

The
People
Centered

PLAY BOOK

Dalberg Data Insights

AI

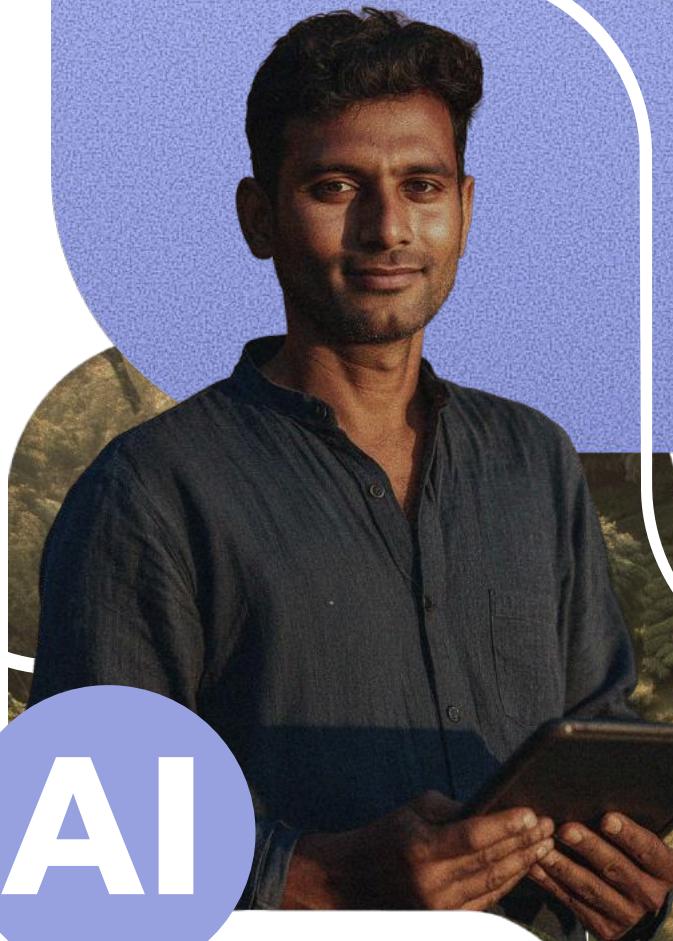


TABLE OF CONTENTS



PAGE		PAGE	
01	EXECUTIVE SUMMARY	13	DEFINE
03	THE 6 PHASES	14	Overview
05	GLOSSARY	15	Assess whether AI is the right solution
06	DISCOVER	16	Translate insights on workflow challenges to AI use cases
07	Overview	17	Assess the Feasibility & Impact of AI Use cases
	HOW TO —————	18	Prioritize AI Use cases
08	Assess your AI readiness	19	DEFINE TEMPLATES
09	Discover a user persona and their workflows	20	DESIGN
11	Define the workflow challenge and the value driver(s)	21	Overview
12	DISCOVER TEMPLATES	22	HOW TO —————
		23	Assess your data needs
			Design user flows and mockups

TABLE OF CONTENTS

PAGE		PAGE	
24	Design scalable & context-aware AI architectures	36	PILOT
25	Evaluate open-source vs proprietary AI tools	37	Overview
26	Assess vendor readiness for Responsible AI	38	HOW TO
27	Define quality and ethical guardrails early	39	Develop a pilot design that sets the right scope & outcomes
28	DESIGN TEMPLATES	39	Monitor performance for iterative improvement
29	DEVELOP	41	SCALE
30	Overview	42	Overview
	HOW TO	43	Key Considerations
31	Understand the technical skills requirements & development process	44	AFTERWORD
32	Select a generative AI model		
33	Build a prototype		
34	Conduct User Acceptance Testing (UAT)		
35	DEVELOP TEMPLATES		



EXECUTIVE SUMMARY

An AI adoption guide without the overwhelm.

This Capsule distills the essentials: core concepts, key frameworks, and practical steps designed for teams working in challenging contexts. Use it to get oriented fast and move from idea to impact with confidence.

Dalberg Design

Developed in partnership with Dalberg Design, whose human-centered design approach helped ensure that the playbook remained rooted in the lived experiences of the people and organizations it aims to serve.

We have seen firsthand how AI solutions can fail when they are not fully localized and contextualized to the users they aim to serve: their workflows, tools, needs, and realities. We have also seen what it takes for AI to create real value, even in challenging settings. These realities do not mean that AI cannot work; they just mean it must be designed differently.

The People-Centred AI Playbook was born out of Dalberg Data Insights' experience working hand in hand with NGOs, social enterprises, governments, and other mission-driven organizations. This work was strengthened through close collaboration with Dalberg Design, whose human-centered design expertise shaped the foundational principles that guide how teams understand users, frame problems, and translate insights into practical design choices. Across health, agriculture, education, financial inclusion, and beyond, we have partnered with teams to design, develop, and adopt AI applications. The message has been consistent: organizations need practical support to move from theory to action, and they do not want to reinvent the wheel.

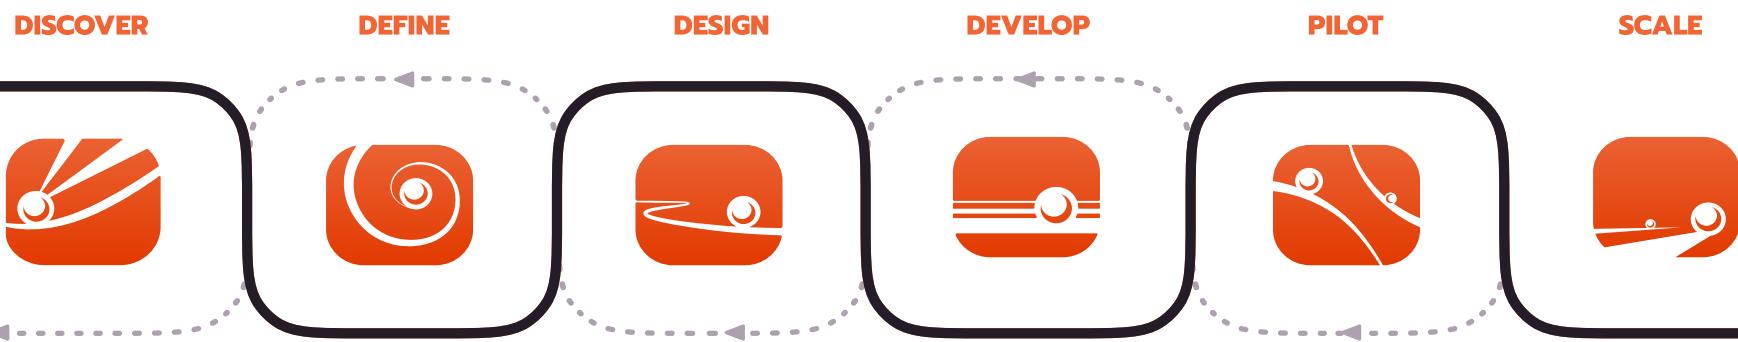
This people-centered playbook is our response. It is designed to give you the essentials of AI adoption in a quick, accessible format, with core concepts, foundational steps, and key frameworks without overwhelming detail. Built for impact-driven teams in low-resource or non-tech-first contexts, it answers the most common questions and provides enough context to get oriented fast. The playbook focuses on helping you get started, from discovering the right problem to designing responsible solutions and launching initial pilots. While we touch on scaling considerations, our deeper emphasis is on laying the right foundation so that when you are ready to scale, you do so responsibly and sustainably. Use this when you need quick answers or want to introduce AI implementation concepts to your team.

What you are reading now is just the beginning. As AI continues to advance and as we learn from teams putting these approaches into practice, this Playbook will expand with it. We do not expect this to be a silver bullet solution, nor do we have all the answers. That is why we are releasing this version early, to invite you to build along with us, share feedback, and help refine and expand this resource as your AI adoption journey unfolds. Your insights and experiences will directly shape what this becomes.



***Throughout this playbook, we use healthcare examples to illustrate key concepts and approaches. However, the frameworks, methods, and principles apply across sectors, including agriculture, education, financial inclusion, and beyond.**

THE 6 PHASES



THE 6 PHASES

Six Phases | Six Questions

PHASE 1 DISCOVER

Is this problem worth solving with AI?

BIG PICTURE

Ground your ambitions in real user needs, organizational realities, and workflow challenges before considering any technology.

OUTCOME

A clear, evidence-based problem statement rooted in what users actually need and what will drive value in your context.

PHASE 2 DEFINE

Which opportunity should we pursue first?

BIG PICTURE

Translate challenges into well-defined AI use cases, assess whether AI is even the right tool for the job, and prioritize which opportunities to pursue based on feasibility, impact, and organizational readiness.

OUTCOME

A shortlist of high-potential AI use cases, ranked by feasibility and impact, with clear criteria for why each one matters and which deserves your resources first.

PHASE 4 DEVELOP

Will this actually work for our users?

BIG PICTURE

Build and test early versions to validate core assumptions with real users, data, and constraints before committing significant resources.

OUTCOME

A tested prototype that meets user needs, passes acceptance checks, and is ready for real-world piloting.

PHASE 3 DESIGN

How do we build this responsibly?

BIG PICTURE

Specify what to build and how: data needs, user flows, architecture, ethical guardrails, and build-vs-buy decisions that fit your reality.

OUTCOME

A solution design that is ethical, contextual, and ready for development, with clear specifications and guardrails that protect users.

PHASE 5 PILOT

Does this work in the real world?

BIG PICTURE

Run a real-world pilot to gather evidence, refine the solution based on actual use, and build trust with users and stakeholders.

OUTCOME

Validated learning about what works in practice, what needs adjustment, and whether the solution is ready to scale.

PHASE 6 SCALE

How do we make this sustainable?

BIG PICTURE

Institutionalize the solution for long-term sustainability and adaptation without compromising quality, trust, or user-centeredness.

OUTCOME

A scale plan covering institutionalization, people and processes, technical maintenance, continuous improvement, and evaluation frameworks.



