

State of AI in Libraries Report

U.S. public libraries | September 2026

Research cutoff: September 16, 2026.

Executive findings

Artificial intelligence is becoming a practical question for public libraries in more than one place. It appears in staff tools, patron questions, and the materials libraries select. A library can decline to generate content with AI and still need to respond to AI-generated books, vendor features, and requests for help evaluating information.

The strongest response connects familiar library responsibilities to specific decisions: selecting credible materials, protecting patron information, supporting staff judgment, and helping people understand what they are reading. The policies and examples in this report show how libraries are beginning that work. They also reveal where more evidence is needed.

- Capacity shapes the conversation. In OCLC's 2026 survey, 54% of small public library directors reported no structured institutional discussion of AI, compared with 33% of medium and 21% of large public library directors. These findings describe respondents, not a census. [1]
- Governance needs an operating procedure. Kenosha connects policy to a documented vetting process; Seattle defines approval and output review. Both offer examples of making responsibility explicit. [4,12-13]
- Collection development deserves sustained attention. The challenge includes evaluating individual titles and addressing material supplied through digital platforms. Kent District Library's account shows why vendor relationships belong in that work. [15]
- Staff education should develop judgment. KU's bias exercises and practitioner resources on book evaluation offer adaptable ways to examine claims, sources, and representation. Their value lies in the questions they help staff ask. [14,16-17]

Recommendation: Give those responsibilities time, ownership, and a record of decisions - the goal is a defensible service or collection decision that the library can explain and revisit. Adoption alone is an inadequate measure of progress.

1. Adoption: distinguish use from implementation

OCLC surveyed 698 U.S. public and academic library directors in March 2026. Its August release reports that 14% of small public library directors used AI daily and 26% did not use it at all. Across the combined sample, 7% reported providing clear organizational direction and expectations. That last figure includes academic libraries. [1]

These measures describe different things. Personal use is not evidence that a library operates an AI reference service. A policy does not show how many staff follow it. Purchasing a product with an AI feature does not show whether that feature is enabled.

Four categories make this activity easier to understand:

- Individual use: staff members experimenting with tools in their work.
- Authorized internal use: approved workflows, named reviewers, and data restrictions.
- Public service use: an implemented service that patrons encounter.
- Public education: instruction about AI, including its limitations and the choice not to use it.

These categories describe different responsibilities rather than stages of maturity. A library could offer excellent AI literacy programming while declining to deploy an AI chatbot.

What can be said about trends

The reviewed sources document policy and programming activity in 2024-2026. They do not provide a consistent annual U.S. public-library adoption series. This report therefore identifies emerging practices without assigning a national growth rate.

A credible annual comparison will require repeated measures with consistent populations and definitions. For now, the evidence is strongest when attached to a specific practice, policy, or institution.

2. Policy and governance: several workable structures

ALA announced adoption of its AI guidance in July 2026. The guidance covers standalone tools and embedded features, supports meaningful human oversight, and recognizes that a documented decision to limit or decline use can be appropriate. This constitutes ALA's professional guidance, but shouldn't be used as evidence, by itself, that local libraries have implemented it yet. [2-3]

Seattle: specify approval and review

Seattle's policy, adopted July 31, 2025, requires an approval process developed by an AI Governance Committee. It expressly includes free tools and AI features added to otherwise approved software. Generative output used officially requires knowledgeable staff review and documentation. Published generated material requires attribution, with an exception for limited editing of human-written material. [4]

Kenosha: turn policy into a vetting process

Kenosha Public Library's August 2024 policy requires staff to coordinate with Computer & Network Services before using an AI tool. It connects the review to an intended use, assigns responsibility, and limits sensitive information to vetted systems within an approved scope. It also calls for a public list of approved tools and their uses. [12]

The linked procedure makes the process concrete. The Head of IT leads; the requester explains the service need; IT staff or a Digital Strategy Librarian considers technical risks. The group develops scenarios, assesses likelihood and severity, and records mitigations and a decision to adopt, reject, or substitute a tool. Its review categories include bias, privacy, security, accessibility, employment, and explainability. [13]

