

Level 2: Applied

All Librarians

Module 12

Making the case to administration

You've tried AI, you've seen the value, and now you need to bring your institution along. This module covers the practical advocacy work: how to frame the case, what data to use, and how to run a pilot that builds institutional confidence.

By Yulia Brusova

20 min read · Reviewed October 2026

AI for Academic Libraries · ai-for-academic-libraries.vercel.app

© 2026 · Licensed under CC BY-NC-SA 4.0

WHAT YOU'LL BE ABLE TO DO

- 1 Present a case for AI adoption using data and framing appropriate for your institutional context
- 2 Draft or contribute to an institutional AI acceptable use policy
- 3 Design a small AI pilot with measurable outcomes
- 4 Identify and frame the business case for an AI tool in terms administration will respond to
- 5 Use the U.S. vs. global AI adoption data to frame urgency without alarmism

The biggest barrier to AI adoption in most academic libraries is not the tools, the professional uncertainty, or the time required to learn. It is the conversation with administration. A librarian may spend months developing a genuine, calibrated AI practice, understanding where the tools add value and where professional judgment is irreplaceable. Bringing the institution along then requires a different skill set entirely. The librarian must translate what they have learned into language that resonates with decision-makers who may be skeptical, risk-averse, resource-constrained, or simply uninformed. This module is about that translation. The goal is not to sell administration on AI but to be the informed, credible voice that helps leadership make good decisions. That requires understanding the data, frameworks, and pilot designs that make the conversation productive.

Understanding how administration thinks about AI: risk, resources, and positioning

Administrative skepticism about AI adoption is rarely ideological. Most library administrators who have not moved on AI are managing a familiar set of pressures. Budgets are limited and must be justified to provosts and boards. Data privacy incidents or policy violations carry consequences for the institution. Staff capacity is already stretched. There is also the reasonable concern that an early commitment to the wrong tool or approach will require costly correction. To make an effective case, the librarian needs to know which of these pressures weighs most on their own administration. The framing that resolves a risk concern differs structurally from the one that resolves a budget concern. Neither is the same as the

framing for an administrator who believes the institution is not yet at a stage where AI is relevant.

The most common framing mistake is leading with capability: describing what AI can do before addressing why the administration should care, in terms it already cares about. An administrator focused on budget constraints does not become more receptive on hearing that AI can draft subject guide descriptions faster. They become more receptive on hearing that AI reduces the staff time needed for recurring writing tasks, and that the savings can be estimated in terms they already track. This reframe, from capability to institutional impact in the administrator's own terms, is the foundational move in effective administrative advocacy.

It is also worth knowing what administrators are hearing about AI through other institutional channels. Since 2024, IT departments, legal offices, provosts' offices, and accrediting bodies have issued AI guidance at an accelerating rate, and much of it frames AI mainly as a compliance and risk management concern. A librarian who arrives knowing what the institution's IT acceptable use policy says, what the provost's office has communicated about AI in academic work, and what legal counsel has flagged as risks is in a much stronger position. One who arrives with library-specific arguments alone is not. This situational awareness signals that the librarian is approaching the conversation as an institutional partner, not an advocate for a particular tool or practice. It positions the library as a resource for the institution's broader AI decision-making, not a department seeking approval for an experiment.

The data that makes the case: Pulse of the Library, ACRL, and the U.S. confidence gap

Administrative conversations are more productive when grounded in external data rather than individual experience. A librarian who says "I have found AI useful in my reference work" is making a personal claim. Another librarian might present data on where academic libraries stand in AI adoption globally, where the U.S. sits relative to international peers, and what professional associations have determined about library AI competencies. That librarian is joining a sector-level conversation, which is the appropriate register for an institutional decision.