CROSS CUTTING ENABLERS

PEOPLE

Success depends on people - building trust, aligning leadership, and equipping teams with the skills and confidence to use AI responsibly. Whether you are mapping user personas, designing guardrails, or scaling a solution, people must remain at the center, engaged, trained, and supported for adoption

EQUITY & INCLUSION

Equity and inclusion ensure that AI systems work for everyone, especially those who are underserved. This requires attention to who is represented in your data, who participates in testing, and who faces barriers such as limited connectivity, literacy, language, or access to devices. Across all phases, teams should look for decisions that may advantage some groups over others and adapt solutions to stay accessible and fair. Keeping inclusion central helps prevent unintended harms and ensures AI delivers value across different needs and contexts.

DATA GOVERNANCE

AI is only as strong as the data it relies on. Effective governance, covering data quality, access, privacy, security, and compliance, runs through all phases. From defining workflow challenges to scaling solutions, strong data governance ensures AI systems are ethical, reliable, and contextually appropriate

GLOSSARY

AI Architecture	The overall design of an AI system, including how components such as data sources, models, and interfaces fit together	People Readiness	The extent to which intended users, staff, and partners are willing, skilled, and motivated to adopt and sustain an AI solution.
AI Readiness	How prepared your organization is to adopt AI, considering factors like data, skills, tools, and leadership support	Problem Statement	A clear, concise description of the specific problem you aim to solve, framed from the perspective of those experiencing it
AI suitability	An assessment of whether AI is the right tool for solving the identified problem, compared to other simpler approaches	Proprietary	AI tools owned and maintained by a company, with usage subject to licensing fees or restrictions
AI Use-Case	A specific way AI could be applied to solve a problem or improve a process in your context	Prototype	An early version of a solution used to test ideas and gather feedback before full development
Contextual adaptation	Adjusting a solution to work well in new settings, cultures, or environments	Quality Assurance	Ongoing checks and processes to ensure the solution meets agreed standards and functions correctly
Diagnostic Tools	Structured assessment instruments or frameworks used to evaluate an organization's current state and readiness for AI adoption	Responsible AI	Designing and using AI in ways that are ethical, inclusive, and aligned with human rights
Ethical Safeguards	Measures built into AI design and deployment to prevent harm, protect rights, and promote fairness	Tech Stack	The combination of software, frameworks, and tools used to build and run a technology solution
Feasibility	How practical it is to implement a solution given the data availability & quality, technical infrastructure, management costs, ethical considerations	UAT (User acceptance testing)	The process of testing a solution with end users to confirm it meets their needs and works as intended
Impact (Value addition)	The positive change a solution creates for the people it serves or the system it operates in	User Persona	A simple, fictional profile that represents a typical user including their needs, challenges, and daily workflows
Institutionalization	Embedding a solution into an organization's regular processes, policies, and structures so it becomes part of "how things are done."	Value Driver	The main benefit a solution delivers, such as saving time, reducing costs, or improving quality of service
MVP (Minimum viable product)	The simplest version of a product that can be released to start generating feedback and learning	Workflow	The series of steps a person or team takes to complete a task or deliver a service
Open Source	AI tools made freely available for anyone to use, modify, and distribute, often with community-driven development	Workflow Challenge	A difficulty or bottleneck in a process that stops users from achieving their goals effectively



PHASE 1
DISCOVER

The People Centered AI Playbook

PHASE 1

DISCOVER OVERVIEW



ABOUT THIS PHASE

The Discover phase helps you lay a strong and contextually relevant foundation for your AI journey by understanding where your organization stands, who your users are, and what real challenges they face in their daily work.



This is where you clarify what problem you're solving, for whom, and why it matters before jumping into solutions.

HOW TO LOCALIZE THIS STEP

This is the phase where your users' perspectives will shape the direction of your AI journey. As you engage with them, focus on making this process feel relevant and respectful to their realities. Think about how they naturally communicate, what influences how openly they share feedback, and whether the way you talk about AI connects to their actual experiences. The goal is to meet users where they are, not where you think they should be.

END GOAL

By the end of this Discover phase, you will have a clear, evidence-based problem statement rooted in what users actually need and aligned with what will drive real value in your context.

You'll leave with

1 AI readiness assessment

A snapshot of your current strengths and priority gaps

2 User personas and workflows

Documented personas with key pain points, behaviors, and workflows

3 Workflow challenges and value drivers

Clear bottlenecks and the value solving them creates

4 Problem statement

A concise, evidence-based statement rooted in user needs

STEP 1

ASSESS YOUR AI READINESS

Assessing your AI readiness means taking stock of your organization's current ability to explore, develop, pilot, and scale AI solutions responsibly. This involves looking at your strategy, people, data, and technology to understand where you have strengths and where gaps exist. In resource-constrained settings, this step is especially important because solutions must account for limited bandwidth, varying device capacity, diverse digital fluency levels, and infrastructure realities. The assessment helps you understand your starting point so you can plan effectively for what comes next.

HOW TO DO IT

SELECT AN ASSESSMENT TOOL

1 **Select an assessment tool** that aligns with your organization's focus areas and capacity. Choose one that addresses the dimensions most relevant to your context: data maturity, ethical considerations, technical infrastructure, or organizational strategy.

GATHER DIFFERENT PERSPECTIVES

2 **Gather diverse perspectives** by involving leadership, program staff, technical teams, and functional teams in completing the assessment. Different roles surface different realities. Leadership may see strengths where frontline staff see challenges. Technical teams may spot infrastructure gaps others miss.

DOCUMENT YOUR RESULTS

3 **Document your results** to create a snapshot of current strengths and priority gaps. Be honest about where you stand today. This baseline provides clarity, not judgment. It shows what you can build on and where you need support.



4

IDENTIFY THE NEXT STEPS

Identify next steps based on where you need to build capacity, seek partnerships, or adjust your approach. Use your assessment to make concrete decisions about training needs, partnership opportunities, or timeline adjustments based on your constraints.

*ASSESSMENT TOOLS

Microsoft AI Readiness Assessment

Business vision, strategy, governance, data, people, infrastructure, model management

GSMA AI Ethics Playbook + Self-Assessment

Fairness, transparency, privacy, security, accountability

Data.org's Data Maturity Assessment (DMA)

Purpose, Practice, People with resource recommendations

CISCO AI Readiness Assessment

Strategy, infrastructure, data, governance, talent, culture

Dalberg's DART

Built for social impact in low-resource settings—strategy, data, capacity, ethics, culture

STEP 2

DISCOVER A USER PERSONA & THEIR WORKFLOW

Understanding the users you are solving for is essential before designing any AI solution. Get to know who you are designing for, what drives their daily work, and where friction slows them down. User personas are archetypes of target user groups built from research and real-world insights. They capture users' context, behavioral drivers, data usage patterns, needs, and motivations. Workflows map the sequence of tasks, decisions, and interactions a user performs to get work done. Mapping workflows helps you surface pain points, inefficiencies, or repetitive steps that AI could automate or improve.

HOW TO DO IT

1 LISTEN & OBSERVE

Talk to target users through interviews, focus groups, or shadowing. Listen for repeated frustrations, workarounds, and what users wish they could do better. The goal is to understand their reality, not validate your assumptions.

2 SPOT PATTERNS

Group patterns that emerge from your research. Who shares similar tasks, barriers, or motivations? These clusters will become the personas you design for. Look for behaviors and contexts, not just demographics.

3 BUILD THE PERSONA PROFILE

Document your results to create a snapshot of current strengths and priority gaps. Be honest about where you stand today. This baseline provides clarity, not judgment. It shows what you can build on and where you need support.



4 MAP THE WORKFLOWS

For each persona, lay out the main steps, touchpoints, and decisions in their daily work. Highlight where things break or slow them down. This workflow map surfaces the pain points, inefficiencies, or repetitive steps that AI could help automate or improve.

5 REFLECT ON ADOPTION

Consider how these users' realities might affect adoption. Will they be excited or hesitant? What will they need to build trust, capability, and motivation? Different personas have different levels of digital confidence, device access, and willingness to adopt new tools.

6 DEFINE THE PRINCIPLES FOR ADOPTION

Use what you have learned about your users' needs, motivations, and context to set clear design principles. These principles help your team stay aligned on how to build AI solutions that people will use. For example, prioritize ease of use, respect local norms, or show immediate value.

Example

PERSONA PROFILE



THE SKEPTICS

“I receive as many as 20 emergency case calls from my CHWs every month. It becomes very difficult for me to cater to these as much as I try.”

An illustrative user persona profile for the community health systems context

Skeptics are committed supervisors who, despite their dedication, are overwhelmed by their heavy workload. Juggling multiple full-time responsibilities alongside CHW supervision, they often struggle to provide adequate support or make timely decisions for their CHWs. The constant struggle to keep up with their workload leaves them with little opportunity for proactive engagement, and they seldom express concerns to their managers, which only adds to their challenges.

Behavioral drivers: They are driven to embrace solutions that reduce their workload.

Data usage: They monitor CHW performance and submit monthly reports using Jamii ni Afya, but only fulfill the minimum requirements of their role.

Needs: They need work optimization to lighten their workload and improve the efficiency of the support they provide to CHWs .

STEP 3

DEFINE THE WORKFLOW CHALLENGE & VALUE DRIVERS

This step is about going deeper than general frustrations to pinpoint the specific bottleneck in a user's workflow and define how solving it creates meaningful value. A workflow challenge is a friction point within a user's specific task or set of tasks that is observable, user-experienced, and grounded in behavior. A value driver is the outcome or benefit that solving this challenge would unlock for users, the system, or both. This could be improved efficiency, accuracy, timeliness, equity, cost savings, or better outcomes.

HOW TO DO IT

1

CHOOSE ONE TARGET FLOW & USER

Start with a concrete workflow performed regularly by a key persona. This keeps the challenge grounded in daily operations, not abstract goals. Focus on a workflow where inefficiencies or breakdowns are visible and frequent.