The procedure proposes meetings in successive weeks and communication after the final meeting. It identifies that timetable as a best-case scenario, not a measured turnaround. The important feature is the connection between a request and a documented decision. A smaller library might adapt the sequence while seeking technical assistance from its system or consortium; this report does not assume that every library has Kenosha's staffing. [13]

Johnson County, Kansas: adapt existing government policy

Johnson County Library's policy, effective September 11, 2025, adopts the county's AI policy with library-specific substitutions and provisions for board approval of significant exceptions. This is a useful example of governance through an existing parent organization. The library's published rule establishes that governance relationship; detailed operating requirements also depend on the county policy. [5]

Provo: make staff and patron boundaries explicit

Provo's policy, approved July 8, 2026, permits specified work support, including brainstorming and program development, while prohibiting AI processing of patron records, circulation data, and confidential information. It assigns staff responsibility for final work and alerts patrons to third-party privacy limitations. Its verification instruction is qualified by "where possible," illustrating why the exact wording of a review requirement matters. [6]

3. Staff concerns and privacy belong in the workflow

WebJunction's February 2025 newsletter poll received 238 responses. Its summary records concerns about creator rights, employment, environmental effects, practical uses, and patron education. This self-selected group is useful for identifying questions staff raise; it cannot establish how common those concerns are among all U.S. public library workers. [10]

Training should make room for those questions. Reviewing an output takes time and expertise; a demonstration that produces text quickly is not a complete productivity measure.

Recommendation: Ask staff to evaluate a real, limited workflow: what improves, what new work appears, what information moves outside the library, and who remains responsible when the output is wrong.

Bias review is a skill staff can practice

The University of Kansas Center for Teaching Excellence provides exercises that compare chatbot responses, vary prompts, examine stereotypes, and check whether cited sources exist. Its materials were developed for higher education and updated in March 2024. They are useful here as teaching methods, not as a current benchmark of particular AI products. [14]

For public-library staff, a short exercise could begin with a fictional request for a community reading list. Compare responses to the same request, inspect whose experiences are represented, and verify suggested titles. Then change one element of the request and discuss what changed. This is an adaptation of KU's comparison approach, not a tested public-library curriculum. [14]

The discussion should end with a work decision: what can be used, what needs correction, and what should be rejected? Save the examples and reasoning so a colleague can understand the review. An AI tool's own assurance that its output is fair should not be the basis for approval.

Privacy review must include embedded features

ALA's guidance calls attention to feature defaults, vendor data collection, retention, training use, subprocessors used by vendors, library data deletion, and exit rights. It recommends continuing review when systems change. That makes privacy an operational responsibility extending beyond the first purchase. [3]

Seattle treats unclear data classification as sensitive until clarified and approved. Provo prohibits specified sensitive data uses and says it will clear public AI-station data where feasible. The latter is a local device practice, not a guarantee that a third-party provider deletes its copies. [4,6]

Recommendation: For a proposed workflow, keep a short written record answering:

- What information will people enter, and what information is prohibited?
- Which organization receives it, and what controls has the library actually checked?
- Who reviews the result, and what authority do they have to reject it?
- How can a patron get human help?
- What event triggers another review or a decision to stop?

These questions support local review; they do not replace applicable legal or municipal requirements. This report does not attempt a state-by-state legal inventory or a technical audit of vendor privacy practices, nor does it claim any specific legal knowledge or advice.

4. Collection development in an AI publishing environment

Collection development brings the AI discussion back to the material a library puts in front of people. The questions concern authority, usefulness, representation, and stewardship of a limited budget. AI provenance may matter under local policy, but a selector still needs to examine the work itself and explain the decision.

