Public Library AI Policy Analysis: A Comprehensive Report

Prepared: March 31, 2026 **Scope:** 15 public library AI policies from the United States and Canada
Method: Structured policy survey instrument (66 questions, 9 sections) applied to each policy; findings synthesized across all libraries

---

Table of Contents

---

1. Overview: The State of Public Library AI Governance

Public libraries across North America are grappling with a new reality: artificial intelligence tools are no longer futuristic. Staff are using them now — to draft marketing materials, answer patron questions, create program flyers, and streamline back-office work. ChatGPT, image generators, AI-enhanced search, and automated chatbots have arrived in library workflows whether or not a policy exists to govern them.

This report analyzes how 15 public libraries have responded to that reality. Between 2023 and 2025, these libraries — ranging from small-town New Hampshire libraries serving a few thousand residents to metropolitan systems in Toronto with hundreds of branches — developed and adopted formal AI policies. We applied the same 66-question survey to each policy, generating 990 individual data points that allow direct comparison across the field.

The results reveal a field that is earnest but uneven. Almost every library studied has done the right things on the basics: they acknowledge both the promise and the risks of AI, they tell staff not to enter patron data into public AI tools, and they require human review of AI outputs. But beyond those fundamentals, the quality and depth of AI governance varies enormously. A handful of libraries have developed genuinely sophisticated frameworks that could serve as models for the field. Many others have produced short, staff-conduct documents that leave significant governance gaps unfilled. And one library — Johnson County Library in Kansas — adopted a county government policy entirely by reference, producing a document that is largely a formality with almost no library-specific substance.

This is not a criticism. These policies were drafted in fast-moving circumstances, often without dedicated legal or technology resources, often by library directors and staff who are experts in library service rather than AI governance. The fact that 15 libraries developed formal policies at all — most before the field had established clear norms — reflects genuine institutional commitment. But understanding where these policies fall short is essential for the next generation of library AI governance.

The 15 Libraries Studied

Library	Location	Adoption Date
Crandall Public Library	Glens Falls, New York	July 23, 2025
DeKalb Public Library	DeKalb, Illinois	September 2024
Hastings Public Library	Hastings, Michigan	October 7, 2024
Hinsdale Public Library	Hinsdale, Illinois	February 25, 2025
Holderness Free Library	Holderness, New Hampshire	December 11, 2023; updated February 10, 2025
Houston County Public Library System	Houston County, Georgia	April 11, 2024
Johnson County Library	Johnson County, Kansas	September 11, 2025
Kenosha Public Library	Kenosha, Wisconsin	Not dated
Naples Library	Naples, New York	March 19, 2025
Oakville Public Library	Oakville, Ontario, Canada	August 21, 2025
Schaumburg Township District Library	Schaumburg, Illinois	August 18, 2025
St. Charles Public Library	St. Charles, Illinois	September 2024
Toronto Public Library	Toronto, Ontario, Canada	January 27, 2025
White House Public Library	White House, Tennessee	Not dated
Wolfeboro Public Library	Wolfeboro, New Hampshire	May 12, 2025

These libraries span urban, suburban, and rural communities; seven U.S. states and two Canadian provinces; communities with populations from a few thousand to millions; and adoption dates ranging from late 2023 to mid-2025.

---

2. The Policy Landscape: Who Has What

Before diving into specific findings, it helps to understand the sheer range of what counts as a "library AI policy." These 15 documents are not all the same type of thing. They fall into roughly four tiers based on depth, sophistication, and coverage.

Tier 1 — Comprehensive Governance Frameworks (3–4 libraries)

Kenosha, Wisconsin; Toronto, Ontario; Oakville, Ontario; and Schaumburg, Illinois have written what can genuinely be called governance frameworks. These documents define their terms, reference specific laws and professional standards, establish formal vetting processes with named criteria, address transparency and disclosure with some specificity, require formal staff training, and deal with subjects like algorithmic bias and (in Schaumburg and Toronto) even environmental sustainability. Kenosha stands apart even within this tier for having built out a complete patron AI literacy program as an appendix.

Tier 2 — Substantive Mixed Documents (3–4 libraries)

Crandall (Glens Falls, NY), Holderness (NH), Wolfeboro (NH), and Naples (NY) have written substantive policies that address multiple dimensions of AI governance with some genuine depth. They are shorter than the Tier 1 documents and have meaningful gaps, but they demonstrate thoughtful engagement with the issues. Holderness is notable for being exceptionally patron-education-focused even though it lacks some internal governance structure. Crandall is notable for its strong legal citations.

Tier 3 — Standard Staff Conduct Documents (6–7 libraries)

DeKalb, Hastings, Hinsdale, Houston County, St. Charles, and White House have written what are essentially employee acceptable-use policies for AI. These documents tell staff what they can and cannot do with AI at work. They are typically a few pages long, they cover the basics competently, and they leave most governance questions unanswered. This is not inherently a failure — a short, clear conduct policy is better than a long, confusing one — but these documents leave substantial governance gaps.

Tier 4 — Administrative Adoption (1 library)

Johnson County Library has done something categorically different: it formally adopted the Johnson County Government AI Policy by reference, with name substitutions. The resulting document is one paragraph long and contains essentially no library-specific AI governance content. Whatever substance exists resides in the county government policy, which was not provided or accessible for this analysis.

This spectrum matters because it shapes what we can fairly ask of each policy. A Tier 3 staff conduct document cannot be evaluated against the same expectations as a Tier 1 governance framework — but both need to answer the most basic questions, and both need to be understood in the context of what their communities actually need.

---

3. Why Libraries Are Developing AI Policies: Framing and Purpose

Nearly every library studied offers a rationale for having an AI policy, and virtually all frame AI as both a promising tool and a genuine risk. This dual framing is universal and important: libraries are not simply adopting AI enthusiastically, nor are they reflexively resisting it. They are trying to thread the needle.

The most commonly cited motivations, in rough order of prevalence, are:

- **Patron privacy and data protection** — mentioned in nearly every policy, usually as the primary concern
- **Ethics and responsible use** — frequently invoked, though rarely defined
- **Service enhancement** — libraries want AI to help them serve patrons better
- **Risk mitigation and security** — protecting library systems and data from breaches
- **Legal compliance** — ensuring AI use doesn't violate applicable laws
- **Productivity and efficiency** — helping staff do their work more effectively

Less commonly cited motivations include **public trust and transparency** (mentioned explicitly in Kenosha, Holderness, Wolfeboro, and Oakville), **community education and AI literacy** (Holderness, Kenosha, Oakville), and **innovation** (Toronto explicitly names it).