2

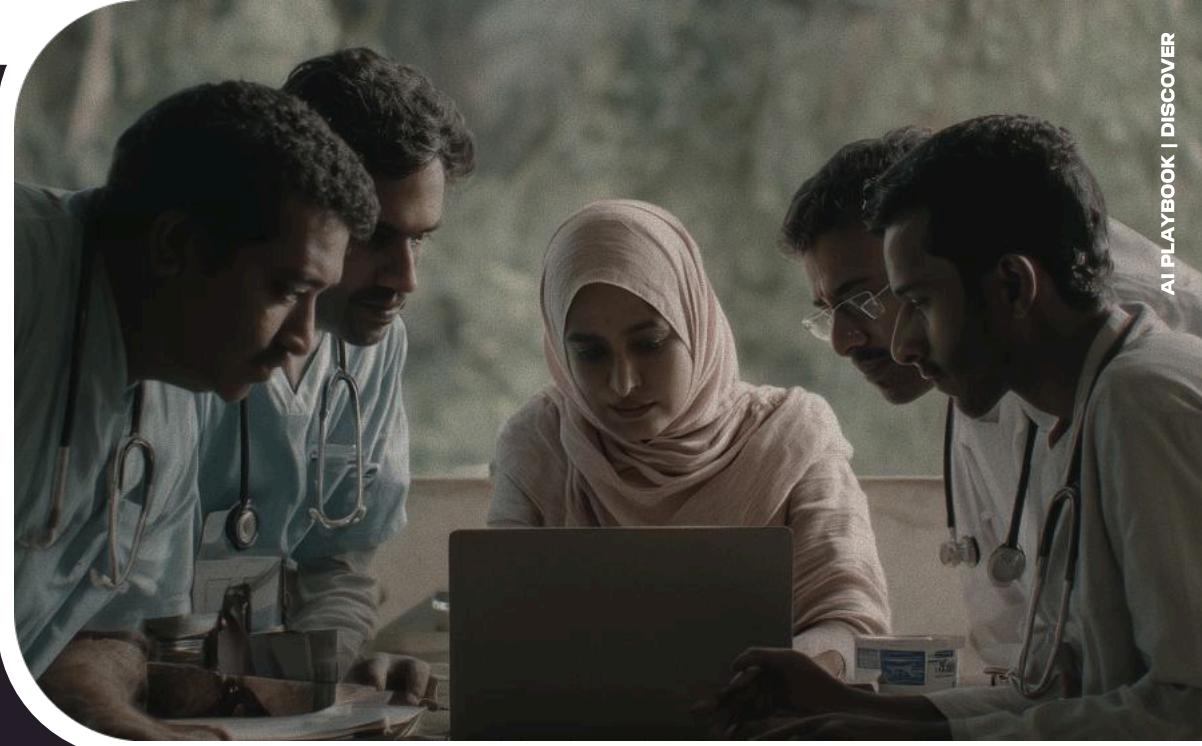
BRING DIVERSE PERSPECTIVES

Facilitate a working session that brings together different voices of the ecosystem. Each sees the problem differently. Prompt each to share what is working, what is not, and where current tools fall short. Capture insights visibly in one place so everyone can see and refine them together.

3

SYNTHESIZE A CLEAR PROBLEM STATEMENT

Use insights from the working session to **craft a short, clear, and specific statement**. Use behavioral language to clearly state what is hard, where the breakdown occurs, and what the impact is. Avoid vague or aspirational language. Be concrete about the problem experienced, when it occurs, who it affects, what users feel, what workarounds exist, and what the costs of these workarounds are.



4

DEFINE THE VALUE DRIVERS

If this challenge is solved, what value will it create? Be explicit about the benefits for the user, the system, or the community. Examples include improved efficiency, better data quality, reduced avoidable complications, timely care, cost savings, or better outcomes.

5

DOCUMENT ANY ASSUMPTIONS OR UNKNOWNs

Note what you still do not know. You may need more evidence or validation before moving to ideation. Acknowledging gaps helps you prioritize what research or validation to do next.

6

Be specific, not aspirational. Vague problem statements lead to vague solutions. If your problem could describe a hundred different situations, it is too broad. Ground it in observable reality: instead of "improve healthcare delivery," say "CHW supervisors spend 5+ hours weekly aggregating fragmented data, delaying care decisions." The test is simple: if someone living the problem would immediately recognize it, you are on the right track. Specificity forces clarity early and makes success measurable.



DISCOVER TEMPLATES

Templates to structure your insights from the Discover phase. These tools guide you through creating user personas and problem statements, with quick tips to help you get started.

USER PERSONA CANVAS

A structured document to capture key details about your target users: who they are, what they do, their challenges, and their context.

TIP

Start with your most representative user group. Use interviews, shadowing, or workshops to understand their exact steps, decisions, and tools.

PROBLEM STATEMENT FRAMING CANVAS

A worksheet to translate user insights into a clear problem statement that identifies who is affected, what challenge they face, and what value solving it creates.

TIP

Keep it short and specific. Aim for one or two sentences reflecting evidence from users, not assumptions. Reference your workflow challenges and value drivers.

Acknowledgment

Dalberg Design

Templates developed in collaboration with Dalberg Design.

SCAN QR CODE



ACCESS TEMPLATES



PHASE 2
DEFINE

PHASE 2

DEFINE OVERVIEW

ABOUT THIS PHASE

The Define phase helps you move from insight to intention, turning observed workflow challenges into concrete AI use cases. This is where you assess whether AI is the right fit or if alternatives would work just as well, apply ethical considerations from the start, and narrow in on the highest-impact, most feasible opportunities.



Rather than jumping straight to building solutions, this phase enables you to frame strong use cases, weigh trade-offs, and define what success looks like.

HOW TO LOCALIZE THIS STEP

This is the phase where you begin to frame concrete opportunities for AI, which means your choices must reflect the realities of your context. As you prepare to define and prioritize use cases, use languages and terminology your stakeholders are comfortable with, account for equity issues that may affect whose problems are prioritized, check policy requirements that could enable or constrain a use case, and consider practical constraints like data availability or connectivity that shape what AI solutions are feasible. The goal is to frame opportunities that are both ambitious and achievable.

END GOAL

By the end of this Define phase, you will have a shortlist of high-potential AI use cases, prioritized by feasibility and impact, ready to move into design.

You'll leave with

1 AI suitability assessment

A clear view of where AI is (and is not) a good fit

2 Defined use cases

Draft descriptions linked to specific pain points and value drivers

3 Feasibility and impact scores

Evaluation across key dimensions with trade-offs identified

4 Prioritized shortlist

Ranked use cases ready to advance into Design

STEP 1

ASSESS WHETHER AI IS THE RIGHT SOLUTION

Not all challenges require or benefit from AI. Start by asking: is AI even the right tool for this challenge? The answer could be that the challenge is better addressed through workflow redesign, basic digital tools, training, or policy changes. Assessing AI suitability means testing whether a problem truly needs AI versus when a simpler solution might work better. In low-resource settings, deploying AI without a clear fit can waste time, introduce risk, or make systems more fragile. This step helps you avoid building AI for its own sake and focus your efforts where it will truly deliver value.

HOW TO DO IT

RECALL YOUR PROBLEM STATEMENT

1 Start by revisiting the workflow challenges you defined in the Discover phase. What exactly are you trying to improve? Keep the workflow challenge and value driver front of mind to ground the assessment.

UNDERSTAND THE AI SUITABILITY FRAMEWORK

2 Use a decision support tool that helps you assess whether a challenge is a strong fit for AI or not. The framework shifts the conversation from "Can we use AI here?" to "Should we?" by guiding you through conditions like whether the task is repetitive, if data exists, or if human judgment is essential.

TEST EACH WORKFLOW CHALLENGE

3 Bring together a cross-functional group and apply the framework to each challenge. For each condition, discuss whether the challenge meets the criteria for a good fit. Different roles will surface different realities: what leadership sees as a strength might be a challenge for frontline staff, and technical teams may identify infrastructure gaps others overlook.

4

FLAG CHALLENGES FOR NON-AI APPROACHES

If a challenge does not meet the threshold for AI, do not force it. Flag it as a good candidate for other interventions such as workflow redesign, digital forms, or future automation when conditions improve.

*AI WORKS BEST WHEN

- Tasks are high-volume, repetitive, or pattern-based
- Relevant data exists or can be collected ethically
- Simpler tools cannot solve it as effectively
- AI offers measurable performance gains
- Technical infrastructure and skills are in place
- Human oversight is feasible
- The environment is safe to test with safeguards





STEP 2

TRANSLATE INSIGHTS ON WORKFLOW CHALLENGES INTO AI USE-CASES

Turn real workflow pain points into well-sscoped, responsible AI use cases. This step helps you move from understanding where users struggle to defining how AI could meaningfully help. A use case describes a specific scenario where AI can support a user to make a decision, take an action, or complete a task more effectively. It specifies what the AI will do, for whom, and under what ethical guardrails. Think of the use case as a focused building block that is smaller than a full solution but sharper than a broad challenge. It is the starting point for building responsibly, validating feasibility, and aligning teams before investing resources.

HOW TO DO IT

1 RECALL THE WORKFLOW CHALLENGE

Start with the workflow challenges you defined in the Discover phase. Identify what kind of intelligence or automation could help. Ask: What is the user trying to achieve but struggling to do? Could AI assist by predicting something, recommending next steps, summarizing or extracting information, generating content, or detecting anomalies or trends?

2 FLESH OUT THE USE CASE DETAILS

Capture who the use case serves, the trigger and workflow touchpoints, required inputs and outputs, constraints, risks, and how success will be measured. Document all of this in one visible place so everyone can see and refine together. This creates a shared reference point that aligns your team around the problem before exploring AI solutions.



3 VALIDATE THE USE CASE WITH YOUR TEAM

Before moving forward, pause to test whether your defined use cases resonate beyond the immediate group that drafted them. Share the draft use cases with a broader circle of stakeholders and invite them to ask: Does this use case truly reflect the challenges and priorities we see in practice? Are the described benefits meaningful for the target users? Are there any risks, blind spots, or unintended consequences we have overlooked?

