HOUSING ACTION ILLINOIS · 2026 HOUSING MATTERS CONFERENCE

The Responsible Innovation Blueprint

Action Toolkit

Aligning AI adoption with your mission

Moving from defensive anxiety to governed leadership

Steven J. Casey

Principal, Kindred Impact LLC

Wednesday, September 30, 2026 · 10:15–11:30 AM
I Hotel, Champaign, Illinois

KINDRED IMPACT LLC
Strategic Philanthropy • Local to Global

WHAT'S INSIDE

A governance framework, an equity screen for housing data, a vendor scorecard, and an 8-week plan, with a fillable template for every phase of the plan.

This is a fillable PDF. Type directly into any shaded field, then save a copy for your organization. Page 2 explains how to get the most out of it.

GETTING STARTED

How to Use This Toolkit

This toolkit turns the session into work your organization can do in the next eight weeks. Every template matches a slide in the deck and a phase of the 8-Week Governance Blueprint.

Type, then save. Click any shaded field to type. Checkboxes and circles click on and off. Fields work in free Adobe Acrobat Reader and most current PDF readers. If your reader will not save your entries, open the file in Acrobat Reader.

Keep a blank original. Save a separate copy for each organization or working group, so the blank version stays ready for the next team.

Work on it together. The templates work best when a small cross-functional group fills them in: leadership, programs, and whoever manages your data.

Short on time? Start with Your First Monday on page 16. Three actions this week are enough to begin.

WHAT'S IN THE TOOLKIT

THE FRAMEWORK · DECK SLIDES 6–7

The Head Chef Rule and the Mission Filter

One rule and one four-stage checkpoint. Every other page in this toolkit puts them into practice.

LINE CHEF · AI

The Line Chef (AI)

Handles the prep work: drafting, summarizing, sorting, transcription. Fast, tireless, and without judgment.

HEAD CHEF · HUMAN

The Head Chef (Human)

Final taste, ethical judgment, and accountability. Decides what leaves the kitchen and signs their name to it.

If the Head Chef doesn't sign off, the plate does not leave the kitchen.

Every AI-assisted intake summary, grant narrative, board memo, or tenant-facing letter needs a named human author.

The Human-in-the-Loop Mission Filter

A four-stage checkpoint for every AI-assisted task, from a grant report to a tenant letter.

STAGE 1

Mission Input

Define the objective and provide clean, appropriate data.

STAGE 2

AI Generation

The first-drafter processes the input into a working draft.

STAGE 3

The Mission Filter

A named human checks bias, equity, accuracy, and tone.

STAGE 4

Verified Impact

The verified version goes out, with the reviewer's name on record.

Stage 3 is the braking system. No output reaches a client, donor, or funder without passing through it.

WORKED EXAMPLE · A GRANT REPORT

Stage 1: Program staff give the AI the actual outcomes data.

Stage 2: The AI drafts narrative language.

Stage 3: The program director checks every number against source data and reviews the tone for dignity and accuracy.

Stage 4: The verified report goes to the funder, and the review is recorded in the Provenance Log (page 14).

Stage 3 is a role, not a vibe. Someone’s name goes next to it.

FILLABLE

The Equity Screen for Housing Data

AI systems learn from historical data, and housing data in America carries the fingerprints of redlining, exclusionary zoning, and biased lending. Run this screen before AI touches intake, outreach, screening, or anything a client will read.

Screening: ______ **Reviewed by:** ______ **Date:** ______

Zip-code proxies

Down-ranking applicants by neighborhood rather than need.

Wealth-biasing

Ranking donors or clients on patterns that track historic exclusion.

Silent disparate impact

Screening tools that quietly reproduce the outcomes you fight.

#	ASK BEFORE AI TOUCHES INTAKE OR OUTREACH	YES	NO	UNSURE	NOTES
1	Does this favor or exclude specific zip codes?	○	○	○	
2	Does it rank people on demographic patterns?	○	○	○	
3	Is the decision explainable to the person it affects?	○	○	○	
4	Does the language preserve their dignity?	○	○	○	
5	Have we tested it against real cases for disparate impact?	○	○	○	

How to read it: a Yes on questions 1–2, or a No or Unsure on questions 3–5, means stop. Fix the input, narrow the use, or escalate to your AI lead before anything goes out.