Three practitioner resources offer complementary perspectives. Althea Christensen's guide organizes warning signs for acquisition work. Illinois Tech presents a shorter introduction for readers. Kent District Library describes how its selectors respond, including through a digital-content vendor. [15-17]

Define the policy before judging the title

Salt Lake County's collection policy, revised May 18, 2026, allows consideration of wholly or partly AI-created materials when they meet its selection criteria. The rule applies across formats. It makes the existing selection standard the basis of the decision. [7]

WebJunction's account of the library's policy work describes a work group, selection guidance, and training development, along with disagreements among staff and training delays. Writing a policy did not remove the need to work through its application. [11]

A local policy should distinguish generated text, illustrations, narration, translation, and limited assistance such as editing. Treating every use as the same can leave staff uncertain about what a rule actually covers.

Recommendation: State which uses affect selection, what evidence is sufficient, and how the library responds when provenance is unknown.

These are policy-design recommendations, not a claim that one position has become the national norm.

Use warning signs to investigate

Christensen, a Senior Library Assistant in Acquisitions and Technical Services at Humboldt County Library, advises examining author information, publication patterns, content, references, and reviews. Her June 2025 guide places particular weight on vague or contradictory information and sources that cannot be verified. [16]

Quinn McKissock's December 2025 Illinois Tech article offers a compact introduction: inspect cover details, investigate the author, and examine unusual publication statements. It cites Christensen's guide, so the two resources should not be counted as independent confirmation of a detection method. [17]

An unfamiliar author, sparse website, self-published edition, or unusual writing style warrants context. None establishes AI authorship by itself. Selectors should distinguish what they can verify about a book from what they suspect about its production. An inaccurate citation is a checkable defect; an impression that prose sounds artificial is a different kind of evidence.

A practical selection and review sequence

The following workflow synthesizes the practitioner resources with ordinary selection responsibilities. It is proposed guidance, not a scoring system validated by this report. [15-17]

- Identify the work. Confirm the edition, format, publisher, creator, and any available disclosure of AI use. Check that descriptive records match the item under consideration.
- Examine authority and substance. Look for relevant expertise, credible review coverage, and a sample substantial enough to assess whether the book delivers what it promises. Verify important factual claims and citations when concerns arise.
- Apply the library's criteria. Record the specific strengths or shortcomings: accuracy, depth, presentation, relevance, demand, and the applicable policy provision. Avoid substituting an unsupported authorship accusation for a selection judgment.
- Decide and document. Acquire, decline, seek more information, or route the item for further review. Record uncertainty where it remains; a decision should be understandable to another selector.
- Respond consistently after acquisition. Apply the library's established review, correction, and weeding procedures when a problem is reported. Explain the decision in terms of the policy and the material.

Digital platforms change what selectors control

Kent District Library's July 2026 article describes patron complaints about brief, poorly written AI-generated summaries on hoopla. KDL asked Midwest Tape to block that content. The account says blocking can use publisher identification or characteristics that suggest AI generation. It documents an intervention, not proof that the platform is free of generated material. [15]

KDL also distinguishes its individually selected Libby licenses from hoopla's pay-per-checkout collection. That distinction matters: the library's responsibility to respond remains, even when its ability to select or suppress individual titles depends on a vendor. [15]

Recommendation: Review the available controls at the collection or platform level. Ask who adds titles, how provenance is recorded, which exclusions the library can apply, and how reported errors or misleading records are handled. Determine whether a change affects one library's view or the shared service. Keep a record of the response rather than assuming a reported item has disappeared everywhere.

Protect credibility without making every purchase an investigation

Libraries need a proportionate process. Established review sources and routine sampling can inform ordinary decisions; conflicting records or substantive errors can trigger closer examination. Time spent investigating is part of the cost of acquiring the material. Shared documentation can reduce repeated work, provided it distinguishes verified defects from suspicion.

The same care applies to patron communication. Staff can explain why a particular work did not meet selection criteria without claiming certainty about how it was written. The purpose is to maintain a collection the library can stand behind while applying its policy consistently.

