Google.org | October 2026

AI Readiness Playbook for Funders

Table of Contents

Quick Start Guide **3**

Letter from Maggie 4

Context 5

Getting Started: A Shared Baseline **6**

Core Definitions 6

Top Use Cases for Funders 8

Module 1: Strategic Foundation **10**

Chapter 1: Responsible AI Use: Move from Policy to Practice 11

Chapter 2: Find Your Starting Point: A 5-Minute “AI Readiness” Survey 13

Chapter 3: Define Your “Why” 15

Module 2: Designing & Implementing an AI Program **18**

Chapter 4: Role-Based Starting Points 19

Chapter 5: Identifying High-Impact Use Cases 20

Chapter 6: Training Design: A Mission-First Model 21

Chapter 7: Activating AI Champions 22

Chapter 8: Creating a Central Hub 24

Chapter 9: Driving Workstream Integration 25

Chapter 10: The individual AI Showcase: Sparking Grassroots Curiosity 28

Chapter 11: The Team AI Showcase: Building Shared Workflows 31

Chapter 12: Measuring Progress and Confidence 34

Module 3: Enabling the Broader Ecosystem **36**

Chapter 13: The Strategic Imperative 37

Chapter 14: Evaluating AI Proposals: A Due Diligence Framework 38

Chapter 15: The Funder’s Role: Shifting the Funding Model 42

Closing **44**

Key Takeaways 45

Resources for Grantees 46

Quick Start Guide

To help you get comfortable with the content before diving into the full roadmap, we have created an interactive Gemini Notebook experience. This tool allows you to converse with the playbook, generating tailored advice and learning paths based on your organization's specific needs.

(Note: AI technology is evolving rapidly, and this playbook reflects the branding of Google AI tools as they were at the time of publication. Readers should rely on Gemini to learn about the latest product offerings, branding, and feature updates).

Try these starter prompts in Gemini Notebook ↓

QUESTION	PROMPT FOR GEMINI NOTEBOOK
A high-level summary	Give me a high-level overview of the three modules in this playbook and the main goal or outcomes of each.
Google.org's key takeaways	What are the top 3 'aha moments' or lessons learned from Google.org's own internal AI journey?
Personalized roadmap	I work at a foundation with the following challenges: [insert your own challenges around AI]. Where should we begin our AI readiness journey?
Quick wins	What are some low-risk, high-impact AI use cases for foundation teams, such as grantmakers or communications staff?
The "Why"	Explain the digital divide and why it is urgent for funders to learn about AI now.

Letter from Maggie



Dear Fellow Philanthropic Leaders,

The social sector is in a precarious position. Nonprofits are being asked to solve some of the world's biggest and most complex problems, but with dwindling resources and support. Many of us have made significant changes to our own operations and funding to help, and as a sector, we are working on new philanthropic models and sources of funding.

AI needs to be a part of this new framework for philanthropy. Not only can it help nonprofits do more with less, but it can also help solve old problems in new, innovative ways. This presents us – Google.org and other funders – with a profound responsibility and opportunity: we must ensure that the benefits of AI extend fully to the social sector. If we do not act deliberately, we risk creating a worse digital divide, where the efficiency, creativity, and scale that AI offers exists in some places, but not in the local communities that can often only be reached by nonprofits.

> To meet this moment, I urge you to make **AI education** for your own internal teams **a top priority**.

From your operations teams to your grant managers: every staff member should understand this technology—not just for their own and the organization's productivity, but so they can confidently and responsibly support the technology in the field.

In the past, technology adoption could live with a single entity like a CTO's team or a systems administrator, but AI is different. Because this technology impacts every function from drafting RFPs to analyzing impact data, we cannot silo this knowledge. There are likely early adopters on your teams already using AI which builds great momentum, but training and support need to be extended to every team member regardless of their role.

This transition requires top-down momentum from the leadership of the organization. It is about organizational change and creating careful strategies for AI adoption and learning that can alleviate the pressure on the social sector at this precarious time.

No single funder or nonprofit can bring the full benefits of AI to the social sector. We must do it together. To help, we've decided to open source how we've navigated these changes as a team. Our goal with this resource is to offer inspiration and practical support that helps bring AI education to your team.

The time is now. I hope you will join us.

Maggie Johnson
Global Head and Vice President, Google.org

Context

Building on our history of funding nonprofit solutions through AI, we recognized a new opportunity: folding AI into all of our work rather than have it exist as a standalone portfolio. This shift was driven by many factors: nonprofits are facing unprecedented demand and budget uncertainty at the same time that AI technology is improving and barriers to adoption are falling.

To meet this moment, we needed to bring AI knowledge to our whole team. From our grantmakers to our operations team, our goal is to establish a baseline of knowledge that empowers everyone to identify where AI can help solve critical problems today.

This playbook captures the key learnings from our own internal AI education program at Google.org. We're sharing it to give you a practical, actionable process for your own teams. We are still learning; AI models are significantly more capable than they were even a year ago. We have adopted a mindset of curious commitment and are investing in “learning to learn” alongside our fellow funders.

We want to demystify the process, share what’s worked for us, and help other funders harness AI to amplify their missions.



Target Audience

This playbook is for philanthropic funders broadly, but with a specific focus on organizations in the mid-to-long-tail of organizational size who may not have dedicated technical strategy teams. However, whether you have a staff of five or five hundred, the principles of change management and "learning by doing" remain the same.

Desired Outcomes

As a result of using the playbook, you will be able to:

- **Grow** your understanding of AI
- **Build** org-wide buy-in for AI capacity building efforts
- **Design** a suite of resources to foster org-wide AI education
- **Collect** feedback and evaluate the impact of new education efforts
- **Engage** in strategies that support grantees in building their own AI capacity

Getting Started
A Shared Baseline

First, we must ground ourselves in a common set of definitions and a fundamental understanding of the nature of AI tools.

Core Definitions

Artificial Intelligence (AI)	Computer programs that can do tasks that normally require human intelligence, like thinking or learning. Recently, advances in machine learning (a specific type of AI) have made AI programs more powerful and more useful.
Generative AI (GenAI)	A specialized type of AI capable of creating net-new content, including text, images, and other media.
Large Language Model (LLM)	An AI model trained on vast amounts of text to understand, synthesize, and generate human-like responses to prompts.
AI Tool	Software powered by AI that can help you with different tasks. Examples include tools for writing, analyzing data, or managing projects.
Autonomous Agent	A system that can pursue a high-level goal and execute multi-step tasks (like sending an email or updating a spreadsheet) autonomously.
Prompt	Text you give to a conversational AI tool to tell it what to do or what kind of output you want. Learning how to write effective prompts ("prompt engineering") is important for getting useful results from generative AI.
Retrieval-Augmented Generation (RAG)	A technical architecture used for source-grounding that allows AI to retrieve facts from your private files before generating an answer.
Human-in-the-loop approach	A way of using AI where humans and AI work together to train, use, and check AI models. This can help ensure accuracy and address ethical concerns.
The Jagged Frontier	The concept that AI capabilities are uneven and predictable, excelling at some hard tasks while failing at some easy ones, requiring humans to test every specific use case.

Transparency	The idea that you should be able to understand how an AI tool works and why it produced a certain result. While full transparency can be difficult, understanding the limitations of AI models is important.
Responsible AI (ACT Framework)	The idea of developing and using AI in a way that is ethical, benefits people and society, and avoids harm. A practice of safe use defined as: Ask (is this task appropriate?), Check (verify output for bias and accuracy), and Tell (maintain transparency about AI use).
Biased data	Data that isn't complete, doesn't accurately represent everyone, or unfairly favors certain groups. If an AI model is trained on biased data, it might also produce unfair or inaccurate results, which is a key concern for nonprofits focused on equity.
Hallucination	When AI produces information that is not true. It's important for nonprofit users to double-check AI-generated content for accuracy.
Source-Grounding	Restricting an AI tool's knowledge base to specific, trusted documents provided by the user to minimize hallucinations and ensure data integrity.



Top Use Cases for Funders

AI excels at solving specific operational challenges. Consider these high-impact applications:

The “Heavy Lifting”
(Data & Analysis)

AI is uniquely positioned to synthesize vast amounts of unstructured information into actionable insights.

- Digesting 50+ page grant proposals or research reports into concise summaries to accelerate due diligence.
- Identifying common themes, challenges, or funding gaps across hundreds of grantee reports simultaneously.

The “First Draft”
(Generative & Creative)

Use AI as a creative partner to overcome the blank page and tailor content for diverse stakeholders.

- Reshaping a single mission “win” into tailored artifacts for different audiences, such as a formal business case for a CEO and a social media post for beneficiaries.
- Brainstorming initial strategic plans or internal announcements by synthesizing meeting notes and core goals.

The “Drudgery”
(Routine Automation)

Leverage AI to tackle repeatable tasks that could cause operational bottlenecks.

- Extracting structured action items, decisions, and follow-up emails from unstructured meeting audio or transcripts.
- Automatically pre-filling standard templates using content from a potential grantee’s proposal.

The “Concierge”
(Information & Personalization)

AI can act as a high-level assistant to help navigate complex internal resources or external landscapes.

- Automating the research of new funding opportunities or monitoring evolving labor markets to align grantmaking with future community needs.
- Creating an internal hub where team members can ask questions of your organization’s trove of historical strategy documents.

Pause for Responsibility

Even in tasks where AI is well-suited, it is critical to always have a human-in-the loop to exercise final judgement.

Inside Google.org

For nearly 20 years, Google.org has lowered barriers to innovation by providing no-cost and discounted technology through **Google for Nonprofits**. We are evolving this model: we don't just provide tools, but also fund the overhead of digital transformation by covering the cost of technical onboarding and training:

Eligible nonprofits can receive **Google Workspace for Nonprofits licenses** for up to 2,000 employees or volunteers. This includes enterprise-grade data and privacy protections for the **Gemini app** and **Gemini Notebook**, which means your organizational data is not used to train Gemini models outside of your domain without permission and is not reviewed by humans. **Gemini in Google Workspace** has also been certified by stringent third-party auditors.

Google.org also covers the operational overhead of deployment of these tools by offering **no-cost 1:1 technical training support and migration guidance** via Google Meet video conferences.

Organizations can get started at google.com/nonprofits and will be invited to technical training sessions upon approval.

MODULE 1

Strategic Foundation

Before advancing to tools or training, start with a solid strategy. Many organizations begin their AI journey by asking, “What tools should we use?” The most successful ones start by asking, “What problems do we need to solve? What untapped opportunities might we pursue?”

This module reframes the conversation from a tech-first approach to a mission-first one, guiding the development of a strategic foundation before tools are introduced.

CHAPTER 1

Responsible AI Use
Move from Policy to Practice

It's important to highlight the importance of responsible AI use while also giving your teams room to experiment across low-risk opportunities. Balancing this tension between safety and innovation requires a shared ethical language rather than just a list of rules.



It’s critical to define what responsible AI means in the context of your organization, and this involves navigating risks and opportunities:

Fairness and Bias: AI-generated outputs may reflect biases present in training data. Make sure to monitor prompts and review outputs to ensure they are fair, representative and do not reinforce harmful stereotypes.

Accuracy and Accountability: AI can unlock creativity and efficiency but can also hallucinate, providing inaccurate information. A “human-in-the-loop” approach requires fact-checking outputs against reliable sources.

Privacy and Security: Make sure you have protocols to limit data input, ensuring that confidential, private or personally identifiable information is never exposed to open AI models.

Inside Google.org

Especially for our work in the social sector, trust and security are the most valuable components of AI tools. That’s why the Google tools that we use (including **Gemini for Workspace** and **Gemini Notebook**) are built with **enterprise-grade data protections**.

Your data is your data: your prompts, uploaded files, and generated responses stay within your organization and are **not used to train** foundational AI models or reviewed by humans without your explicit permission.

Built-in compliance: our AI solutions are among the first to attain comprehensive safety and security certifications, including **SOC 1/2/3, ISO 27001**, and **GDPR** compliance.

Safe grounding: tools like **Gemini Notebook** use **source-grounding** which means the AI only answers based on the specific, trusted documents you provide, significantly reducing the risk of hallucinations.

As a first step toward articulating your AI policy, start by operationalizing your values. As a company, Google has established AI Principles that govern all of our work across the company. At Google.org, we found that starting with a guiding conversation was more effective than a strict compliance checklist to bring those principles to life for our team.

We call this the "Principles in Practice" model.
Use the Principles in Practice Workshop Guide to run this session with your team.

Helpful Resources

Once you have aligned on your principles, consult these industry resources to learn any new ideas for inclusion.

1. **For AI Policy:** Fast Forward Nonprofit AI Policy Builder
   A free and easy tool to help nonprofits craft a custom AI policy.

2. **For Data Classification:** CISA Traffic Light Protocol (TLP)
   Adopt this globally recognized cybersecurity standard to help your team instantly classify documents (from 'Public' to 'Sensitive') without needing to memorize complex data laws.

3. **For Bias & Design:** Google PAIR (People + AI Research) Guidebook, “Feedback & Controls” section
   Leverage this framework to design 'human-in-the-loop' workflows that prevent automation bias and ensure every AI output remains under human oversight.

4. **For "Responsible AI" Basics:** Google for Nonprofits: Responsible AI Guide
   Share this primer with your wider team to establish a shared baseline on the core risks of AI (such as data privacy and hallucination) specifically tailored for the social impact sector.

Teams can get stuck in circular debates because they lack a shared language for navigating trade-offs around safety and innovation. To address this, we utilize a structured ethics inquiry deck: a set of prompts designed to uncover blind spots in high-stakes projects.

The Ethics Inquiry Deck
Use this resource when a project requires deeper scrutiny or involves high-risk communities. This method is adapted from Google DeepMind (GDM) to help you stress-test AI and data projects.



Inside Google.org

We leaned on Google’s AI Principles and ran workshops to translate them into our daily work. We use three pillars to guide our decision-making:

1. **Bold Innovation:** We will use AI to solve massive problems and scale our impact.
2. **Responsible Development:** We will never trade speed for safety.
3. **Collaborative Progress:** We will not build in a silo. We empower the ecosystem.

CHAPTER 2

Find Your Starting Point
A 5-Minute “AI Readiness” Survey

In order to design an effective skill development initiative, you’ll need to understand the baseline tool use and use-cases across your org. The most informed way to do this is to run a short, anonymous survey using this AI Readiness Survey Template to map your organization across three dimensions:

- **Frequency** (are they using it?)
- **Confidence** (do they know how to use it?)
- **Mission Readiness** (can they apply it to the mission?)

If you prefer a customized AI Readiness Survey, use this AI prompt as a starting point:

"Act as an organizational change management expert. Draft a 5-minute anonymous AI Readiness Survey for a philanthropic foundation. The survey should measure: (1) Frequency of general AI use, (2) Current confidence levels, (3) Operational use cases (like writing, meeting summaries, or data analysis), (4) Reported technical, administrative, and security barriers, and (5) Specific learning goals tied to grantee success. Ensure the questions are written in an encouraging, non-technical, and highly approachable format."

How to Read the Results

When you get the results, look for the "Confidence Gap" and the "Ecosystem Void."