BEFORE IT GOES OUT · ACCURACY CHECK

- ☐ Every number and factual claim is traced to a source document. Anything not in the source data is treated as a hallucination and corrected or removed.
- ☐ A named reviewer has signed off in the Provenance Log.

Result: ○ Clear to proceed ○ Proceed after fixes ○ Stop and escalate

Governing AI well is fair-housing work. These are the same patterns you already fight in lending, appraisal, and zoning.

FILLABLE

The 8-Week Governance Blueprint

From ad-hoc, ungoverned AI use to a formalized, mission-aligned policy in eight weeks, not eight months. Most of it is naming decisions and writing them down, not buying new technology.

Organization: ______ **Blueprint lead:** ______ **Start:** ______

WEEKS 1–2

Discovery & Alignment

- [ ] Form a cross-functional team across leadership and programs.
- [ ] Survey staff on the AI tools already in use.
- [ ] Define your non-negotiables in writing.

Template: Shadow Kitchen Inventory, p. 6

Owner: ______

Dates: ______

WEEKS 3–5

Drafting the Guardrails

- [ ] Build the three-tier acceptable use policy.
- [ ] Name reviewers by role in the policy, by name in practice.
- [ ] Share the draft with your liability carrier.

Template: Acceptable Use Policy, Reviewers, Request Form, Vendor Scorecard, pp. 7–11

Owner: ______

Dates: ______

WEEKS 6–7

Pilot & Board Approval

- [ ] Run a two-week beta in one department and note the friction points.
- [ ] Score any tools used in the pilot.
- [ ] Bring the policy to your board, framed as protecting the mission.

Template: Board Brief, p. 12

Owner: ______

Dates: ______

WEEK 8+

Launch & Training

- [ ] Roll out organization-wide as empowerment, not restriction.
- [ ] Publish your AI Use Statement.
- [ ] Put a 90-day review on the calendar.

Template: AI Use Statement, Provenance Log, 90-Day Review, pp. 13–15

Owner: ______

Dates: ______

BY WEEK 8 YOU SHOULD HAVE

- [ ] **A signed AI use policy**
- [ ] **A named human review team**
- [ ] **A board that approved it**

FILLABLE

Shadow Kitchen Inventory

Staff are already using AI, with or without a policy. Find out where before you write one. The goal is a clear picture, not blame: the people using these tools are solving a real capacity problem.

Step 1. Ask. Send a three-question survey to all staff: Are you using any AI tools for work? Which ones? For what tasks? An anonymous survey gets more honest answers.

Sent:

Replies:

Step 2. List what you find.

AI TOOL	WHO USES IT	WHAT FOR	DATA IT TOUCHES	RATING
EXAMPLE *Otter.ai*	*Intake staff*	*Transcribing intake interviews*	*Client names, income, immigration status*	*Red*
				○ G ○ Y ○ R
				○ G ○ Y ○ R
				○ G ○ Y ○ R
				○ G ○ Y ○ R
				○ G ○ Y ○ R
				○ G ○ Y ○ R

Green: public or internal, non-sensitive information. **Yellow**: drafts that reach funders, boards, donors, or tenants. **Red**: client records, HMIS data, or case files in a tool with no contract protecting them.

Step 3. Flag the biggest risk.

Highest-risk tool:

Red-Face Test: if this tool leaked client data tomorrow, what would the headline be?

One stop-work step we will take this week:

Step 4. Name your limits. Write down the non-negotiables your policy will build on.

Example: We will never enter client records or HMIS data into a public AI tool.

1.
2.
3.

FILLABLE

Three-Tier Acceptable Use Policy

Sort every AI task into one of three tiers, so staff know before they start whether a task is approved, needs sign-off, or is off-limits.

Organization: | **Policy owner:** | **Effective:**

GREEN *Approved*

Low-risk tasks using public or internal, non-sensitive information. Staff may proceed with approved tools.

From the deck: *Summarizing public meeting notes, drafting internal agendas.*

Our examples	Approved tools

YELLOW *Supervisor sign-off*

Anything that reaches a funder, board member, donor, or tenant. A named reviewer checks it before it goes out.

