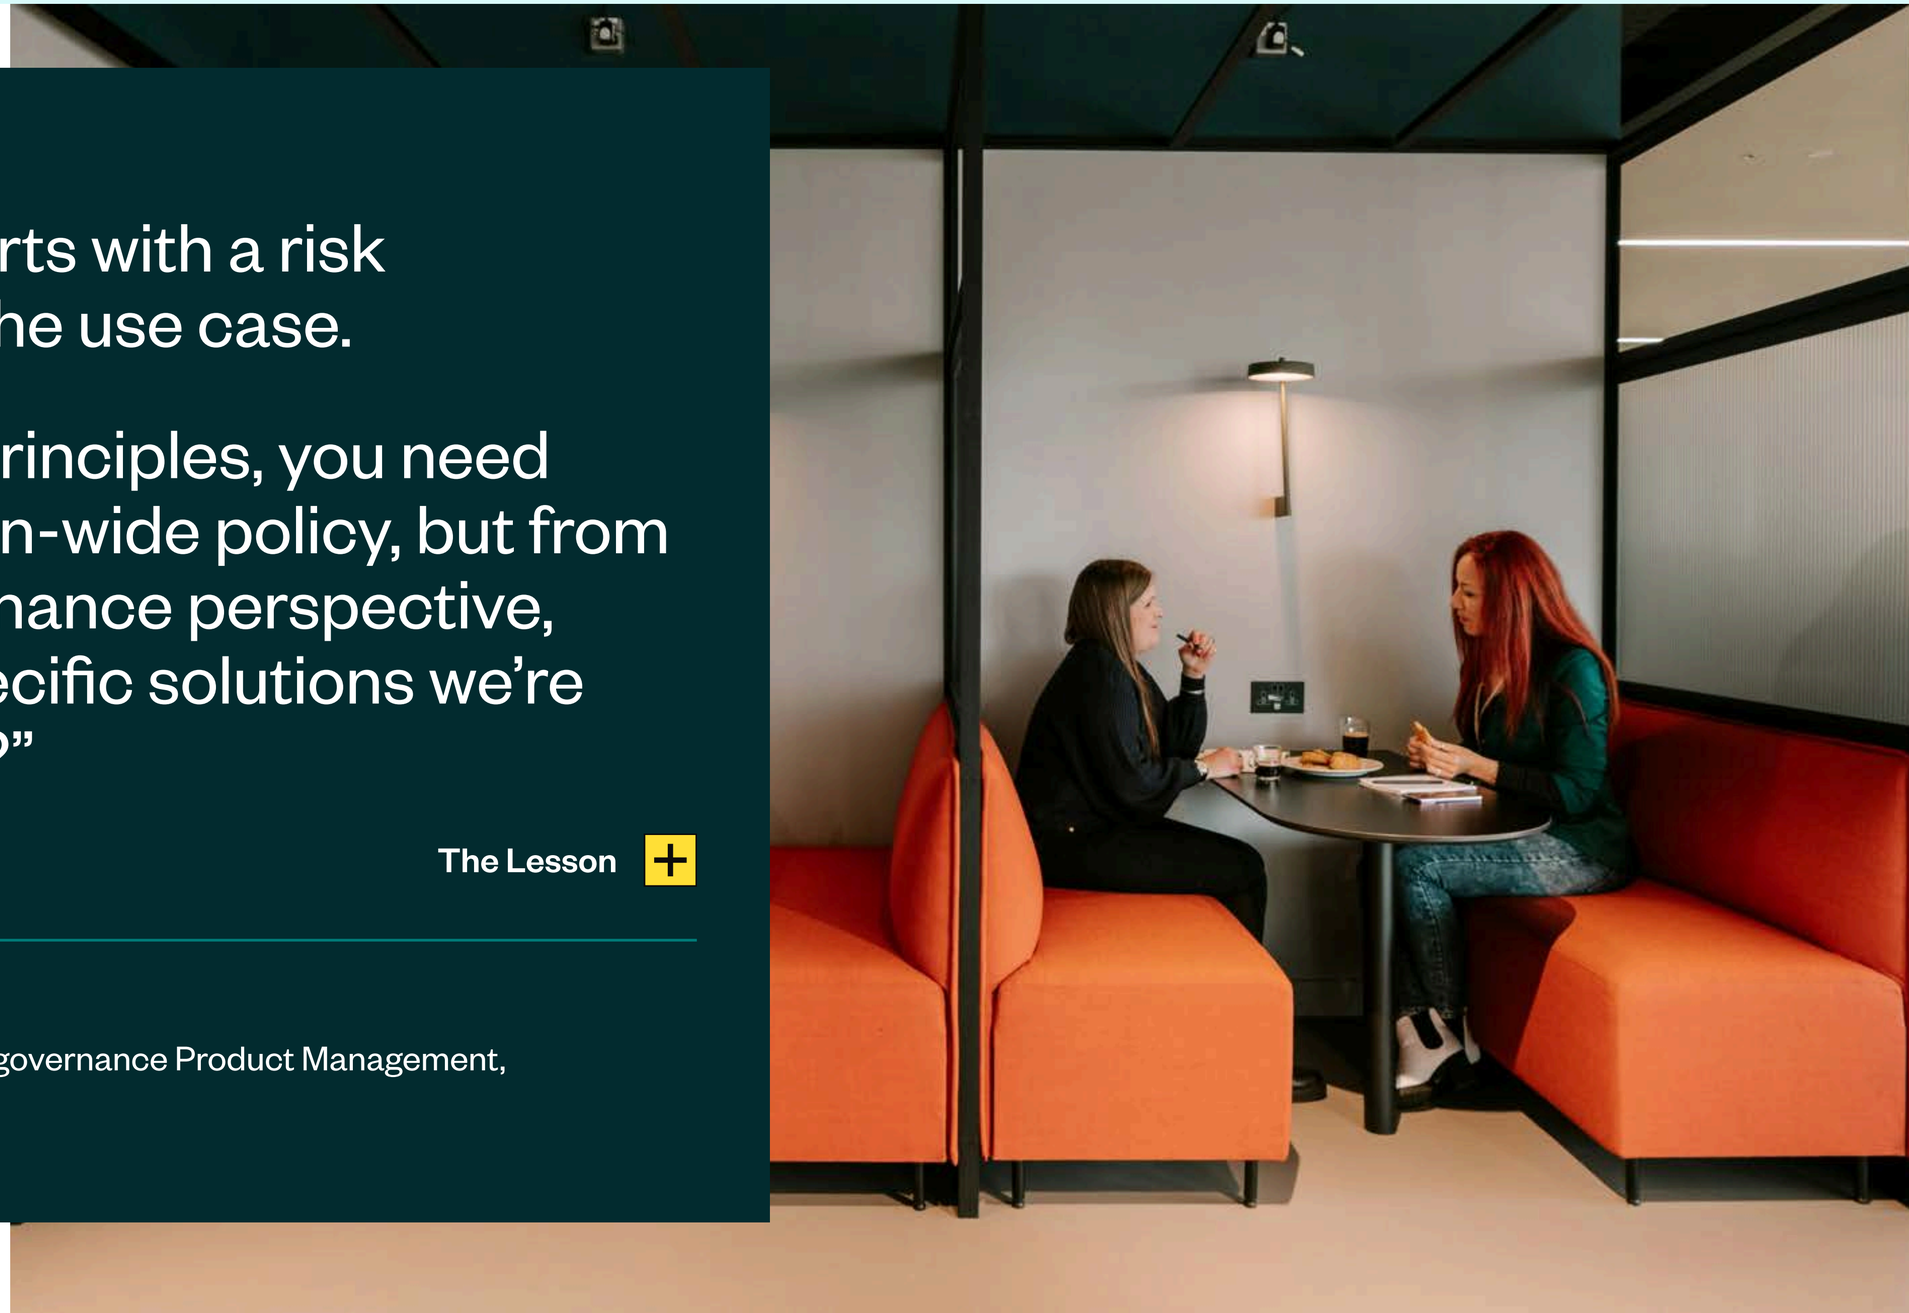

"Governance starts with a risk assessment of the use case.

You need your principles, you need your organization-wide policy, but from a risk and governance perspective, what are the specific solutions we're going to deploy?"

The Lesson

Heather Gentile
Director of watsonx.governance Product Management,
IBM Data and AI