4 CONFIRM ALIGNMENT & PREPARE FOR ASSESSMENT

Once validated, consolidate your shortlisted use cases and agree as a team on which ones to carry forward for assessment. At this point, you are not yet deciding feasibility or impact, only confirming that the use case is well-framed, resonates broadly, and is ready to be stress-tested for viability.



A good use case is concrete, user-driven, and shaped collaboratively. It defines what AI will do, for whom, and under what safeguards before any building begins.

STEP 3

ASSESS THE FEASIBILITY & IMPACT OF AI USE-CASES

Before you decide which AI use case to pursue, you need to know two things:

1. Can it realistically be built (feasibility)
2. Will it create enough value if it succeeds (impact).

Feasibility is not only about technology, it also depends on whether people are ready and able to adopt, use, and sustain the solution. This step walks you through structured ways to evaluate both feasibility and impact so your organization invests in AI opportunities that are not only exciting but also viable, meaningful, and centered on people.

HOW TO DO IT

1

TAILOR YOUR ASSESSMENT CRITERIA

Avoid using a generic list. Select the feasibility and impact dimensions most relevant to your context. For feasibility, you might include data availability, infrastructure, people readiness, management cost, and ethical safeguards. For impact, anchor criteria in your defined value drivers such as efficiency, equity, cost savings, or improved outcomes.

2

RECALL YOUR USE-CASE

Choose one well-defined use case from the previous step. Keep the workflow challenge and value driver front of mind to ground the assessment.

3

ASSESS FEASIBILITY

For each feasibility dimension, assign a score on a 1 to 5 scale. Be explicit about what "1" (very low) and "5" (very high) mean in your context. Document the reasoning behind each score, especially where opinions differ. The discussion itself often reveals blind spots.



4

ASSESS IMPACT

For each feasibility dimension, assign a score on a 1 to 5 scale. Be explicit about what "1" (very low) and "5" (very high) mean in your context. Document the reasoning behind each score, especially where opinions differ. The discussion itself often reveals blind spots.

5

DOCUMENT RESULTS FOR DISCUSSION

Capture scores and the reasoning behind them, especially where opinions differ. The discussion itself often reveals blind spots. Keep notes ready for use in the next step: prioritization across multiple use cases.



Feasibility is more than technical readiness. It includes whether people will adopt the solution and whether your organization can sustain it. Assess both feasibility and impact honestly to invest in what can actually succeed.

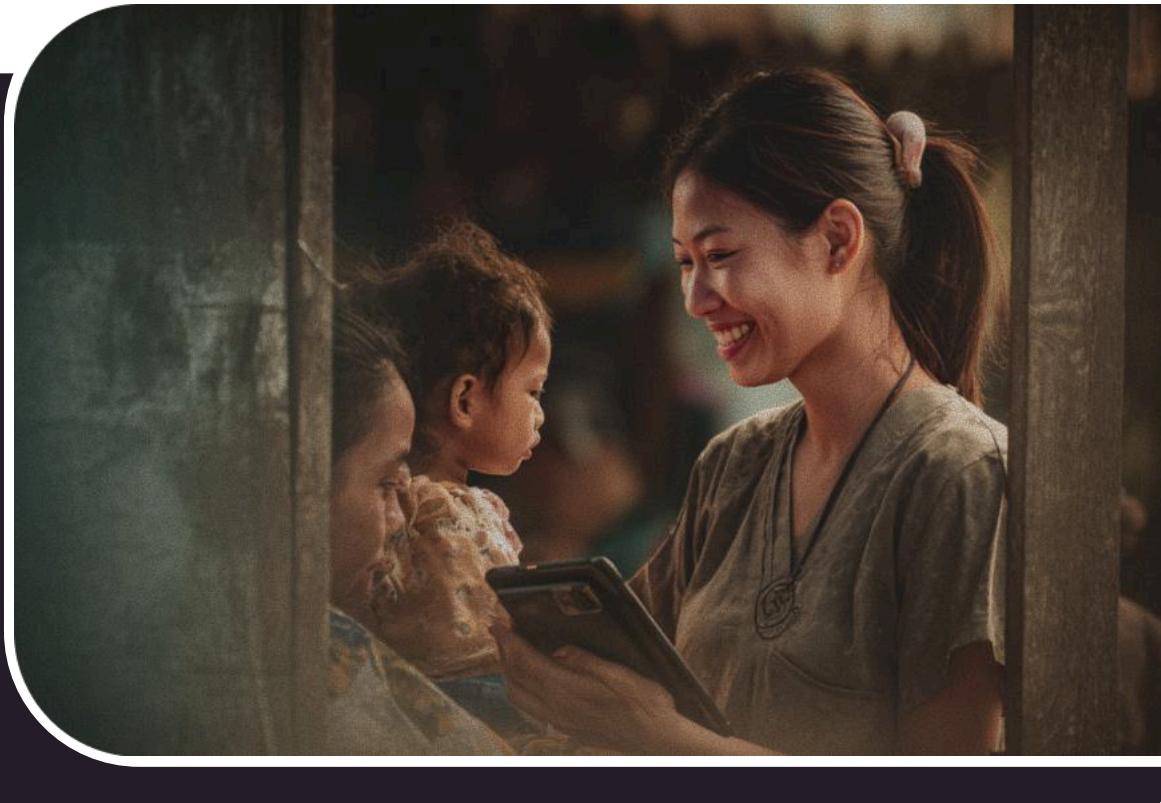
STEP 4

PRIORITIZING AI USE CASES

You have scored your use cases, now it is time to decide which to pursue. Prioritization is about turning scores into strategy. Focusing scarce resources on the use cases that offer the best balance of value and feasibility. In resource-constrained contexts, trying to build everything at once dilutes impact and overwhelms teams. Prioritization ensures you start with the most promising opportunities, delivering early wins and building momentum for long-term AI adoption.

HOW TO DO IT

- PLOT THE SCORES**
Create a simple 2x2 grid with feasibility on the X-axis and impact on the Y-axis. Plot each use case to see which are worth pursuing.
- VISUALIZE PATTERNS & TRADE OFFS**
Visualize patterns and trade-offs. Discuss what the distribution shows. Are ideas clustered? Any outliers? Are many strong-value ideas blocked by low feasibility, suggesting data or people-readiness gaps?
- FACILITATE A CO-WORKING SESSION**
Facilitate a co-working session. Invite technical, program, and user-side teams to challenge assumptions. Ask: What would move a promising but low-feasibility idea to the right? What risks or ethical flags change positioning?



- 4** **AGREE ON PRIORITIES & RECORD WHY**
Select 1 to 2 use cases to advance into Design. Note the decision rule, for example, prioritize cases above threshold on both axes. Capture a short decision log: rationale, diverging views, dependencies, and what must be true before building.

- 5** **REVISIT LATER**
Keep the matrix updated. A low feasibility case today might shift to ready in a year with better infrastructure or data. Add non-selected ideas to a backlog with unlock conditions, such as data pipeline, policy clearance, user training, and set a review date.



Feasibility is more than technical readiness. It includes whether people will adopt the solution and whether your organization can sustain it. Assess both feasibility and impact honestly to invest in what can actually succeed.

DEFINE TEMPLATES

Templates to structure your insights from the Define phase. These tools help you translate challenges into clear, well-scoped AI opportunities and shortlist what to take forward into design.

USE CASE DEFINITION CANVAS

A structured document to capture key details about your AI use case: who it serves, what AI does, how it works, what value it creates, and what safeguards are needed.

TIP

Keep your use case focused and concrete. Avoid vague ideas or overly ambitious concepts. Start by drafting multiple use cases if needed, but refine them into one or two that are specific, user-driven, and feasible in your context.

2X2 PRIORITIZATION MATRIX

A simple 2x2 grid to map your use cases based on feasibility and impact. This visual tool helps you compare options and identify which opportunities are ready now versus which may need to wait.

TIP

Beyond which use cases are in the "high-high" quadrant, document why other use cases fell where they did and what conditions would make postponed ideas viable in the future.

SCAN QR CODE



TEMPLATE DOWNLOAD



PHASE 3 DESIGN

The People Centered AI Playbook

PHASE 3

DESIGN OVERVIEW

ABOUT THIS PHASE

The Design phase is where ideas become concrete plans. After choosing your best AI opportunity in the Define phase, you now figure out how to actually build it. This means answering key questions: What data do you need? How will users interact with it? What tools will you use? What could go wrong, and how will you prevent it? By the end of this phase, you will have a clear blueprint that your team can follow to build, test, and launch the solution.



This is where you turn your chosen opportunity into a workable plan that fits your team, your users, and your constraints.

HOW TO LOCALIZE THIS STEP

This is the phase where design choices must reflect the realities of your context. As you prepare to shape your solution, use languages and interaction styles that reflect your users' everyday experience, factor in digital literacy levels when shaping interfaces, check policy or vendor rules that may limit tool choices, and account for infrastructure realities like connectivity and device types when making design decisions.

END GOAL

By the end of this Design phase, you will have a clear plan with data requirements, tool choices, and success criteria that guide your team as they build.

You'll leave with

1 Data requirements

What data you need, what you already have, and what is missing

2 User interaction designs

Visual sketches showing how people will use the solution

3 System diagram

A simple drawing showing how information flows through your solution

4 Tool and vendor decisions

Documented choices about which tools to use and why

5 Success criteria and safety measures

Clear definitions of what good looks like and how to prevent harm

STEP 1

ASSESS THE DATA NEEDS FOR A SPECIFIC USE CASE

AI runs on data, but for many organizations, it is the bottleneck. You cannot design a viable solution without first knowing if the right data exists, what state it's in, and whether you can use it responsibly. This step helps you move from wishful thinking ("we'll figure out the data later") to grounded design choices. Many AI projects fail because either the required data does not exist, the available data is incomplete, messy, or biased, or the team underestimates the time and skills needed to prepare it. In low-resource settings, these challenges are magnified: paper records, fragmented systems, sparse data, language diversity, and infrastructure barriers. Understanding your data needs early prevents wasted effort and ensures you design responsibly.