Two motivations that are conspicuously rare in most documents: explicit grounding in the library's public service mission, and acknowledgment that AI poses labor/workforce challenges for library employees. Several policies invoke the library's mission and values by reference (especially Holderness, Oakville, Toronto, and Schaumburg), but most staff conduct documents never mention the library's purpose or the communities it serves — they read as internal HR policies rather than public-service governance. And while Crandall's policy notably acknowledges "the impact of AI on CPL as an employer" as a concern, none of the policies develop substantive guidance on the workforce dimension of AI adoption.

Most policies explicitly acknowledge that AI is evolving rapidly and commit to regular review cycles. Kenosha and Crandall commit to review at least every six months. Holderness calls its policy "a living document." Toronto notes that "the technological and regulatory environment of AI rapidly develops." However, a few policies — most notably Hinsdale and St. Charles — do not characterize AI as evolving, which raises questions about how those policies will adapt as the landscape changes.

Almost all policies treat AI as exceptional technology warranting special rules rather than just another digital tool covered by existing IT policies. The fact that they wrote dedicated AI policies at all signals this belief, and many reinforce it explicitly. The exception is Naples, which takes an interesting middle-ground position: some of its rules are AI-specific, but AI use is also required to comply with the library's existing Code of Ethics, EEO, and conduct policies.

---

4. Who These Policies Govern: Scope and Actors

One of the most important and most overlooked questions in any policy is: *who does this actually apply*

to? The answer varies more than you might expect.

Staff Are Universal; Everyone Else Is Optional

All 15 policies apply to library employees. That is, of course, the minimum. But the similarities largely end there.

Board members and trustees are addressed by only two policies (Holderness, by implication; Oakville, explicitly). For every other library, it is unclear whether elected or appointed trustees are expected to follow the same AI rules as paid staff.

Volunteers are explicitly covered by only Toronto and Oakville. Libraries often rely heavily on volunteers, and this gap could create real privacy and data-protection risks if volunteers are using AI tools without governance oversight.

Interns and students are covered explicitly by Naples and Oakville. Naples has the broadest scope definition in the group, covering "all employees, interns, whether paid or unpaid, anyone who is (or is employed by) a contractor, subcontractor, vendor, consultant, or anyone providing services in our workplace who has access to AI technologies." This is notably comprehensive — most other policies could leave contractors and vendors operating in the library's environment without any AI policy obligation.

Patrons and the public present the most interesting scope question. Most libraries simply don't address this: patron use of AI on library computers or through library services is never mentioned. Some libraries (Oakville, Toronto, Schaumburg) explicitly exclude patron use from their policy's scope, noting that existing computer use, internet, and conduct policies already govern patron activity. Holderness and Crandall touch on patron computer use but don't create a comprehensive framework for it. **No policy studied attempts to govern patron use of personal AI tools they bring to the library.**

The fact that most of these policies are silent on patron use is understandable — libraries generally do not control what software patrons access, and there are serious intellectual freedom concerns about attempting to do so — but it means there is a significant gap between AI governance for staff and for the communities libraries serve.

Defining Key Terms: A Field Still Finding Its Language

Most of the stronger policies define AI itself, but many simpler documents do not. Among the definitions provided:

- **Crandall** provides four separate definitions: Artificial Intelligence, Large Language Models, Generative AI, and AI Systems & Tools, plus a definition of AI Hallucinations.
- **Kenosha** mirrors this structure and adds the same five definitions.
- **Toronto** provides a formal machine-learning-based definition of AI: "A machine-based system that, for explicit or implicit objectives, infers from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments."
- **Oakville** goes furthest with definitions, defining not just AI but also Generative AI, Public AI, Private AI, Responsible AI, Personal Information, and Confidential Business Information.

Several policies — including DeKalb, Houston County, Hastings, Hinsdale, Wolfeboro, St. Charles, and White House — never define AI at all. They use the term throughout but leave it to readers to infer what counts as AI. As more tools gain AI features, this creates ambiguity about what the policy actually covers.

Defining sensitive data is equally inconsistent. Toronto leads the pack with eight specific sub-categories of Personal Information drawn directly from MFIPPA (Ontario's privacy law). Oakville provides formal definitions of both Personal Information and Confidential Business Information. Naples provides an illustrative list (passwords, dates of birth, social security numbers, banking information, protected health information). Everyone else uses terms like "patron data" or "sensitive information" without formal definition — which means staff must make judgment calls about what information is protected.

5. Legal Grounding: How Policies Connect to the Law

A policy that says "comply with applicable laws" without naming any laws is weaker than one that cites specific statutes — both because it provides less guidance to staff and because it signals less institutional investment in understanding the legal landscape.

The spread on legal citations is dramatic.

At one end are the deeply grounded policies:

- **Crandall (NY)** cites four New York State statutes: Technology Law §103-e (which provides the statutory definition of AI the policy adopts), Civil Service Law §80 (effective July 1, 2025, relating to employee rights), General Business Law §1102, and Election Law §12-106. It also references the Library Bill of Rights Article VII, pending state senate and assembly bills, and "state and federal government guidelines." This is the most comprehensive state-law grounding in the study.
- **Kenosha (WI)** is the most internationally grounded document in the study: it cites Wisconsin Statute 43.30 (Public Library Records), the NIST AI Risk Management Framework (a federal tool), the EU Artificial Intelligence Act, the ALA Library Bill of Rights Article VII, and the Wisconsin Governor's Task Force on Workforce and Artificial Intelligence.
- **Toronto (ON)** cites Ontario's Municipal Freedom of Information and Protection of Privacy Act (MFIPPA) as the foundational privacy law, Ontario's Enhancing Digital Security and Trust Act (2024), Ontario's Working for Workers Four Act (2024, which requires disclosure of AI use in hiring), the City of Toronto Digital Infrastructure Strategic Framework, and the NIST AI Risk Management Framework.
- **Schaumburg (IL)** cites the Illinois Library Records Confidentiality Act (75 ILCS 70/1, et seq.) — the specific Illinois library patron privacy statute. It also quotes ALA's privacy statement directly.
- **Houston County (GA)** cites Georgia Confidentiality Law, PINES Policy (the Georgia library consortium's policy), and a specific federal court case — Civil Action No. 22-1564 — to support its position that AI-generated images are not copyrightable.
- **Wolfeboro (NH)** relies entirely on state sources: the New Hampshire Code of Ethics for the Use and Development of Generative Artificial Intelligence and Automated Decision Systems, the NH Statewide Information Security Manual, and NH Privacy Laws. Its entire governance framework is adapted from state government rather than library-sector guidance.
- **Oakville (ON)** references the Government of Canada's Responsible Use of AI guidance, the United Nations Resolution on Artificial Intelligence, and the Urban Libraries Council's AI Leadership Brief.