Learn more

A Look into IBM's AI Ethics Governance Framework

IBM AI Risk Atlas

IBM Responsible AI Maturity Assessment

Responsible Technology Board

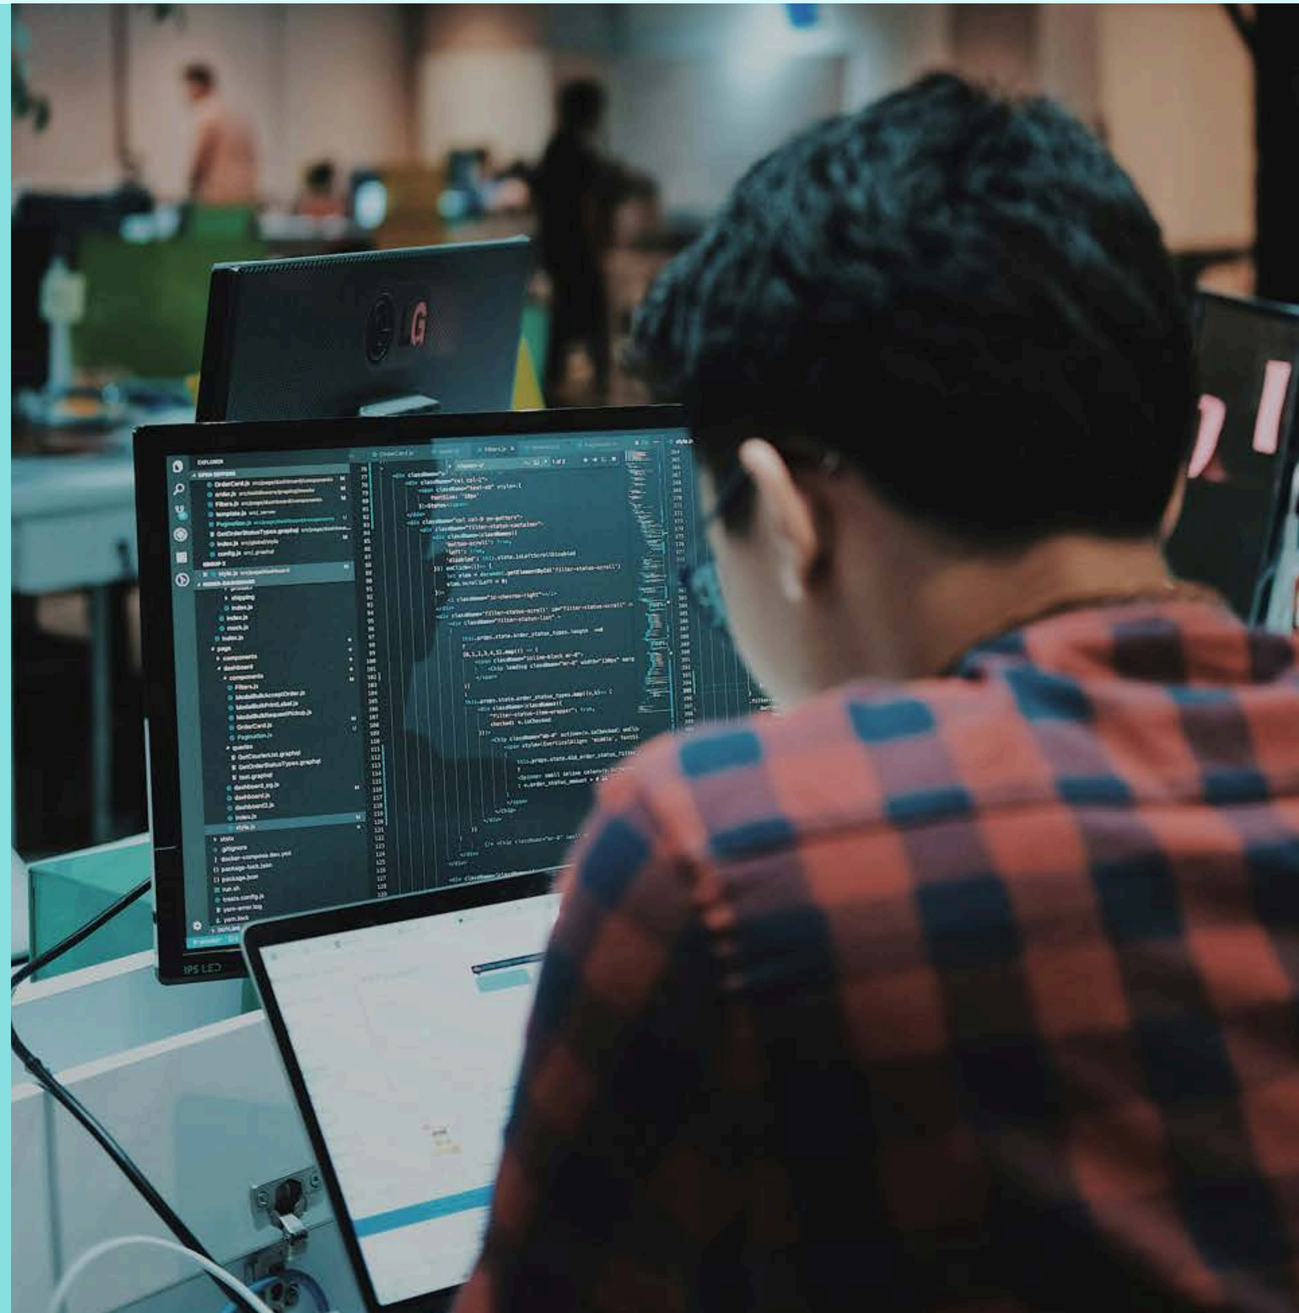

The Industry Exchange Network *(IXN)*

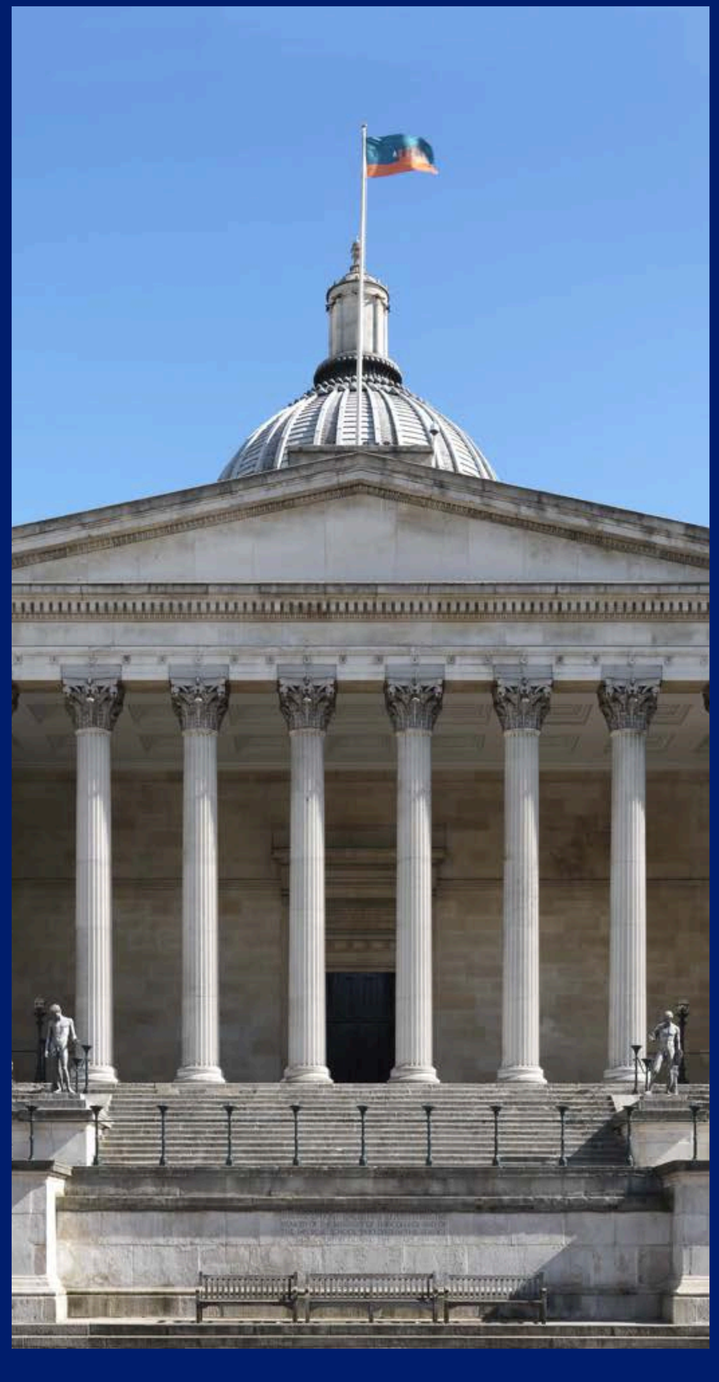

Lessons from University College of London's Decade of Connecting Students, Tech Companies, and Nonprofits.

Since 2011, the Industry Exchange Network (IXN) at University College London has connected thousands of students with 300+ organizations through a structured methodology for real-world problem solving.

Created by Professor Dean Mohamedally and Dr. Graham Roberts with initial partners including IBM, IXN cited as a best practice in the UK government's Topol Review (2019) and endorsed by the Foreign & Commonwealth Office.

The Core Innovation

Students work on real projects as 25% of their coursework, with formal agreements protecting all parties and requirements ensuring sustainability.

By the Numbers

Learn more about the Industry Exchange Network

Learnings for Nonprofits Partnering with Universities from the Industry Exchange Network

Define Your Specific Outcome

The Challenge

Organizations may approach AI capacity building without a clear understanding of how they will define success.