From the deck: *Grant narratives, board memos, donor and tenant communications.*

Our examples	Who signs off

RED *Prohibited*

Sensitive client data in any public or non-contracted tool, and any eligibility, allocation, or waitlist decision made by AI. AI can draft and summarize; a person decides who gets housing.

From the deck: *Client intake records, HMIS data, case files in any public tool. Eligibility, allocation, or waitlist decisions made by AI.*

Our examples	Governed alternative, if any

Red does not mean "no AI." It means no public, non-contracted tools for sensitive data. There is almost always a governed path to the same benefit, such as an enterprise tool under a contract that bars the vendor from training on your data.

FILLABLE

Named Human Reviewers

Name reviewers by role in the policy, so it survives staff turnover, and by name in practice, so everyone knows who signs.

OUTPUT TYPE	REVIEWER ROLE (IN POLICY)	CURRENT REVIEWER (NAME)	BACKUP
Grant narratives and funder reports			
Board memos			
Donor communications			
Tenant- and client-facing letters			
Outreach and social media			

Before the board sees it

- [ ] Draft shared with our liability or cyber insurance carrier. Ask whether your coverage addresses AI use. **Date:**
- [ ] Draft reviewed by legal counsel, if available. **Date:**
- [ ] Staff input gathered from the Shadow Kitchen survey and reflected in the tier examples. **Date:**

Send an interim memo now

While the full policy is drafted, send staff a one-paragraph memo that sets the no-go zones. Edit the model below, then send it this week.

Model: Until our AI policy is adopted, please do not enter client records, HMIS data, case files, or donor personal information into any AI tool. Anything AI-assisted that goes to a funder, board member, donor, or tenant needs sign-off from your supervisor first. Questions go to [AI lead].

Memo sent on: **Sent by:**

 FILLABLE

AI Tool Request Form

The front door for every new AI tool. The person asking fills in the top; the AI lead or policy owner completes the decision. Anything that will touch Yellow or Red data also goes through the Vendor Scorecard on page 10.

Requested by: ________ **Department:** ________ **Date:** ________

1. What tool is being proposed? ________ **Cost:** ________

2. What problem does it solve?

Name the task and the time it takes today.

____

3. What data will it touch?

____

4. Which tier is the data? ◯ Green ◯ Yellow ◯ Red

Green: public or internal, non-sensitive. Yellow: reaches funders, boards, donors, or tenants. Red: client records, HMIS data, or case files, or any eligibility, allocation, or waitlist decision.

5. Who is the reviewer? ________

6. What is the exit plan?

How we would stop using it, and how we get our data back or deleted.

____

7. Red-Face Test. If this tool leaked client data tomorrow, what would the headline be, and what one step would prevent it?

Headline: ________

Prevention: ________

DECISION · AI LEAD OR POLICY OWNER

◯ Approved ◯ Approved with conditions ◯ Vendor Scorecard first ◯ Not approved

____

Decided by: ________ **Date:** ________

Add approved tools to your Shadow Kitchen Inventory (page 6), so the list stays current.

FILLABLE

Traffic-Light Vendor Scorecard

Run every AI-enabled contract through this before you sign: HMIS add-ons, donor CRMs, intake chatbots, transcription tools. New tools arrive here from the AI Tool Request Form on page 9.

Tool / vendor: ____ **Reviewed by:** ____ **Date:** ____

What it will do: ____ **Data it will touch:** ____

#	CRITERION	WHAT TO LOOK FOR	YES	PARTIAL	NO
1	Your data stays yours	The contract bars the vendor from using your data to train or "improve" its models, or you have opted out in writing. Read the data usage section of any agreement.	○	○	○
2	Audit rights in writing	You can review the vendor's security practices, such as a current SOC 2 report, and you know where your data is stored.	○	○	○
3	Written exit and data-export plan	You can export your data and have it deleted when you leave, on timelines stated in the contract, including backups.	○	○	○
4	Human review built in	A named person can check output before it reaches anyone. Nothing is sent automatically.	○	○	○
5	Accuracy and bias testing	The vendor can show accuracy and disparate-impact testing, or you can test it against your own cases first using the Equity Screen (page 4).	○	○	○

STOP	Do not use	A No on criterion 1, 2, or 3: the vendor trains on your data, there are no audit rights, or there is no exit plan.
CAUTION	Use with controls	Any Partial, or a No on criterion 4 or 5. Requires enhanced human oversight. Document your safeguards below.
GO	Safe to use	Yes on all five. Still subject to your tiers and named reviewers.

