

AI Privacy

The decision facing the board

What information may staff use with AI, for which purposes, and under whose approval?

What boards need to know

An AI policy needs examples staff can apply during ordinary work. Drafting a program announcement from public information and summarizing an identifiable patron's reference conversation should not receive the same automatic approval.

NIST's Generative AI Profile identifies privacy risks including unauthorized disclosure and the possibility of connecting data back to individuals. It recommends evaluating third-party AI risks and monitoring generated content for sensitive information. These risks warrant reviewing the particular use, data, and product rather than treating every AI tool as equivalent. [Source: NIST AI 600-1, section 2 and actions GV-6.1-009 and MP-4.1-001.](#)

ALA's vendor guidance recommends collecting only the user data needed for a specific task and addressing retention and permitted use in agreements. Apply those questions when reviewing AI services. [Source: ALA, Data Integrity and Security and Agreements sections.](#)

Questions to ask

- Which uses are approved, and how will staff recognize an unapproved use?
- Could an uploaded document or connected account expose information beyond the intended task?
- What do the applicable terms say about storage, human access, model training, and deletion?
- Who checks the output and handles a suspected disclosure?
- Can staff accomplish the task with public or fictional information, or without AI?

A practical next step

Ask the director to bring three ordinary workflows to a policy discussion: one suitable for approved AI assistance, one requiring further review, and one that should remain outside the tool. Document the reason for each boundary and who can authorize a change.

Illustrative discussion: A staff member wants help responding to a reference question. A proposed safer starting point is to create a fictional example that captures the research need without copying the patron's correspondence. Removing a name alone should not be treated as proof that the original message is safe to share.