- **The Confidence Gap**
  You will likely see many people answering "Frequently" (Question 1) but "Somewhat Confident" (Question 3). This means your team is using AI, but they might be using it without training or security protocols. Your training must focus on confidence, not just access.
- **The Ecosystem Void**
  Look closely at Question 7. If very few people select "How to assist beneficiaries," you have a strategic opportunity. It means your team views AI as a desk tool, not a mission tool. Your launch communications should explicitly link AI skills to grantee success.
- **The Efficiency Plateau**
  When analyzing your readiness survey, look for the “Efficiency Plateau.” This occurs when an organization successfully reaches high individual tool usage but stalls before achieving team-scale impact. Your goal is to identify this plateau in your data and use it to transition your team from being "AI-Adopted" to "AI-Adaptive":
  - **AI-Adopted (ad hoc individual shortcuts):** In your survey results, this emerges as high individual usage but low collaborative scaling. You will see staff reporting frequent use of unvetted, free tools for isolated, one-off tasks (such as drafting an email or summarizing a single document) with no shared prompt libraries, little-to-no team-wide standardization or collaboration, and a lack of clear data safety protocols.
  - **AI-Adaptive (cohesive, governed workflows):** This is the maturity level your program aims to build. In the survey, look for signals that staff report collaborating on shared, enterprise-grade solutions (like custom Skills or structured prompt templates), demonstrate a unified understanding of organization-wide data privacy boundaries, and actively build documented human-in-the-loop checkpoints directly into standard team operating procedures.

This Readiness survey will provide you with foundational data to compare to your AI Progress survey, which we recommend sending after the organization has completed all training initiatives.

Key Strategy

Measure confidence, not just usage. High usage does not equal high competence. A team member might be using AI daily but doing so inefficiently or unsafely. Your goal is to identify the gap between "I experiment with this tool" and "I can rely on this tool for critical work."

Inside Google.org

- When we surveyed our team, we found that most of our employees were using AI tools on a weekly basis. Impressive!
- However, many did not feel "confident" choosing the right tool for a task.
- This insight completely changed our training strategy. We realized we didn't need to evangelize (convince people to try it); we needed to *stabilize* (teach them to use it well). We shifted our approach from "Intro to AI" to specific, workflow-based deep dives. Learn more in Module 2's Training Design section.

CHAPTER 3

Define Your “Why”

AI capacity building is not just an IT project, it’s a core organizational change initiative. Your "why" must be tied directly to your foundation's mission.

Define Your “Why”

1. **Calibrate**
   Before identifying problems you hope to solve with AI, it’s important to gauge where AI tools excel. Consider the “Jagged Frontier” concept. AI capabilities are uneven: it excels at some tasks (like synthesizing 500 pages of text) but can fail at basic math or fact-checking. View roles as bundles of tasks and help your team distinguish between them:
   - **Augmentable tasks:** drafting, summarizing, brainstorming, coding and pattern matching.
   - **Human-only tasks:** final decision-making, high-stakes relationship management, and nuance/empathy.



2. **Identify Bottlenecks**
   Bring the team together to identify critical strategic goals and the bottlenecks hindering progress that align with AI's capabilities. Position AI as a potential tool to solve these existing business problems.
   - Instead of “We don't know what is happening in Education and AI.” (This isn’t something AI can directly solve.)
   - Try: “It takes our team two weeks to synthesize new research and landscape trends in Education, creating a lag in our strategic response.” (AI can help automate research and extract core takeaways from dense articles instantly.)

Pause for Responsibility

As you brainstorm, remember: **Efficiency should never come at the cost of equity.** AI is excellent at processing information, but needs our human context and gutcheck. Use the ACT framework to ensure your goals prioritize professional integrity over pure automation.

- **Ask**
  Is this strategic goal safe to pursue with AI? Avoid goals that would require inputting sensitive, proprietary, or confidential donor data (PII) into non-enterprise tools.
- **Check**
  If AI supports an information-dense friction point, at what specific stage will a human expert apply their judgment? AI should only inform the process, never control final high-stakes decisions like funding allocations.
- **Tell**
  How will we maintain transparency with the communities we serve? Commit to a clear trail that explains when AI-assisted analysis has influenced strategic shifts.

3. **Dual-Purpose Brainstorming**
   Facilitate a session centered on two key opportunities:
   - **Internal Goals (Back-Office Solutions):** How can AI unlock time for deeper, more strategic work? How can AI boost creativity or bring to bear new information that could enable progress? These are productivity investments—tools that help your team (and grantees) operate more efficiently.
   - **Ecosystem Goals (External-Facing Solutions):** How can AI directly serve our mission? How does becoming a strong AI user make the organization a better grantmaker or partner? These are programmatic investments—tools that interact with the community.
     - *Note:* Most organizations start with Back-Office Solutions to build confidence, but true impact requires a strategy for supporting Beneficiary-Facing innovation.

4. **Draft the Statement**
   Combine these insights into a dual-purpose statement using this simple template:
   - We are building our AI capacity in order to...
   - **[Internal Goal]:** e.g., Reduce administrative burden to focus on relationship building.
   - **[Ecosystem Goal]:** e.g., Better evaluate technical proposals to ensure capital supports scalable solutions.

Inside Google.org

For us, we've created initiatives to ensure our whole team are confident AI users to better support our nonprofit partners, build stronger AI portfolios, and help the entire social sector harness this technology.

Our "why" is dual-purpose:

- **Internal Efficiency:** Unlocking new insights and time for deep work.
- **Ecosystem Enablement:** We believe that to be good partners, we must understand the tools our grantees are using. We build AI capacity internally so we can recognize, evaluate, and support innovation across the sector. This fluency allows us to act as strategic matchmakers: when a grantee surfaces a problem, we can help them determine if AI can be a part of a meaningful solution.

Securing Buy-In (Framing the 'Why')

Whether you're pitching this to your leadership or you are the leadership aligning your board and executive peers, this dual-mission framing is key. You can pitch the AI initiative by highlighting the risk, the opportunity, and closing with a clear commitment.

- **The Risk:** If the private sector adopts AI and the social sector does not, a massive digital divide will emerge.
- **The Opportunity:** By becoming AI-fluent, the team can better evaluate proposals, offer better technical assistance, and ensure your capital supports truly scalable solutions.
- **The Commitment:** Every team member must dedicate time for the initial workshop and experimentation. Mandate that everyone include an AI proficiency goal in their official performance expectations for the year.

Key Strategy

This initiative must be relevant to your team's existing workflows. If AI is an extra, optional thing, it will be the first thing to drop when people get busy. For example, identify an upcoming, high-volume deliverable (e.g., Q3 board presentation, grant cycle reporting) and make that a training target. Call out how AI tools can augment the tasks involved for creating the deliverable. And, by formally integrating it into performance goals, you signal that cultivating these skills is a core priority.

Build Inspiration and Momentum

Your why will be aspirational, so you'll need to provide a "pull" to get your team excited about your ambitious goal.

For inspiration and tangible demonstrations of how AI tools can solve massive, sector-wide challenges, reference the **Social Good AI Hub** (https://www.thesocialgoodhub.ai/) by Google.org and IDInsight.

Internally, we built inspiration and momentum internally by sharing role-specific examples (our "aha" moments) in a simple spreadsheet that everyone could add to. (Learn more about Creating a Central Hub.) Managers reminded team members to share inspiring use cases and to build on peers' experiences.

Inside Google.org

A few of our effective "aha" moments:

- **Communications**
  Use **Gemini Notebook** to draft communications in alignment with our global and regional narratives and legally approved language.
  - Use **Gemini Canvas** to transform event transcripts into an interactive, visually engaging presentation. This makes complex information accessible and impactful.
- **Grantmakers**
  Use **Gemini** and/or **Gemini Deep Research** to analyze 50 grantee reports to identify common themes and challenges, then systematically evaluate them.
  - Use **Gemini Notebook** to synthesize 40+ page proposals into a business review template designed for faster leadership decision-making.
- **Operations**
  Use **Google Workspace with Gemini** to get a first draft of a complex formula in Sheets or map out a 6-month project plan for a new initiative.
- **Analysis**
  Use **Gemini** and **Google Forms with Gemini** to analyze AI Readiness Survey & Progress Survey results, then use **Gemini Canvas** to build an interactive results dashboard (complete with AI-generated insights).
- **Ecosystem**
  Use **Google AI Studio** to ingest hundreds of grantee AI use cases and build a structured repository that would become the **Social Good AI Hub**, transforming scattered files into a searchable learning tool for the entire sector.

MODULE 2

Designing & Implementing an AI Program

Now that you have established your strategic “why” and considered your ethical guardrails, the next step is execution. This module outlines the practical design of an AI capacity building program.

We recommend a role-specific, workflow-integrated approach delivered in four phases:

1. **Map Your Team.** Use your AI Readiness survey data (from Module 1) to segment your team into specific personas based on their role and confidence levels.
2. **Discover High-Impact Use Cases.** Before you design the training, you must identify what to train on. Conduct a "Task Audit" to identify bottlenecks (like summarizing 50+ grantee reports) that are a good fit for AI augmentation.
3. **Customize the Curriculum.** While off-the-shelf training is a helpful starting point, we recommend tailoring it to your team's specific needs where possible. You can use our 90-Minute Workshop Guide as a foundation, plugging in your own high-impact use cases and tools to ensure the session speaks directly to your team's actual pain points.
4. **Activate Champions.** Deploy the training and simultaneously deputize power users to sustain momentum after the session ends.

CHAPTER 4

Role-Based Starting Points

AI adoption is not linear. To develop skills effectively, you must map use cases to your team's specific personas. We recommend a dual-persona approach:

By Functional Role (What they do)

- **Grantmakers**
  The learning themes here are two-fold. First, **operational efficiency** (summarizing proposals, analyzing landscape reports). Second, **strategic evaluation**. Grantmakers must learn to assess the technical feasibility of AI components in grantee proposals.
  - *Potential Goal:* Shift time from manual synthesis to strategic decision-making, and build the technical confidence to assess the feasibility of AI components in grantee proposals.
- **Comms & Advocacy**
  Focus on drafting, editing, and translating complex ideas into different formats.
  - *Potential Goal:* Accelerate content production by using AI for first drafts and reformatting, allowing the team to focus on narrative strategy and stakeholder alignment.
- **Operations**
  Focus on data structuring, policy drafting, and soft-coding (e.g., generating formulas or scripts).
  - *Potential Goal:* Streamline workflows by augmenting high-volume administrative tasks.

By Adoption Stance (How they feel)

- **The Experimenter**
  Uses AI occasionally when prompted but lacks a routine.
  - *Goal:* Move them to consistent weekly usage.
- **The Power User**
  Proactively identifies new opportunities and shares prompts.
  - *Goal:* Deputize them as peer teachers.
- **The Skeptic**
  Concerned about data privacy, "workslop" (low-quality output), or mission alignment.
  - *Goal:* Address specific pain points with high-value, low-risk use cases (“aha” moments) to build trust. Reinforce your AI policy through these trainings to make sure the trade-offs feel well considered.

Key Strategy

Once you’ve mapped these groups, use the data to customize your workshop materials. Create a simple one-pager for each persona on how they can leverage AI. Before you host the workshop, distribute these “cheat sheets” and give the team time to practice so no one feels intimidated or starts from zero.



CHAPTER 5

Identifying High-Impact Use Cases

Now that you've mapped your team's personas, you must also identify the problems you need to solve. This can help your training feel less generic and more tailored to your organization's specific needs.

1. Ask teams to list **3 high-volume tasks** that feel draining or repetitive, e.g. manually copying data from PDF grant reports into Sheets.
2. Check if these tasks fall within **current AI capabilities**. For example, AI excels at synthesis, drafting and pattern matching but shouldn't be used for high-stakes decision making.
    - Good fit: "Summarize these 10 proposals against our rubric."
    - Bad fit: "Decide which of these 10 proposals gets funding."
3. Select the use cases that offer the most **team-level relief**, not just personal productivity or convenience.
    - Personal win: "Help me write this email faster."
    - Team win: "Analyze our entire portfolio's quarterly reports to identify challenges across the ecosystem."

Pause for Responsibility

Never trust AI completely. When vetting a new administrative or grantmaking use case, run it through the ACT checklist:

- **Ask**
  Am I following our internal data classification rules? For example, using AI to summarize a grant agreement is AI-assisted, but entering Protected Health Information (PHI) from a grantee's program data is prohibited.
- **Check**
  How will we verify that the AI hasn't flattened community nuance or hallucinated a metric? You must commit to independently verifying all facts and names before the output moves to the next workstream.
- **Tell**
  How will you disclose that an AI summarized a grantee's proposal? Maintain transparency about tool use to protect grantee trust.

CHAPTER 6

Training Design
A Mission-First Model

By framing your AI workshop as a strategic mission-alignment exercise, you empower your team to proactively upgrade their capacity and deepen their support for grantees. Use our 90 Minute Workshop guide to host customized sessions for your team.

As you design your workshops, prepare your team for the Agentic Era. AI training is no longer just about prompt engineering for first-draft emails. It is about learning to oversee autonomous systems that execute multi-step workflows. Staff must be trained as architects and evaluators who monitor agentic actions, rather than just writers.

Key Strategy

Frame the workshop not just as “learning new tools,” but as building mission capacity.

Inside Google.org

Our AI Readiness survey results showed that our team’s #1 barrier was **knowing which tool** to use for which task. But more importantly, we realized that high internal usage didn’t automatically translate to use cases that support the ecosystem.

We shifted our training examples. Instead of just showing how to draft an email or other basic use cases, we showed tasks that were tied to creating impact, like analyzing a grantee report for common themes across a portfolio or using **Gemini Deep Research** to kickstart a landscape analysis of a new problem space. This subtle shift legitimized the use of AI tools, proving that **AI could deepen their understanding of their beneficiaries** rather than distancing them.



CHAPTER 7

Activating AI Champions

You do not need a large technical team to run an effective development program. In fact, peer-to-peer learning is often more effective than top-down training because it feels relevant instead of prescriptive.

To scale this work without needing to hire new staff, you must identify and empower a network of AI Champions, or internal enthusiasts who act as force multipliers within their specific teams.

Key Strategy

While this is a peer-led model, it cannot be purely volunteer work. Formally recognize and protect the time of your Champions (even just 10-15%). If you treat AI leadership as a professional development opportunity, they will thrive instead of burn out. This opportunity must be reinforced top-down from leadership and communicated consistently as a key priority for the organization.

Commit Capacity
Secure temporary support for the launch phase.

- **Program Lead** (1 person at ~20%): A project manager to coordinate the surveys, training and resources.
- **Tiger Team** (2-3 people at ~10%): A cross-functional group to pressure-test material and lead workshop discussion sessions.

Identify Champions
Don’t just tap people with the most technical backgrounds. Look for the Champion mindset. Champions tend to be power users who are not just proficient, but proactive and generous:

- **Typical user:** “I use AI to finish my work faster.”
- **Champion:** “I figured out how to use AI to augment and improve this task; here’s the prompt so you can do it too or share how to make it even better.”

Enable Influence
Dedicate 5-10 minutes of each team meeting for a Champion to demo one workflow that saved them time that week.



Characteristic	The Experimenter	The Power User
Usage	Uses AI tools occasionally or when prompted	Uses AI tools daily or weekly as a standard part of their workflow.
Initiative	Follows suggestions on where to apply AI.	Independently identifies opportunities to apply AI to new tasks.
Mindset	Cautious and hesitant, unsure of AI's capabilities for complex tasks.	Confident and capable, leveraging AI to tackle complex challenges.
Collaboration	Consumes AI tips shared by others.	Proactively shares their own successful prompts, use cases, and tips.

Inside Google.org

We defined the "Power User" to set a cultural standard. We made it clear that collaboration was a key metric of success. (Reference the table above.)

Peer validation was critical for our team to overcome skepticism. When a grantmaker saw a peer using AI to analyze a grant report, the hurdle of trust disappeared and the hope for time-savings became tangible.

CHAPTER 8

Creating a Central Hub

One of the biggest risks in a skill development initiative is that the learning happens in silos. A grantmaker in Climate might figure out the perfect prompt to summarize grantee reports while a colleague in Education struggles with the exact same task. To prevent this knowledge loss or double work, move from individual discovery to collective intelligence in the form of a central hub.