The Clarivate Pulse of the Library survey is the most widely cited benchmark for library AI adoption. The 2026 findings, based on 1,876 library workers surveyed in April and May 2026, are particularly useful for administrative framing. About seven in ten libraries are exploring or implementing AI, but only about one in six has reached moderate or active implementation;

among academic libraries, that share rises to 22%. In administrative terms, this distribution means the field is moving from experimentation toward deployment. Institutions not yet evaluating AI are not holding a stable position relative to their peers; they are falling behind a moving field. The 2026 data also shows budget constraints cited by 52% of respondents as a concern in adopting or scaling AI, level with lack of expertise. That reframes the budget conversation. The question shifts from "should we spend money on AI" to "every library facing this decision is navigating the same resource pressure, and here is how the ones moving forward are addressing it."

The U.S. confidence gap deserves specific attention in administrative conversations at American institutions. The 2025 Pulse data showed that only 7% of U.S. librarians reported optimism about AI, compared to 27% to 31% among their counterparts in Asia and other global regions. The 2026 survey did not repeat that question, but it shows the gap persisting in practice: 27% of U.S. libraries are at initial implementation or beyond, compared to 43% in Asia and 52% in Europe. To use this data without creating unproductive alarm, make the framing specific. Institutions that build library AI capacity now will be better positioned to support the AI-literate graduates that employers and graduate programs already expect. Those that defer the decision are not avoiding a risk. They are accepting a different one: falling behind in a capability that is becoming a professional baseline.

The ACRL AI Competencies for Academic Library Workers, published in October 2025, is the professional authority anchor for any administrative conversation. The framework defines what library workers need to know about AI across four competency areas: ethics, knowledge, evaluation, and application. It is the field's answer to the question "what does responsible AI engagement look like for academic librarians?" Citing it signals to administration that the library is following a professionally sanctioned development pathway, not experimenting on its own. For example, a librarian might say "the ACRL framework identifies [specific competency] as a foundational expectation for academic library professionals, and this pilot is designed to build that competency." That grounds an institutional request in the field's own standards, a very different conversation from requesting approval for a personal interest in new technology.

Two additions strengthen this evidence base for two audiences the survey data does not reach. The first is the administrator who discounts vendor research on principle, a defensible position given who pays for it. Leo S. Lo's study "Evaluating AI literacy in academic libraries," published in *College & Research Libraries* in 2024 and awarded the 2025 CALA Jing Liao Award for Best Research, surveys U.S. library employees directly and is peer-reviewed. The training-gap argument therefore does not have to rest on a vendor's survey of its own customers. For example, a librarian told that the Clarivate numbers are marketing can make

the same case from a peer-reviewed source in the field's flagship journal, and the objection falls away. The second addition is for the administrator worried about spending money badly rather than spending it at all, which at a community college is usually the real concern. CalMatters reported in March 2026 on California community college districts that spent substantially on chatbot platforms that then gave wrong answers drawn from outdated FAQ libraries. An account like this persuades more than an optimistic one, because it shows that the librarian making the request understands how these projects fail and has designed around it. Naming a documented failure, then explaining what the proposed pilot does differently, turns advocacy into risk management, and risk management is the register administrators already think in. Whether a particular tool is worth its cost at this institution remains a local judgment that no national survey settles.

Framing for different types of administrative skepticism

The same proposal needs different framing for different audiences. A librarian who can identify which concern weighs most on their administration, and adjust accordingly, will be far more effective than one who presents a single argument to every listener. This audience awareness is not manipulation. It is the professional practice of meeting people where they are, the same skill librarians apply in reference consultations and instruction sessions.

For risk-averse administration, lead with risk mitigation rather than capability. The argument runs like this. AI adoption without policy creates more institutional risk than adoption with policy. A structured pilot with clear guidelines is less risky than the current situation, in which individual staff may already be using AI tools without institutional oversight or data protection guardrails. For example: "Our staff are likely already using AI tools informally. A pilot with defined parameters, approved tools, and data protection guidelines gives us visibility into that use and a framework for managing it responsibly, which is a better position than not knowing what's happening." This framing does not require the administrator to be enthusiastic about AI. It requires only that they recognize structured adoption as more manageable than unmanaged individual adoption. Present the ARL Guiding Principles and the ACRL AI Competencies framework as the professional authority behind the proposed approach, not as documents the administrator needs to read in detail.