5. Public learning and implementation outcomes

Palo Alto: learning through a shared project

Palo Alto City Library's February 24, 2025 recap describes a middle-school AI hackathon held on December 22, 2024. Students worked in small teams with mentors, discussing AI and developing ideas related to community and environmental problems. The event is an earlier example reported in 2025. Its recap documents activity but does not supply a validated learning assessment. [8]

Programs tied to familiar community needs

WebJunction's June 2025 roundup describes Santa Clara's job-search instruction, Allen County's genealogy demonstration, Rocky River's scam-awareness program, and Pottsboro's grant-writing class. These examples are reported by WebJunction; their original event records were not all independently retrieved for this report. They illustrate subject choices, not proven effectiveness or current recurring offerings. [9]

Measure what the service is meant to accomplish

Recommendation: For a future case study, document the starting problem, people served, staff preparation, costs, review requirements, and observed results. A patron-learning program might assess whether participants can verify a generated claim or identify information they should not submit. An internal workflow might compare total staff time, including corrections, with its previous process.

6. Strategic outlook

These are evidence-informed editorial predictions for the coming 12-18 months, not measured trends or statistical forecasts.

Embedded AI will make product changes a governance issue

Seattle's explicit treatment of embedded features and ALA's review guidance suggest that approval will increasingly concern changes within existing products. Watch for revised contracts, feature inventories, and documented decisions to enable or disable functions. Confidence: moderate; the direction is plausible, but this review does not measure the rate of vendor change. [3-4]

Shared support will matter for libraries with limited capacity

The size differences in OCLC's findings suggest a need for support that reduces the work required to evaluate tools and establish local expectations. Watch for regional training, shared vendor reviews, and reusable procedures. Confidence: moderate about the need; uncertain about whether funding and staffing will meet it. [1]

Public education will remain a distinct service opportunity

The documented programs connect AI to existing interests such as jobs and genealogy. That suggests a continuing role for practical instruction even in libraries that limit internal AI use. Watch for repeat offerings, participant feedback, and assessed learning. Confidence: moderate; the examples demonstrate activity, not demand everywhere. [8-9]

Collection work will require closer vendor engagement

KDL's intervention suggests that local collection decisions will increasingly involve platform settings, disclosure, and vendor responses. The practical question is whether libraries can act on their policies across all formats. Watch for clearer provenance fields, documented reporting routes, and transparent limits on local control. This is a moderate-confidence prediction about the direction of work, not an estimate of how much generated content collections contain. [15]

What leaders can do now

Recommendation: Begin with one clearly described service problem and decide whether AI is appropriate to it. Set the data boundary and review responsibility before a pilot. Evaluate the full workload and service result, including errors and accessibility. Record a decision to continue, revise, or stop.

For associations and researchers, a useful next contribution would be comparable public-library-only evidence: institutional use separated from personal use, consistent size definitions, frontline staff perspectives, and measured outcomes. For vendors, the practical contribution is documentation that lets libraries make those decisions without guessing.

Methods and limits

This report uses targeted desk research completed September 16, 2026. Sources include U.S. public-library policies and procedures, library-authored accounts, association guidance, survey summaries, and practitioner teaching resources. KU and Illinois Tech contribute adaptable educational material rather than evidence of public-library adoption. Sources were chosen for relevance and inspectability, not through a systematic-review protocol. The source set overrepresents organizations that publish detailed material online. It is not a representative sample of public libraries.

OCLC's survey was fielded March 3-12, 2026 to 5,670 directors, with 698 responses; reported response rates were 16% for public libraries and 9% for academic/research libraries. The published summary does not provide subgroup denominators or size thresholds. Its aggregate uncertainty statement cannot be applied to individual size subgroups. [1]

Policy provisions are reported as written requirements, not audited behavior. Event recaps establish reported activity, not causal impact. WebJunction's newsletter poll provides qualitative context, not population estimates. This report offers no national adoption percentage, longitudinal growth calculation, return-on-investment estimate, or claim that AI improves service quality across public libraries. These are all excellent topics for future research, however.

Recommendations and predictions represent this report's synthesis. No interviews or technical audits were conducted. Codex assisted with research, drafting, and layout under my editorial direction. This report incorporates supplied notes and resources.