Start Simple
Don’t over engineer it! The goal is low friction and ease of use. The best hub is one that your team already knows how to use, like a simple shared spreadsheet or running document.

Useful Entries
Steer users away from generic entries like “used AI for emails.” Every entry in your hub should contain four specific components:

- **Task:** What specific problem were you solving?
- **Tool:** Which specific tool did you use?
- **Prompt:** Copy and paste the exact prompt that worked.
- **Impact:** Quantify the win. (e.g. “saved 30 min”, “improved clarity”, “caught an error”)

Create a “Toolshed”
In addition to a section to add use cases, your hub should have a “toolshed” section or tab. This acts as your organization’s source of truth for governance. It should list:

- **Approved tools:** Which tools are approved and safe to use with internal data?
- **Use status:** Is the tool for piloting, general availability, or restricted use?
- **Access:** A direct link to where the team can access or request license for the tool.

Driving Adoption
A hub can easily become a digital parking lot if not actively managed. Here are a few ideas to keep it alive:

- Have your AI Champions select one high-impact entry from the hub and demo it for 2 minutes during your regular staff meeting.
- Offer public recognition or another small incentive to recognize the best prompt or use case of the month.
- Ensure managers are leading by example and adding their use cases to the hub to encourage reports to do the same.

Key Strategy

Collect problems, not prompts. Don't just ask your team to share cool prompts. Ask them to share the problem they were solving. A repository of problems solved is more valuable than a repository of prompts.

Inside Google.org

We built a **Central AI Use Case Hub** using a standard spreadsheet accessible to the entire organization. We kept the barrier to entry low, asking team members to simply log the **Task, Tool, Prompt and Impact**.

We also found that adding a “Contributor” column was vital. It added peer validation because if a Grant Lead vetted a prompt, others were more likely to trust it and try it themselves, aiding in those “aha moments”!

CHAPTER 9

Driving Workstream Integration

Individual adoption is a great start, but real ROI comes when you embed AI into the “hand-offs”, or the points where work passes from one team to another.

To unlock systemic efficiency, you need to identify high-volume, cross-functional workflows where AI can reduce friction.

Identify Processes for Integration

Don’t try to overhaul everything at once. Look for processes that suffer from bottlenecks, for example where one person or team has to read, summarize or reformat a massive amount of information before the next person or team can act.

Common Candidates

- **Grant Reporting:** Grantee → Grant Lead → Leadership
- **Board Prep:** Multiple Teams → Strategy Team → Board Deck
- **Communications:** Technical Report → Comms Team → Public Blog Post

Example 1

Grant Reporting

The Old Way

A grantmaker spends a week reading 50 individual PDF reports, manually copying data into a spreadsheet, and writing a summary for leadership.

The Integrated Way

- **Ingest:** Upload all 50 PDFs into a secure tool like **Gemini Notebook**.
- **Structure:** Use a standard prompt to extract key metrics (e.g., "Lives impacted," "Budget utilization") into a structured table.
- **Synthesize:** Ask the tool to "Identify the top 3 shared challenges across this portfolio."
- **Review:** Use Gemini Notebook’s built-in citations to verify data points or claims against the original source documents. This “human-in-the-loop” check ensures the AI hasn’t generalized away community nuance or hallucinated a metric.
- **Result:** The grantmaker spends their time analyzing the implications of the data, rather than just finding it.

Example 2

Strategic Matchmaking

The Old Way
A grantee asks for help with a specific challenge (e.g., marketing strategy) or requests technical assistance. Grantmaker schedules a call, brainstorms based on their personal memory, and sends a few links.

The Integrated Way

1. **Context:** The grantmaker uses a custom AI tool like a **Gemini skill.** When they go in to create a skill, they provide custom text instructions for AI and upload their organization’s source of truth documents (partner directory, strategy documents, any best practices) to serve as the knowledge base as well as any other relevant information.
2. **Action:** Upload the grantee’s proposal and any other supporting resources and prompt the tool: "Based on [Grantee Name]'s target audience and budget, brainstorm 5 low-cost marketing strategies and identify 3 potential technical assistance partners."
3. **Review:** The grantmaker reviews the output through the lens of organizational voice and feasibility and performs a “human-in-the-loop” check:
    - Do these partners have proven experience in the grantee’s specific region?
    - Does the proposed tone of the marketing plan align with the grantee’s mission and community values?
    - Does the plan avoid any high-risk tactics that might lead to low-quality engagement?
4. **Result:** The grantmaker shows up to the meeting with a tailored menu of options and resources, transforming a check-in into a strategic working session.

Pause for Responsibility

Ensure the drive for efficiency doesn't compromise grantee privacy (Input) or misrepresent their impact (Output).

- **Input Check**
  Before uploading files, you must classify the data. Does this document contain PII (names, addresses of beneficiaries) or confidential intellectual property? If the tool is not approved for sensitive data, do not upload it.
- **Output Check**
  AI summarization often erases complexity. Does this summary strip away the community's voice or generalize their struggle? Verify quotes and metrics against the source file before sharing it.



Inside Google.org

Our grantmaking teams built custom **Gemini skills** to act as dedicated thought partners for grantee proposals or new initiatives.

Set up: We supplied the tools with specific ecosystem strategy documents, evaluation criteria and goals.

Workflow: Grantmakers upload a proposal or brief and engage the tool to:

- **Analyze** the proposal's strengths and weaknesses against our documentation.
- **Test** the technical feasibility and alignment with our goals.
- **Draft** targeted follow-up questions for the potential grantee.

Guardrails: We frame these skills as tools, not decision-makers. They surface critical insights and blind spots, but our human grantmakers always check the AI's work and make the final decisions. Users also fine tune the Skills as they go, both from the skills' ability to learn through continued conversation and as the users notice blind spots.

CHAPTER 10

The Individual AI Showcase: Sparking Grassroots Curiosity

To build genuine excitement, we launched a lightweight, low-barrier submission window to celebrate how team members were integrating AI into their daily work. The initiative was designed to show that AI adoption doesn't require deep coding or overly technical skills, but rather a problem-solving mindset. The outcome was beneficial to everyone: we not only celebrated those demonstrating curiosity and exploration, but we also built more AI use cases to socialize across the entire organization.

Our core philosophy with this showcase is to celebrate and reward various levels of implementation across the entire learning curve—from "just getting started" to seasoned experts—because every step forward makes the entire organization more impactful.

Establish the Categories and Evaluation Process

To encourage maximum participation while ensuring high-quality evaluations, we divided the judging into two parallel tracks: **Peer Voting** for organization-level engagement with all entries, which sparks inspiration, and **Expert Evaluation** for technical and functional depth. This dual-track approach yields five distinct options for winning.

We encourage you to tailor these categories to what works best for your specific goals and matches the size of your staff (e.g., smaller teams may want to simplify to just the peer-voted categories, while larger organizations can support the full specialized rubric).

1. **Peer-Voted Categories (Broad Engagement)**

    - **Wow Factor:** The most creative, highly innovative, or technically impressive use of AI to solve a complex challenge.
    - **Daily Driver:** The most practical, efficiency-focused workflow hack that saves time on repeatable, daily tasks.

2. **Expert-Evaluated Categories (Deep Functional Alignment)**

These categories are paired with an evaluation rubric to support the judges. We provide an example prompt to generate a starting point for your rubric in the upcoming Launch Resources section.

- **Operations:** Identifies and implements an AI solution that directly enhances internal operational efficiency (e.g., simplifying grant reporting).
- **AI Consultant:** Helps a nonprofit or funder partner identify where AI could be a meaningful part of a solution for a social impact challenge.
- **AI Application:** Uses specific AI technology and any internal support offerings to address a complex social sector challenge.

Design the Right Incentives

We recommend motivating teams with incentives that build cultural momentum and peer recognition, such as:

- **All-Teams Spotlights:** Feature winning submissions and workflows during org-wide All Teams meetings to celebrate creativity publicly.
- **Direct Leadership Exposure:** Invite winners to present their AI workflows directly to executive leadership or the board of directors, establishing them as internal strategic thought partners.
- **Professional Development Support:** Offer specialized training, certifications, or attendance at tech-for-good conferences as rewards.

Generative Prompts for Launch Resources

To make this launch low-lift for your operations team, you can leverage AI to generate your forms for submission and voting, and an evaluation rubric for the expert-evaluated categories. Our team created simple submission and peer voting forms in Google Forms, and leveraged Gemini to build an evaluation rubric that human experts used to begin reviewing submissions.

Here are three ready-to-use prompts:

Generate the submission form structure:

"Act as an organizational development specialist. Draft a clear structure for a Google Form designed to capture AI showcase submissions. Organize the form into four clear sections: (1) Essentials (Project Title, Lead, Team, and selected category), (2) The Pitch (1-2 sentences on the Problem, AI Solution, and Impact), (3) Category-Specific Criteria (which internal supports were leveraged, which workflow was improved, or how a partner was supported), and (4) Responsibility & Scale (explaining Human-in-the-Loop checks, novelty, and safe use of data if relevant)."

Generate the peer voting form structure:

"Act as a user experience and Google Forms expert. Draft a short, engaging peer-voting form structure to capture votes for our AI showcase. The form should allow voters to select one winner for the peer-voted categories: 'Wow Factor' (the most creative or impressive entry) and 'Daily Driver' (the most practical, time-saving entry), alongside a brief, optional comments section for peer feedback."

Create the evaluation rubric:

"Act as an AI change management expert. Create a detailed scoring guide based on a 40-point rubric. Detail the specific scoring thresholds for three tiers: Catalyst Champion (35-40 pts), Strong Contender (25-34 pts), and Emerging Innovator (under 25 pts). Break down how judges should award points for domain alignment (up to 20 pts) and the three responsibility pillars: Safe/Responsible (10 pts), Novelty (5 pts), and Scalability (5 pts)."

- Reminder: the AI-generated evaluation rubric is just a tool that serves as a starting point for your expert human judges! This is the point system we used, but you are welcome to customize the prompt to meet your needs.

Key Strategy

Do not wait for someone to build a complex, custom application before celebrating them. A five-minute workflow optimization that saves a grantmaker two hours every week is far more valuable to organizational culture than an over-engineered solution that no one else can replicate. Focus on celebrating these highly transferable "aha" moments.

Inside Google.org

When we launched our inaugural individual showcase, we wanted to prove that AI was a tool for everyone, not just our technical leads. One of our winning entries came from an operations team member who demonstrated how they used Gemini to instantly generate complex spreadsheet formulas that previously required hours of troubleshooting.

Seeing a peer share a tangible "daily driver" was highly motivating across all our teams, especially non-technical. The energy really took off across the organization because it was paired with active, top-down support. When our leadership team (including our managers and VP) championed the initiative and allocated dedicated time for learning, it gave everyone permission to play without the fear of falling behind.

CHAPTER 11

The Team AI Showcase: Building Shared Workflows

Once individual confidence is established and teams have a shared knowledge baseline from the recent workshop and prototyping session, a worthwhile progression is to drive shared workflows across core operations. The Team AI Showcase model brings teams or departments together to collaborate on automating or streamlining highly impactful processes. Over a multi-month period, teams move through structured phases to build, test, and implement their solutions in real-world scenarios, presenting their measurable efficiency gains and workflows at the end.

Team AI Workshops as a Launchpad

We found that a team-level AI Workshop was helpful to kick off our team showcase. We required our organization to run the AI Workshops for Team Enablement as the official kickoff for the Team Showcase. These workshops are designed to build on top of the previous 90-Minute AI Workshop, which focused more on individual experimentation and early confidence.

By dedicating focused, structured time to whiteboarding bottlenecks and building rapid prototypes during the workshop, teams walked out of the workshops with an improved problem-solving mindset and a set of tools to put to the task. This hands-on momentum helped assist with tool discernment, adoption friction and built confidence around using prioritized tools as a team (rather than solely as individuals).

Team Showcase Resources and Guardrails

Instead of acting as restrictive barriers, clear guardrails serve to accelerate, empower, and de-risk the Team AI Showcase. They provide teams with a safe, predefined runway to experiment confidently, knowing exactly where the data safety boundaries lie so they can focus entirely on problem-solving.

To establish this supportive environment, we recommend deploying three resources:

1. **Submission Guidance Sheet:** outlines clear parameters, evaluation criteria, and milestones for the Showcase.
2. **Approved Tool Directory:** outlines approved AI tools or platforms and their specific capabilities, data use policies by tool (as relevant), and guidance on which tool is best for which task.
3. **Showcase Project Tracker:** a shared database (can be as simple as a spreadsheet) to log ideas, prioritize solutions, and share progress updates.
    a. Access Tip: Giving everyone access to this tracker builds a sense of shared progress, breaks down silos, and encourages collaboration across teams who might be addressing a similar challenge!

By leading with these resources, you encourage cohesive AI use. Rather than leaving staff to navigate the landscape on their own, this approach guides them toward secure, pre-vetted environments that protect organizational integrity.

This proactive framing prevents the friction and safety risks of fragmented tool use, which can unknowingly expose sensitive foundation data. Providing a secure runway reinforced by clear parameters and a tool directory ensures that team collaboration remains uninterrupted, secure, and fully scalable.

Team Showcase Phases: From Concept to Implementation

To keep teams aligned, focused, and moving forward without feeling overwhelmed, we structured the three-month operational sprint into four clear, sequential periods. This lifecycle guides teams from initial ideation to final organization-wide scale.

1. Brainstorm, Align & Submit

This phase is about focusing momentum. Intact teams align on a practical problem and formally log their intent.

- **Brainstorm & Choose (Weeks 1–4)**: In regular team meetings, teams revisit their workshop notes, discuss operational challenges, and select 1–2 high-priority ideas that are high impact enough to warrant the extra effort and focus.
- **Categorize & Submit (Week 5):** Teams enter their idea into the central tracker, detailing a 1–2 sentence problem statement, the proposed tool, and its focus classification (Whole Org, Multi-Team, or Team-Specific).
- **The "Green Light" (Week 6):** Team Managers or department leads review submissions and issue a quick confirmation to activate the building phase. They can also connect teams exploring similar ideas to reduce duplication of efforts.

---

2. The Building Sprint

This phase is a focused execution window designed to move teams from concept to working prototype.

- **Build (Week 7):** We recommend aligning this phase with dedicated calendar blocks. (We planned ours during a set “No Meeting Week.”) Team members set aside a few focused hours to collaborate, write system instructions, configure their custom LLMs or agents, and build a functional "Version 1" of their workflow.

---

3. The Test Phase

This phase moves the tool from the sandbox into daily operations to build trust and check limits.

- **Test (Weeks 8–12):** Teams casually integrate the "Version 1" prototype into their day-to-day work tasks. This period is critical for answering: Does this actually save time? Does it require better prompting? Are the human-in-the-loop safeguards working? Teams use this time to gather feedback, update monthly tracking logs, and iterate on their designs.

---

4. Implement, Showcase & Scale

This phase is about celebration, integration, and scaling successful solutions.

- **Team Showcase (Week 13):** Teams submit their final project status updates. The Showcase wraps up with an organization-wide event where teams present their "Before vs. After" workflows, and celebrate accomplishments.
- **Implementation & Integration:** Move tested, refined AI workflows from the pilot stage to permanent, team or organization-wide implementation. Integrate the successful tools into standard operating procedures and share the finalized templates, system prompts, or custom agents/tools across the organization (even in a new tab of your Showcase tracker).

