

# What Boards Need to Know About AI

A practical governance briefing for library trustees, directors, and public-sector leaders

**AI is not just a software question. It is a governance question: who may use it, what data is off limits, how outputs are verified, and how public trust is protected.**

## Prepared by Robin Hastings

Library technology consultant, author, teacher, and AI policy practitioner

June 2026 | [rhastings.net](https://rhastings.net)

# Executive Summary

Boards do not need to become AI technicians. They do need to make sure the library has a defensible way to use AI without compromising privacy, accuracy, equity, staff judgment, or community trust.

|                                   |                                                                                                                      |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Board responsibility</b>       | Set direction, approve policy, ask risk-aware questions, and ensure AI use aligns with mission and values.           |
| <b>Operational responsibility</b> | Directors and staff translate policy into procedures, training, approved tools, vendor review, and local workflows.  |
| <b>Public responsibility</b>      | Libraries should help patrons understand AI strengths, limits, privacy risks, misinformation, bias, and ethical use. |

## The essential board-level position

A good AI policy should protect people without freezing innovation. It should encourage responsible experimentation while drawing bright lines around confidential data, automated decisions, surveillance, high-risk outputs, vendor data use, and staff accountability.

**Think of AI policy as the seatbelt, not the brake pedal. It makes movement safer.**

## The Human-AI-Human test

Use this simple test for board conversations: a human frames the task, AI assists with a bounded activity, and a human verifies, contextualizes, and owns the final decision or public output.

|                                                 |                                                                    |                                                                                   |
|-------------------------------------------------|--------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Human frames</b><br>Purpose, context, limits | <b>AI assists</b><br>Drafting, sorting, summarizing, brainstorming | <b>Human verifies</b><br>Accuracy, privacy, bias, local fit, final responsibility |
|-------------------------------------------------|--------------------------------------------------------------------|-----------------------------------------------------------------------------------|

# Seven Things Boards Need to Know

## 1. AI is already in the organization.

Staff may be using public chatbots, vendor tools may be adding AI features, patrons may ask for help, and AI may be embedded in search, translation, transcription, or productivity systems. No policy does not mean no AI. It means undocumented AI.

## 2. Privacy is the highest-stakes issue.

Prompts, uploaded files, chat transcripts, and outputs are data governance decisions. Patron records, personnel files, donor records, protected institutional records, and sensitive community information should not be entered into unapproved systems.

## 3. Accuracy requires human verification.

AI output often sounds polished before it is correct. Policies should require higher verification standards for public information, legal or medical topics, disciplinary matters, policy work, and anything affecting patron rights or public trust.

## 4. Staff need permission and boundaries.

A policy should say what is allowed, what requires approval, and what is prohibited. This creates confidence for responsible use instead of leaving innovation to the boldest staff member in the building.

## 5. Vendors matter as much as chatbots.

AI features are increasingly embedded in platforms libraries already license. Boards should expect vendor review before enabling features that process, retain, analyze, or train on library or patron data.

## 6. Patron education is part of the mission.

Libraries should not only manage internal AI use; they should help patrons understand AI literacy, misinformation, bias, privacy, employment and school use, and ethical participation in AI-shaped information environments.

## 7. AI policy must stay alive.

AI policies should be reviewed more often than many traditional policies. The durable policy can stay principles-based; procedures, tool lists, and staff guidance can change more frequently.