Specific Enough to Succeed

- Reduce client intake time from 3 days to under 24 hours through automated document classification
- Predict service demand 2 weeks ahead to optimize staff scheduling
- Automate initial screening to handle 80% of inquiries without significant human intervention

Get Your Data House in Order

Professor Dean Mohamedally's Reality Check

"All want innovation but they don't have their data house in order."

The Data Readiness Checklist

- ☐ **Data Exists**
  You actually have the data you think you have
- ☐ **Data is Accessible**
  It's not locked in paper files or legacy systems
- ☐ **Data is Documented**
  Someone knows what each field means
- ☐ **Data is Clean Enough**
  A critical mass of your data is clean enough to work with or you have a viable path to getting there
- ☐ **Privacy is Clear**
  You know what can be shared under what conditions
- ☐ **Someone Owns It**
  A person in your organization can make decisions about data access

Focus on Strategic Advancement, Not Innovation for Innovations Sake

Professor Dean Mohamedally's Guidance

"Charities shouldn't just shoot for innovation but strategic ways to advance their work."

Strategic Questions to Answer Before Launching an AI Innovation Project

- ☐ **Mission Alignment**
  How does this help us serve more people or serve them better?
- ☐ **Sustainability**
  What's our plan when students graduate and the project needs maintenance?
- ☐ **Capacity Reality**
  Can our team actually use and maintain this solution?
- ☐ **Opportunity Cost**
  What won't we do because we're focusing resources here?

The Litmus Test

Can you complete this sentence?

"We will know this project succeeded when [specific metric] changes from [current state] to [desired state], enabling us to [mission impact]."

Learn more

How to Craft a Data Cleanliness Policy

IBM SkillsBuild Data Fundamentals

IBM SkillsBuild Data Analytics Fluency Pathway

How Funders Can Catalyze Responsible AI Across the *Social Impact Sector*

Funders have a critical role to play in developing the shared infrastructure and organizational capacity to ensure the benefits of AI are shared broadly and the potential harms are minimized.

Strategy 1

Support the Development of Public Purpose AI

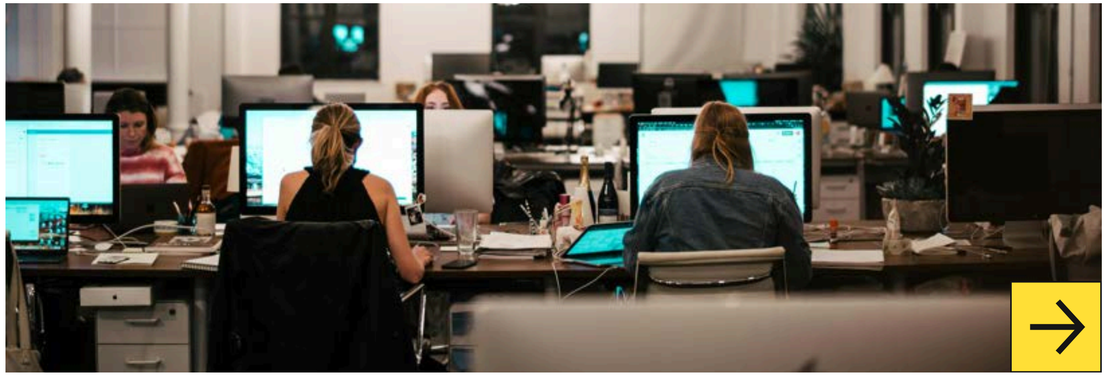

Strategy 2

How Funders Can Support AI Literacy and Capacity Building

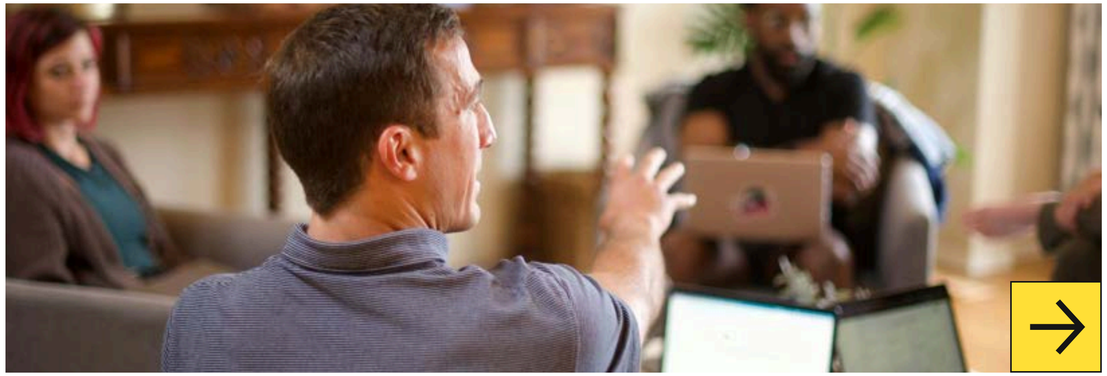

Strategy 1

Support the Development of Public Purpose AI

Individual nonprofits lack the resources to develop the AI tools needed to serve their organizations. The issue isn't just technical capacity within a single organization, it's about creating new institutional fabric that aggregates talent, compute resources, and data across organizations to achieve scale.

What the Intermediary Layer Looks Like

Philanthropy has the opportunity to create tools and institutions that sit between individual nonprofits and big tech, aggregating:

- **AI Talent:** Shared data scientists and engineers that multiple nonprofits can access
- **Compute Resources:** Pooled funding for expensive cloud computing and AI processing
- **Data Infrastructure:** Secure platforms for organizations to share and analyze data collectively
- **Standardized Tools:** Solutions built once for the sector, not recreated by each organization