At the other end are policies that reference no specific laws at all: **DeKalb, Hastings, and White House** simply say they comply with "applicable laws" or "copyright and privacy laws" without naming any. **Hinsdale** makes no external references whatsoever.

For **patrons and community members**, this matters because specific legal citations signal that the library has done the homework to understand its actual legal obligations and is not just paying lip service to compliance. For **staff**, specific citations answer the "why?" question that helps them understand why a rule exists and how seriously to take it.

A notable gap across all policies: only one policy (Toronto) explicitly addresses a law regulating AI *disclosure* to employees about AI use in hiring (Ontario's Working for Workers Four Act). As AI is increasingly used in workforce management, this will become a significant compliance issue for libraries as employers, not just as service providers.

---

6. What Staff Are Allowed to Do: Permitted Uses

One of the most practical questions any AI policy must answer is: what can staff actually use AI for? The approaches taken by these 15 libraries span a spectrum from a formal approved-use list to complete

silence.

Three Approaches to Permitted Uses

Approach 1: An exclusive approved list. Only **Houston County** explicitly frames its permitted uses as a definitive, closed list. Its "Approved Uses/Actions" section, paired with a "Prohibited Uses/Actions" section, creates a clearer governance picture than most: if a use isn't on the list, it isn't approved. The approved list includes in-house text creation, text editing for grammar and clarity, non-sensitive Q&A exchanges, and basic photo editing (not AI image generation). This conservatism is intentional — the library is more concerned about risks than about expanding use cases.

Approach 2: An illustrative examples list. Most policies provide examples of acceptable uses with language like "including but not limited to" or "the following list is not exhaustive." Libraries using this approach include Hinsdale, St. Charles, Hastings, White House, and Holderness. This gives staff some guidance without creating the bureaucratic overhead of approving every new use case, but it also creates ambiguity — staff may not know whether a use that isn't on the list is acceptable.

Approach 3: General permission with guardrails. The more comprehensive policies — Kenosha, Crandall, Toronto, Oakville, Schaumburg, and Naples — grant broad permission to use AI across library functions while imposing specific guardrails (vetted tools only, no patron data, human review required). This approach trusts staff with professional judgment and avoids the administrative burden of maintaining an exhaustive list.

Johnson County sits outside all three approaches: the policy simply doesn't address permitted uses at all.

What Specific Uses Are Most Commonly Permitted?

Among the policies that do name specific use cases:

- **Text drafting, editing, and proofreading** — permitted by most policies that enumerate uses (Hinsdale, St. Charles, Kenosha, Houston County, Hastings)
- **Brainstorming and ideation** — generating ideas for programs, content, marketing (Hinsdale, St. Charles, Kenosha, White House, Hastings)
- **Marketing and communications** — creating promotional materials, social media posts, flyers (DeKalb, Hinsdale, Hastings, St. Charles, Kenosha, White House)
- **Research assistance** — looking up information, especially on behalf of patrons (Hinsdale, Holderness, Kenosha, Hastings, Houston County)
- **Data analysis** — analyzing library usage data, informing decisions (DeKalb, Kenosha, Schaumburg, Wolfeboro, Hastings)

Less commonly permitted: translation (almost never explicitly named), image and video generation (addressed with restrictions in Kenosha and White House), and search enhancement (only Holderness and Kenosha specifically address search tools).

An important pattern: most permitted uses are *staff productivity* uses (drafting, editing, brainstorming). Far fewer policies explicitly address AI for *patron service* — that is, using AI to help patrons in real time. Only Kenosha has a dedicated section on AI use in reference interactions, establishing specific protocols for how staff should disclose AI use and verify AI-generated answers before sharing them with patrons.

Conditions on Permitted Use

Virtually every policy attaches conditions to any AI use:

- **Human review of AI outputs before use** — required by 13 of 15 policies (Johnson County and Naples do not explicitly require it, though Naples prohibits using AI "as a final work product")
- **Only approved or vetted tools** — required by the stronger policies; absent from simpler ones that don't establish a vetting process
- **No patron or sensitive data** — the most universal data restriction; addressed in virtually every substantive policy
- **Attribution or disclosure** — required by several policies in specific contexts (see Section 9)

7. What Staff Are Forbidden from Doing: Prohibited Uses

Prohibited use provisions are where policies most clearly reveal their underlying concerns. Almost every policy includes some prohibitions, though their breadth and specificity vary enormously.

The Universal Rule: Don't Put Patron Data Into AI

This is the one rule that virtually everyone has: staff may not enter patron data, personally identifiable information, or sensitive library data into AI systems — at least into unapproved or public AI tools.

The specific framing varies:

- Most policies (Crandall, Kenosha, Hinsdale, Hastings) prohibit uploading confidential data to *unvetted* or *public* AI systems, but allow it on vetted/approved systems when necessary.
- Houston County and White House state **absolute prohibitions** — "under no circumstances" — without qualification by tool type.
- St. Charles and Hinsdale prohibit entering patron data into "AI" generally (applying to all tools, not just unapproved ones).
- DeKalb's prohibition is about unauthorized data *collection* by AI systems, rather than staff entering data — a subtle but meaningful difference in framing.

The protected data named most frequently includes patron records, employee information, and personally identifiable information. Toronto goes furthest, defining Personal Information with eight MFIPPA-grounded subcategories including health/medical information, financial transactions, identifying numbers, and more. White House and Naples, by contrast, simply say "personal information of any individual" without definition.

Employment Decisions: AI Shall Not Hire or Fire

Prohibiting AI from making employment decisions is common among the more substantive policies. Libraries that include this prohibition include Crandall, Hinsdale, Kenosha, St. Charles, and Naples. The typical language bars using AI for "selecting or narrowing down potential hiring candidates, or for disciplinary actions." Bias in hiring algorithms is frequently cited as the reason.