Technical Support Infrastructure

To help teams along in case they get stuck in any phase of the showcase process, we recommend supplying a support system during the sprint. Depending on the structure and resources available to your organization, you can tap internal or external experts for help:

- **Internal AI Champions or Tiger Team:** peer mentors embedded in each team to guide colleagues and troubleshoot tool use. You can elect new Champions or rely on your existing cohort selected earlier.
- **External AI Consultants:** provide deep expertise on security integrations, data checkpoints, and enterprise suites.

Inside Google.org

When we launched our Team AI Showcase, we used these three focus levels to categorize and track projects based on how they scaled across teams or the whole organization:

- **Whole Org:** Universally applicable across the organization; can be centralized for broad adoption.
- **Multi-Team:** Relevant to daily work across multiple teams or regional offices; can be scaled beyond a single group.
- **Team-Specific:** Highly tailored to a specific team, role, or unique task; not scalable across the broader organization.

These classifications helped us determine how broadly to share or implement each solution, and determine the level of cross-functional communication required for each project.

CHAPTER 12

Measuring Progress and Confidence

You can’t improve what you don’t measure. To evaluate the success of your program, you must measure both confidence and integration.

We recommend deploying a brief follow-up survey using the AI Progress Survey template. By comparing these results against your baseline AI Readiness Survey (from Module 1), you can perform an analysis to quantify your team's growth.

If you prefer to create your own progress survey, use this AI prompt as a starting point and attach your AI Readiness Survey as a resource:

> "Act as a professional learning and evaluation expert. Draft a brief, anonymous AI Progress Survey designed to measure shifts from our baseline AI Readiness Survey. To ensure our data sets align for a clean delta analysis, map the questions directly to the categories established in our baseline survey. The survey must capture: (1) Changes in tool usage frequency and integration levels compared to the starting baseline, (2) Current confidence levels contrasted with our baseline, (3) Measurable progress on previous barriers identified in the first survey, (4) Behavioral changes (such as moving from simple drafting to complex, multi-step workflows), and (5) Future training needs. Formulate the questions so we can easily compare the two data sets to show measurable growth and program attribution."

Modular Assessment: Customize to Match Your Timeline

At Google.org, we found that measuring progress is not based on linear milestones. We actually deployed our progress survey before we kicked off our showcase initiatives. This allowed us to measure the baseline impact of our initial upskilling and establish a clear, data-backed rationale for launching the showcases.

Because every organization's upskilling journey has its own rhythm, we encourage you to adopt a modular approach to this measurement. You have full discretion to deploy this survey when it feels most meaningful and actionable for your team.



Key Markers

What to look for

- **The Confidence Shift:** Look for a decrease in "Not at all confident" and a corresponding increase in "Somewhat" or "Very Confident."
- **The Integration Leap:** Look for movement from "Experimenting" (using AI for isolated tasks) to "Integrated" (using AI for core workflows). This proves your training was relevant to their actual jobs.
- **Barrier Removal:** Review the "Barriers" question. If "Lack of knowledge" drops but "Tools are buggy" rises, congratulations! Your team has moved from uncertainty to informed critique.

Sharing Your Results (Dashboard)

For an impact dashboard, meaningful metrics are those that prove behavior change, not just activity.

- **The Confidence Delta:** The percentage increase in users rating themselves "Confident" or "Very Confident" between the baseline and progress surveys. (e.g., "Confidence rose from 30% to 65%.")
- **The Integration Rate:** The percentage of users who moved from "Experimenting" to "Integrated" in their workflows. This proves the training transferred to real work.
- **Barrier Drop-off:** The specific percentage decrease in "Lack of knowledge" as a reported barrier.
- **Net Promoter Score (NPS) for Training:** Derived from the "Impact" question ("To what extent do you agree..."), showing the net percentage of "Strongly Agree/Agree" vs. "Disagree."

Inside Google.org

When it came time to analyze our own progress, we moved beyond simple spreadsheet formulas. We used Gemini to conduct a delta analysis between our baseline and progress surveys, asking it to identify specific cohorts where confidence had stalled.

Our workflow:

- **Analyze:** We uploaded both datasets and prompted **Gemini** to "Identify the top 3 shifts in user behavior and flag any correlation between specific workshops and increased confidence."
- **Verify (“Human in the loop”):** We asked **Gemini** to "Show your work" by citing the specific survey rows it used for its conclusions. We then spot-checked its findings against the raw data to ensure the "Confidence Delta" was real and not a hallucination.
- **Strategize:** Once verified, we asked **Gemini** to "Draft 3 strategic recommendations for our 2026 roadmap based on these confirmed gaps."
- **Visualize:** Finally, we used **Gemini Canvas** to build a custom dashboard that surfaced only the most meaningful metrics (like the Integration Rate and Confidence Delta) prioritizing the story of growth over a dense wall of charts.

MODULE 3

Enabling the Broader Ecosystem

You’ve built the skills internally, and now it’s time to apply them to your portfolio. The goal of the final module is to **transform your team’s AI fluency into another lever for impact** that they can use in support of their external partners. We’ll guide you on how to use your technical intuition to evaluate proposals with confidence, helping to bolster the social sector’s readiness to innovate and ensuring these powerful tools serve the public good.

CHAPTER 13

The Strategic Imperative

We operate on a simple belief: You cannot effectively fund what you do not understand.

As AI accelerates in the private sector, there's a risk that the social impact sector (often constrained by financial and time resources) will be left behind. This could create a new digital divide where organizations solving society's most critical problems lack access or the ability to build the most powerful tools.

As funders, we can serve as a bridge. Our internal AI programs are not just about our own productivity; they are the prerequisite for deploying resources that ensure the benefits of AI are distributed equitably.



From Funder to Partner

Building your team's AI capacity allows them to help be a matchmaker between the benefits of these technologies and nonprofit challenges. When you understand the capabilities and limitations of these tools, you can:

- **Navigate feasibility** and help grantees distinguish between AI hype and reality.
- **Mitigate risk** and act as a sounding board for ethical considerations.
- **Build empathy** through the understanding of "jagged frontier" tasks, helping set realistic expectations about what AI can augment versus where human judgment is still essential.
- **Facilitate connections** by identifying opportunities for redeployment or partnership, leveraging existing tools and proven use cases rather than starting from scratch.

Inside Google.org

We anchor our own strategy in the core principles outlined in our white paper, Investing in AI for Good (SSIR). Another helpful resource is Patrick J. McGovern Foundation's A Practical Guide: Shaping Strong Tech Proposals.

We advocate for a **Mission-First, Tech-Second** approach, deliberately ensuring AI is not treated as a separate vertical or standalone portfolio. Instead, we view it as a tool that sits alongside legal, policy, and programmatic levers to drive progress.

When evaluating proposals, we use this framework to force a rigorous **problem-first assessment**, ensuring that technology serves the mission goals rather than letting the AI hype drive the strategy.

CHAPTER 14

Evaluating AI Proposals
A Due Diligence Framework

You can use this framework (adapted from Patrick J. McGovern Foundation's AI for Good Diligence Guide and our SSIR white paper, Investing in AI for Good) to evaluate proposals, ensuring the technology is necessary, feasible, and safe. As you connect with nonprofits, some may already be experts in AI and some may be learning alongside you.

1. **Categorize the Application (Operational vs Programmatic)**
   Before applying the diligence framework, determine the primary beneficiary of the AI solution. This distinction determines your path for risk tolerance and evaluation criteria:

- **Back-Office Solutions:** These tools improve nonprofit operations (e.g. drafting grant reports, analyzing donor data).
  - Look for: Efficiency gains, data privacy for donor/employee info, and cost savings that can be reinvested into the mission.
  - Consideration: Don't trap these in overhead. As we will note in the Shifting the Funding Model section, these are essential operational capabilities.
- **Beneficiary-Facing Solutions:** These tools interact directly with the community to assist beneficiaries (e.g. service triage), where AI is used for mission impact rather than just internal tasks.
  - Look for: "Human-in-the-loop" workflows where human judgment engages to remedy any unintended consequences, and rigorous testing to ensure training data is representative.
  - Consideration: The risk of hallucination or bias here can cause direct harm or inequity. Verify that the organization has a plan to monitor risk and explicitly designates someone who is responsible for ethical oversight.

2. **Vet the Problem, Not the Tool**
   Beware of solutions looking for a problem. Some organizations might start with a desire to use a specific AI tool and work backwards to find a use case. If an AI tool improves an outcome, that doesn't automatically mean it is a strong candidate for funding. You must assess if the technical intervention is truly novel and worthwhile.

- Key questions: Did the organization start with a validated user need or did they start with a specific tool in mind? What is the current approach to this problem, and does the AI use provide an outsized benefit?
- What to look for: Evidence of user research or user testimonies that predates the technical solution. If users haven't been engaged yet, the project plan should explicitly account for community inclusion in the design phase and better yet, be co-designed with the community.

3. **Responsible AI & Ethics**
   Ensure the organization is leveraging AI use with accountability, operational responsibility and a “human-in-the-loop” approach:

- Who in the organization is responsible for ethical oversight of the tool or solution? Does that organization have AI principles?
- Is there a designated workflow for high-stakes decisions? Ensure there are safeguards where human judgment engages to identify and remedy unintended consequences.

4. **EU AI Act & Transparency Compliance (Article 50):**
   Grantees operating in Europe, or serving European constituents, must have an explicit plan to label AI-generated content and disclose automated systems. Failure to build these transparency markers directly into the technical architecture poses immediate legal and compliance liabilities.

- Does the proposal involve public-facing chatbots, synthetic media generation, or automated triage systems? If yes, how will they visually disclose AI involvement to users and embed digital watermarks in generated assets?

5. **Maintenance & Sustainability**
   AI solutions require ongoing care, so it’s crucial the organization identifies who will build, maintain and sustain the solution over time.

- Does the organization have a plan for the technical roles required? Aim to have the long-term maintenance team members involved from the beginning to avoid a hand-off gap (where a solution risks being forgotten once external developers leave).
- Is there a plan for ongoing upgrades and operational continuity once funding ends?

6. **Interrogate the Data (Inputs)**
   Data integrity is essential for an AI tool or solution’s utility and safety. It’s key to understand the profile of the data that is training the model as an input.

- Does the organization **own or have the right to utilize the data** required to build the solution?
  Was it collected with consent from the individuals it represents?
- Is the training data representative enough to **avoid bias**? E.g. a machine learning model used for social services triage must be trained on diverse datasets to avoid automating historical exclusions.

7. **Stress-test the Tech (Feasibility)**
   Push for specificity. Look for thorough plans that include the following:

- For Generative AI, **mitigate hallucinations**: how will they prevent the model from making things up? Look for architectures like Retrieval-Augmented Generation (RAG) which grounds answers in proprietary information (like a verified policy resource) rather than relying solely on the training data.
- For Predictive AI, **seek explainability**: if a model is making a recommendation (e.g. credit scoring or health assessment), is it explainable? The organization should articulate how they will communicate to the beneficiary that a machine is making the assessment.

8. **Infrastructure & Cost Realism**
   AI projects risk failure when teams often underestimate the recurring costs of running a model.

- Is the budget realistic about computing costs? Many proposals account for the cost of training a model but not the cost of inference (cost per query when the model is actually being used). The Google Cloud pricing calculator is a helpful tool to estimate costs.
- As user volume scales, inference costs increase. Does the organization have a plan to keep these within budget? Ensure they have mapped out estimated costs for ongoing expenses, e.g. cloud infrastructure and API licenses.

9. **Post-Deployment Evaluation**
   Evaluation doesn't end at launch; it begins there.

- Does the team have a plan to monitor for "data drift" (where incoming data no longer matches training data)? This can cause model degradation over time.
- Evaluate "time-to-value". Is the AI solution actually saving human time, or are teams spending their time overseeing and correcting the AI's errors?
- For Generative AI, look for protocols to score outputs for adherence to factual sources rather than just fluency.

10. **Ecosystem Integration & Open Innovation**
    We believe in collaborative progress over building in silos.

- Has the organization explored adapting or redeploying existing open-source tools instead of building from scratch?
- Will the organization open-source their code, share their training datasets (if ethical), or provide an API so the broader social sector can benefit from this investment? As a principle, Google.org generally requires all grantee tools be open sourced.

Pause for Responsibility

Data is not neutral. It is a reflection of history, and history is often biased.

- **Question:** Don't just ask if the model works. Ask who it works for.
- **Safeguard:** Challenge proposals to demonstrate how their data represents marginalized groups. If the training data excludes a specific population, the resulting tool will likely fail them. We must ensure our capital closes the digital divide rather than automating inequality.

CHAPTER 15

The Funder's Role
Shifting the Funding Model

To enable the ecosystem, we must account for the unique lifecycle of software and AI enablement in how we deploy capital.



Reclassify "Overhead" to “Program Costs”

AI is not a one-time expense; it is an operational capability that requires ongoing funding. Grantmakers must place emphasis on sustainability by considering licenses, compute power, and model maintenance as essential to the project’s success.

Action: Classify technology costs as Program Costs, not Overhead, to ensure sustainability. Ensure proposals include estimates for these expenses.

Fund the "Pivot"

AI development is experimental, unlike traditional program delivery. Traditional, rigid milestone grants often fail here.

Action: Structure grants with flexibility for prototyping and pivoting in the first 12-18 months. Encourage organizations to test ideas with low-fidelity prototypes and ensure budgets are structured to allow resources to be reallocated as their teams learn what works.

Support "Boring" Infrastructure

The most critical part of an AI project is often the least visible. Acknowledge that successful AI projects require unglamorous investments in data cleaning, storage, and cloud infrastructure before any modeling can happen.

Action: Proactively fund the foundational data layer. Recognize that investments in data standards and security certifications are not just administrative expenses: they are prerequisites for scalable, safe AI solutions.

Invest in Technical Human Capital

Successful execution often requires specialized skills that are expensive and can be difficult to recruit in the social sector.

Action: Ensure budgets realistically account for the cost of technical talent. Beyond funding, consider how your organization can provide "Technical Assistance" through employee volunteering or pro-bono consulting to help grantees bridge the talent gap.

Inside Google.org

For nearly 20 years, Google.org has lowered barriers to innovation by providing no-cost and discounted technology through Google for Nonprofits. We are evolving this model: we don't just provide tools, but also fund the overhead of digital transformation by covering the cost of technical onboarding and training:

- Eligible nonprofits can receive **Google Workspace for Nonprofits licenses** for up to 2,000 employees or volunteers. This includes enterprise-grade data and privacy protections for the **Gemini app** and **Gemini Notebook**, which means your organizational data is not used to train Gemini models outside of your domain without permission and is not reviewed by humans. **Gemini in Google Workspace** has also been certified by stringent third-party auditors.
- Google.org also covers the operational overhead of deployment of these tools by offering **no-cost 1:1 technical training support and migration guidance** via Google Meet video conferences.

Organizations can get started at google.com/nonprofits and will be invited to technical training sessions upon approval.

Closing

Building your AI capacity is not a singular training event; it's a fundamental shift in how we achieve social impact. By moving through these modules, you have begun to transition your organization to active, strategic partners who can help lead the social sector through its digital transformation.

This journey requires a dual-purpose mindset. AI skill development can unlock new insights and time for deep work internally, but also ensures we can recognize, evaluate, and support innovation across the ecosystem.
If the private sector adopts AI tools while the social sector does not, a digital divide will emerge. As funders, our own fluency can act as a bridge while we work to ensure the benefits of AI are distributed equitably.



Key Takeaways

As you finalize your roadmap, keep these three major principles at the center of your strategy:

1. **Mission First, Tech Second**
   Don't just ask "What tools should we use?" but rather "What problems do we need to solve?" AI is not a standalone portfolio; it is a tool that should fold into all of your work to solve problems and deliver impact at scale. When identifying use cases look for opportunities where AI doesn't just expedite admin work but actively assists beneficiaries.

