

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 June 2026

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, not the professional uncertainty, and not the time required to learn. It is the conversation with administration. A librarian who has spent months developing a genuine, calibrated AI practice (who understands where the tools add value and where professional judgment is irreplaceable) often finds that bringing her institution along requires a different skill set entirely: the ability to translate what she has learned into language that resonates with institutional 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, and to understand the data, frameworks, and pilot design approaches that make that 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 institutional pressures: limited budgets that must be justified to provosts and boards, risk exposure from data privacy incidents or policy violations that have consequences for the institution, staff capacity that is already stretched, and the reasonable concern that an early commitment to the wrong tool or approach will require costly correction. In order to make an effective case for AI adoption, the librarian needs to understand which of these pressures is most operative for her specific administration, because the framing that resolves a risk concern is structurally different from the framing that resolves a budget concern, and neither is the same as the framing that

addresses the 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 who opens a conversation thinking about budget constraints does not become more receptive by hearing about AI's ability to draft subject guide descriptions faster. She becomes more receptive when she hears that AI reduces the staff time required for recurring writing tasks, and that the time savings are estimable in terms she already tracks. Such a reframe (from capability to institutional impact in the administrator's own terms) is the foundational move in effective administrative advocacy.

It is also worth understanding what administrators are hearing about AI from other institutional channels. IT departments, legal offices, provosts' offices, and accrediting bodies are all generating guidance about AI at an accelerating rate in 2025, and much of that guidance frames AI primarily as a compliance and risk management concern. The librarian who arrives at an administrative conversation aware of what the institution's IT acceptable use policy says, what the provost's office has communicated about AI in academic work, and what the institution's legal counsel has flagged as risk areas is in a substantially stronger position than one who arrives with library specific arguments alone. Such situational awareness signals that the librarian is approaching the conversation as an institutional partner rather than an advocate for a particular tool or practice, and it positions the library as a resource for the institution's broader AI decision making rather than 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 they are 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. A librarian who presents data showing where academic libraries are in AI adoption globally, where the U.S. sits relative to international peers, and what professional associations have determined about library AI competencies is participating in a sector level conversation, which is the appropriate register for an institutional decision.

The Clarivate Pulse of the Library survey is the most authoritative current benchmark for library AI adoption. The 2025 findings are particularly useful for administrative framing: 67%

of libraries globally are exploring or implementing AI, up from 63% in 2024, and the proportion of libraries at active implementation stages has tripled over the same period. Such a rate of change is meaningful in administrative terms; it indicates that institutions which are not yet evaluating AI are not maintaining a stable position relative to their peers, but are falling behind a moving field. Furthermore, the Pulse data identifies budget constraints as the primary barrier to AI adoption, cited by 62% of respondents, which reframes the budget conversation 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 shows that only 7% of U.S. librarians report optimism about AI, compared to 27-31% among their counterparts in Asia and other global regions. U.S. libraries are also significantly behind in implementation stages (14-16% at initial implementation or beyond, versus 37-40% in Asia and Europe). In order to use this data effectively without creating unproductive alarm, the framing should be specific: institutions that build library AI capacity now will be better positioned to support the AI literate graduates that employers and graduate programs are already expecting. Those that defer the decision are not avoiding a risk; they are accepting a different kind of risk, the risk of falling behind in a capability that is becoming a professional baseline.

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

Framing for different types of administrative skepticism

The same proposal requires different framing for different audiences, and the librarian who can identify which concern is most operative for her administration and adjust accordingly will be substantially more effective than one who presents a single argument to every listener.

Such audience awareness is not a form of manipulation; it is the professional practice of meeting people where they are, which is the same skill librarians apply in reference consultations and instruction sessions.

For risk averse administration, the most effective framing leads with risk mitigation rather than capability. The argument structure is: AI adoption without policy creates more institutional risk than adoption with policy, and a structured pilot with clear guidelines is less risky than the current situation where 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." Such a framing does not require the administrator to be enthusiastic about AI; it only requires her to recognize that structured adoption is more manageable than unmanaged individual adoption. The ARL Guiding Principles and ACRL AI Competencies framework should be presented as the professional authority backing the proposed approach, not as documents the administrator needs to read in detail.

For cost focused administration, the most effective framing quantifies 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 tasks, that is approximately 125 hours per year per full time librarian. At a median librarian salary of \$65,000, and accounting for benefits at a standard institutional multiplier, thirty minutes per day represents a meaningful recoverable capacity: time that can be redirected to instruction, reference, or collection work that serves institutional priorities directly. Such a calculation does not require precise data; it requires a defensible estimate that makes the time value of AI assistance visible in institutional cost terms. Furthermore, the cost of AI tool subscriptions, typically \$20-25 per user per month for professional tiers, is modest relative to even a conservative estimate of recovered staff time, and the comparison should be made explicit rather than left for the administrator to calculate independently.