St. Charles has the most expansive list of prohibited employment uses: hiring, discipline, suspension, termination, promotion, demotion, salary changes, and performance reviews. Naples uniquely extends the prohibition to AI *surveillance* of employees and AI use in *disability-related inquiries* — two areas with significant legal implications that no other policy addresses.

Most policies are silent on this entire dimension — DeKalb, Hastings, Houston County, Holderness, Toronto, Oakville, Wolfeboro, Schaumburg, White House, and Johnson County do not address employment-decision uses at all. As AI-driven HR tools become more prevalent, this is a growing gap.

Standout Prohibitions Worth Noting

Several libraries have prohibitions that appear in no other policy studied:

- **Kenosha** explicitly prohibits using AI to "deceive, manipulate, mislead, misrepresent, or falsify" — a values-based prohibition that goes beyond data protection into professional ethics.
- **Kenosha** also prohibits text-to-image AI for generating *original* images, while allowing AI-generated images as components within larger original works. This is a nuanced distinction that reflects both copyright concerns and questions about authenticity.
- **White House** prohibits purchasing AI-written books for the collection — a collection development prohibition that no other policy includes. It also prohibits AI-generated photos in marketing materials, making it the only policy (along with Houston County) to specifically address AI in visual content for library communications.
- **St. Charles** requires that staff *not erase AI prompt history*. The rationale is employer monitoring —

the library retains the right to review what staff asked AI to do. This is a notable employment relations dimension that most policies ignore.

- **Crandall** commits to disabling "desktop AI tools that passively monitor patron or staff behavior" — framing the library as an active technical defender of patron privacy against AI surveillance, not just a passive rule-setter.
- **Naples** explicitly prohibits using AI to "surveil or gather information regarding employees and other individuals" — one of very few policies to name employee surveillance as a prohibited use.

How Prohibitions Are Enforced

Almost all prohibitions in all 15 policies rely on staff behavior and professional responsibility for compliance. Technical controls — software that blocks certain tools, network filtering, audit logs — are described in only a handful of policies and only briefly. Kenosha mentions anonymization techniques (hashing, k-anonymity, differential privacy) in its data protection section, and Wolfeboro mentions encryption and access controls, but neither describes a comprehensive technical enforcement infrastructure. The prevailing model is: here are the rules, staff are responsible for following them, and violations may result in discipline.

---

8. Keeping Humans in Charge: Oversight and Output Quality

Perhaps the most fundamental question in AI governance is: who is responsible when AI produces something wrong, harmful, or misleading? Every substantial policy studied answers the same way: a human must review AI outputs before they are used, and that human bears responsibility for the result.

Human Review: Near-Universal but Variably Defined

14 of 15 policies require human review of AI outputs in some form. (Johnson County's policy contains no operational guidance.) The strength of this requirement ranges considerably:

- **"Always review and validate AI generated results"** (Hastings, Hinsdale — both use nearly identical language, suggesting a common template source)
- **"Library staff utilizing AI and ML tools must review all output before publishing or using any generated materials"** (Crandall, Kenosha — again, nearly identical language)
- **"No AI-generated work product will be used or published without oversight from a TPL staff person"** (Toronto)
- **"AI-generated content is reviewed for accuracy and compliance with Library standards"** and **"reviewed by Library personnel for accuracy and appropriateness before publication or use"** (Schaumburg)
- **"Staff must fact check all AI generated information for accuracy and bias"** (White House)
- **"final publications must have human review and approval"** (Wolfeboro)

These all say essentially the same thing, but the framing matters. Schaumburg explicitly ties review to *compliance with Library standards* as well as accuracy. White House explicitly mentions checking for *bias* alongside accuracy. Wolfeboro requires review and *approval* — a stronger gate than simply review.

Naples is the outlier among substantive policies: it prohibits using AI "as a final work product," which implies human intervention, but never explicitly requires human review as a separate step. Accuracy verification is not mentioned.

The Human-in-the-Loop Model

The concept of "human-in-the-loop" — keeping a person involved at critical decision points in any AI-assisted process — is functionally present in most policies but rarely named. **Kenosha is the only library to give this model a formal name:** it calls it the "Human-AI-Human model" and describes it explicitly: "human staff members initiate interactions, utilize AI as a tool to assist or enhance the service, and then interpret and communicate the AI's output with a human touch."

Toronto, Schaumburg, Oakville, and St. Charles all describe the same principle without naming it: AI assists decision-making but does not replace professional judgment. Crandall references a New York State pending law that addresses "any system making decisions without human intervention" — implicitly invoking the human-in-the-loop concept through its legal citations.

Hallucinations, Inaccuracy, and Bias: What Are Staff Warned About?

A policy that tells staff to review AI outputs but doesn't explain *why* — what kinds of errors AI makes — is less useful than one that prepares staff for specific failure modes. The policies vary significantly on this.

Hallucinations (AI generating confident but completely fabricated content) are explicitly named by only five policies: Crandall, Kenosha, White House, White House, and Crandall (both define the term), plus White House ("due to hallucinations and other AI pitfalls"). This is surprisingly low given how important hallucination awareness is for anyone using AI for reference or content creation.

Inaccuracy is addressed more broadly — most policies that require human review at least imply inaccuracy as a concern, and many name it explicitly.

Bias in AI outputs is named by about two-thirds of policies, usually in the context of requiring staff to check outputs for biased or discriminatory content. Kenosha goes furthest, requiring staff to "assess the potential for bias or errors" in AI outputs, and providing patron-education materials that explain how datasets reproduce societal biases.

Kenosha is the only library to warn about all four categories — hallucinations, inaccuracy, bias, and misleading outputs — and to require cross-referencing AI-generated information against authoritative sources (academic journals, government websites). This level of output quality management is not matched by any other library in the study.

When Does Human Review Actually Matter?

Most policies require review before "use or publication." But what does that mean in practice? If a staff member uses AI to generate a draft email, does every sentence need to be fact-checked? If AI helps create program flyers, does the color scheme need verification?

The more thoughtful policies recognize that not all AI uses carry the same risk of harm:

- **Kenosha** distinguishes between uses where disclosure and cross-referencing are required (original published works, reference interactions with patrons) and uses where they are not (grammar correction, brainstorming that doesn't become content).
- **Schaumburg** requires review specifically for "work product substantially generated by AI" — acknowledging that minor AI assistance doesn't need the same scrutiny as major AI generation.
- **Holderness** focuses human oversight requirements specifically on cataloging, where accuracy is especially critical.