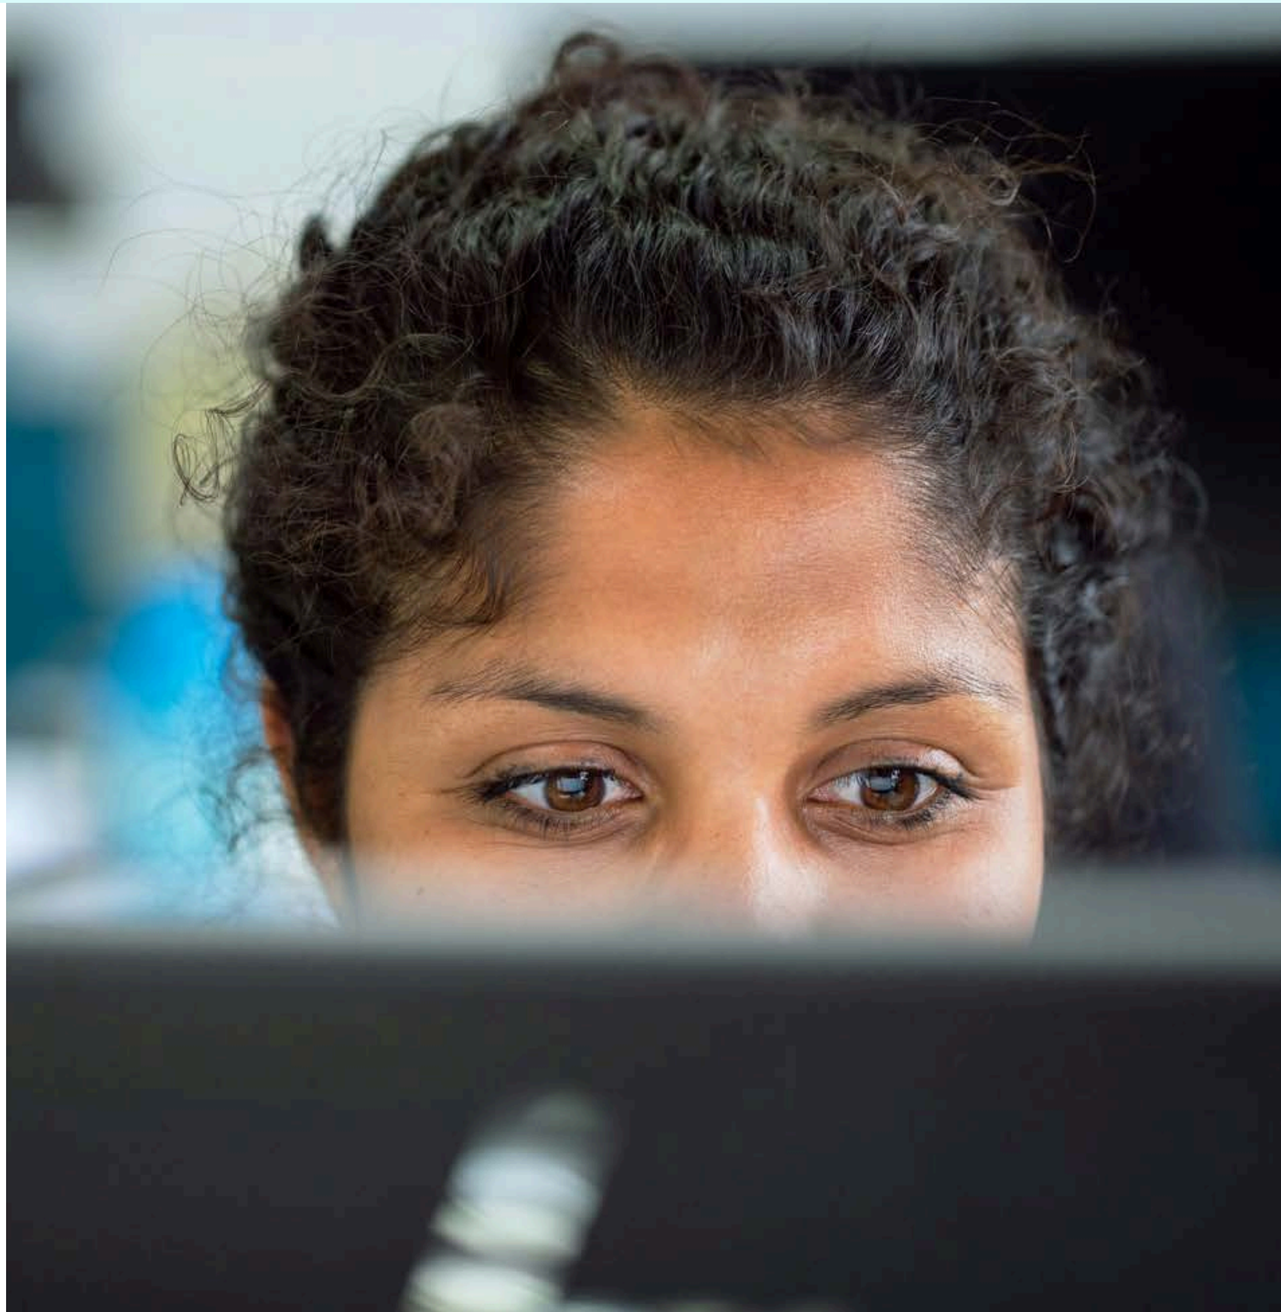

"One nonprofit organization can't realistically hire an entire AI team. There is a class of challenges where philanthropy can play a unique role in building an intermediary layer of institutions that support nonprofits to build public purpose AI."

Vilas Dhar
President, Patrick J. McGovern Foundation

Real-World Examples

IBM Impact Accelerator

The IBM Impact Accelerator builds and implements AI and other technical solutions for nonprofit, governmental, and academic organizations. The program offers:

- Up to two years of intensive support
- Access to IBM watsonx AI platform and other leading technologies
- Dedicated technical mentorship from IBM subject matter experts
- Shared learning network connecting nonprofits facing similar challenges

Real-World Examples

Grant Guardian

Patrick J. McGovern Foundation's (PJMF) Grant Guardian demonstrates this model in practice. Rather than having each foundation build its own financial analysis AI, PJMF created a free tool now used by 189+ philanthropies that:

- Standardizes financial due diligence across the sector
- Uses Anthropic's Claude to extract data from 990s and audits
- Saves hours per grant review while improving consistency
- Is provided free as public infrastructure for philanthropy

Strategy 2

How Funders Can Support AI Literacy and Capacity Building

We stand at a critical technological juncture. Just like the adoption of the internet and mobile phones, AI is increasingly being integrated into our daily work.