For the administration that believes the institution is not yet at a stage where AI is relevant, typically expressed as "our faculty aren't asking about this" or "our students seem to be managing," the most effective reframe is student outcomes rather than operational efficiency. The argument is not that students are currently struggling, but that the AI literacy competencies they 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. Such a reframe shifts the conversation from "does the library need this" to "are we preparing our students for the environment they are entering", which is a question with 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 that is 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 that is difficult to report on and easy to question. In order to design a pilot that builds institutional confidence rather than extending ambiguity, four design 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." Such a definition is specific enough that the pilot can actually be evaluated: everyone knows what the tool is, what it is being used for, what the data protection constraint is, 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 identifying, before the pilot begins, what evidence will constitute success. For reference email drafting, the relevant measures might include time spent per email response before and after AI assistance, assessed by the librarians themselves; email response quality, assessed through patron follow up rates or explicit patron feedback; and any incidents requiring the pilot to be modified or paused, documented as they occur. Such prespecified measures mean the end of pilot report can present data rather than impressions, which is the difference between an evaluation that administration can act on and one that requires further clarification before a decision can be made.

Low stakes means starting with internal administrative work rather than patron facing services. The appropriate starting point is 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. Such a sequence (internal first, patron facing after evidence of reliability) is also the most credible sequence for administrative purposes, because it demonstrates that the library is managing AI adoption incrementally and responsibly rather than deploying tools directly to patron services before their performance is understood.

Furthermore, a pilot in the internal administrative domain produces outcome data that is easier to measure and report than patron facing outcomes, where attribution of changes to AI assistance is more complex.

The defined review point, a specific date or milestone at which the pilot outcomes are evaluated and a recommendation is made, is what converts a pilot from an indefinite experiment into a decision generating process. 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." Such a structure gives administration a timeline, a clear deliverable, and an explicit decision point: all of which reduce the openness that makes experimental programs difficult to oversee.

Budget framing: presenting the cost of AI tools and training

Budget requests for AI tools and training are most effective when they are presented as investment proposals with estimable returns rather than as requests for resources to cover 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 tool subscription request as an experiment is implicitly accepting that the outcome is uncertain and the administration's skepticism is reasonable. A librarian who presents it as an investment (with a stated expected return, a measurement approach, and a comparable cost reference) is engaging the administration as a partner in a resource allocation decision that has a defensible rationale.

Professional AI tool subscriptions for individual librarians typically cost between \$20 and \$30 per user per month for the tiers that provide meaningful professional functionality such as Claude Pro, ChatGPT Plus, and their equivalents. For a library team of five, a full year subscription commitment is approximately \$1,200-\$1,800 per year. Such a figure should be presented 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 a midrange librarian salary of \$60,000 plus benefits, that recovered capacity has an institutional value that substantially exceeds the subscription cost. Such a comparison does not require the administration to accept the time savings estimate as precise; it requires only that the estimate is reasonable and that the comparison makes the scale relationship between cost and potential return visible.

Training investment requires a separate and adjacent argument. The Clarivate Pulse data is directly relevant here: among libraries where AI literacy is part of formal training or

onboarding, 28% are already at moderate or active implementation, compared to just 8% where there is little to no institutional focus on AI training. Such a gap in implementation 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 rather than an optional professional development add on. For example, requesting funding for three librarians to attend a professional development workshop on AI applications in library work can be framed as: "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." Additionally, the ACRL AI Competencies framework provides specific competency targets against which training can be designed and evaluated, which is a level of specificity that supports a professional development budget request more effectively than a general statement about the value of 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 generating general institutional AI guidance (covering academic integrity, data security, and vendor contracts), but the library's specific obligations around patron privacy, professional disclosure, and collection related AI use are typically not addressed in general institutional guidance. In order to ensure that the policy framework governing library AI use reflects library professional standards rather than only institutional risk management priorities, the library needs to be a proactive participant in policy development rather than a passive recipient of guidance written by offices that do not share the library's 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 address governance, algorithmic equity, and the librarian's role as a responsible mediator of AI systems: foundational orientations that should inform any institutional policy. The ACRL Competencies define what professional AI engagement looks like across four areas, providing a vocabulary for policy provisions that reflect 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, which is a stronger position in institutional policy conversations and provides a defensible basis for the policy's specific provisions.

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

If the library drafts the policy rather than waiting for IT or legal to draft it, the library owns the framework. Such ownership has practical consequences: a policy written by library professionals will include provisions that reflect the FERPA obligations, the patron privacy commitments, and the professional disclosure norms that library practice requires. A policy written by IT or legal without library input may be technically sound but professionally insufficient: addressing data security without addressing the professional obligation to disclose AI use, for example, or treating library patron data as equivalent to general institutional data without recognizing the specific confidentiality protections that library records carry. Furthermore, a library that proactively brings a draft policy to administration signals institutional maturity: it is demonstrating that the library is not asking for permission to experiment without a plan, but proposing a responsible framework that the institution can adopt with confidence.

PRACTITIONER NOTE

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 such conversations, a straightforward package tends to work well: the Pulse of the Library data showing where academic libraries stand globally, the 7% U.S. optimism figure alongside peer country comparisons, a single page description of the pilot to run, and a draft policy outline. What works in such conversations 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 the concern the administrator is already carrying rather than 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 2025 data is the most authoritative external benchmark: 67% of libraries globally are exploring or implementing AI, and the U.S. confidence gap (7% optimism versus 27-31% in peer regions) 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-\$1,800 per year for a five person team, 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. (2024). *Pulse of the library 2024*. <https://doi.org/10.14322/pulse.of.the.library.2024>
3. Clarivate. (2025). *Pulse of the library 2025*. <https://clarivate.com/pulse-of-the-library/>

ACRL AI Competencies covered

Ethical Considerations

Use & Application

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