Most other policies apply a single blanket requirement without distinguishing high-risk from low-risk use cases. While a blanket rule is simpler to follow, it may in practice lead staff to either ignore the requirement (when it feels too burdensome for minor uses) or apply excessive scrutiny to everything (wasting time on review that isn't warranted).

---

9. Being Honest About AI: Transparency, Disclosure, and Attribution

Should patrons know when library staff used AI to help answer their questions? Should library materials disclose when AI helped create them? Should library websites list which AI tools the library uses? These questions of transparency and disclosure are among the most contested and least resolved in this policy set.

The short answer: most policies require disclosure in some contexts, but very few require it broadly.

Disclosure to Patrons During Service Interactions

When a reference librarian uses an AI tool to help answer a question, should they tell the patron? Only three libraries unambiguously require this:

- **Hinsdale** requires staff to cite the AI tool used "in the same manner as other sources of information" — treating AI disclosure as equivalent to citing any reference source
- **Hastings** requires staff to "ensure patrons understand that AI was used" when providing AI-generated results
- **Kenosha** requires staff to "disclose AI-generated answers in reference interactions to maintain transparency and trust" and specifically states that "staff will disclose AI use verbally"

Most other policies are silent on this question, even when requiring human review and accuracy verification. Crandall, DeKalb, Houston County, Napls, Schaumburg, St. Charles, Toronto, Oakville, and White House do not require staff to tell patrons when AI was used in preparing their answers. This is a significant transparency gap — patrons may not realize they are receiving AI-assisted information, and they may not know to apply additional critical evaluation.

Disclosure in Published Content

More policies require disclosure when AI-generated content is published or distributed:

- **Holderness** requires "clearly label AI-generated or AI-assisted content, including images, text, and recommendations, in both physical and digital spaces" — the broadest labeling requirement in the study
- **Kenosha** requires that "all original work in which AI plays a role in its creation will link to a statement indicating the use of AI" — providing a formal attribution statement as an appendix
- **Oakville** requires that "all content produced or assisted by AI will be clearly attributed"
- **Schaumburg** requires that "any work product substantially generated by AI must be clearly identified"
- **Toronto** requires staff to "clearly identify all work products that have been substantially generated using AI"
- **White House** requires that "staff must cite any AI related source they do use" for pamphlets, brochures, and presentations
- **Wolfeboro** requires that "all AI generated information should be clearly marked before distribution"

Naples requires "properly citing or disclosing when AI has been used as a resource for a final work product" but does not specify any format for doing so.

Several policies require no disclosure at all: **Crandall, DeKalb, Houston County, and St. Charles** have no disclosure requirement for published AI-assisted content.

The Most Transparent Policy: Kenosha

Kenosha's approach to transparency goes further than any other library in this study. In addition to requiring verbal disclosure during reference interactions and a link to an AI usage statement on published works, Kenosha commits to **maintaining a public list of approved AI tools on its website**, including the name, date added, and description of use. This gives patrons and the community a way to understand concretely which AI tools the library has vetted and uses. No other library studied makes this level of public commitment to transparency about its AI tool inventory.

Copyright and AI-Generated Content

The question of whether AI-generated content can be copyrighted — and who owns it — is legally contested and practically important for libraries. Most policies either ignore it or address it only superficially.

Houston County has the most substantive copyright position, explicitly stating that AI-generated images "are not legally considered intellectual theft" are "still considered to be derivative, and therefore not copyrightable by law" and citing a specific federal court case (Civil Action No. 22-1564) to support this position. Whether this legal interpretation is current is a matter of ongoing litigation, but the fact that the library took a position is noteworthy.

Crandall notes flatly that "AI cannot own copyright."

Most other policies acknowledge copyright as a concern and tell staff to respect copyright laws, but do not address the ownership status of AI-generated content or explain what that means in practice for library materials.

---

10. Learning and Teaching: Staff Training and Patron Education

AI tools require learning — both to use them effectively and to use them responsibly. Governance without education produces policies that sit in drawers. And libraries, as institutions whose fundamental purpose includes connecting people with information and skills, have a particular responsibility not just to govern staff AI use but to help their communities understand AI.

The policies studied reveal a dramatic gap between these two responsibilities.

Staff Training: Required by Less Than Half

Of the 15 libraries studied, **seven require staff training on AI:**

- **DeKalb** requires training on "ethical AI practices, privacy protection, and bias mitigation before proceeding" with any AI use
- **Holderness** requires "ongoing, role-specific training for staff on AI technologies, with a focus on both operational knowledge and ethical considerations"
- **Kenosha** requires role-specific staff training with a detailed curriculum (see below)
- **Oakville** requires "appropriate training on ethical AI use, privacy protection, and responsible application of AI tools"
- **Schaumburg** requires training on "the implications of AI for privacy, security, and ethics, as well as on how to assist patrons in understanding AI-generated content"
- **St. Charles** requires training on "the ethical use of AI, privacy protection, and how to assist patrons in using AI tools"
- **Toronto** requires "ongoing training for staff on the use of AI in the workplace" covering cybersecurity, privacy, and how to help customers understand AI content

The remaining eight libraries — Crandall (encouraged but not required), White House (self-directed, not formally required), and Hastings, Hinsdale, Houston County, Johnson County, Naples, and Wolfeboro (no training mentioned at all) — leave staff to figure things out without institutional support.

What Training Covers

Training topics named across all policies that require training:

- **Ethics and responsible use** — mentioned by most training policies (DeKalb, Holderness, Kenosha, Oakville, St. Charles)
- **Privacy protection** — mentioned by most (DeKalb, Kenosha, Oakville, Schaumburg, St. Charles, Toronto)
- **Security and cybersecurity** — Kenosha, Schaumburg, Toronto
- **Bias and fairness** — DeKalb, Kenosha
- **Evaluation and verification of AI outputs** — Kenosha (most explicit)
- **Copyright and intellectual property** — Kenosha
- **AI hallucinations and inaccuracy** — Kenosha
- **The AI vetting and approval process** — Kenosha

Kenosha's training curriculum is by far the most comprehensive — it covers history and current use of AI, general information literacy, the library's AI system vetting process, in-depth understanding of biases and hallucinations, data privacy, copyright and fair use, bias in datasets, credibility of AI outputs, and how to distinguish human from AI-generated works. No other library comes close to this level of training specification.