HOW TO DO IT

1

RECALL YOUR AI USE CASE

Here you translate your AI use case into the actions an AI model must perform (e.g., classify, extract, predict, rank, generate, summarize). This drives the data requirements. What exactly are you trying to improve? Keep the workflow challenge and value driver front of mind to ground the assessment.

2

DEFINE THE DATA YOU WILL USE

What inputs are needed? What outputs should the model generate? Be specific: text, audio, images, structured fields. Understanding the type of data shapes everything from storage to model selection. Document what quality and quantity is needed for the model to learn or generate something useful.

3

AUDIT WHAT YOU ALREADY HAVE

Check internal systems, paper logs, partner datasets. Where does this data live today? What condition is it in? What is missing? Without mapping how work is really captured, you risk overestimating availability and designing around gaps you didn't see coming.



4

ASSESS IMPACT

Is the data too sparse, inconsistent, biased, or inaccessible? What risks could it introduce? Data may underrepresent certain groups due to connectivity gaps, incomplete records, or collection biases. Outputs based on incomplete data risk reinforcing existing inequities or simply failing in practice.

5

DEFINE YOUR DATA STRATEGY

Will you clean, label, synthesize, or collect new data? Who will do it, and with what resources? Think about whether you'll partner with organizations to digitize records, use human-in-the-loop validation, prototype with public datasets, or apply transfer learning on smaller sets. These approaches can bridge gaps and make pilots safer.



You do not need perfect data to start. Map what you actually have, identify gaps honestly, and build iteratively with small steps. Prototype with public datasets, test manual processes, or use human-in-the-loop approaches to bridge gaps. The goal is to design around your data reality, not wishful assumptions.

STEP 2

DESIGN USER FLOWS & MOCK-UPS

Once you've clarified your data needs, the next step is to make your AI concept real. This is where ideas move from theory to experience. By designing user flows and mock-ups, you help the team see how people will actually use the solution and where AI makes a difference. These quick visualizations reveal what users see, do, and feel as they interact with the system. They also surface design gaps, accessibility barriers, or technical limits that could slow down development later. In settings where devices, bandwidth, or literacy levels vary, this step is especially useful for grounding your design in real conditions.

HOW TO DO IT

1

MAP THE FULL JOURNEY

Begin by outlining the steps a person takes to complete their task today. Include what triggers each action, what tools or people they interact with, and where frustration often appears. Mark points where AI could remove repetition, increase accuracy, or improve decision-making. The goal is to understand the current experience deeply before changing it.

2

SKETCH END-TO-END FLOWS

Once you understand the journey, draw it out. Use simple shapes, arrows, and notes to show how information moves from one step to another and how people interact with it. Mark where AI might appear, such as offering suggestions or simplifying an action. Keep it light so your team can add, remove, or rearrange ideas easily. These sketches show the whole picture of how humans and AI will work together and where the flow can stay smooth and natural.



3

BUILD QUICK MOCK-UPS

Turn your sketches into simple screens or storyboards that look close to real use. Use short phrases, icons, and feedback messages that feel authentic. Focus on showing the most important interactions, such as what users tap, read, or respond to, and how the AI reacts. You do not need polished visuals at this stage. What matters is testing how the flow feels and making sure every step communicates clearly.

4

TEST & REFININE WITH USERS

Share your mock-ups with a few people who match your intended audience. Ask them to walk through a task and explain what they think is happening. Watch where they hesitate or appear uncertain. Listen to their feedback and note if the flow feels too complicated or unclear. Use what you learn to adjust the layout, wording, or sequence so the experience feels natural and supportive.



Prototype early and share often. A quick sketch that users understand saves more time than a polished interface built on guesswork. Keep each iteration focused on trust, ease, and clarity between human and machine.

STEP 3

DESIGNING SCALABLE & CONTEXT AWARE AI ARCHITECTURE

Once your user experience is clear, focus on how your AI solution will operate behind the scenes. This step turns ideas into a system that is stable, adaptable, and ready to grow. A scalable and context-aware design ensures that your AI remains reliable across different environments, user groups, and resource conditions. Thinking about scale early helps prevent costly redesigns, while considering context ensures the solution stays inclusive and relevant. Together, these principles make your system strong, efficient, and future-ready.

HOW TO DO IT

1. DEFINE THE SYSTEM'S CORE PARTS

Connect each function in your system to a clear step in the user journey. Every action should serve a real moment in the experience. This keeps technical choices grounded in user needs. When design and architecture move together, the system feels cohesive and easier to manage as it grows.

2. PLAN FOR FLEXIBILITY AND GROWTH

Identify your main data sources and show how information travels between them. Note who manages it, where it's stored, and what requires extra protection. Mapping the flow prevents duplication, improves accuracy, and helps teams spot risks early. It also supports privacy, security, and responsible use.

3. PROTECT PRIVACY & MANAGE ACCESS

Decide who can view, change, or share different types of information. Create straightforward permission levels that keep data safe without blocking everyday work. Store only what is needed and explain why. When users understand how their data is handled, trust becomes part of the system's design.

4. OUTLINE THE PROCESS FLOW

Draw the path from input to output, showing where automation starts and people step in. Add checkpoints for review so the system can be tested and improved. A clear process makes it easier to monitor quality and explain how results are created.

5. DESIGN THE INTERACTION LAYER

Plan how users experience feedback from the AI. Decide when guidance appears, how it sounds, and how errors are handled. Use consistent cues that show what the system is doing. Clear, predictable communication helps users feel supported and confident.

6. PLAN FOR GROWTH & CHANGE

Expect the system to evolve. Use modular parts that can be updated or replaced easily. Leave room for new users, tools, or languages. Document connections so future teams can extend the system without disruption. Designing for growth now keeps it flexible and future-ready.



STEP 4

EVALUATE OPEN-SOURCE VS PROPRIETARY TOOLS

The tools you select determine how easily your system can grow and how well it performs over time. This step helps you compare options and choose those that fit your goals, values, and capacity. Whether the tools are open-source or proprietary, focus on flexibility, reliability, and long-term support. Taking time to evaluate carefully ensures your solution stays stable as your team, data, and needs evolve.

HOW TO DO IT

1 CLARIFY YOUR PRIORITIES

Begin by outlining what the system must achieve. Identify essential features such as language support, processing speed, and data handling. Separate must-haves from nice-to-haves. This clarity narrows options and keeps attention on what matters most. A focused list prevents distraction from flashy but unnecessary tools.

2 COMPARE A RANGE OF OPTIONS

Review open-source and proprietary tools side by side. Evaluate performance, usability, and integration with what you already have. Consider security, transparency, and cost over time rather than up-front price. Record findings in a shared document so the reasoning behind each choice remains clear.

3 ASSESS RISK & LONG TERM FIT

Examine the sustainability of each tool. Check update frequency, community strength, and vendor reliability. Look for detailed documentation and responsive support. Tools that evolve with your context reduce maintenance stress and future costs. Prioritizing durability ensures your system remains reliable and current.



4 RUN SMALL PILOT TESTS

Test your top options in realistic conditions. Use sample data or short workflows to observe speed, accuracy, and stability. Include both technical and non-technical teammates to understand usability. Small pilots reveal hidden challenges early and guide smarter final decisions.

5 DECIDE & RECORD YOUR RATIONALE

When testing is complete, make your choice transparent. Document which tools were chosen, why, and what trade-offs were accepted. Note risk mitigation or backup plans. Keeping this record supports accountability and simplifies future upgrades or audits.



Pick tools your team can learn, explain, and fix. Simpler and well-supported always beats complicated and fragile.

STEP 5

ASSESS VENDOR READINESS FOR RESPONSIBLE AI

When partnering with external vendors, confirm they meet strong ethical and operational standards. This step ensures that everyone contributing to your AI solution upholds responsible practices that protect users and sustain trust. Assessing readiness early prevents costly risks and helps your team choose partners who are transparent, reliable, and aligned with your values. Responsible collaboration turns compliance into shared accountability and builds the foundation for long-term, trusted partnerships.

HOW TO DO IT

1 REVIEW GOVERNANCE & POLICIES

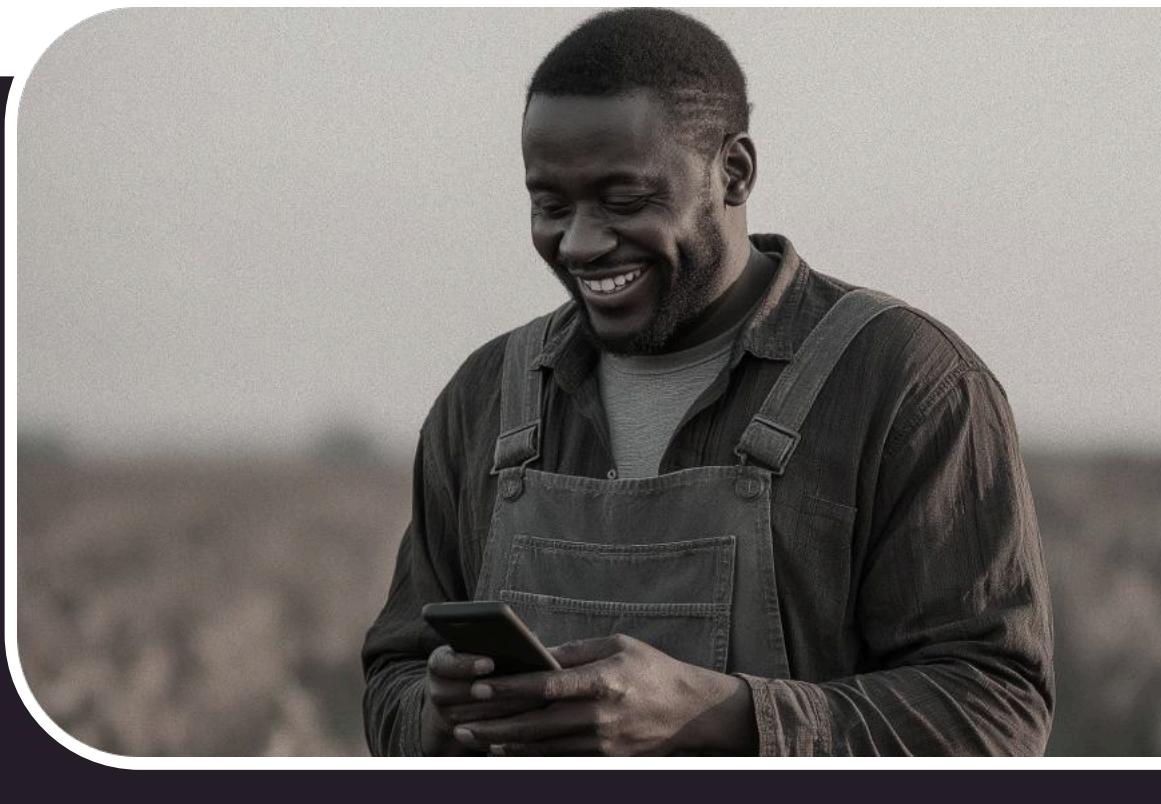
Request written details on how vendors handle data collection, consent, and retention. Check how they manage user rights and deletion. Reliable partners document their approach clearly and follow legal and ethical standards. Keep this material on file to confirm that vendor practices align with your own organizational commitments to safety and privacy.