The Red-Face Test: if this tool leaks client data tomorrow, can I look my board in the eye? Worst-case headline:

____

Decision: ○ GO ○ CAUTION ○ STOP

Safeguards we will document (if CAUTION):

____

This scorecard is general guidance, not legal advice. Have counsel review contract terms before you sign.

FILLABLE

Practice: Spot the Flags

Sample clauses modeled on common enterprise AI service agreements. Mark each one Green (mission-safe), Yellow (negotiate), or Red (high risk), then compare your read with the notes at the bottom.

3.1 Customer Content. "Customer" (the Nonprofit) owns all input data. However, by using the Service, Customer grants "Provider" (the Vendor) a worldwide, non-exclusive, royalty-free license to use, host, store, and reproduce Customer Content to improve the Service and train Provider's underlying machine learning models, including models used for other customers.	◯ G	◯ Y	◯ R
3.2 Output Ownership. Provider hereby assigns to Customer all rights, title, and interest in and to the Outputs generated by the Service. Customer acknowledges that Outputs may not be unique and that the Service may generate the same or similar Outputs for other users.	◯ G	◯ Y	◯ R
8.1 No Warranty. The Service is provided "as is." Provider makes no warranties that the Outputs will be accurate, complete, or unbiased. Provider shall not be liable for any decisions made by Customer based on Outputs, including clinical, financial, or legal advice.	◯ G	◯ Y	◯ R
8.2 Indemnification. Customer shall indemnify and hold Provider harmless against any third-party claims arising from the Customer's use of the Outputs, including claims of defamation, copyright infringement, or violation of privacy rights.	◯ G	◯ Y	◯ R
12.1 Right to Audit. Provider maintains internal logs of model performance. Customer may request a summary of these logs no more than once per calendar year. Provider is not required to disclose specific training data sets, algorithmic weights, or logic used to produce specific Outputs, as these are protected trade secrets.	◯ G	◯ Y	◯ R
12.2 Data Portability. Upon termination of this Agreement, Customer may download their raw Input Data for 30 days. After 30 days, Provider will delete all Customer Content from primary servers but may retain anonymized fragments in backup archives indefinitely.	◯ G	◯ Y	◯ R

ONE READ · COMPARE AFTER YOU MARK YOURS

3.1 Red: your data trains models used for other customers. Ask for a no-training clause.

3.2 Green: you own the outputs. Just don't assume they are unique.

8.1 Yellow: common "as is" language, and the reason human review is mandatory.

8.2 Red: you carry the vendor's privacy risk. Push for mutual indemnification.

12.1 Yellow: limited audit rights. Ask for a current SOC 2 report.

12.2 Yellow: 30 days to export is workable; keeping fragments indefinitely is not. Ask for a deletion certificate and a limit on backup retention.

 FILLABLE

Board Brief

A one-page memo that secures formal authorization and fiduciary oversight. Frame it as protecting the mission, not restricting staff: boards approve protection; they hesitate on restriction.

To: Board of Directors **From:** **Date:**

Subject: Adoption of our AI Use Policy

1. Summary In two or three sentences: what we are asking the board to do, and why now. Try opening with: "We are not adopting AI. We are governing the AI already in use."

2. What we found What the Shadow Kitchen Inventory showed: which AI tools staff already use, and the risks.

3. What the policy does Our three tiers, named reviewers, and vendor scorecard, in plain language.

4. What the pilot showed The two-week beta: which department, what worked, what we changed.

5. Resources Staff time, any license costs, and who leads. Say so if no new budget is needed.

6. How we will measure it Two measures the board can see at the 90-day review.

For example: share of external AI-assisted documents with a named reviewer on record; staff hours returned to direct service.

7. THE ASK · SELECT ONE

- ◯ **For approval:** Adopt the AI Use Policy, effective
- ◯ **For review:** Send feedback on the draft policy by
- ◯ **For information only:** No action needed at this meeting.

FILLABLE

AI Use Statement

A short public paragraph for your website, annual report, and grant applications. Edit the model language into your own words, then publish it in week 8.