Patron AI Literacy: The Field's Most Significant Gap

While staff training has room for improvement, the gap in patron AI literacy is more striking. Libraries are institutions built around the idea that an informed public makes better decisions. AI is reshaping how people find information, make decisions, interact with institutions, and understand the world. Yet most library AI policies contain nothing about helping the public navigate this shift.

Of the 15 policies studied, **only two have developed meaningful patron AI literacy programs:**

Holderness has the most patron-education-focused policy in the study. Its "AI Literacy and Public Education" section commits to: offering regular workshops and resources on AI topics; developing programs that are "accessible to all ages" including intergenerational programs; empowering patrons to "critically understand and assess AI tools they encounter in everyday life"; and framing the library explicitly as a community educator about AI. Holderness is the only library in this study that explicitly states it is positioning itself as a community AI educator — which is perhaps remarkable for a small New Hampshire library serving a rural community.

Kenosha has a detailed appendix on patron AI literacy with specific learning targets for community members, including: understanding the history and current uses of AI; data privacy and security in the AI context; copyright and fair use; recognizing bias in AI-generated content; evaluating the credibility of AI outputs; and distinguishing human-created from AI-generated works. Kenosha also commits to coordinating with local educational institutions to share best practices.

Several other libraries mention patron education indirectly: Crandall mentions "Education and awareness for library users" as a concern but provides no program. Oakville and Toronto note the library's role in promoting digital and AI literacy but develop no programs. Schaumburg trains staff to "assist patrons in understanding AI-generated content" but describes no patron-facing programs.

The majority of libraries in this study — DeKalb, Hastings, Hinsdale, Houston County, Johnson County, Naples, St. Charles, White House, and Wolfeboro — have no patron education component at all. Their AI policies are entirely inward-facing documents about staff conduct.

This is arguably the most significant missed opportunity in public library AI governance as a field. Libraries occupy a unique position as trusted public institutions with reach across communities, including communities that are least likely to encounter high-quality AI literacy education elsewhere. A staff conduct document that says nothing about this opportunity represents an incomplete response to the challenge AI poses.

---

11. Accountability: Records, Reporting, and Consequences

Good governance requires not just rules, but ways to know whether rules are being followed and what happens when they are not.

Records Retention: A Near-Universal Gap

Only one library — Kenosha — explicitly addresses records retention for AI-generated or AI-assisted content. Kenosha's policy states that "all content published by the Kenosha Public Library through the use of AI" must be retained "for the purposes of public records retention as may be required by applicable law." This matters because public libraries are typically subject to government records retention requirements, and if AI generates a social media post, a reference answer, or an internal report, that content may have records retention implications.

All other libraries are silent on this. If AI generates content that later becomes relevant to a legal dispute, an audit, or a public records request, most libraries studied would have no policy-level guidance about how to handle it.

Documentation of Approved Tools: Only Kenosha Goes Public

Kenosha is also the only library that commits to maintaining a **public list of AI tools approved for staff use** — including the name, date added, and description of each tool — on the library's website. This serves multiple purposes: it gives staff a clear reference point, it gives the public transparency about

what tools their library is using, and it creates an accountability record for how the library's AI use has evolved over time.

Several other libraries maintain or imply an internal list of approved tools (Hastings's TEC_6 list, White House's director-maintained list, St. Charles's management-maintained list), but none of these are public.

Incident Reporting: Minimal

Staff are required to report AI-related incidents, violations, or concerns in only three policies:

- **Naples**: "Any concerns, incidents, or violations of the use of AI technologies under this policy must be reported immediately to the Director for investigation and resolution"
- **Oakville**: Staff must "Report any AI-related security or privacy concerns immediately to supervisor"; Leaders must "Report policy violations or concerns to EMT [Executive Management Team]"
- **White House**: "Staff are to report any violations of this policy to the director who will investigate the matter further and determine recommended correction"

All other libraries state that staff should follow the policy and that consequences exist for violations, but provide no structured mechanism for raising concerns or reporting incidents.

Disciplinary Consequences: Named but Not Specified

Four libraries state explicitly that policy violations can result in discipline including termination:

- Kenosha: "Violations of this policy will be investigated and may result in disciplinary action, up to and including termination of employment"
- Oakville: "Failure to adhere to the ethical and responsible guidelines for AI use as outlined in this policy may result in disciplinary action, up to and including termination of employment"
- Naples: "Violation of this policy may result in disciplinary action, up to and including termination of employment"
- White House: director "will investigate the matter further and determine recommended correction"

No policy specifies what specific violations warrant what specific consequences. The spectrum from "policy violation" to "termination" is enormously wide, and leaving the entire range to discretionary investigation creates both legal exposure and uncertainty for staff.

---

12. Cross-Cutting Themes

Several important themes emerge across the survey dimensions that don't fit neatly into a single section.

Theme 1: Privacy Dominates; Other Values Are Underrepresented

Every single policy studied addresses patron privacy as a core concern, and most address it explicitly as a priority. This reflects the library profession's longstanding commitment to patron confidentiality. It is the field's strongest consensus.

But privacy is not the only value at stake in AI governance. **Equity and fairness** — the question of whether AI treats all community members fairly, or whether AI-assisted library services disadvantage people who don't know how to use or evaluate AI — is addressed explicitly by fewer than half the policies (DeKalb, Holderness, Kenosha, Schaumburg, Toronto, Oakville, Wolfeboro). **Intellectual freedom** — the concern that AI tools may filter, shape, or limit what information people access — is addressed by even fewer (Toronto, Holderness, Oakville, and Kenosha, typically by reference to existing policies or frameworks). **Environmental sustainability** — the fact that AI requires massive energy consumption — is addressed by only two policies: **Toronto and Schaumburg**. Both commit to considering environmental impact when choosing AI tools, which is ahead of the rest of the field but far from universal.

Theme 2: All Policies Rely Primarily on Staff Behavior, Not Technology

Across all 15 policies, the primary mechanism for ensuring compliance is staff professional responsibility.

Staff are told what to do, what not to do, and trusted to follow those rules. Technical controls — software that enforces data restrictions, audit logs that record AI use, monitoring systems that flag prohibited inputs — are described in only a handful of policies and at a high level of abstraction.