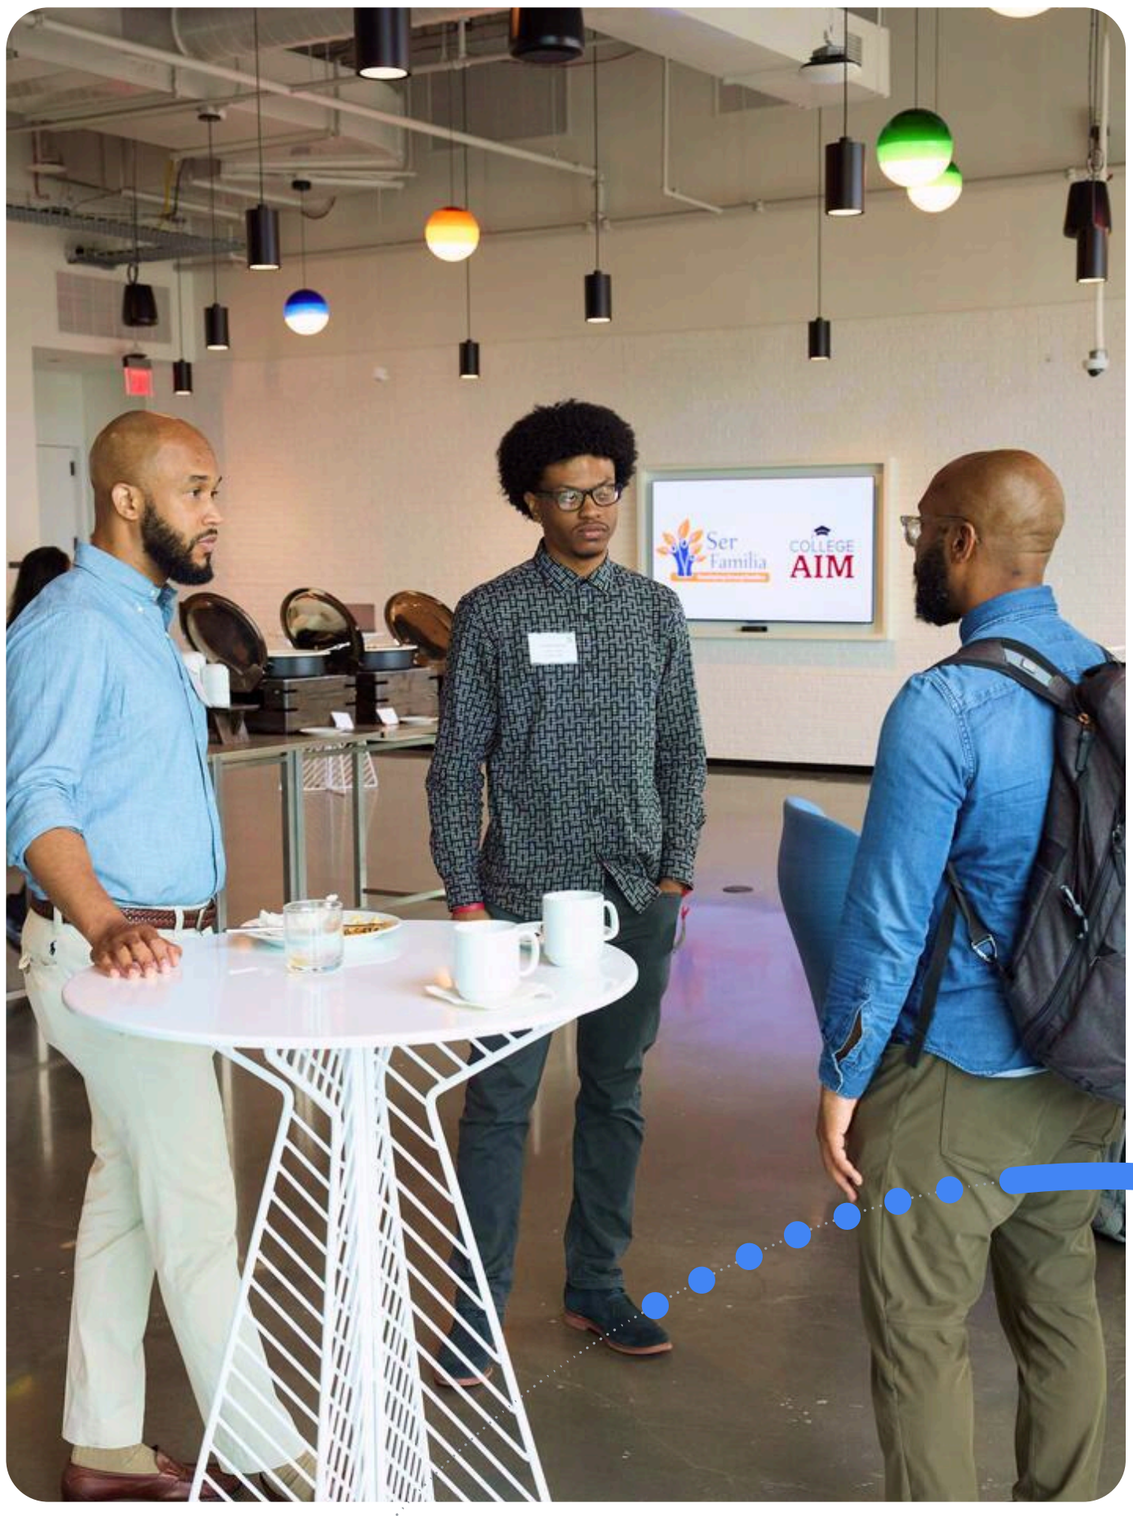
2. **Never Trade Speed for Safety**
   Efficiency should never come at the cost of equity. Establish a "human-in-the-loop" standard where human judgment is non-negotiable for high-stakes decisions. Use the **ACT** framework (**A**sk, **C**heck, **T**ell) to ensure that while AI may inform the process, it never controls the final outcome, maintaining integrity and trust.

3. **Empower the Ecosystem**
   Don't build in a silo. By making progress toward building your own technical intuition, you can become a better partner and help grantees navigate feasibility, navigate risk, and distinguish between hype and reality. Share your learnings and tools openly with the sector to help other funders and grantees harness this technology for the public good.

Inside Google.org

We used a multi-AI-tool approach to research, draft, and refine this playbook. Early in the process we created a **Gemini Gem** (whose capabilities are now mostly available in **Gemini skills** to help customize AI) and added our AI Strategy documents, the results of our workshops and surveys, Google's AI Principles, and our core organizational goals as knowledge sources. The tool served as a continuity engine and ensured that every module remained aligned with our high-level outcomes for this resource. It helped us generate initial outlines and pinpointed specific internal resources to cite.

When it was nearly complete, we used **Gemini Notebook** to pressure test the material and identify any logic gaps. And, since **Gemini Notebook** only draws from the sources uploaded, it was the perfect tool to help compile our **Glossary**.



Resources for Grantees

Reports and White Papers

- The Philanthropic Reset: How Philanthropy Can Lead in the Age of AI (Fast Forward, Google.org)
- AI in Action: Accelerating Progress Towards the Sustainable Development Goals (Google)
- AI in the nonprofit sector: Insights for unlocking the productivity potential of AI for social impact (Google.org)
- The 2025 AI for Humanity Report (Fast Forward)
- The Playbook on AI for Humanity (Fast Forward)
- Realizing the Reset: Turning AI Adoption into Impact (Fast Forward, Google.org)
- The 2026 AI for Humanity Report (Fast Forward)

Google Resources

- Google.org Resources for Funders and Nonprofits
- Google for Nonprofits Prompt Library
- Google for Nonprofits Responsible AI help center article
- Google AI Essentials Specialization (Coursera)

Resources

Resource List by Module

MODULE 1

AI Principles

Google's AI Principles

The foundational ethical charter that governs Google’s own development pipeline, serving as a robust template for drafting your own responsible AI guidelines.

Principles in Practice Workshop Guide

Move from policy to practice in one hour with this ready-to-use agenda that aligns your team’s daily AI habits with your organization’s mission.

AI Ethics: Ethics Inquiry Deck

Adapted from Google DeepMind’s internal research, this tool equips non-technical funders with the right questions to ask to stress-test AI and data projects.

AI Policy: Fast Forward Nonprofit AI Policy Builder

A free and easy tool to help nonprofits craft a custom AI policy.

Data Classification: CISA Traffic Light Protocol (TLP))

Adopt this globally recognized cybersecurity standard to help your team instantly classify documents (from 'Public' to 'Sensitive') without needing to memorize complex data laws.

Bias & Design: Google PAIR (People + AI Research) Guidebook, “Feedback & Controls” section

Leverage this framework to design 'human-in-the-loop' workflows that prevent automation bias and ensure every AI output remains under human oversight.

"Responsible AI" Basics: Google for Nonprofits: Responsible AI Guide

Share this primer with your wider team to establish a shared baseline on the core risks of AI, such as data privacy and hallucination, specifically tailored for the social impact sector.

Starting Point Survey: AI Readiness Survey Template

A ready-to-deploy survey designed to baseline your team’s current AI maturity, uncover hidden barriers, and tailor your AI skill development roadmap to actual needs.

Ecosystem Use Cases: Social Good AI Hub (IDinsight & Google.org)

An AI-powered repository that transforms hundreds of grantee AI use cases into a structured, searchable learning tool for the philanthropic sector.

MODULE 2

Training Design: 90 Minute Workshop Guide

A step-by-step facilitation guide designed to move teams from general awareness to role-specific application through interactive exercises like "Task Audits" and "Prompt-Critique-Refine" relays.

AI Progress Survey: AI Progress Survey template

A follow-up evaluation tool designed to measure the impact of your program by tracking shifts in employee confidence, barrier removal, and workflow integration over time.

MODULE 3

Investment Approach: Investing in AI for Good (SSIR)

This article outlines a "mission-first, tech-second" strategic framework, arguing that funders should view AI not as a separate vertical but as a tool alongside policy and programming to accelerate social impact.

Diligence Guide: A Practical Guide: Shaping Strong Tech Proposals (Patrick J. McGovern Foundation)

This guide equips grantmakers with a comprehensive framework of diligence questions to assess the problem definition, data integrity, and scalability of proposed AI interventions.

Grantmaker Diligence: AI Proposal Assessment Tool (Fast Forward)

This tool can help teams estimate computing costs based on need.

Impact & Evidence: AI Evidence Playbook (J-PAL)

A reference for donors and policymakers mapping concrete pathways for AI in development, emphasizing randomized evaluations and real-world delivery design

Computing Cost Calculator: Google Cloud pricing calculator

This tool can help teams estimate computing costs based on need.

Readiness and Progress Surveys

TEMPLATE

The Funder AI Readiness Survey

Copy this directly into Google Forms or your survey tool of choice

1. **Frequency**

   How often do you use AI tools for your own productivity or for the work that you share with others?

   a. Never
   b. Very Rarely (once per month)
   c. Rarely (twice per month)
   d. Occasionally (once per week)
   e. Frequently (twice per week)
   f. Always (daily)

2. **Integration**

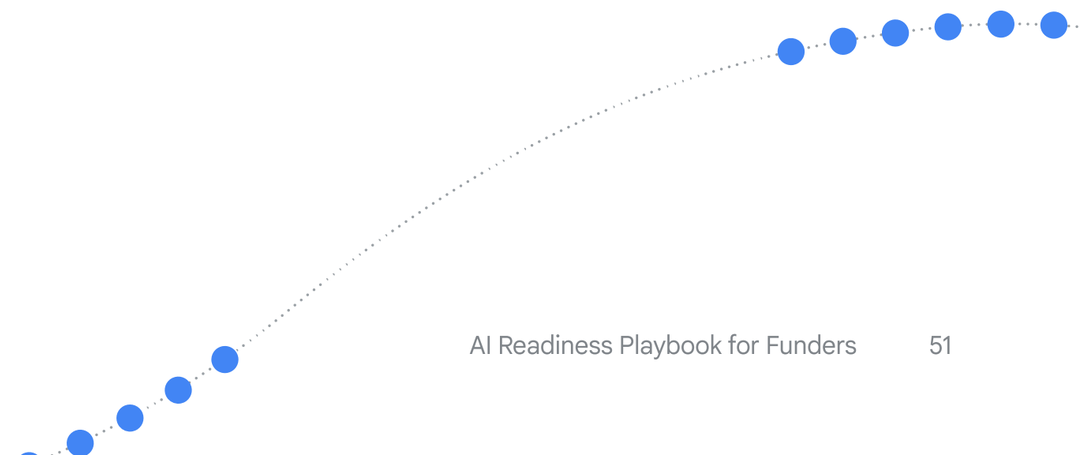
   To what extent have you used AI-powered tools?

   a. Experimenting (Low): I am figuring out how to use them in ONE or SOME of my workflows.
   b. Experimenting (High): I am figuring out how to use them in MANY of my workflows.
   c. Integrated (Low): I have successfully integrated AI into ONE or SOME of my workflows.
   d. Integrated (High): I have successfully integrated AI into MANY of my workflows.

3. **Confidence**

   Overall, how confident are you right now using AI tools for work?

   a. Very confident
   b. Confident
   c. Somewhat confident
   d. Not at all confident



4. Use Cases

How are you currently using AI in your role? (Select all that apply)

a. Grant making & Evaluation: Summarizing proposals, analyzing grantee reports, or landscape analysis.
b. Writing Support: Drafting/refining emails, memos, or project docs.
c. Brainstorming: Idea generation, problem-solving, or strategy development.
d. Meeting Management: Note-taking, summaries, and action items.
e. Data Analysis: Summarizing datasets, identifying trends, or finding information quickly.
f. Accessibility: Generating captions, transcripts, or screen-reader-friendly documents.
g. Operations: Automating repetitive tasks or managing project timelines.
h. Audio/Visuals: Generating images, videos, or slide presentations.
i. Coding: Writing, testing, or debugging code (if applicable).
j. User Support: Answering FAQs or personalizing interactions.
k. Other

5. Tooling

Which tools are you currently using? (Customize this list based on your organization's approved stack)

a. Enterprise Chat: (e.g., Gemini for Workspace, ChatGPT Enterprise)
b. Research Tools: (e.g., Gemini Notebook, Perplexity)
c. Creative Suites: (e.g., Adobe Firefly, Midjourney)
d. Coding Assistants: (e.g., GitHub Copilot)
e. None / I don't know

6. The Barrier Matrix

To what extent have any of the following been barriers?
(Rate as: Not a barrier, Somewhat, Moderate, Extreme)

a. Lack of time to learn/experiment.
b. Lack of knowledge on how to use AI effectively.
c. Unsure how AI can specifically help my role.
d. Concerns about data privacy/security.
e. Concerns about accuracy/reliability (hallucinations).
f. Lack of clear internal guidelines or policies.

7. Learning Goals

What do you hope to gain from our AI skill development program? (Select all that apply)

a. The "Ecosystem" Goal: How to assist grantees/beneficiaries with GenAI solutions. [Crucial for Strategy]

b. The "Ethical" Goal: Understanding bias, privacy, and responsible use.

c. The "Efficiency" Goal: Automating repetitive tasks and saving time.

d. The "Strategic" Goal: Understanding social impact use cases for GenAI.

e. The "Technical" Goal: Prompt engineering (writing effective prompts).

f. The "Analysis" Goal: Summarizing data and identifying trends quickly.