2 EXAMINE TECHNICAL PRACTICES

Ask vendors how they design, test, and monitor their systems. Look for evidence of regular checks for bias, error, and misuse. Responsible vendors share their methods openly and welcome feedback. Clear documentation and a willingness to discuss limitations signal maturity and integrity.



If a vendor cannot show proof of responsible practice, they are not ready. Choose partners who demonstrate both good intent and clear evidence of action.



3 CHECK INCIDENT RESPONSE PLANS

Learn how vendors handle failures or breaches. Ask who leads the response, how users are informed, and what steps ensure quick recovery. Strong vendors communicate transparently and have clear escalation timelines. Knowing this plan in advance prepares your team to coordinate effectively when something goes wrong.

4 SCORE & DOCUMENT READINESS

Create a simple checklist covering ethics, security, and transparency. Use it consistently for all vendors so comparisons are fair. Keep notes on strengths, weaknesses, and follow-up actions. A standardized review turns vendor assessment into a transparent, repeatable process rather than a one-off judgment.

STEP 6

DEFINE QUALITY & ETHICAL GUARDRAILS EARLY

Before development begins, decide what quality and responsibility mean for your AI system. Guardrails give your team direction when trade-offs arise and keep the system fair, transparent, and safe. Setting these standards early helps balance innovation with care and ensures decisions align with your organization's mission and the needs of real users. When everyone shares a common definition of quality, progress becomes faster, safer, and easier to explain.

HOW TO DO IT

- 1 **SET FUNCTIONAL ASSESSMENT CRITERIA**
Define what the system must do to meet its core purpose. Include how it supports users, completes tasks, and handles errors. Functional criteria create shared clarity on when the system is delivering real value.
- 2 **AGREE ON PERFORMANCE BENCHMARKS**
Establish measurable targets such as accuracy, speed, or reliability. Define acceptable limits for mistakes and specify when human review is required. Clear benchmarks guide testing and help teams decide when the system is ready for release.
- 3 **INCLUDE USER-EXPERIENCE STANDARDS**
Identify expectations around clarity, tone, and ease of use. Account for language, literacy, and cultural context. Designing for inclusivity makes the system feel welcoming and reduces confusion or misinterpretation.



- 4 **SET OPERATIONAL EXPECTATIONS**
Prepare for real-world challenges by including goals for uptime, offline function, and resilience in low-resource settings. Reliable systems plan for constraints instead of assuming ideal conditions.
- 5 **DEFINE ETHICAL BENCHMARKS**
Capture principles like fairness, transparency, and human oversight. Specify how these values will be tested and maintained. Ethical benchmarks make responsibility measurable and visible rather than abstract.



Protect people first. A system that is safe, fair, and transparent will always outlast one that focuses only on performance.

DESIGN TEMPLATES

Templates to guide your work in the Design phase. These tools help you plan how the AI solution will function, define data needs, and ensure quality and ethics are built in from the start.

PRODUCT REQUIREMENTS DOCUMENT (PRD)

A shared document linking your system design, architecture, and quality guardrails. It aligns technical, ethical, and user needs before development begins.

i TIP

Build it alongside your User Flow Diagram. Keep it simple and update as design decisions evolve.

BUY VS BUILD CHECKLIST

A quick guide to compare open-source and proprietary tools. It helps balance cost, flexibility, and long-term support.

i TIP

Complete it as a team. Discuss trade-offs and document why each choice works for your context.

USER FLOW DIAGRAM

A visual map showing how users move through your solution and where AI steps in.

i TIP

Sketch by hand first. Focus on key actions before refining digitally.

SCAN QR CODE



TEMPLATE DOWNLOAD

Acknowledgment

Dalberg Design

The [User flow diagram](#) was developed in collaboration with Dalberg Design.





PHASE 4
DEVELOP



PHASE 4

DEVELOP OVERVIEW

ABOUT THIS PHASE

The Develop phase turns ideas into real, working systems. It focuses on assembling the right team, selecting an appropriate Generative AI model, building an initial prototype, and testing it with real users through User Acceptance Testing (UAT). This is where plans take shape and concepts become products that can be validated in practice. Development bridges the gap between design and delivery, ensuring what you build is functional, inclusive, and ready for real-world use.



This phase is about building responsibly, aligning talent, tools, and testing so that solutions are practical and ready for scale.

HOW TO LOCALIZE THIS STEP

This is the phase where the solution begins to take shape, so it must reflect local realities. Match your development approach to the skills your team or partners already have, and choose models that can handle local languages or non-standard inputs. Prototype on the devices people already use, paying attention to performance under low bandwidth or offline settings. Run UAT sessions in real environments, not labs, to capture issues such as literacy, connectivity, and workflow differences early.

END GOAL

By the end of the Develop phase, you will have a working prototype that demonstrates your solution's core function, tested in real conditions, and documented with clear performance and maintenance notes.

You'll leave with

1 Team Alignment

A clear structure and workflow connecting technical and non-technical contributors.

2 Model Selection

A chosen Generative AI model suited to your context and user needs.

3 Prototype and Testing

A functioning prototype tested with users and informed by UAT feedback.

4 Path to Scale

A sustainable setup with next steps for pilot and scale-up.

STEP 1

UNDERSTAND THE TECHNICAL SKILL REQUIREMENTS & DEVELOPMENT PROCESS

Turning an AI use case into a real, working solution requires both the right expertise and a realistic development plan. This step helps you identify the skills needed, assess your team's capacity, and plan for gaps before building begins. It ensures your system is not only created but maintained, scaled, and improved over time. In resource-constrained settings, mapping skills early avoids projects collapsing after prototypes are built and ensures development remains sustainable and grounded in real capabilities.

HOW TO DO IT

1 MAP SKILLS ACROSS THE AI JOURNEY

List the types of expertise you'll need from start to finish. Some stages need technical builders, others need people who understand data or can keep things running. Identifying these gaps early means you can plan how to fill them.

2 LIST KEY ROLES & RESPONSIBILITIES

Translate needed skills into clear roles. Define who leads model design, manages data pipelines, and ensures testing and feedback. Clarity in roles avoids overlap and confusion. When each person knows what success looks like in their area, teams can collaborate smoothly across technical and non-technical functions.

3 ASSESS YOUR CURRENT CAPACITY

Evaluate your team's readiness. Ask if they can write production-ready code, manage data preparation, and build usable interfaces. Identify where skills are strong and where external help is needed. Honest assessment now saves time later by making resource decisions visible and actionable.



5 PLAN FOR LONG-TERM SUSTAINABILITY

Assign ownership for retraining, monitoring, and maintenance early. Ensure documentation and handovers are clear before scaling. Sustainable systems depend on shared responsibility rather than individual heroics.

4 FILL THE GAPS STRATEGICALLY

Few teams have every skill in-house. Partner with universities, local tech-for-good groups, or global networks for specialized support. Outsource short-term tasks like data labeling while training your own team. Strategic partnerships build momentum without overextending internal staff.

Keep teams lean and roles clear. Build only the skills needed for the next phase, not for every possible future.

***Use the skills matrix template at the end of this phase to support you in this step.**



STEP 2

SELECT A GENERATIVE AI MODEL

Selecting the right Generative AI model is one of the most critical technical decisions in development. This step helps you evaluate available models, understand their capabilities and limits, and choose one that aligns with your goals, data quality, and infrastructure. The right model balances accuracy, cost, speed, and accessibility. In environments where computing resources or data are limited, choosing an adaptable, lightweight model ensures your solution performs reliably and can scale responsibly without overburdening local systems.

HOW TO DO IT

1. DEFINE SUCCESS & PREPARE TEST DATA

Start by clarifying what success looks like for your use case. Build a small “golden dataset” of real examples that reflect both common and edge cases. This will act as your reference for testing and scoring model performance consistently across candidates.

2. SELECT MODELS THAT FIT YOUR CONTEXT

Look for AI tools (free or paid) that work with your language, your type of data, and your tech setup. If internet is unreliable, prioritize tools that can work offline or with slow connections. Pick tools that can handle the same real-world conditions your users face.

3. RUN CONTROLLED COMPARISON TESTS

Test shortlisted models using identical prompts and conditions. Measure accuracy, speed, and cost per output. Invite two reviewers to rate results for quality and clarity. These tests highlight trade-offs and reveal which model performs best under real-world pressures.



4. ESTIMATE COST & MAINTENANCE EFFORT

Look beyond performance to understand what it takes to sustain the model. Account for hosting, API fees, retraining schedules, and staff time. A model that performs well but costs too much to maintain is not sustainable.