This is understandable for small and medium libraries that lack the technical infrastructure to implement sophisticated controls. But it does create a governance model that depends heavily on individual staff judgment and is difficult to audit. If a staff member inadvertently enters patron data into an AI tool, most libraries have no system that would catch it.

Kenosha stands apart by describing specific technical anonymization techniques (hashing, k-anonymity, differential privacy) and by requiring that Computer & Network Services conduct a security evaluation of every tool before use. **Wolfeboro** describes encryption and access controls. But these remain exceptions rather than the norm.

Theme 3: Vetting Processes Exist in the Stronger Policies, But Criteria Are Rarely Specified

Most substantive policies require that AI tools be approved before use. But the criteria for what makes a tool approvable are rarely spelled out. Privacy and security are mentioned most frequently as vetting criteria. Bias/fairness, accessibility, and legal compliance are less consistently named. Data retention and vendor practices (what does the vendor do with data entered into their system?) are named only by the more sophisticated policies.

Toronto goes furthest: before any tool can be approved, it must undergo three separate assessments — an IT Security Assessment, a Privacy Impact Assessment, and a Human Rights AI Impact Assessment. This three-gate vetting process is unique in this study and significantly more rigorous than any other approach.

Kenosha requires coordination with Computer & Network Services and has a formal AI Vetting Request form and process — the most defined vetting workflow of any U.S. library in the study.

Theme 4: The Canadian Libraries Are More Comprehensive

Comparing U.S. and Canadian libraries is instructive. The two Ontario libraries — Toronto and Oakville — are among the most comprehensive policies in the study on virtually every dimension. Both explicitly state that patron use is out of scope (and direct patrons to existing policies). Both have formal vetting processes tied to government IT standards. Both require mandatory attribution for published AI content. Both have clear definitions and formal accountability structures.

This likely reflects differences in the governance environment: Ontario's MFIPPA is a well-established privacy framework that gives libraries clear legal grounding, and the Town of Oakville and City of Toronto have municipal IT governance structures that library policies can align with. U.S. libraries typically operate in a more fragmented legal environment and may lack a municipal parent organization with established AI governance standards.

Theme 5: Policy Age Doesn't Necessarily Predict Quality

The oldest policy in this study — Holderness (December 2023) — is one of the most thoughtful, particularly on patron education. The newest policies (Toronto's effective January 2025, Johnson County's September 2025) range from comprehensive to nearly empty. Quality correlates more with institutional investment and the specific focus areas each library prioritized than with when the policy was written.

That said, there is a general trend toward greater comprehensiveness in the 2025 policies compared to the 2024 ones, which may reflect lessons learned as the field developed shared norms.

---

13. Standout Policies Worth Learning From

Kenosha Public Library (Wisconsin): The Field Leader

Kenosha's policy is, by every measure applied in this analysis, the most comprehensive public library AI policy in this study. Key features that other libraries should consider adopting:

- **The Human-AI-Human model**: A named, explicit framework for how staff should think about AI in service interactions — initiate with human judgment, use AI as a tool, close with human interpretation
- **A public list of approved tools**: Transparency about what tools the library uses, when they were added, and for what purpose
- **Required cross-referencing of AI outputs**: Staff must check AI-generated information against authoritative sources — not just review it, but verify it
- **Prohibition on using AI to deceive or manipulate**: A values-based rule that goes beyond data protection
- **Prohibition on text-to-image AI for original images**: A nuanced position that distinguishes AI image components from AI-generated original works
- **Record retention**: The only library that explicitly addresses public records obligations for AI content
- **A full patron literacy program as an appendix**: With specific learning targets, staff training curriculum, and community partnerships
- **The most extensive legal and professional framework citations**: Wisconsin library privacy law, NIST AI RMF, EU AI Act, ALA Library Bill of Rights

Kenosha's weaknesses are mainly omissions shared across the field: no environmental sustainability section, no employment-impacts discussion, no community consultation process.

Toronto Public Library (Ontario): Most Rigorous Pre-Deployment Process

Toronto stands out for its three-gate vetting process — every AI tool must pass an IT Security Assessment, a Privacy Impact Assessment, and a Human Rights AI Impact Assessment before deployment. This is more rigorous than any U.S. library in this study. Toronto also leads the group in data definitions (MFIPPA-grounded Personal Information with eight subcategories), scope clarity (explicit about what is and isn't covered), and legal grounding. Its sustainability section (Section 5) commits to considering environmental impact — one of only two libraries to address this. Toronto's main weaknesses: no patron-facing AI literacy program, disclosure requirements that are vague in form, and some accountability gaps.

Holderness Free Library (New Hampshire): Champion of Patron Education

Holderness is, by a wide margin, the most patron-education-focused policy in this study. It explicitly frames the library as a community AI educator, commits to offering AI literacy workshops, and designs programs for "all ages" including intergenerational programs. Its stated purpose — "to enrich, inform, and empower our community" — is the clearest statement of library mission in any AI policy studied. Holderness was also the first policy adopted in this study (December 2023) and among the first public library AI policies in the nation.

Holderness's weaknesses include limited vetting processes, no employment-decision prohibitions, and an absence of formal accountability structures.

Crandall Public Library (Glens Falls, New York): Best Legal Grounding

Crandall's policy is notable for using the actual New York State statutory definition of AI (from NYS Technology Law §103-e), citing multiple state statutes relevant to AI use, and explicitly acknowledging the impact of AI on the library as an employer — a dimension most policies ignore. It also explicitly commits to disabling passive monitoring AI tools on patron computers, framing the library as a technical defender of patron privacy.

Oakville Public Library (Ontario): Most Values-Comprehensive

Oakville explicitly names all 11 of its library values in the AI policy — privacy, intellectual freedom, equity, belonging, transparency, human-centered service, lifelong learning, curiosity, access, professional judgment, and community trust. This grounding in the library's complete value system is unique in this study and represents best practice for connecting governance documents to institutional identity.

Oakville is also one of only two libraries to explicitly exclude patron use from scope and direct patrons to existing policies.

Houston County Public Library System (Georgia): Best Copyright Position

Houston County's decision to take an explicit legal position on AI image copyright — citing a specific federal court case to explain why AI-generated images are not copyrightable by the library — is a model for how libraries can engage with specific legal questions rather than making vague references to "applicable laws." Whether the specific position remains current is a legal question, but the approach of citing authority is right.