8. Demographic

What represents your primary function?

a. Program / Grantmaking

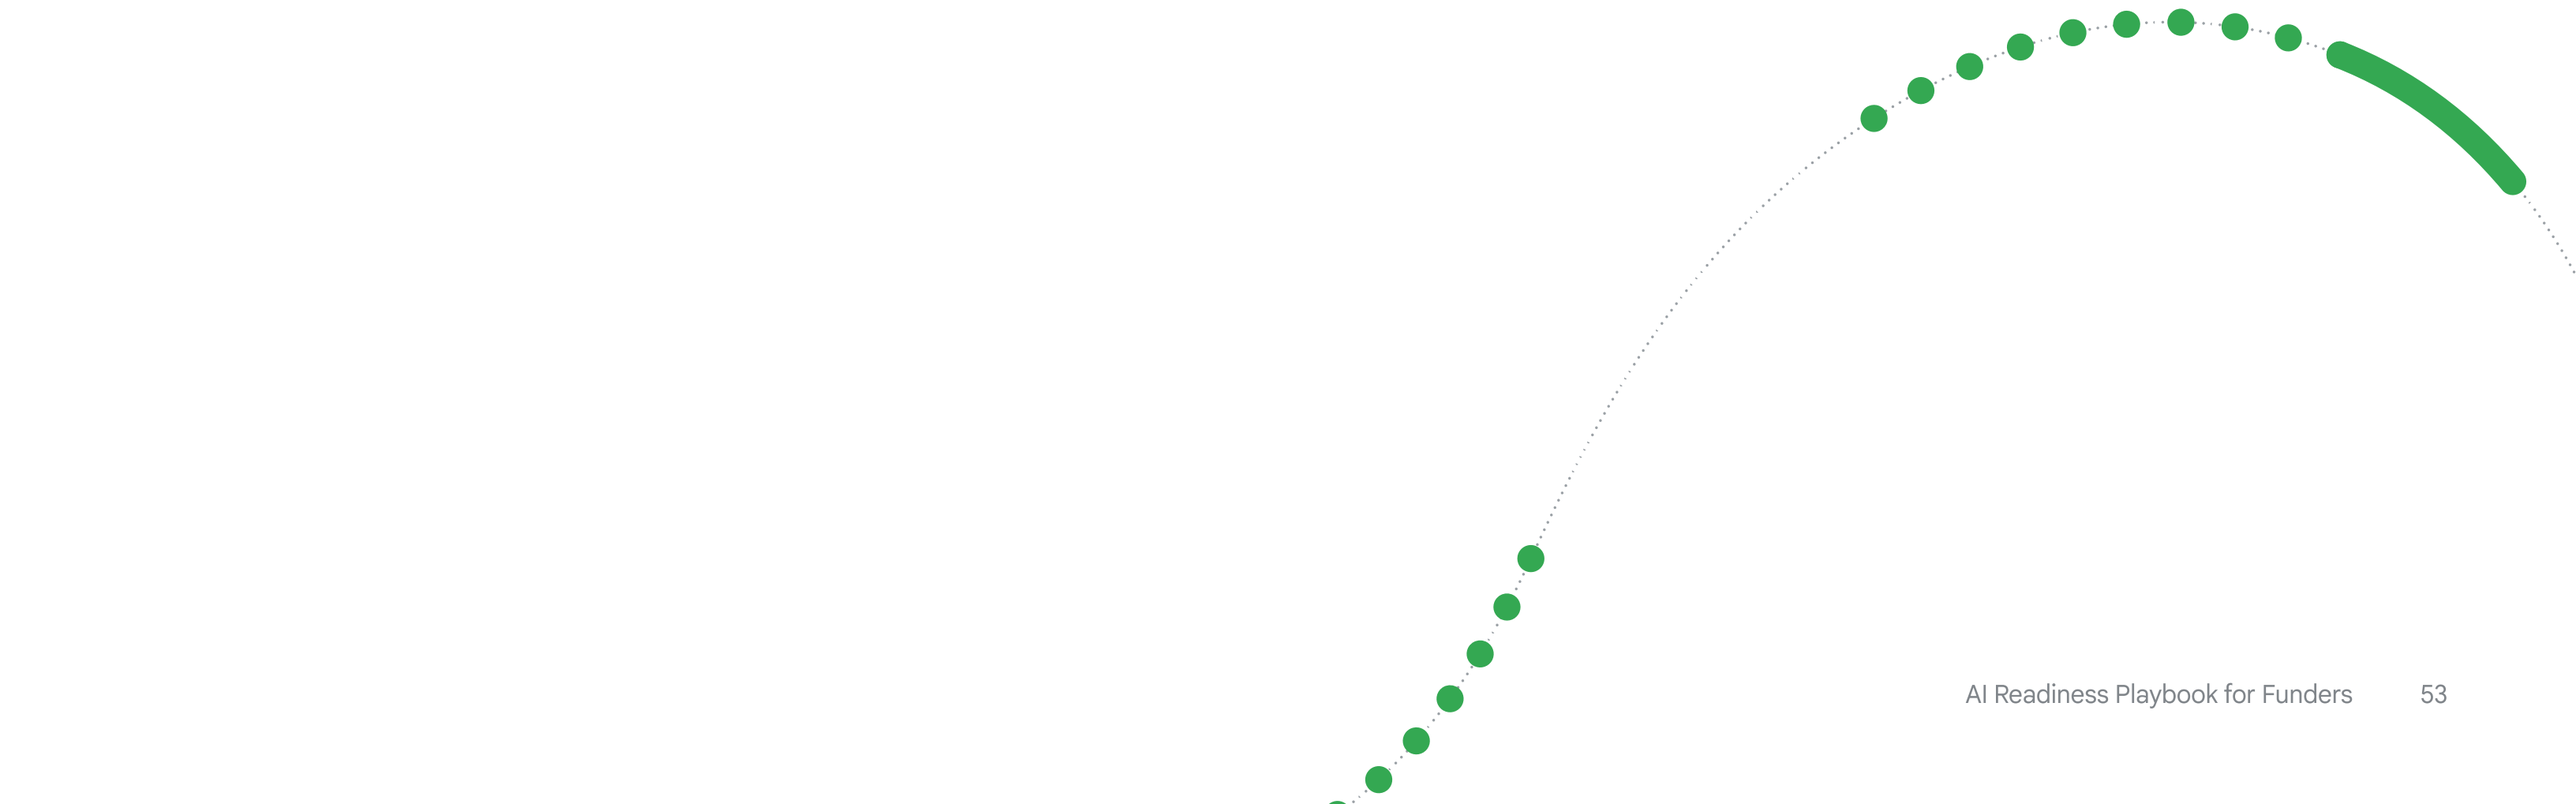
b. Operations / Finance / Legal

c. Communications / Advocacy

d. Strategy / Leadership

e. Research / Evaluation

f. Template: The Funder AI Progress Survey



TEMPLATE

The Funder AI Progress Survey

Copy and paste this into your survey tool. Customize the bracketed text to match your specific tools and program name.

1. Frequency

How often do you use AI tools for your own productivity or for the work that you share with others?

a. Never
b. Very Rarely (once per month)
c. Rarely (twice per month)
d. Occasionally (once per week)
e. Frequently (twice per week)
f. Always (daily)

2. Integration Level

To what extent have you used AI-powered tools for your own productivity or for the work that you share with others?

a. Experimenting (Low): I am figuring out how to use them in ONE or SOME of my workflows.
b. Experimenting (High): I am figuring out how to use them in MANY of my workflows.
c. Integrated (Low): I have successfully integrated AI into ONE or SOME of my workflows.
d. Integrated (High): I have successfully integrated AI into MANY of my workflows.

3. Use Cases

How are you currently using AI in your role? (Select all that apply)

a. Writing Support: Drafting/refining emails, memos, or project docs.
b. Summarization: Summarizing documents to grasp key insights.
c. Meeting Management: Note-taking and meeting summaries.
d. Brainstorming: Creativity, idea generation, and problem-solving.
e. Research: Landscape analysis and information retrieval.
f. Data Analysis: Analyzing datasets and identifying trends.
g. Task Management: Prioritizing tasks or managing timelines.
h. Accessibility: Generating captions, transcripts, or screen-reader-friendly docs.
i. Audio/Visuals: Generating images, videos, or presentations.
j. User Support: Answering FAQs or personalizing stakeholder interactions.
k. Coding: Writing, testing, or debugging code.
l. Automating Repetitive Tasks

4. Tooling

Which AI tool(s) are you currently using in your workflows? (Select all that apply)

a. Enterprise Chat (e.g., Gemini Enterprise, ChatGPT Enterprise, Copilot)

b. Research & Notebook Tools (e.g., Gemini Notebook)

c. Creative Tools (e.g., Image or Video generation tools)

d. Data Analysis Tools

e. Coding Assistants

f. None

5. Confidence

Overall, how confident are you right now using AI tools for work?

a. Very confident

b. Confident

c. Somewhat confident

d. Not at all confident

6. Barriers

To what extent have any of the following been barriers to using AI in your workflows? (Rate as: Not a barrier, Somewhat, Moderate, Extreme)

a. Lack of time to learn/experiment.

b. Lack of training on how to use AI effectively in my role.

c. Unsure how AI can help my specific role.

d. Concerns about data privacy/security.

e. Concerns about accuracy/reliability (hallucinations).

f. Lack of clear internal guidelines or policies.

g. Tools are buggy or inconsistent.



7. Program Attribution

Which of the following resources or efforts have you participated in or used? (Select all that apply)

a. [Your Organization's] AI Strategy Workshops

b. Peer-to-peer sharing sessions

c. Central Resource Hub

d. Team-level AI use case discussions

e. I have not used any of these resources.

8. Impact

To what extent do you agree with this statement? "The AI skill development efforts this quarter have improved my ability to use generative AI tools effectively or impactfully in my work."

a. Strongly agree

b. Agree

c. Neutral

d. Disagree

e. Strongly disagree

9. Behavioral Change

What new or different problems are you using AI tools to solve that you were not solving 3 months ago? (e.g., "I moved from just drafting emails to analyzing qualitative data from grantee reports.")

a. [Open Text Field]

10. Evolution

Which statement best describes the evolution of your AI use over the last 3 months?

a. Integration: My use has expanded significantly; AI is deeply integrated into multiple parts of my daily workflow for both simple and complex tasks.

b. Expansion: I have expanded my use from simple tasks to include more complex, core parts of my role (e.g., strategy, planning).

c. Stasis (High): No change; I was already using AI tools for complex, core parts of my role.

d. Stasis (Low): No change; I am still using AI tools only for simple/experimental tasks (e.g., drafting emails).

e. Non-Use: No change; I am still not using AI tools regularly.

11. Future Needs

What future AI training or support would be most valuable to you? (Select top 3)

a. Advanced Prompt Engineering

b. Using AI for Data Analysis & Insights

c. AI for Project Management & Workflow Automation

d. AI for Research & Synthesis

e. Assessing the Quality & Reliability of AI Content

f. How to Solve Social Impact Challenges Using AI

g. Keeping up with new tools & capabilities

12. Demographics

Which function represents your primary role?

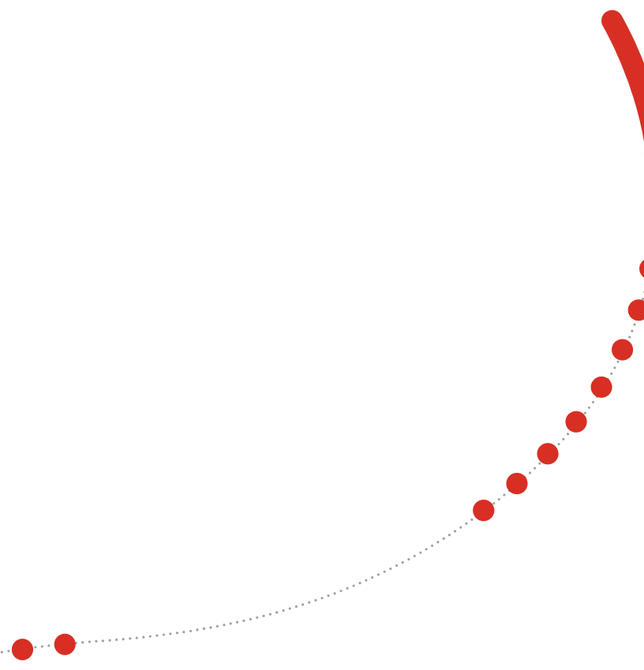
a. Program / Grantmaking

b. Operations / Finance / Legal

c. Communications / Advocacy

d. Strategy / Leadership

e. Research / Evaluation



Principles in Practice Guide

GUIDE

The "Principles in Practice" Workshop

Use this agenda to facilitate a 60-minute team session.

Step 1: Select Your Anchor Principles

Choose 3-4 existing organizational values or adopt standard AI principles (like Google’s "Bold Innovation","Responsible Development", “Collaborative Progress”).

Step 2: The "Workflow Audit"

Ask your team to map these principles to their specific daily tasks using the "If/Then" format.

- **The "Bold" Challenge (Innovation):**
  - Prompt: "Where can AI help us solve a 'Bold' problem for our grantees?"
  - Example Output: "If we use AI to pre-fill grant applications using data from a non-profit's website, then we can save them 10 hours of redundant data entry."
- **The "Responsible" Challenge (Safety):**
  - Prompt: "Where does AI threaten our principle of 'Responsibility'?"
  - Example Output: "If we use AI to draft policy briefs, then a single hallucinated statistic could destroy our credibility. Rule: Every AI statistic must be manually verified."
- **The "Collaborative" Challenge (Ecosystem):**
  - Prompt: "How does our internal use enable 'Collaborative Progress'?"
  - Example Output: "If we figure out a safe way to use AI for legal reviews, then we must open-source that policy to our grantees so they can use it too."

Step 3: Adopt Simple Mental Models (The Toolkit)

Once your team understands the principles, give them simple tools to make decisions in real-time.
We use three core frameworks to keep our work safe without slowing down.

Actionable Checklist

- ☐ **The Anchor:** Have you reviewed the 3 Pillars (Bold, Responsible, Collaborative) with your team?
- ☐ **Tool Audit:** Explicitly list which tools are "Approved" vs. "Banned."
- ☐ **Data Red Lines:** List 3 specific data types that must never touch AI.

Ethics Inquiry Deck

Ethics Inquiry Deck

Use this agenda to facilitate a 60-minute team session.

Origin

Adapted from Google DeepMind (GDM) Ethics Deck of Cards, co-developed by Google DeepMind and The Markkula Center for Applied Ethics

Purpose

To move ethical debates from abstract feelings to concrete analysis. You can use it during annual planning to ensure funding priorities align with your core values, as you review grant proposals to identify risks or mitigation needs, or as a resource for your grantees, empowering them to run internal ethics sessions.

This is not a checklist of right answers. It is a catalog of the right questions. The deck is organized into three strategic categories, allowing you to pull the right card for the specific challenge you are facing.

Cards

Download the Ethics Deck of Cards

1. **Considerations for Assessment (“Net Impact")**
   Use these cards to evaluate the viability of a new proposal or strategy. This category forces a rigorous look at the balance between hope and reality. It moves beyond general optimism to specific accounting of impact.
   - **Benefit Specificity:** Who specifically benefits, and is that benefit distinct from the business case?
   - **Risk Reality:** If the technology works perfectly, could it still cause harm?
   - **Mitigation Feasibility:** Are the proposed safety measures theoretical or actually implementable?

2. **Frameworks ("Values")**
   Use these cards to align on why you are making a decision. Different stakeholders often operate from different moral frameworks without realizing it. One may value "speed" (Utilitarianism) while another values "fairness" (Justice). These cards help teams explicitly name their values.
   - **Justice:** How are benefits and burdens distributed among different communities?
   - **Rights & Duties:** What rules must we follow regardless of the outcome?
   - **The Common Good:** How does this work contribute to the long-term health of society?