5. DECIDE & DOCUMENT YOUR RATIONALE

Choose the model that best meets your goals and resource realities. Record why it was selected, its known limits, and how you will retrain or upgrade it over time. This documentation supports transparency and prepares the team for scaling.



Choose the simplest model that reliably meets your needs. Complexity does not equal quality and often increases cost, maintenance, and barriers to scale.

STEP 3

BUILD A PROTOTYPE

Building a prototype transforms your model and design concepts into something tangible. This step allows your team to test assumptions, gather user feedback, and refine the experience before full-scale development. A good prototype demonstrates core functionality without the need for complete integration. It helps teams identify usability issues, performance limits, and trust gaps early. In environments with limited time or resources, rapid prototyping encourages creativity, reduces cost, and prevents expensive mistakes later in the process.

HOW TO DO IT

1

DEFINE THE CORE PURPOSE OF THE PROTOTYPE

Focus on testing a few high-value features, not the entire product. Identify what you need to learn, such as how users interact with AI feedback or how the model performs under real conditions. Setting a narrow scope keeps the build efficient and purposeful.

2

PLAN YOUR PROTOTYPE STRUCTURE

Outline how the prototype will function from end to end. Sketch user screens, AI interactions, and input flows. Use low-code or no-code tools where possible to save time. A clear structure helps the team visualize dependencies and simplifies coordination between design and engineering.

3

BUILD & CONNECT THE MAIN COMPONENTS

Assemble the minimum set of elements that bring the idea to life. This includes connecting the AI model, input interface, and data handling. Keep the build modular so it can be replaced or extended easily. Prioritize usability and clarity over advanced features.



TEST WITH REAL USER EARLY

4

Run quick sessions with small user groups who represent your audience. Observe how they navigate the prototype, where they hesitate, and how they interpret AI responses. Feedback from just a few users can reveal critical improvements before you invest in a full build.

5

ITERATE & DOCUMENT LEARNINGS

Update the prototype based on feedback. Record what worked, what failed, and what still needs testing. These insights become the foundation for your next development phase and help new team members understand key decisions.



Build light, test fast, and learn early. A prototype's goal is to teach your team what to improve, not to deliver a finished product.

STEP 4

CONDUCT USER ACCEPTANCE TESTING

User Acceptance Testing (UAT) is where your AI solution meets real users. It tests whether the system performs as expected, feels intuitive, and supports people in achieving their goals. This step validates functionality and usability while revealing gaps that might not appear during internal testing. UAT ensures your solution is both effective and trusted. By grounding feedback in real use, you build confidence that the product is ready for wider deployment or refinement.

HOW TO DO IT

1

DEFINE WHAT SUCCESS LOOKS LIKE

Identify the key behaviours or outcomes that will prove your system is working as intended. Set measurable success criteria, such as accuracy, response time, or ease of use. These become the benchmarks against which user feedback is compared.

2

PLAN YOUR TEST SESSIONS

Recruit participants who represent your real users and design short, focused sessions. Prepare realistic scenarios that mirror everyday conditions. Keep the setup lightweight, ensuring people can test using their own devices when possible. Testing in authentic settings reveals insights lab tests may miss.

3

GUIDE OBSERVE & LISTEN

Introduce the system clearly, then step back and watch how users interact with it. Note where they hesitate, ask questions, or misinterpret AI feedback. Encourage them to describe what they expect to happen next. Honest observation uncovers both usability issues and opportunities for improvement.



4

CAPTURE & PRIORITIZE FEEDBACK

Collect feedback through short surveys, voice notes, or group discussions immediately after each session. Categorize issues by severity and frequency. Address critical issues first, especially those affecting trust, accessibility, or task completion.

5

DOCUMENT OUTCOMES & NEXT STEPS

Record key findings, what was fixed, and what remains unresolved. Summarize what changes are required for the next development cycle. Well-documented UAT results help align teams, track accountability, and guide decisions about scaling or further testing.



Treat UAT as a learning exercise, not an exam. Every challenge users face is a clue to making your system stronger, clearer, and easier to trust.



DEVELOP TEMPLATES

Templates to support your work in the Develop phase.

SKILLS MATRIX TEMPLATE

A tool to map the technical and non-technical skills needed for development. It helps teams identify strengths, skill gaps, and partnership needs before building begins.

TIP

Update it regularly as roles evolve. Use it to plan hiring, training, or external support.

SCAN QR CODE



TEMPLATE DOWNLOAD



PHASE 5 PILOT

The People Centered AI Playbook

PILOT OVERVIEW

ABOUT THIS PHASE

The pilot phase tests the solution in real conditions with a defined, representative sample group. Its purpose is to validate usability, feasibility, and early outcomes through focused evaluation and feedback loops. This phase helps teams learn how the solution performs in the hands of real users and whether it delivers measurable value. By running a lean, time-boxed pilot, you gather the evidence needed to decide whether to iterate, pause, or prepare for scale.



This phase is about learning quickly, validating real-world performance, and using evidence to guide scale decisions.

HOW TO LOCALIZE THIS STEP

This is the phase where design meets reality, so each pilot should reflect local conditions and user contexts. Use familiar communication channels such as WhatsApp, SMS, or in-person facilitators to collect feedback. Keep participation light with short, clear tasks that respect users' time and literacy levels. Partner with local organizations to ensure cultural fit, data consent, and smooth implementation. Test in realistic environments so you can uncover practical challenges like connectivity, device sharing, or language barriers early.

END GOAL

By the end of the Pilot phase, you will have tested your solution in realistic settings, gathered reliable evidence on usability and performance, and decided whether to iterate, scale, or pause based on the results.

You'll leave with

1 Pilot design & Scope

A structured plan outlining sites, sample group, roles, and evaluation approach.

2 Evidence of Usability and Feasibility

Data and user feedback showing how the solution performs in context.

3 Iterative Improvement Plan

A documented list of issues, fixes, and design updates for the next phase.

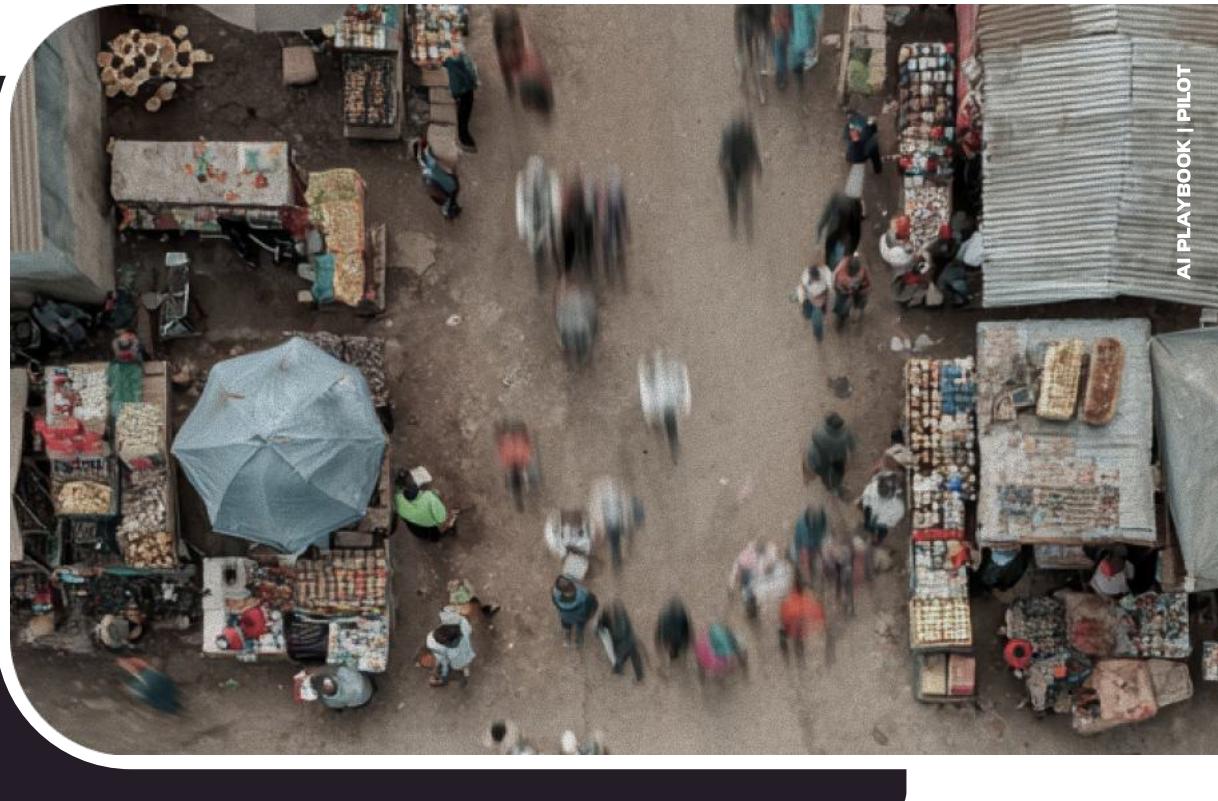
4 Decision path for Scale

A clear recommendation on whether to advance, adapt, or pause based on pilot findings.

STEP 1

DEVELOP A PILOT DESIGN THAT SETS THE RIGHT SCOPE & OUTCOMES

A pilot sets the boundaries for learning before scale. It defines where, with whom, and for how long your solution will be tested to validate usability, reliability, and early outcomes. By focusing on clear decision gates, lean indicators, and representative users, a well-scoped pilot turns limited resources into credible evidence that guides whether to iterate, expand, or refine your approach.



HOW TO DO IT

1

FRAME THE PILOT PURPOSE & BOUNDARIES

Start by defining what the pilot must prove. Identify specific questions it should answer and the features or workflows being tested. Clarify what success looks like and how long the pilot will run. Outline the decision points that determine whether the solution continues, adapts, or stops at midline and endline.

2

DEFINE ROLES & RESPONSIBILITIES

Map partners and team members who will support implementation. Assign clear roles for management, data collection, training, and feedback review. Develop a short RACI to clarify who is responsible, accountable, consulted, and informed. Include leads for ethics, consent, and communications to ensure oversight and transparency.