For funders, this moment is about providing activation energy during the transition period. After that, most nonprofits will have AI integrated as a core part of their organizational infrastructure.

For Every $1M in Nonprofit Operating Budget

> "Nonprofit organizations achieving AI fluency, where LLMs are used as frequently as excel, can reasonably see 10-15% efficiency gains — equivalent to a $100,000 unrestricted grant for a $1 million nonprofit."
>
> – Rich Leimsider, Director of AI for Nonprofits, Fund for the City of New York

Result: More resources for mission delivery

This Applies to Every Funding Area

- **Education funders:** Automated administrative tasks mean more teacher-student interaction
- **Health funders:** AI-enabled case management frees clinician time for patients
- **Environmental funders:** Data analysis that took weeks now takes hours

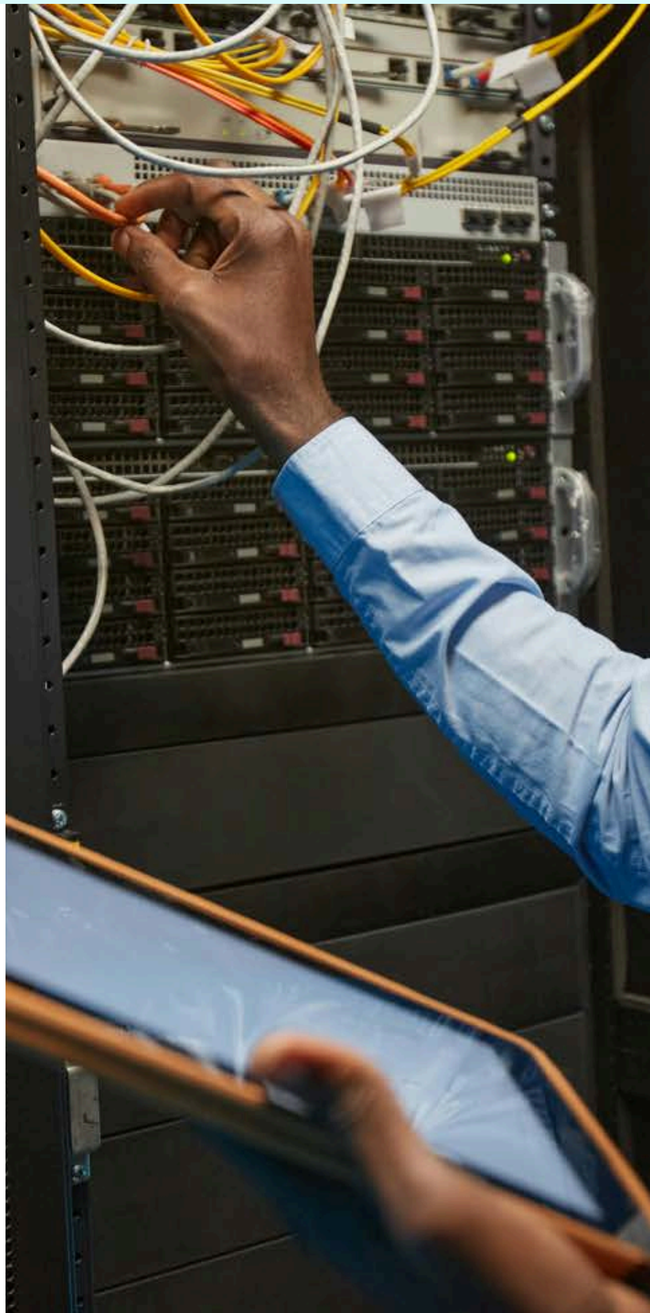

"Nobody needs a grant anymore to realize the internet is important.

This is that same moment. Philanthropy shouldn't expect to invest in basic AI literacy for more than the next 3 to 5 years."

