

AI POLICY IMPLEMENTATION CASE STUDY

From AI anxiety to operational practice

Implementing human-centered AI policy in libraries

ROBIN HASTINGS

Human-centered technology strategy for libraries and knowledge organizations

A policy implementation system, not a policy file

Robin Hastings designed and delivered a practical method for libraries responding to AI already present in staff work, patron questions, and vendor platforms. The method paired durable policy language with guided drafting, staff training, tool governance, and a path from informal experimentation to accountable practice.

3 x 90	14	10
workshop sessions	registrations	records marked attended

THE PROBLEM

Libraries did not need another abstract debate about whether AI would arrive. It was already showing up through staff experimentation, patron use, search and office tools, and vendor features. The operational gap was shared guidance: what work AI may assist, what information it may never receive, where human judgment must remain decisive, who reviews tools, and how staff learn to apply the policy.

THE EVIDENCE BOUNDARY

This case study describes the design and documented delivery of the method. The program was designed to leave participating libraries with a policy draft and implementation plan. Available records do not verify how many institutions completed board approval or adopted a final policy; those outcomes are not presented as measured impact.

Constraints became design requirements

The implementation had to work across organizations with different sizes, staffing levels, governance cycles, technical capacity, and appetite for experimentation. It also had to remain useful as product names and features changed.

Concern	Design response
Staff fear or uncertainty	Frame policy as permission with boundaries, then practice realistic decisions.
Patron privacy	Make data minimization and approved-tool rules non-negotiable.
Accuracy and credibility	Require verification of facts, citations, calculations, and claims.
Leadership accountability	Name owners, approval paths, escalation points, and review dates.
Vendor-created risk	Move AI questions into procurement, renewal, and product review.
Fast-changing tools	Keep principles durable; move product details into procedures.

The policy could not become a catalog of current tools. Durable principles belonged in policy. Tool-specific instructions, examples, and approval steps belonged in procedures and training.

Human-AI-Human

The workflow places professional judgment before and after AI assistance. The final human step is not ceremonial; it is where responsibility lives.

1	2	3
DEFINE	ASSIST	REVIEW AND OWN
A person sets the purpose, task, context, constraints, and acceptable inputs.	AI produces options, drafts, patterns, summaries, or other bounded assistance.	A person checks facts, privacy, bias, tone, accessibility, and consequences, then decides what to use.

THREE SAFEGUARDS

- **Privacy:** Keep patron, employee, credential, and confidential organizational data out of unvetted systems.
- **Accuracy:** Verify facts, citations, calculations, and claims before use.
- **Judgment:** Keep meaningful decisions and accountability with people.

None works alone. An accurate answer can still violate privacy. A private answer can still be wrong. A carefully edited answer can still represent an inappropriate delegation of judgment.

Policy becomes real through four connected layers

01	POLICY Purpose, scope, principles, responsibilities, and boundaries.
02	PROCEDURES Tool review, approved uses, documentation, incidents, and escalation.
03	TRAINING Practice applying the policy to daily decisions and patron support.
04	REVIEW Feedback, owners, vendor changes, incidents, and scheduled reassessment.

Risk-based guardrails

LOW RISK	REVIEW REQUIRED	PROHIBITED
Brainstorming, rewriting, generic drafts, search-term expansion.	Public-facing content, analysis, recommendations, substantial AI-generated text.	Confidential data, decisions about people, deception, unreviewed high-stakes output.

From template to implementation

Define the foundation

Map current use; compare short and long templates; draft purpose, scope, and guiding principles.

Make operational decisions

Draft approved and prohibited uses, patron guidance, data protection, and tool governance.

Make the policy usable

Peer-review language, identify training needs, connect related policies, and name approval and rollout steps.

DOCUMENTED TRANSFERABILITY

The framework was also adapted into a single-session ByWater Solutions webinar and later distilled into a 19-slide conference presentation. That progression demonstrates a method designed to travel from hands-on institutional work to executive briefing and professional education.

A 90-day path from draft to practice

DAYS 1-30 Discover and define • Inventory staff, patron, and vendor uses. • Name an accountable owner and review group. • Draft purpose, scope, principles, and data boundaries.	DAYS 31-60 Test and prepare • Test realistic scenarios. • Review privacy, retention, training-use, and deletion terms. • Create tool-request and escalation procedures. • Train supervisors and early adopters.	DAYS 61-90 Publish and reinforce • Move through the approval path. • Run hands-on staff exercises. • Publish tool and disclosure guidance. • Set feedback channels and review dates.
--	---	--

Staff training mirrors the decisions the policy requires

- Compare an AI answer with trusted sources and document what needed correction.
- Rewrite a prompt to remove personal or confidential data.
- Decide whether a task is low risk, requires review, or should be prohibited.
- Review a vendor's privacy and data-retention terms.
- Practice disclosing substantial AI assistance in public-facing material.
- Work through a data incident or misleading-output scenario.

Sample provisions

These excerpts are starting points from Robin Hastings's workshop templates. They must be adapted to local policy, law, staffing, and governance.

AI will be used to support human judgment, not replace it. Human oversight is required in all meaningful decisions.

No confidential or personally identifiable patron data may be entered into non-approved AI systems.

AI outputs must always be reviewed before use.

This policy is considered a living document.

Why the wording matters

The provisions establish accountability and durable boundaries without attempting to predict every tool. Procedures can then specify approved systems, request paths, settings, and review dates.

RESULTS

Documented delivery signals

- Three 90-minute workshop sessions were delivered on August 4, 7, and 14, 2026.
- The registration record contains 14 registrations; 10 records are marked as having attended workshops.
- Participants received an integrated workbook, three session decks, short and long policy templates, guided drafting prompts, and implementation-planning tools.
- The second workshop session ran the full 90 minutes, a qualitative signal that discussion and exercises sustained engagement.
- A separate ByWater Solutions webinar was delivered successfully on August 12, 2026.
- The material was subsequently distilled into a 19-slide conference presentation covering human oversight, risk-based guardrails, vendor review, training, and a 90-day implementation path.

Lessons learned

- 1 Policy publication is the beginning of implementation. Staff need examples, practice, feedback, and reinforcement.
- 2 Durable policy and adaptable procedure are different jobs. Tool names age quickly; principles should not.
- 3 Privacy, accuracy, and judgment must be taught together.
- 4 A usable policy makes ownership visible: who approves, where questions go, and when review happens.
- 5 Honest evidence builds more credibility than inflated metrics.

This work combines policy expertise, systems design, adult learning, facilitation, change management, and practical execution. The result is not only a policy template. It is a repeatable way to translate public-service values into daily decisions about AI.

SOURCE RECORD

Robin Hastings, *Creating Generative AI Policies: A Guide for Public and Academic Libraries*; CALL workshop project index, workbook, policy templates, and Session 1-3 decks; AI Policies for ByWater project index and presentation; Human-Centered AI - MPLA/NLA 2026 presentation; HQ end-of-day records for August 7 and August 12, 2026.