# A Board-Level Risk Map

Boards should not try to evaluate every prompt or product. They should make sure leadership has a consistent method for classifying risk and deciding when safeguards are required.

| Risk area                 | What can go wrong                                                                                         | Board-level safeguard                                                                                                  |
|---------------------------|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| Privacy + confidentiality | Patron, staff, donor, student, or personnel data gets entered into a system the library does not control. | Adopt data-use rules; require approved tools for sensitive workflows; ask for vendor retention and training-use terms. |
| Accuracy + authority      | AI-generated content is wrong, outdated, biased, or fabricated but appears credible.                      | Require human verification and risk-based review before public use.                                                    |
| Equity + accessibility    | Tools produce biased, exclusionary, inaccessible, or culturally inappropriate results.                    | Require accessibility and bias review in tool approval and staff training.                                             |
| Labor + accountability    | AI quietly changes staff expectations, job duties, or performance evaluation without governance.          | Prohibit AI-only personnel decisions; keep human responsibility explicit.                                              |
| Vendor + procurement      | AI features are enabled inside existing platforms without privacy, security, or opt-out review.           | Require AI review language in procurement and renewal workflows.                                                       |
| Public trust              | Community members do not know when AI affects library services, content, or decisions.                    | Set transparency standards and maintain a human point of contact.                                                      |

**Board question:** What information should never leave library-controlled systems without review?

# The Five Decisions Boards Should Make

## 1. Adopt a clear AI policy scope.

Cover staff-facing, patron-facing, administrative, instructional, and vendor-embedded AI. Avoid naming only one product.

## 2. Define prohibited uses.

At minimum: confidential data in unapproved tools, AI-only employment or disciplinary decisions, unauthorized surveillance, discriminatory uses, and unverified high-risk public outputs.

## 3. Require a tool approval pathway.

Small libraries can use director approval and a simple checklist. Larger organizations may need IT, administration, legal, accessibility, and privacy review.

## 4. Fund staff and patron training.

A policy without training is a decorative binder. Staff need practical scenarios, verification practice, privacy rules, and shared vocabulary.

## 5. Set a review cycle.

Review the policy at least annually, and consider semi-annual review while tools, vendor features, and legal expectations remain volatile.

## Ten questions trustees can ask without becoming technologists

- |                                                                            |                                                                                            |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| 1. Where is AI already being used by staff, patrons, or vendors?           | 6. Can patrons opt out of AI-enabled services or recommendations?                          |
| 2. What data are staff forbidden to enter into public or unapproved tools? | 7. What training do staff receive before using AI in library work?                         |
| 3. Who approves new AI tools or AI features in existing platforms?         | 8. How are bias, accessibility, and language issues evaluated?                             |
| 4. How are AI outputs verified before public use?                          | 9. What vendor assurances do we require about retention, training, deletion, and security? |
| 5. When do we disclose AI use to patrons or the public?                    | 10. When will this policy come back to the board for review?                               |

# A 30-60-90 Day Path Forward

|                      |                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>First 30 days</b> | Inventory staff uses, patron questions, vendor AI features, and current policy gaps. Issue interim guidance: no confidential data in unapproved tools; human review required for public output. |
| <b>Next 60 days</b>  | Draft or adapt an AI policy. Map AI language to privacy, internet use, personnel, collection development, procurement, patron behavior, and technology policies.                                |
| <b>Next 90 days</b>  | Approve policy, publish staff guidance, train staff with real scenarios, create a public-facing patron education plan, and schedule the first review.                                           |

## Suggested board motion language

The Board directs the Library Director to develop and maintain AI use guidance, including an AI policy, approved-use procedures, staff training, patron education resources, and vendor review practices, with policy review at least annually and sooner when major legal, technical, or service changes require it.

## Selected sources and evidence base

- Robin Hastings, AI Policy in Libraries materials and forthcoming ALA Special Report; Human-AI-Human policy framework.
- NIST, Artificial Intelligence Risk Management Framework, 2023.
- American Library Association, Library Privacy Guidelines and Privacy: An Interpretation of the Library Bill of Rights.
- American Library Association, Library Bill of Rights, Article VII on privacy and confidentiality.
- Federal Trade Commission, guidance on AI privacy, confidentiality commitments, and deceptive AI claims.
- UNESCO, Recommendation on the Ethics of Artificial Intelligence.

## About this series

The Hastings Executive Briefing Series translates emerging library technology issues into board-ready governance language. Each briefing is designed for leaders who need practical oversight questions, not vendor hype or technical fog machines.