Rich Leimsider
Director of AI for Nonprofits, Fund for the City of New York

Lessons from the *IBM Impact Accelerator*

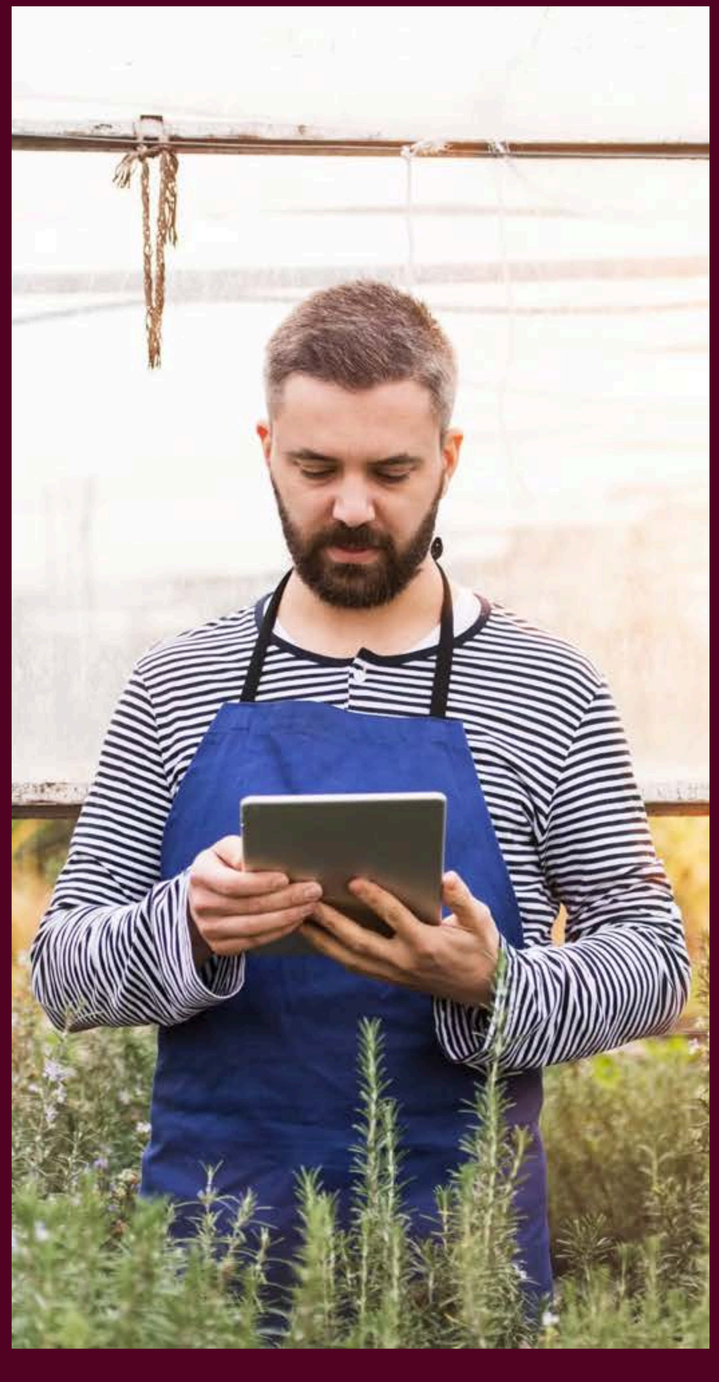

Begin with the end user in mind.

The IBM Impact Accelerator demonstrates a critical lesson: successful AI projects begin with the end user, not the technology.

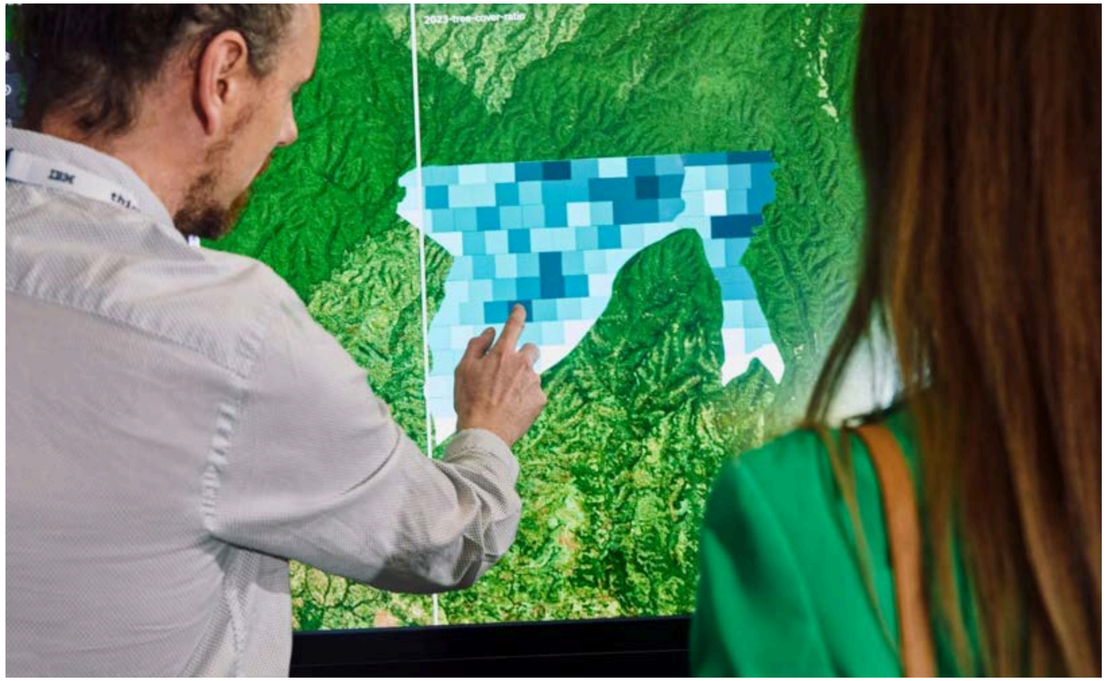

The IBM Impact Accelerator is a pro bono social innovation program that applies IBM technologies, such as hybrid cloud and AI, and expertise to enhance and scale solutions, helping populations facing environmental, economic, or social stress. IBM selects five organizations annually for cohorts that receive up to two years of support. Participants gain access to technical mentorship and enterprise-grade technology, which is configured and implemented to meet specific but public challenges.

Key Lessons from the IBM Impact Accelerator

01

Have a Clear Implementation Pathway

Organizations need more than a good idea — they need boots on the ground. The most successful participants arrive with implementation partners ready, whether that's community health centers, extension services, or local governments.

02

Embrace User-Centered Design

When Mass General Brigham initially envisioned a dedicated app for extreme heat warnings, user research revealed a clear reality: vulnerable populations face fierce competition for screen time among all target users. **The solution?** Simple text messages integrated with existing health communications.

03

Build Beyond Your Organization

IBM and its collaborators deliberately develop tools to meet the needs of core stakeholders. Mass General Brigham's heat warning system isn't just for their patients — it's designed for community health centers who could never afford to build it themselves.

04

Be Comfortable with Evolution

Being comfortable changing course or being comfortable with varying levels of impact is essential. First movers face headwinds, and the path from vision to implementation rarely runs straight.

Mass General Brigham

Extreme Heat Risk Tool

"[The Extreme Heat Risk Tool] is aimed at scaling up innovative action in cities around the world.

Together, Mass General Brigham and IBM will develop an AI tool for healthcare systems and community health centers confronting health risks from extreme heat."

Learn more

My FarmWell

Mobile App with AI-Powered Chatbot Advances Sustainable Farming in the UAE

"Farmers in the United Arab Emirates now have a smarter way to manage water and irrigation thanks to the My FarmWell app. Developed in collaboration with the University of Sharjah through the IBM Impact Accelerator and leveraging IBM's watsonx, the app empowers farmers with practical, actionable insights, as well as demonstrating how AI can transform agriculture and support sustainable development at scale."

Learn more

Sustainable Energy for All (SEforALL)

Open Building Insights

"IBM and Sustainable Energy for All (SEforALL) announced a new, publicly available artificial intelligence (AI)-powered solutions to inform more sustainable urban development for cities and communities around the world, enabling decision-makers and policymakers to map urbanization and identify energy and infrastructure needs for communities in developing regions."