For cost-focused administration, quantify time savings in terms the institution already tracks. Even conservative estimates are useful. If AI assistance saves each librarian thirty minutes per day on drafting, communication, and administrative writing, that is approximately 125 hours per year per full-time librarian. At the median librarian salary of about \$64,000 (U.S. Bureau of Labor Statistics, May 2024), with benefits at a standard institutional multiplier,

thirty minutes per day is meaningful recoverable capacity. That time can be redirected to instruction, reference, or collection work that serves institutional priorities directly. The calculation does not require precise data. It requires a defensible estimate that makes the time value of AI assistance visible in institutional cost terms. AI tool subscriptions, typically \$20 to \$25 per user per month for professional and team tiers, are modest relative to even a conservative estimate of recovered staff time. Make that comparison explicit rather than leaving the administrator to calculate it.

Some administrations believe the institution is not yet at a stage where AI is relevant, typically saying "our faculty aren't asking about this" or "our students seem to be managing." For them, the most effective reframe is student outcomes rather than operational efficiency. The argument is not that students are struggling now. It is that the AI literacy competencies students will need after graduation are being developed now, at institutions that are preparing them, and that the library has a professional role in that preparation. This reframe shifts the question from "does the library need this" to "are we preparing our students for the environment they are entering?" That question carries clearer institutional stakes for an administration concerned with graduate employment outcomes and institutional reputation.

Designing a pilot that builds institutional confidence

The most common pilot design error is scope ambiguity: a pilot defined well enough to start but not well enough to evaluate. An administration that approves a "six month AI pilot" without a defined scope, a specified outcome measure, and a clear review point has approved a vague experiment. It is hard to report on and easy to question. To design a pilot that builds institutional confidence rather than extending ambiguity, four characteristics are essential: bounded scope, measurable outcome, low stakes, and a defined review point.

Bounded scope means one tool, one use case, one team or individual, and a defined time period. For example: "A thirty day pilot in which the reference team uses Claude for drafting reference email responses, with all output reviewed before sending and patron data limited to anonymized inquiry descriptions." That definition is specific enough for the pilot to be evaluated. Everyone knows the tool, what it is being used for, the data protection constraint, and when the evaluation will occur. An unbounded pilot such as "let's try AI for various things over the next six months" cannot be reported on meaningfully, because there is no baseline and no defined outcome to measure.

Measurable outcome means deciding, before the pilot begins, what evidence will count as success. For reference email drafting, the measures might include time spent per email response before and after AI assistance, as assessed by the librarians themselves. They

might also include response quality, assessed through patron follow-up rates or explicit patron feedback, and any incidents that required modifying or pausing the pilot, documented as they occur. With prespecified measures, the end-of-pilot report can present data rather than impressions. That is the difference between an evaluation administration can act on and one that needs further clarification before a decision.

Low stakes means starting with internal administrative work rather than patron-facing services. Start with the work that would cause the least institutional exposure if a problem occurred: drafting internal communications, improving subject guide content that is reviewed before publication, or generating lesson plan structures that instructors evaluate before use. This sequence, internal first and patron-facing only after evidence of reliability, is also the most credible one for administrative purposes. It shows the library managing AI adoption incrementally and responsibly, rather than deploying tools to patron services before their performance is understood. A pilot in internal administrative work also produces outcome data that is easier to measure and report than patron-facing outcomes, where attributing changes to AI assistance is more complex.

The defined review point, a specific date or milestone at which outcomes are evaluated and a recommendation is made, turns a pilot from an indefinite experiment into a process that produces a decision. For example: "At the end of the thirty day period, we will present the reference team's documented time measures, any quality or incident observations, and a recommendation for whether to expand, modify, or conclude the use of AI for email drafting." This structure gives administration a timeline, a clear deliverable, and an explicit decision point. Each of these reduces the open-endedness that makes experimental programs hard to oversee.