Organization:

1. Our intent

Model: We use AI tools to [draft grant reports and summarize meeting notes]. Our goal is not to replace human connection, but to buy back staff time for [helping families find and keep stable homes].

2. Our data promise

Model: We protect the privacy of the people we serve. We do not enter client records, HMIS data, or other personally identifiable information into public AI tools, and we use sensitive data only in tools approved under our policy.

3. Our human review

Model: AI is a tool, not a decision-maker. A named staff member, such as our [Program Director], reviews every AI-assisted document for accuracy, bias, and tone before it reaches a client, funder, or partner. AI assists; humans author.

4. Our transparency

Model: We label materials created substantially with AI, and we will explain how we use AI to anyone who asks.

Publish it: ☐ Website ☐ Annual report ☐ Grant applications ☐ Staff handbook

Test it before you publish. Swap drafts with a colleague from another organization and ask: if you were a tenant or donor reading this, would you trust us? A working example is on page 18.

FILLABLE

Provenance Log

A running record of which tool touched what data, verified by whom. It is the cheapest artifact to keep and the most persuasive when a funder asks how you know an AI-assisted report is accurate. Keep the live version in a shared spreadsheet with these columns.

DATE	TOOL + VERSION	OUTPUT	DATA USED (TIER)	VERIFIED BY	CHECKS PERFORMED	RESULT
EXAMPLE *9/30/26*	*[Approved tool, version]*	*Q3 funder report narrative*	*Outcomes data, no PII (Yellow)*	*[Name], Program Director*	*Every figure checked against HMIS export; tone check*	*Passed*

FOR EACH TOOL IN THE LOG

- ☐ Vendor scorecard on file (page 10), with the contract reference noted.
- ☐ Input data stripped of PII, or kept inside a tool approved for sensitive data.
- ☐ A next review date set to recheck accuracy as the tool or its terms change.

When a program officer asks, “How do you know this is accurate?” this log shows which tool was used, who signed off, and what they checked.

FILLABLE

90-Day Policy Review

Put this on the calendar the day you launch. The tools change quickly, and the review keeps your policy from going stale.

Review date: ______ **Led by:** ______ **Attendees:** ______

1. Paper trail

- [ ] Every external AI-assisted output has a Provenance Log entry.
- [ ] No output reached a client, funder, or partner without a named reviewer's sign-off.
- [ ] No Red-tier data went into a public or non-contracted tool.

2. Tools and vendors

- [ ] We reran the Shadow Kitchen survey and added new tools to the inventory.
- [ ] We rescored any vendor whose terms or features changed.

3. Quality and equity

- [ ] We tracked how often outputs needed significant correction.
- [ ] We ran the Equity Screen on at least one live use, with no unresolved flags.
- [ ] We reviewed any complaints or near-misses.

4. Policy fit

- [ ] We collected the friction points staff reported.
- [ ] We added or moved tier examples where staff were unsure.

DECISION

○ **Keep the policy as is** ○ **Update the policy** ○ **Escalate to the board**

Next review date: ______

If a tool fails the review: stop using it, ask the vendor to delete your stored data and confirm in writing, and document what happened for the board.

FILLABLE

Your First Monday

Three actions to take this week. Your 8-week goal: a documented governance framework ready for grant proposals and your board.

- [ ] **Identify the AI tools already active in your organization.**

- [ ] **Draft one sentence of your AI Use Statement.**

- [ ] **Name your human-in-the-loop reviewer.**

The Pocket Decision Guide

Three questions before you paste data into a prompt or send an AI-assisted draft.

1. **Data.** Which tier is this? If it includes client records, HMIS data, or case files, stop: no public tools.
2. **Review.** Who is the named Head Chef who will check it before it goes out? No name, no send.
3. **Trust.** If the person this is about saw how it was made, would it honor them? If not, write it yourself.

AI MISSION FILTER

Stop and check:

DATA	Is this non-sensitive?
REVIEW	Who checks facts and tone?
TRUST	Does this honor our mission?

AI assists. Humans author. Every external output gets a name in the Provenance Log.

Screenshot this card for your phone.

> **"You aren't asking for permission to experiment. You're demonstrating you've already governed the risk. That is leadership."**
>
> **Your organization doesn't need to be perfect. It needs to be governed.**

FILLABLE

Team Drill: The Screening Tool Request

A 30-minute exercise for your leadership team or AI working group. Run it before you approve any new AI tool, or as practice at a staff retreat.