Learn more

InfoSaneamento

Instituto Yarandu

"In Brazil, millions still lack access to basic sanitation, affecting health and community well-being. Through the IBM Impact Accelerator, IBM, Instituto Yarandu, and BioSaneamento created InfoSaneamento — a web platform that uses IBM Cloud, watsonx.ai, and open-source tools to analyze local data and guide better water and waste infrastructure.

By empowering local leaders with data-driven insights, the project helps advance Brazil's progress toward universal sanitation."

Learn more

Begin your *AI Journey*

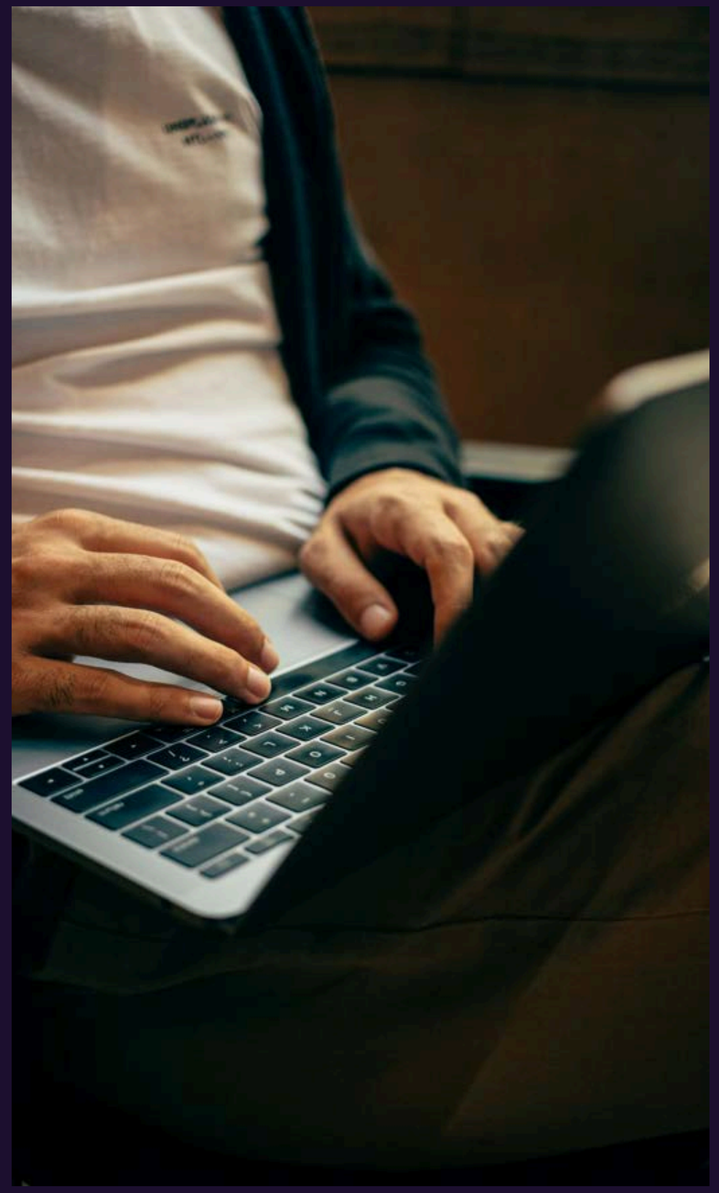

Take the AI competency assessment

Acknowledgements

Mobo Akiode
Paul Biddinger
Rafaella Bonogli Bruno
Raissa Caselas
Vilas Dhar
Heather Gentile
Leslie Hernandez
Michael Jacobs
Rich Leimsider
Sara Link
Lydia Logan
Mariana Paál Martinato
Professor Dean Mohamedally
Justina Nixon-Saintil
Dr. Olugbemiga A. Olatidoye
John Thomas

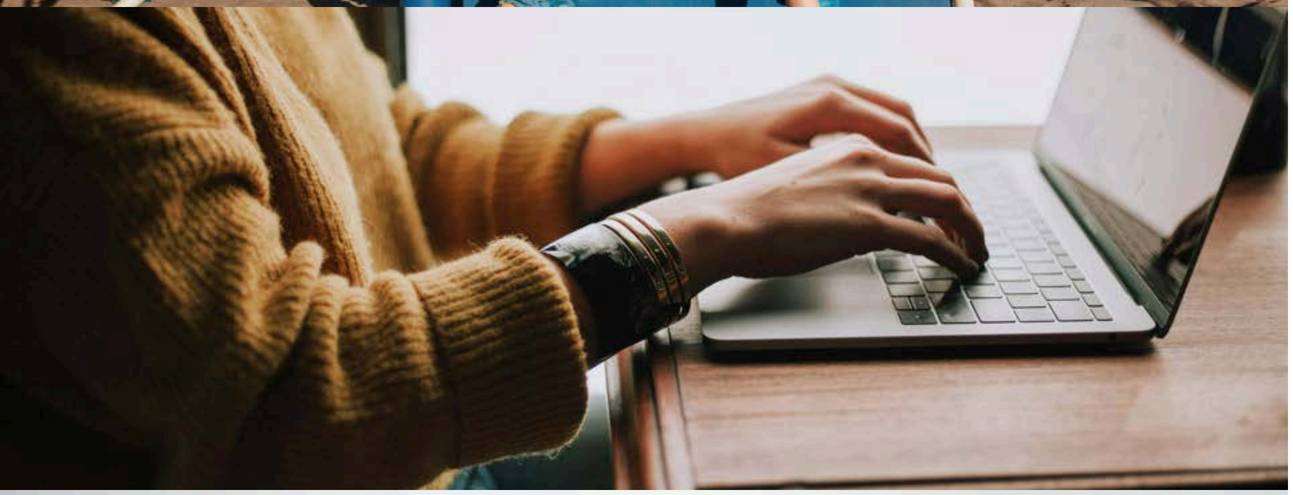

NATION SWELL

In collaboration with

IBM

Responsible Use of AI *for Social Impact*