Budget framing: presenting the cost of AI tools and training

Budget requests for AI tools and training work best when presented as investment proposals with estimable returns, not as requests to fund an experiment. The distinction matters in administrative terms. An experiment consumes resources to find out something unknown; an investment commits resources in anticipation of a specified return. A librarian who presents an AI subscription request as an experiment implicitly accepts that the outcome is uncertain and that the administration's skepticism is reasonable. One who presents it as an investment, with a stated expected return, a measurement approach, and a comparable cost reference, engages the administration as a partner in a resource decision with a defensible rationale.

Professional AI tool subscriptions typically cost between \$20 and \$25 per user per month. The institutional team plans, Claude Team and ChatGPT Business, now cost about the same

as the individual Claude Pro and ChatGPT Plus plans. Request the team plans for any pilot that touches patron inquiries, because they exclude conversations from model training by default and the individual plans do not. For a library team of five, a full-year subscription is approximately \$1,200 to \$1,500. Present that figure alongside a time savings estimate. If each team member saves thirty minutes per day on drafting and administrative tasks, the annual recovered capacity across five librarians is approximately 625 hours. At the median librarian salary of about \$64,000 plus benefits, that capacity has an institutional value well above the subscription cost. The administration does not have to accept the estimate as precise. It only has to find the estimate reasonable, and the comparison makes the scale relationship between cost and potential return visible.

Training investment requires a separate, adjacent argument, and the Clarivate Pulse data is directly relevant. In 2025, among libraries where AI literacy was part of formal training or onboarding, 28% were already at moderate or active implementation, a far higher share than among libraries without such programs. In 2026, only 13% of libraries made AI literacy part of formal training, while 30% reported little to no institutional focus on it. This gap in outcomes between institutions that invest in training and those that do not is an evidence-based argument for training as a prerequisite to effective adoption, not an optional professional development add-on. For example, a request to fund three librarians at a professional development workshop on AI in library work can be framed this way: "The institutions that are successfully adopting AI are the ones that have invested in building the professional capacity to do it well. This training is the foundational step that makes the rest of the pilot program viable." The ACRL AI Competencies framework also provides specific competency targets for designing and evaluating training. That specificity supports a professional development budget request far better than a general statement about staying current.

Drafting an institutional AI policy: why the library should lead

Most academic institutions do not yet have an AI acceptable use policy that addresses library-specific functions, data, and professional obligations. IT departments and legal offices are issuing general institutional AI guidance on academic integrity, data security, and vendor contracts. That guidance typically does not address the library's specific obligations around patron privacy, professional disclosure, and AI use in collections work. The policy governing library AI use should reflect library professional standards, not only institutional risk management priorities. For that to happen, the library must take an active part in policy development rather than passively receiving guidance from offices that do not share its professional context.

The ARL Guiding Principles for Artificial Intelligence, published in 2024, and the ACRL AI Competencies framework, published in October 2025, together provide the professional authority foundation for any library AI policy. The ARL Principles cover equitable access and AI literacy, awareness of bias, transparency in the AI technologies libraries license, human involvement in consequential decisions ("no human, no AI"), user privacy, and the copyright and licensing terms that protect scholarly use. These foundational orientations should inform any institutional policy. The ACRL Competencies define professional AI engagement across four areas, giving policy provisions a vocabulary that reflects professional practice rather than only risk avoidance. A library AI policy that references both documents is anchored in nationally recognized professional standards rather than institutional convention. That is a stronger position in policy conversations, and it gives the policy's specific provisions a defensible basis.

A minimum viable library AI policy covers four provisions. First, which tools are approved for staff use, with a brief rationale for each and the data security basis for including it. Second, what patron data can and cannot be entered into AI tools. Specifically, patron records, personally identifiable inquiry information, and any data protected under FERPA or library confidentiality law cannot go into commercial AI tools without explicit institutional authorization for that data category. Third, disclosure requirements when AI is used in patron-facing work. AI-assisted communications, guides, or finding aids must be reviewed by a professional before publication, and significant AI assistance in patron-facing content must be noted in appropriate institutional documentation. Fourth, who is responsible for reviewing AI output before it is published or used with patrons. This is a professional accountability provision rather than a technical one.