---

14. Common Gaps and Missed Opportunities

Across all 15 policies, certain topics are consistently absent or underdeveloped. Libraries developing or revising AI policies should consider these areas carefully.

1. Environmental Sustainability

Only Toronto and Schaumburg address the environmental impact of AI. AI systems require enormous amounts of energy and water for training and operation. Libraries that care about sustainability — as many do — have not yet integrated this value into their AI governance. Procurement decisions about AI tools could include environmental criteria (energy use, corporate sustainability commitments, data center locations) that none of these policies require.

2. Labor and Workforce Impacts

Remarkably, only Crandall's policy even gestures toward the question of AI's impact on library employment. The concern that AI could change job responsibilities, reduce positions, or create new skill requirements for library workers is real and pressing — and entirely absent from almost every policy studied. This may reflect the fact that AI policies are written by library leadership rather than in consultation with frontline staff, or it may reflect that this conversation is still too politically sensitive. But workforce impact is a governance question, not just a labor relations question, and it belongs in AI policy.

3. Accessibility

AI-generated content can create new accessibility barriers (AI image generators don't automatically produce alt text; AI-generated video may lack captions) or new opportunities (AI can help generate plain-language summaries of complex documents). Only a handful of policies mention accessibility, and none develop specific accessibility requirements for AI-generated content.

4. Community Consultation

None of the 15 policies describe any process for consulting with the community about AI use. In the public library sector, where service responds to community need and community trust is foundational, this seems like a significant gap. How should a library decide whether to deploy a patron-facing AI chatbot? Who gets input on what AI tools are used in children's programming? These questions of community governance are not addressed.

5. AI Embedded in Existing Products

Libraries buy hundreds of products — cataloging systems, databases, discovery layers, learning platforms — that are increasingly incorporating AI features. None of the policies studied, with the partial exception of Kenosha (which has a dedicated appendix on subscription products), provide clear guidance on how to handle AI that comes bundled into tools the library already uses. Kenosha's appendix honestly states: "We cannot guarantee that vendor-provided AI tools meet the Library's internal standards" — which is the right starting point for an honest conversation about a real governance challenge.

6. Patron Use of AI in the Library

Most policies simply don't address what happens when a patron uses their own AI tools — on their own device — while in the library. Should the library's wireless network terms of service address AI? Should staff be trained to help patrons who are being misled by AI outputs? These questions are left to separate computer use and conduct policies that predate AI and may not address it well.

7. Records Retention

Only Kenosha addresses records retention for AI content. As AI-generated materials accumulate in library records, and as government records obligations apply to more of what libraries produce, this will become a compliance issue. Libraries should develop guidance now rather than waiting for an audit or legal dispute to force the question.

8. Specific Disciplinary Clarity

Every policy that mentions discipline does so with "up to and including termination" — a maximum without any guidance on proportionality. Staff would benefit from understanding what kinds of violations are serious and what kinds are correctable. A framework distinguishing inadvertent errors, negligent violations, and intentional misconduct would make the accountability structure more credible and more fair.

---

15. Key Takeaways for Practitioners

For library directors, trustees, and staff members who are thinking about AI policy — whether developing a first policy or revising an existing one — the following observations from this analysis are offered.

What Every Library Needs (The Non-Negotiables)

1. **A clear prohibition on entering patron data into unapproved AI tools.** This is the field's consensus minimum, and it is right. Patron data protection is the core legal and ethical obligation that no library can afford to ignore.
2. **A requirement for human review of AI outputs before use.** AI makes mistakes — it fabricates information, amplifies bias, and produces plausible-sounding nonsense. Every piece of AI-assisted content should have a human checkpoint before it reaches patrons or the public.
3. **Some form of tool approval or vetting process.** Staff should not be using whatever AI tool they happen to encounter without any institutional review. Even a simple process — request to director, basic privacy check — is better than none.
4. **At least minimal staff education.** A policy without training produces compliance by chance. Staff need to understand what AI can and cannot do, what the risks are, and what the rules mean in practice.

What Most Libraries Should Add

1. **Patron disclosure in reference interactions.** When AI helps generate an answer a staff member provides to a patron, that patron deserves to know. Only three libraries currently require this. It is a matter of respect for patron autonomy and professional honesty.
2. **Named legal and professional references.** Citing specific laws (the relevant state library confidentiality statute, any applicable privacy laws) grounds the policy in real obligations and tells staff why the rules exist, not just what they are.
3. **Employment decision prohibitions.** As AI-driven HR tools proliferate, libraries need explicit rules about not using AI to select candidates, make discipline decisions, or conduct performance reviews. The risk of discrimination is real and legally significant.
4. **At least a starting point on patron AI literacy.** Even a single line committing to AI literacy programming, or a paragraph on how staff should assist patrons who encounter AI-generated

content, would represent progress toward what libraries' public service mission demands.

What the Field Leaders Are Doing That Others Should Consider

1. **A formal model for patron-facing AI interactions** (like Kenosha's Human-AI-Human model) that gives staff a mental framework for when and how to use AI in serving patrons.
2. **A public list of approved AI tools** that gives the community transparency about what AI the library is using and accountability if tools are added or removed.
3. **Cross-referencing requirements** — not just "review" AI outputs, but check them against authoritative sources before relying on them.
4. **Multiple vetting gates** (security, privacy, and bias/human rights assessments, as in Toronto) before deploying AI tools.
5. **A patron education program** with specific learning targets, not just a vague commitment to digital literacy (as in Kenosha and Holderness).
6. **Environmental sustainability criteria** in AI tool procurement (as in Toronto and Schaumburg).

A Final Word on Process

The best AI policy is not the longest one or the most comprehensive one — it is the one that gets followed. A short, clear policy that staff have been trained on and that is reviewed regularly is better than a comprehensive policy that sits unread. Most of the policies in this study will need to be revised within a year or two: the technology is moving too quickly, and the legal landscape is shifting rapidly in multiple jurisdictions.

The libraries that have built their policies around clear values (privacy, patron trust, human oversight), connected them to specific legal obligations, and committed to regular review and staff training are best positioned to keep their governance relevant as the landscape changes. That is ultimately the lesson from the best work in this field: good AI governance is not a document, it is a practice.

This analysis is based on the 15 library AI policies as they existed at the time of survey completion (2026-03-31). Policy documents and the laws they reference change over time. Libraries should verify the current status of any policy or legal citation before relying on it.

Survey instrument and completed surveys are available in the project working files.