The request. Your intake team wants to pilot an AI tool that scores applicants for rental assistance and suggests a waitlist order. The vendor says it predicts “likelihood of a successful tenancy.”

The complication. The model is proprietary. A peer agency using it noticed that applicants from a few zip codes rarely score near the top, regardless of need.

Split into groups of four. Each person takes one role.

The Innovator Intake Manager	Efficiency and capacity	*“We have a long waitlist and two intake staff. This could save weeks.”*
The Ethicist Program Lead	Equity and bias	*“Are we rebuilding redlining with better software?”*
The Guardian Data or IT Lead	Privacy and security	*“Where does applicant data go, and did applicants consent?”*
The Strategist Executive Director	Reputation and mission	*“If an applicant learned how they were scored, would they trust us?”*

Decide as a group

1. Transparency: can we explain a score to the applicant it affects? ○ Yes ○ No

2. Equity: run the Equity Screen (page 4). What did it flag?

3. Vendor: what does the Vendor Scorecard (page 10) say? ○ GO ○ CAUTION ○ STOP

4. Red-Face Test: if this tool were the headline of our annual report, would we be proud or defensive?

COMMITTEE DECISION

○ **Approve** ○ **Reject** ○ **Approve with conditions**

A note on screening tools: tools that score or rank housing applicants can raise fair housing concerns. Involve counsel before piloting any tool that influences who gets housing.

CREDITS · DECK SLIDE 19

Sources and Acknowledgments

The frameworks in this toolkit were developed by Kindred Impact LLC for the “AI for Impact” workshop series. Several concepts build on the publicly available frameworks below.

PRIMARY FRAMEWORK

Kindred Impact LLC (Steven J. Casey)

Head Chef / Line Chef model · Human-in-the-Loop Mission Filter · Shadow Kitchen · Equity Screen for Housing Data · AI Tool Request Form · Traffic-Light Vendor Scorecard · 8-Week Governance Blueprint · AI Use Statement · Board Brief · Provenance Log · Pocket Decision Guide

COMPANION RESOURCE

ANB Advisory Group and NTEN

Afua Bruce and Rose Afriyie, *AI Policy Template: AI Governance Framework for Nonprofits* (September 2024). Data-sensitivity tiering, overarching AI principles, and general-use guidelines draw on this template. Download the full template from NTEN’s AI for Nonprofits resources at nten.org, or from anbadvisory.com.

ADDITIONAL FRAMEWORKS AND RESEARCH

1. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1. January 2023. doi.org/10.6028/NIST.AI.100-1. Data provenance, transparency, and privacy practices.
2. U.S. Department of Agriculture, Food and Nutrition Service. *Framework for State, Local, Tribal, and Territorial Use of Artificial Intelligence for Public Benefit Administration*. April 29, 2024. fns.usda.gov/framework-artificial-intelligence-public-benefit. Equity and access principles. Issued under Executive Order 14110, which was rescinded in January 2025; check the page for its current status.
3. Microsoft. Responsible AI principles. microsoft.com. The fairness principle as adapted in the ANB/NTEN template.
4. The HIPAA Journal. “Is ChatGPT HIPAA Compliant?” hipaajournal.com/is-chatgpt-hipaa-compliant. Generative AI and protected health information; the article is updated as vendor offerings change.
5. Frontegg and the Cloud Security Alliance. Guidance on multi-factor authentication, least privilege, just-in-time privilege, and access reviews, as cited in the ANB/NTEN template.

AI USE STATEMENT FOR THIS TOOLKIT

AI assists. Humans author. This toolkit was drafted and designed with help from Claude, an AI model made by Anthropic, working from Kindred Impact’s workshop materials. Steven J. Casey reviewed the content and is accountable for it. The toolkit contains no AI-generated images.

KINDRED IMPACT LLC
Strategic Philanthropy • Local to Global

Steven J. Casey, Principal, Kindred Impact LLC
steve@kindredimpactllc.com · kindredimpactllc.com · 773.960.8242

Ask about the full two-day “AI for Impact” executive workshop for your board and leadership team.

Thank you to Housing Action Illinois and the 2026 Housing Matters Conference. ·