3. **Methods of Inquiry ("Process")**
   Use these cards to stress-test your execution plan. These are active exercises designed to break groupthink and challenge assumptions before they become expensive mistakes.
   - **The Pre-Mortem:** A forward-looking exercise where the team imagines a project has failed and works backward to find the cause.
   - **Expanding the Ethical Circle:** A prompt to identify and center the voices of those most marginalized by the technology (often the "non-users").
   - **The "Terrible People" Test:** A red-teaming exercise to anticipate how malicious actors might weaponize the tool.

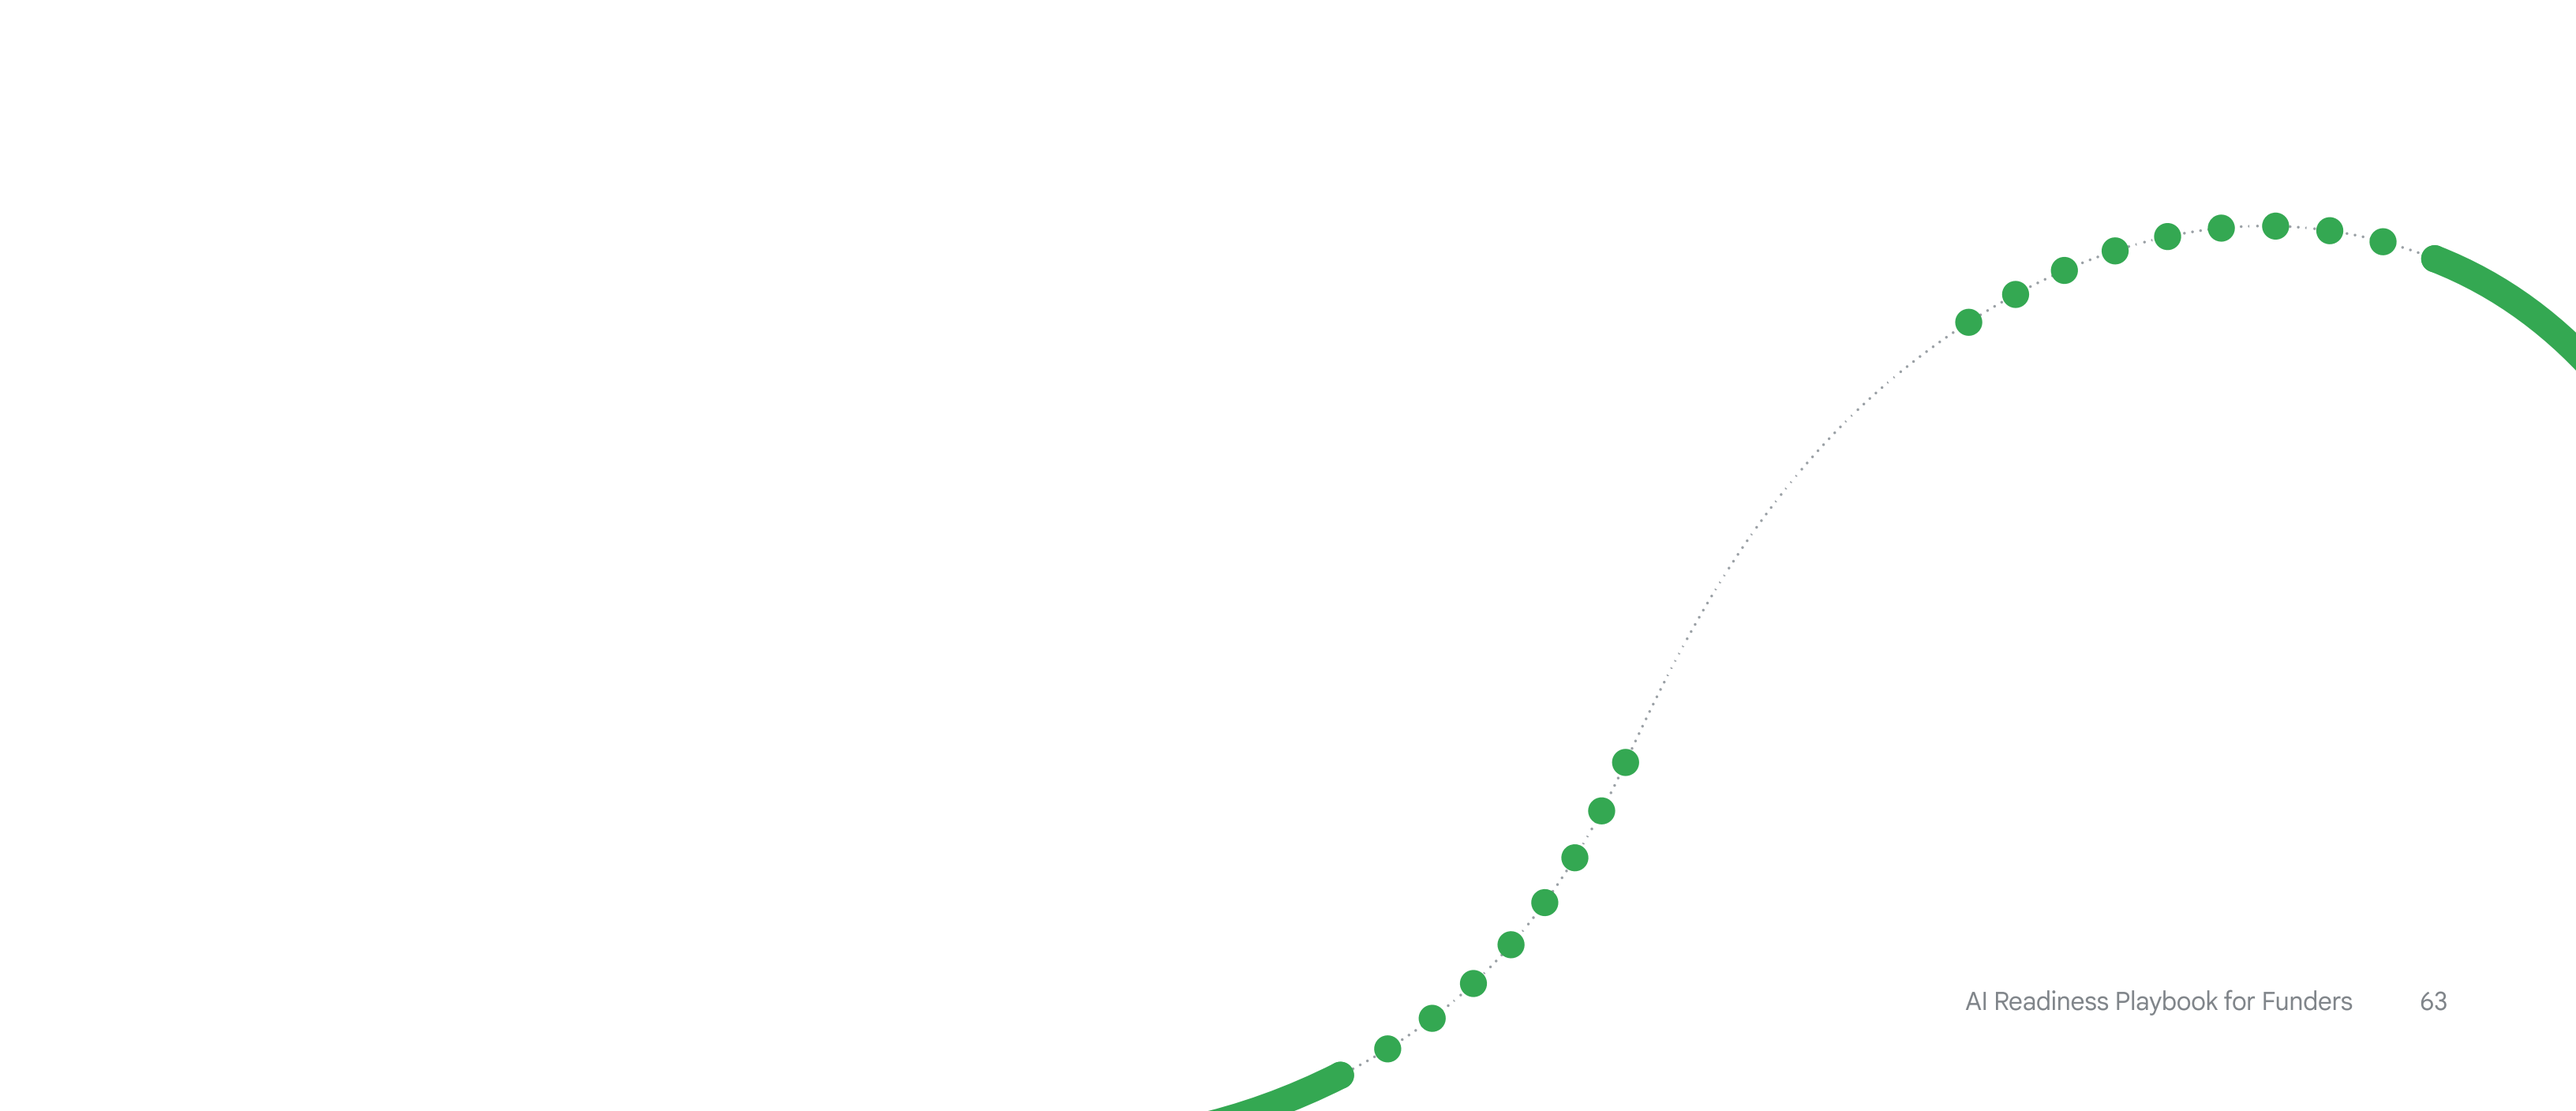
90 Minute Workshop

The 90 Minute Workshop

We recommend a discussion-based session that moves participants from personal utility to strategic application.

Workshop Outcomes

- By the end of this session, participants will be able to analyze diverse opinions on AI to pinpoint the specific professional values or risks (such as privacy, efficiency, or quality) driving those views, to foster an open dialogue around AI and AI tool use.
- Articulate the strategic importance of AI capacity building for the social sector.
- By the end of this session, participants will be able to evaluate their daily tasks to determine the specific role AI should play, such as ideation partner, drafter, or critic, to best enhance and support human decision-making.
- Effectively apply genAI tools to org-specific workflows.



Part 1: Pulse Check, Evolution of AI, Tool Comparison (15 min)

Activity:
“Hopes & Hurdles” pulse check. Instead of jumping right into slides, start with a prompt to gauge the room’s sentiment around AI.

- **Set up:**
  Ask everyone to share (via chat or live): what is one hope they have for AI in their work and what is one hurdle (fear/hesitation) they feel.
- **Prompts:**
  - "I see a lot of ‘time savings’ as a hope, but ‘accuracy’ as a hurdle. Who has experienced a moment where the tool was confidently wrong (hallucinated)? How did that impact your trust?"
  - "If we fast-forward two years, and the private sector has fully adopted these tools but our grantees haven't, what does that inequality look like in practice?"
  - “Which of these hurdles do you think our grantees are also facing right now?”
- **Content:**
  Briefly cover the evolution of AI, framing it through the lens of the digital divide.
  - While the private sector leverages AI for speed, the social sector risks falling behind. However, properly applied, AI is also the most powerful tool we’ve seen for dismantling that very disadvantage. We are developing our AI skills to turn this potential threat into a capacity-building asset that allows grantees to close the efficiency gap rather than letting it grow.
- **Content:**
  General tool comparison. Present a table like this to help teams understand the logic of tool selection.

Tool Category	Conversational Generalist	Grounded Researcher	Autonomous Agent
Google Example Tools	Gemini	Gemini NotebookPro	Gemini Enterprise, Workspace Studio
What it Knows	The entire public internet (broad but can hallucinate)	Only your uploaded, trusted documents (pinpoint accuracy)	Your internal systems (Drive, Salesforce, Calendar)
Primary Goal	Creativity: brainstorming, drafting, summarizing web data	Analysis: fact-finding, deep synthesis, and citation-backed Q&A	Action: executing multi-step tasks across different apps
Best Used For	“Help me draft a warm email to this new partner.”	“Summarize these 50 grantee reports and find common trends.”	“Monitor this folder for new proposals and draft a summary email.”

Part 2: Grantmaker Toolkit (30 min)

Objective
Build technical intuition by evaluating how workflows can leverage AI to ensure accuracy, safety and human accountability.

Content
Introduce three filters that determine what type of an AI tool is appropriate for a specific task:

- **Data Source -** Does the task require public web data (broad and creative) or private internal resources (grounded and secure)?
- **Goal -** Is the goal high creativity (ideas) or high accuracy (verifiable facts)?
- **Mandate -** Establish the "human-in-the-loop" standard immediately. AI is a reasoning engine, not a truth database! It requires human architects to evaluate all outputs, especially those affecting grantee trust or funding decisions.

Activity
Capability categorization. Provide participants with a pre-filled version of this table. Identify a few tasks that are relevant to your team. Have them brainstorm one or two new tasks to add to the table.

Task	Challenge	AI Capability (Best Fit)	Recommended Tool Type
Due diligence	Need to analyze lengthy proposals against a rubric.	Summarization and grounded reasoning: extracts facts from a closed set of data.	**Grounded Researcher:** Gemini Notebook uses only your uploaded sources to prevent hallucinations and provide citations.
Operations management	Must identify top challenges in a region.	Pattern matching: finding themes across vast unstructured text.	**Conversational Generalist:** Gemini and Gemini Deep Research (to identify sources), then Gemini Notebook for low-creativity analysis.
Portfolio analysis	Need to keep the team up to date on multiple grants and initiatives each week.	Send automatic updates to the team on grant statuses.	**Autonomous Agent:** utilize Gemini Enterprise or Workspace Flows to act on your behalf across internal Workspace.
[New Task...]			

Discussion
Use these types of questions to help teams build their own technical intuition to any future tasks.

- **Tool Selection:** "If you had to choose between Gemini and Gemini Notebook for this specific regional analysis, which is the 'right' choice and why? (Focus on the trade-off between high creativity vs. high accuracy)"
- **The Risk of Wrong:** "If the AI makes a subtle error here, what is the impact on our mission or our trust with partners?"
- **Source Integrity:** "Does the AI need to see the whole internet for this, or just a few specific sources we can upload? Why does that distinction matter for safety?"
- **Human Accountability:** "How do we document the human expert's active intervention, shaping: correcting, or approving AI suggestions, to ensure that the final course of action was a human choice, not an automated default?"

Part 3: Mission Application (40 min)

Activity
This section uses a "Choose Your Own Adventure" model for discussion-focused workshops. The facilitator should select one, two or all of these activity paths based on your team culture:

- **Analytical:** Teams that thrive on data, objective evidence, and "why" before "how."
- **Process-heavy:** Teams that value clear frameworks, sequence, and established workflows.
- **Collaborative:** Teams that prioritize group consensus, open dialogue, and shared ideation.

Activity A (Analytical Teams)

Task Audit. Cross-functional groups (Ops, Programs, Comms) ready to break down complex projects into AI-augmentable elements.

1. **Setup**
   Ask participants to select one major project or recurring deliverable from their current to-do list (e.g., a board presentation or a grant cycle report).
2. **Prompts**
   a. “If you haven’t done so already, break your project down into individual elements (e.g., data retrieval, synthesis, drafting, final review).”
   b. “Which elements align with AI's core strengths, e.g. drafting from scratch, transforming content, summarizing resources, or identifying patterns?”
   c. “Which elements require high-stakes judgment, deep human context, or empathy that must remain human-only?”
3. **Activity outputs**
   a. A breakdown of your workflows showing where specific AI capabilities can accelerate or improve the process.
   b. Clear list of where human oversight is required for accountability and ethics.

Activity B (Process-Heavy Teams)

"Aha" Moment Discovery Lab. Best for teams looking to solve specific operational bottlenecks or who are feeling overwhelmed by information-dense workflows.