3

SPECIFY THE PILOT APPROACH

Select 1 to 3 representative sites that reflect real-world conditions, such as connectivity, device access, and user diversity. Plan how users will be recruited, trained, and supported throughout. Build in phases for onboarding, implementation, evaluation, and handover. Use short, practical methods such as community facilitation or SMS prompts to capture local feedback.

4

DESIGN A LEAN EVALUATION FRAMEWORK

Define the key learning questions the pilot will answer, such as “Does the system perform reliably?” or “Do supervisors feel more confident in their work?” Select a few measurable indicators linked to your value drivers. Track only what is essential, such as completion rate, uptime, or satisfaction. Document how evidence will be collected and ensure consent processes are followed.

5

PLAN FOR DOCUMENTATION & LEARNING CAPTURE

Create an evidence pack that includes logs, screenshots, anonymized notes, or short case summaries. Record what worked, what didn’t, and what needs adjustment. Keep it concise but complete, ensuring findings can be shared easily with partners or funders.

6

REVIEW & FINALIZE THE PLAN

Share the pilot plan with all key stakeholders for feedback. Consolidate comments, confirm ownership, and secure final approval before launch. A shared and signed-off plan creates alignment and accountability across all teams.

STEP 2

MONITOR PERFORMANCE FOR ITERATIVE IMPROVEMENT

Monitoring during the pilot helps teams track how the solution performs in real conditions and adjust before problems grow. It turns observation into improvement by revealing where systems slow down, users get stuck, or assumptions break. This step focuses on collecting data that explains both performance and experience. Regular check-ins and rapid fixes keep learning active, ensuring the solution stays responsive and grounded in real use rather than theoretical design.

HOW TO DO IT

1. DEFINE WHAT TO MONITOR & WHY

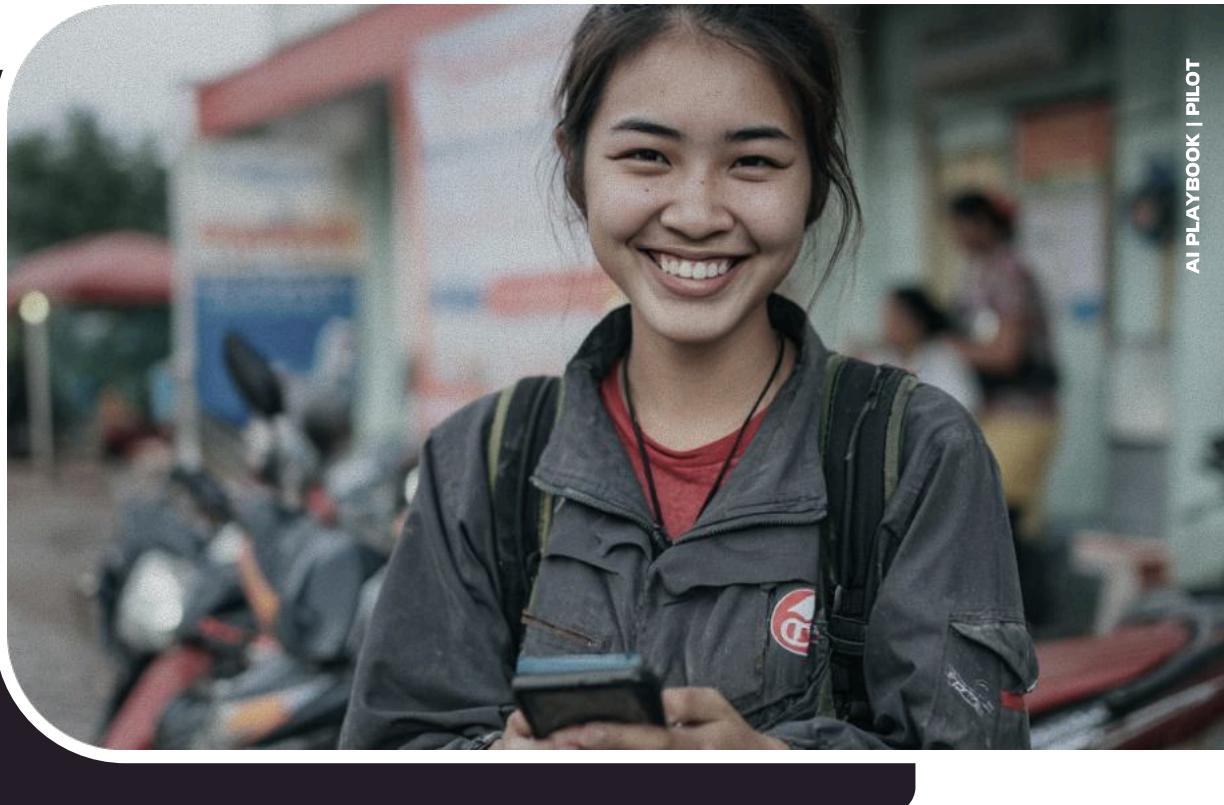
Start by identifying the indicators that reflect success, risk, or progress. Include both technical metrics, such as uptime and response time, and user measures like ease of use or satisfaction. Link every data point to a clear question: what does it tell us, and what decision will it inform?

2. BUILD A SIMPLE MONITORING SYSTEM

Use tools that your team can manage easily, such as shared spreadsheets, dashboards, or short daily forms. Keep data entry minimal and consistent. Ensure everyone understands what to record and when. Regular updates create a real-time view of performance and prevent surprises at review meetings.

3. COLLECT FEEDBACK CONTINUOUSLY

Set up clear feedback loops with users, field teams, and partners. Use short interviews, WhatsApp chats, or short survey links to gather insights quickly. Capture both what is working and what feels confusing or frustrating. Feedback collected in the moment is more accurate than reflections gathered later.



4. REVIEW & ADAPT REGULARLY

Hold short weekly or biweekly reviews to discuss what the data is showing. Prioritize critical issues first, especially those affecting trust, accuracy, or safety. Agree on quick fixes and assign owners. Use the pilot's flexibility to make immediate improvements instead of waiting until the end.

5. DOCUMENT LESSONS AS YOU GO

Record what you changed, why it was changed, and how it affected performance. This creates a learning trail that strengthens institutional memory and supports future scaling. Summarize key findings at midline to share with stakeholders and refine your next testing cycle.



Monitor lightly but often. Continuous observation and small adjustments deliver more progress than heavy reporting done too late.

TOOL

PILOT PHASE

LEAN SCORECARD

Use the scorecard at midline and endline to capture results. Focus on three to five indicators that matter most, such as adoption, reliability, and user trust. Keep scoring simple — red for gaps, yellow for progress, and green for readiness.

TIP

Before scoring, align your team on what each colour means in practice. Spend 5–10 minutes agreeing on what “red,” “yellow,” and “green” look like for each indicator so scoring stays consistent across midline and endline.



SCALE

The People Centered AI Playbook

PHASE 3

SCALE OVERVIEW

ABOUT THIS PHASE

The Scale phase focuses on expanding what works while maintaining quality, trust, and usability. It is the stage where a proven concept becomes a sustainable system that reaches more users. Scaling is about consistency, not speed. It ensures what worked in a small test continues to perform as teams, data, and demands grow. This phase helps you build strong foundations for governance, partnerships, and capacity so that success can last beyond the initial project team.



This phase is about growing intentionally, protecting quality, and ensuring long-term impact.

HOW TO LOCALIZE THIS STEP

Scaling must be shaped by local capacity and existing systems. Adapt your delivery model to align with national and regional processes rather than copying a single design. Empower local partners with the skills, documentation, and autonomy to maintain the solution independently. Use infrastructure and communication channels that already exist. Work closely with communities and regulators to ensure compliance, trust, and continuity. Scaling succeeds when it strengthens the ecosystem around it, not just the product itself.

END GOAL

By the end of the Scale phase, your solution will have expanded responsibly, with strong systems, clear ownership, and ongoing user trust. The focus shifts from testing to sustaining, ensuring the work continues to grow in value and reach.

KEY CONSIDERATIONS

These considerations help you maintain quality, manage risk, and build sustainability as your solution grows.

1 INSTITUTIONALIZATION

Scaling succeeds when your solution becomes part of daily systems and decision-making, not a one-off project. Integrate it into existing workflows, reporting structures, or policies to ensure accountability and long-term ownership.

2 PROCESS & PEOPLE

Equip and empower the right people to sustain and evolve your solution. Define roles for maintenance, training, and leadership. Support teams through capacity building, change management, and incentives that align with daily work.

3 MAINTENANCE & CONTINUOUS DEVELOPMENT

AI systems require ongoing care. Plan for updates to models, data, and interfaces as needs evolve. Build routines for quality checks, bug fixes, and user experience improvements to maintain reliability and trust over time.

4 CONTEXTUAL ADAPTATION

What works in one setting may not in another. As you scale across geographies or user groups, re-examine assumptions, language, and data flows. Adapt the design so the solution remains relevant, inclusive, and easy to use.

5 EVALUATION & EVIDENCE FOR SCALE

As solutions mature, shift from activity tracking to impact evaluation. Use proportionate, credible methods to validate performance, equity, and cost-effectiveness. Share findings transparently to strengthen learning across teams and partners.



Scaling responsibly means building the capacity to evolve. Keep embedding reflection and accountability as you grow.

The
People
Centered

AI **PLAY BOOK**

Building with AI is not a one-time effort. As organizations adopt solutions, we all learn more about what works, what does not, and why. The lessons come from practice, from getting it wrong, and from listening to the teams doing this work every day.

At Dalberg Data Insights, we know this playbook will get better only if it reflects more than our own experience. Share your feedback by scanning the QR code below. Tell us what worked, what did not, and what you would change. Your experience will directly shape the next version.

We are committed to refining this resource based on what we learn together. The goal remains the same: to help organizations design and adopt AI that starts with people, responds to real needs, and works in the contexts where it is meant to be used. That takes all of us.

The
People
Centered

AI

PLAY BOOK

Dalberg Data Insights