If the library drafts the policy instead of waiting for IT or legal to draft it, the library owns the framework. That ownership has practical consequences. A policy written by library professionals will include the FERPA obligations, patron privacy commitments, and professional disclosure norms that library practice requires. A policy written by IT or legal without library input may be technically sound but professionally insufficient. It might address data security but not the professional obligation to disclose AI use, or treat library patron data like general institutional data without recognizing the specific confidentiality protections library records carry. A library that brings a draft policy to administration also signals institutional maturity. It shows the library is not asking permission to experiment without a plan; it is proposing a responsible framework the institution can adopt with confidence.

PRACTITIONER NOTE

The administrative conversations that prove most impactful are often not with direct supervisors but with senior academic leadership: deans, provosts, and academic affairs offices that have been waiting for someone to bring them a governance framework rather than a request. In these conversations, a straightforward package tends to work well. It includes the Pulse of the Library data showing where academic libraries stand globally, the U.S. implementation gap alongside peer country comparisons, a single-page description of the proposed pilot, and a draft policy outline. What works is not framing the request as "we want to adopt AI" but as "we want to build the institutional infrastructure to do this responsibly, and here is what that looks like." That framing resolves a concern the administrator already carries instead of introducing a new priority, and it consistently moves the conversation forward.

KEY TAKEAWAYS

- Administrative skepticism about AI adoption is typically about risk, resources, and institutional positioning, not ideology. Effective advocacy identifies which concern is most operative and frames accordingly.
- The Clarivate Pulse of the Library 2026 data is the most widely cited external benchmark: only about one in six libraries has reached moderate or active AI implementation, and the U.S. gap (27% at initial implementation or beyond versus 43% to 52% in Asia and Europe, after just 7% U.S. optimism in 2025) reframes the conversation from experimentation to institutional positioning.
- The ACRL AI Competencies for Academic Library Workers (October 2025) and the ARL Guiding Principles (2024) provide the professional authority anchors for any administrative proposal, grounding a local request in nationally recognized professional standards.
- A well-designed pilot has four required characteristics: bounded scope (one tool, one use case, defined period), a prespecified measurable outcome, low institutional stakes to start, and a defined review point that generates a recommendation rather than extending ambiguity.
- AI tool subscriptions are most effectively presented as investment proposals with estimable returns: \$1,200 to \$1,500 per year for a five-person team on plans that keep patron inquiries out of model training, compared to hundreds of hours of recoverable staff capacity, a cost-return comparison that makes the scale relationship visible without requiring precision.

- The library should draft the institutional AI policy rather than waiting for IT or legal to do so without library input; a policy written by library professionals will include patron privacy, professional disclosure, and FERPA provisions that general institutional guidance is unlikely to address adequately.

References

APA 7TH EDITION

1. Association of College and Research Libraries. (2025, October). *AI competencies for academic library workers*. American Library Association. <https://www.ala.org/acrl/standards/ai>
2. Clarivate. (2026, September). *Pulse of the library 2026*. <https://doi.org/10.14322/pulse.of.the.library.2026>
3. Clarivate. (2025). *Pulse of the library 2025*. <https://clarivate.com/pulse-of-the-library/>
4. Lo, L. S. (2024). Evaluating AI literacy in academic libraries: A survey study with a focus on U.S. employees. *College & Research Libraries*, 85(5), 635–657. <https://crl.acrl.org/index.php/crl/article/view/26409>
Winner of the 2025 CALA Jing Liao Award for Best Research.
5. California colleges spend millions on faulty AI systems. (2026, March). *CalMatters*. <https://calmatters.org/education/higher-education/college-beat/2026/03/college-ai-chatbot/>
6. U.S. Bureau of Labor Statistics. (2025). Librarians and library media specialists. In *Occupational outlook handbook*. <https://www.bls.gov/ooh/education-training-and-library/librarians.htm>
Median annual wage \$64,320 (May 2024).

ACRL AI Competencies covered

Ethical Considerations

Use & Application

Sub-competencies: 1.4, 1.5, 4.2 · ACRL AI Competencies (2025)