The book-evaluation guides provide indicators for investigation, not tested measures of detection accuracy. Historical descriptions of AI capabilities or prices in teaching resources are not treated as current product claims.

References

All sources accessed September 16, 2026. Numbered references support the adjacent claims in the report.

- [1] OCLC Research. August 18, 2026. What 698 Library Leaders Are Telling Us About AI.
<https://www.oclc.org/en/learn/perspectives/insider-insights-what-698-library-leaders-are-telling-us-about-ai.html>
- [2] American Library Association. July 16, 2026. ALA Council adopts Guidance on the Use of Artificial Intelligence in Libraries.
<https://www.ala.org/news/2026/07/ala-council-adopts-guidance-use-artificial-intelligence-libraries>
- [3] American Library Association. Guidance on the Use of Artificial Intelligence in Libraries. Current web text reviewed.
<https://www.ala.org/tools/standards-and-guidelines/guidance-use-artificial-intelligence-libraries>
- [4] Seattle Public Library. Adopted July 31, 2025. Library Use of Artificial Intelligence.
<https://www.spl.org/about-us/policies/library-use-of-artificial-intelligence>
- [5] Johnson County Library. Effective September 11, 2025. Artificial Intelligence Policies, ARM 50-40-10.
https://www.jocolibrary.org/plans-policies/arm_50-40-10/
- [6] Provo City Library. Approved July 8, 2026. Artificial Intelligence Use Policy. Publisher text retrieved through indexed search; direct page retrieval failed. <https://provolibrary.gov/artificialintelligenceusepolicy>
- [7] Salt Lake County Library. Revised May 18, 2026. Collection Management Policy, page 3.
<https://www.slcolibrary.org/policies/pdfs/Collection-Management-Policy.pdf>
- [8] Palo Alto City Library. February 24, 2025. AI Wonder World Hackathon Program Recap.
<https://library.cityofpaloalto.org/blogs/post/ai-wonder-world-hackathon-program-recap/>
- [9] WebJunction. June 19, 2025. How public libraries are teaching AI and digital literacy skills.
<https://www.webjunction.org/news/webjunction/ai-programming-for-patrons.html>
- [10] WebJunction. March 25, 2025. AI poll results: What we learned.
<https://www.webjunction.org/news/webjunction/ai-poll-results.html>
- [11] WebJunction. February 26, 2026. Considering AI in your collection policy: A practical framework for libraries.
<https://www.webjunction.org/news/webjunction/ai-collection-policy-framework.html>
- [12] Kenosha Public Library. Approved August 13, 2024. AI Policy, particularly pages 2-4.
<https://mykpl.info/wp-content/uploads/sites/120/2024/08/AI-Policy-V1-Final.pdf>
- [13] Kenosha Public Library. Artificial Intelligence Vetting Process. Procedure linked from its published policy; current document retrieved September 16, 2026. https://docs.google.com/document/d/1HbldZyYCFilsneo8Gwa-_Mvg-XldJ2De_XFy_AdyqUo/edit
- [14] University of Kansas, Center for Teaching Excellence. Updated March 2024. Helping students understand the biases in generative AI. <https://cte.ku.edu/addressing-bias-ai>
- [15] Abby Lynn and Gwennan Lawcock. July 8, 2026. Can You Spot an AI-Generated Book? KDL's Collection Experts Weigh In. Kent District Library. <https://kdl.org/blogs/post/can-you-spot-an-ai-generated-book-kdls-collection-experts-weigh-in/>
- [16] Althea Christensen, Humboldt County Library. Updated June 10, 2025. Tips for spotting AI books. File titled AI book spotting.docx; supplied by Robin and retrieved through Google Drive.
https://docs.google.com/document/d/1uL_BKwIzjdRkj83d7EoT1LUQKo2y9cnr/edit
- [17] Quinn McKissock. December 5, 2025. Spotting AI-Generated Books. Illinois Institute of Technology Library. Publisher page retrieved directly after web-tool retrieval failed. <https://guides.library.iit.edu/scholcomm/blog/spotting-ai-generated-books>