1. **Setup**
   Ask the group to identify one shared process that is notoriously slow or information-dense (e.g., "Summarizing board notes," "Finding a document in Drive," "Getting a contract signed")
2. **Prompts**
   a. “In this process, where do we experience the most significant bottleneck?”
   b. "Which element of this bottleneck aligns with AI’s core strengths (e.g. drafting a starting point, transforming content into a new format, summarizing vast amounts of information, or analyzing patterns)?”
   c. “If AI supports the information-dense friction points, how does that enhance our team’s ability to perform human-only work, like our ability to deploy capital and support partners?”
   d. "What is our verification step to ensure the AI output is accurate, unbiased and grounded in our values before it informs a final action?”

3. **Activity outputs**
    a. A list of processes that can be augmented with AI
    b. Ideas for which steps in those processes should be “human-only”

Activity C (Collaborative Teams)

"Prompt-Critique-Refine" Relay. Best for building shared confidence with the tools through iterative problem-solving.

1. **Setup**
   Ask a volunteer to share their screen with an LLM tool open (Gemini). Instead of prepared notes, ask the group to suggest a nuanced writing or synthesis task they do weekly (drafting a donor thank-you, summarizing a messy internal brainstorming session, or responding to a common grantee inquiry). Collect 3–5 "messy" bullet points or a short paragraph from the team to act as the live raw data.
2. **Step-by-Step**
    a. Identify a small, annoying task. Have the group describe the "hidden context" (e.g., "This donor is very formal" or "This grantee is over-capacity").
    b. Have the volunteer paste the raw data into the AI with a simple, low-effort instruction like "Summarize this" or "Draft an email."
    c. Facilitate a critique of the output. Compare the result against the human context identified in Step 1.
    d. Invite the team to shout out "tuning" instructions one by one. The volunteer enters each change (e.g., "Make it warmer," "Limit to 3 sentences") to see the output improve in real-time.
    e. Ask the AI to play "Devil's Advocate." Prompt it: "What are three ways a reader might misunderstand this draft?" to find remaining blind spots.
    f. Discuss the real-world implications of using only the AI version vs. the "human-steered" version.
3. **Activity outputs**
    a. A visual Before & After demonstration showing how professional steering transforms a generic draft into a mission-aligned tool.
    b. A crowd-sourced Team Tuning List of 3–5 custom rules or instructions the team can use to steer AI in future tasks.

Content (closing this section for all selected activities):

When we identify which tasks are augmentable and which are best handled with our unique expertise and human-centric approach, we shift from being a “drafter” to an “architect.” With this mindset shift, teams can focus their energy on high-value judgment, empathy and strategy.

Part 4: Closing (5 min)

- **Set up**
  Ask participants to scan next week's calendar for one low-stakes, high-drudgery task to run as a specific experiment (e.g., "Drafting that recurring agenda").
- **Prompt**
  Ask participants to share their tasks with the group.
- **Content (closing)**
  Launch a central hub (simple spreadsheet) for the team to record which prompts worked ("Wins") and which failed ("Lessons"), turning individual trials into collective intelligence.

Team Enablement AI Workshops

Team Enablement AI Workshops

Once grassroots excitement is established through the individual showcase, the next step is to activate your organization by team. To scale AI training across your organization, we recommend a hybrid facilitation model. We split our training across two days, but you could space them out more or group them together:

1. **Central Facilitation (Session 1):** The entire organization gathers for a unified, central training led by an internal AI Champion or facilitator to build tool confidence and set safety guardrails.

2. **Manager-Led Breakouts (Session 2):** Department teams break out into localized sessions facilitated by their respective managers to map operational bottlenecks and prototype custom solutions.

Note: The 90-Minute AI Workshop introduced earlier is an individual-focused training session focused on establishing a shared technical baseline, whereas the Team Enablement AI Workshops outlined here are designed for intact teams to shift the focus from personal productivity to collaborative optimization.

Electing an In-House Facilitator
While external consultants or facilitators could be helpful to run an effective AI enablement session, you do not need to find an AI expert to lead Session 1. In fact, peer-to-peer change management is highly effective because it feels relatable rather than prescriptive.

When selecting an in-house facilitator, focus on these qualities:

- **Learn in Public:** Someone who is comfortable showing vulnerability and learning alongside their peers, managers, or senior leaders.

- **Embrace Troubleshooting:** Someone who doesn't get flustered when an AI tool gives an unexpected answer, but instead embraces curiosity, laughs off the errors, and troubleshoots live with the room.

- **Champion the Mission:** Someone who feels passionately about convening the organization around a shared mission, using the workshop to deepen your team's collective impact in the field.

Principles to Design an Enablement Workshop
We found these core principles helpful when designing our own training:

- **Mandate Light Pre-Work:** To prevent "first-hour friction" (such as forgotten passwords or uncreated accounts), require participants to spend 10–15 minutes logging into approved enterprise tools and completing a basic baseline reading about the tools' capabilities before arriving.

- **Adopt a "Less Talking, More Practicing" Approach:** Keep presentation slides to an absolute minimum. Keep tool introductions brief and focus on step-by-step, handheld labs where participants copy, run, and modify pre-written prompts.

- **Champion the Mission:** Frame the prototyping session as a fun, low-risk experiment. Let teams know that failing to build a perfect tool is completely fine—the primary goal is to build muscle memory around tool discernment and team collaboration.

Dual-Session Architecture

This curriculum is structured to move your organization from foundational tool awareness to active, collaborative prototyping.

Session 1: Foundations & Tool Confidence (~90 minutes)

- Focus this session on introducing your team to four essential tool layers of the modern enterprise stack:
  - **Conversational Generalists** (e.g., Gemini Advanced)**:**
    Focus on drafting, restructuring complex text, and reformatting.
  - **Grounded Researchers** (e.g., Gemini Notebook): Focus on deep document analysis, fact-finding, and citation-backed Q&A using closed, private datasets.
  - **Developer Playgrounds** (e.g., Google AI Studio): Introduce system prompts, temperature settings, and structured prompt engineering without writing code.
  - **Agentic Orchestrators** (e.g. Gemini skills): Walk through how to configure custom agentic workflows (e.g. specialized, persistent persona assistants) using the tool of your choice.
- Close the session by presenting the room with a few hypothetical operational challenges and have them debate which tool handles each task best based on speed, security, accuracy, and tool capabilities. Participants should be able to leave this session having successfully executed prompts across these layers, understanding how they can operate these tools responsibly.

Leverage Managers
While central AI Champions are ideal for teaching general tool mechanics, managers and team or department leads are essential for operational translation. They bring the local context needed to turn abstract AI concepts into daily productivity or creativity gains. The managers' mandate for the sessions should be to:

- **Lead the Gap Audit:** Guide the team through a structured whiteboard brainstorm to map process friction.
  - **Example Tasks:** Upload unstructured grantee narrative reports and impact spreadsheets, leveraging AI's code execution to clean data, normalize metrics, and instantly synthesize portfolio-wide trends; or ingest a grantee's proposal, financial audits, and board bios into a secure workspace, querying the AI to screen the project against funding guidelines and draft a 1-page due diligence memo.
- **Filter for Feasibility:** Push for simple, elegant prompt sequences or custom tools (like Gemini Skills) that can be realistically built in an hour.
- **Establish Human Oversight:** Ensure the team documents clear "human-in-the-loop" checkpoints for every prototype.

To set up your leaders for success, provide them with a light preparation roadmap beforehand:

- **Session Preview:** Share a high-level overview of the two sessions in advance so managers can understand their roles and prepare accordingly.
- **Pre-Evaluation:** Have managers pre-assess their team's tasks leading up to the sessions to identify workflows that fit the capabilities of your approved AI tools (e.g., summarizing recurring grantee check-ins, extracting financial line items from audit PDFs, standardizing board updates, or building an interactive knowledge base of past strategy docs).

Example Prompt: If you have an enterprise AI tool (e.g. Gemini Enterprise), managers can use a prompt to assist with task identification, such as

> "Help me assess the workflows that my team manages, and identify tasks that fit the capabilities of approved AI tools based on process efficiency, automation, and blue-sky opportunities."

Session 2: Hands-on Workflows & Prototyping (~90 minutes)

Respective managers lead their teams in translating Session 1 concepts into functional, localized workflow solutions.

- **Strategic Framing** (10 mins): Manager kicks off the breakout by sharing how they personally use AI, highlighting their support, and setting a creative tone.
- **Macro Team Whiteboard Brainstorm** (20 mins): Managers facilitate a whiteboard session with their teams to map efficiency leaks across three structured streams:
  - Process Efficiency: Pinpointing administrative bottlenecks, meeting clutter, and reporting drag.
  - Workflow Automation: Highlighting repetitive document pipelines, email sequences, or data translations.
  - Blue Sky Opportunities: Auditing legacy bottlenecks traditionally assumed to be unresolvable due to lack of time or resources.
  - Output: A prioritized whiteboard listing 3–5 team-specific operational bottlenecks.
- **Micro Pod Prototyping Sprint** (60 mins): The team splits into self-organized micro pods (2–4 people). Each pod selects exactly one bottleneck from the whiteboard, commits to a single approved enterprise tool, and builds a working, early-stage prototype (such as a custom agent/persona, a structured system prompt, or a shared sheet integration).
- **Output:** A working, early-stage prototype ready to be tested in real-world workflows.

Pause for Responsibility

During the Pod Prototyping phase, require every group to explicitly define their human oversight checkpoint. If a pod prototypes a prompt to summarize grantee proposals or draft partner follow-ups, they must document:

- **Verification**
  How will the human team member verify that the AI's output is accurate, unbiased, and free of hallucinations?
- **Judgment**
  At what specific stage does the human expert apply their unique professional judgment to make the final, high-stakes decision?

Glossary

Glossary

ACT Framework	A mnemonic for responsible AI use defined as: Ask (is this task appropriate?), Check (verify output for bias and accuracy), and Tell (maintain transparency about AI use). This framework ensures that while AI may inform a process, it never controls final high-stakes decisions.
Agentic AI (AI Agents)	AI systems capable of pursuing high-level goals, executing multi-step workflows, and taking actions across multiple applications rather than just responding to single text prompts. **Funder Context:** Moves teams from passive "search-and-draft" prompting to active tool orchestration (e.g., automatically extracting data from a folder of grantee PDFs, structuring metrics into a spreadsheet, and drafting a summary email).
AI-Adaptive	An advanced organizational maturity stage where intact teams deploy cohesive, governed, and collaborative workflows with shared prompt templates, clear data boundaries, and documented human oversight. **Funder Context:** The ultimate goal of internal upskilling. Signals that an organization has moved beyond isolated personal shortcuts to systemic, department-wide capacity building.
AI-Adopted	An early organizational maturity stage where individuals frequently use AI tools for ad-hoc, isolated personal tasks, but without team standardization, shared libraries, or formal governance.
AI Champion	An internal enthusiast who acts as a force multiplier within a specific team. Unlike a typical user who uses AI to finish work faster, a Champion proactively identifies opportunities to apply AI to new tasks and shares prompts and workflows with colleagues to scale learning.
AI Tiger Team	A dedicated, cross-functional cohort of peer mentors embedded across departments to troubleshoot prompts, guide colleagues, and lead grassroots workflow reviews.
Artificial Intelligence (AI)	Computer programs that can do tasks that normally require human intelligence, like thinking or learning.
Autonomous Agent	A system that can pursue a high-level goal and execute multi-step tasks (like sending an email or updating a spreadsheet) autonomously.
Back-Office Solutions (Internal Goals)	AI tools designed to improve nonprofit operations and internal productivity, such as drafting grant reports or analyzing donor data. **Funder Context:** These should be viewed as essential operational capabilities rather than overhead. They are often the entry point for organizations to build capacity and confidence.

Beneficiary-Facing Solutions (Ecosystem Goals)	AI tools that interact directly with the community to assist beneficiaries (e.g., service triage chatbots or personalized learning tools). **Funder Context:** These tools carry a higher risk profile regarding hallucination and bias. They require strict "Human-in-the-Loop" workflows and ethical oversight to prevent direct harm to vulnerable populations.
Biased Data	Data that isn't complete, doesn't accurately represent everyone, or unfairly favors certain groups. **Funder Context:** If an AI model is trained on biased data, it will produce unfair results. This is a critical concern for nonprofits focused on equity, as history is often biased, and data reflects history.
Confidence Gap	A phenomenon observed in survey data where team members report high frequency of AI usage but low confidence in their ability to use it effectively or safely. AI skill development programs should target this gap rather than just access.
Data Drift	A technical risk where incoming real-world data no longer matches the data the model was trained on. **Funder Context:** This causes model performance to degrade over time. Diligence should confirm the grantee has a plan to monitor for this post-deployment.
Digital Divide	The risk that the efficiency, creativity, and scale offered by AI will exist in the private sector but not be fully leveraged by the social sector. Funders educate themselves to act as a bridge to ensure these benefits are distributed equitably.
Ecosystem Void	A strategic gap identified in readiness surveys where teams use AI for internal desk work (administrative tasks) but are not yet applying it to mission-facing work to assist beneficiaries.
Efficiency Plateau	The behavioral gap in an organization's AI adoption curve where individual tool usage reaches a peak but team-scale operational transformation fails to materialize.
Generative AI (GenAI)	A specialized type of AI capable of creating net-new content, including text, images, and other media.
Hallucination	When AI produces information that is not true. **Funder Context:** It is critical for users to double-check AI-generated content for accuracy against reliable sources.
Human-in-the-Loop	A workflow where humans and AI work together to train, use, and check AI models. **Funder Context:** This is a non-negotiable standard for high-stakes decisions (e.g., funding or hiring) to ensure accuracy, accountability, and ethical oversight.

Inference	The process of a model processing a query when it is being used. **Funder Context:** Distinct from the one-time cost of training a model, inference costs are recurring. As user volume scales, inference costs increase, which must be accounted for in budget sustainability.
J-Curve of Adoption	The temporary dip in team productivity and confidence that occurs when an organization transitions to new workflows or systems before long-term efficiency gains are realized.
Jagged Frontier	The concept that AI capabilities are uneven and unpredictable; AI may excel at some hard tasks (like synthesizing 500 pages of text) while failing at easy ones (like basic math), requiring humans to test every specific use case.
Large Language Model (LLM)	An AI model trained on vast amounts of text to understand, synthesize, and generate human-like responses to prompts.
Overhead vs. Program Costs	A funding classification distinction. **Funder Context:** Funders are encouraged to reclassify technology costs (licenses, compute power, maintenance) as "Program Costs" rather than "Overhead" to ensure the sustainability of AI solutions.
Prompt	The text you give to a conversational AI tool to tell it what to do or what kind of output you want. Learning to write effective prompts is known as "prompt engineering".
Responsible AI	The practice of developing and using AI in a way that is ethical, benefits people and society, and avoids harm. This includes navigating risks such as fairness, bias, accuracy, privacy, and security.
Retrieval-Augmented Generation (RAG)	A technical architecture used for source-grounding that allows AI to retrieve facts from private, trusted files before generating an answer. **Funder Context:** This significantly reduces the risk of hallucinations compared to using open internet models.
Shadow AI (Tool Fragmentation)	The unapproved use of consumer-grade, unvetted AI tools by employees for work tasks, often using personal logins and exposing sensitive organizational data to public training models.
Source-Grounding	Restricting an AI tool's knowledge base to specific, trusted documents provided by the user to minimize hallucinations and ensure data integrity.
Transparency	The principle that one should be able to understand how an AI tool works and why it produced a certain result. **Funder Context:** This includes disclosing when AI has been used to draft content or analyze data to maintain trust with grantees and communities.