

Adaptive AI Audit Guide

An institutional planning tool for the strategy phase of any adaptive AI adoption

This guide turns the four EDUCAUSE AI Ethics principles most relevant to adaptive learning systems — Transparency, Nondiscrimination, Respect for Autonomy, and Accountability — into a practical audit framework. Use it before procurement, during deployment reviews, and whenever a vendor materially updates its models.

How to Use This Guide

When to run the audit

Before procurement, annually thereafter, and whenever the vendor announces a material model update. A vendor's answer at procurement may not hold six months into deployment.

Who participates

Institutional decision-makers with authority over procurement, learning strategy, and accountability for outcomes — typically including provosts, deans, directors of learning, and institutional effectiveness teams.

What to document

Capture institutional positions, escalation paths, review thresholds, and accountability decisions before deployment and revisit them on a recurring cadence.

What this guide is not

This is not a vendor feature checklist. It is a governance and learning strategy audit focused on whether institutional oversight is possible.

Transparency

Why it matters

Adaptive AI systems make instructional decisions that institutions must be able to investigate and explain.

Institutional Self-Audit Questions	Vendor Questions
• Do we know how the system arrives at recommendations and interventions? • Can we trace a specific decision back to the logic that produced it? • Are model updates documented and communicated clearly?	• Walk us through how the system would recommend X for learner profile Y. • What audit logs or decision trails are available on request? • How are material model changes communicated to institutions?

Contract & SLA Language to Require

- ☐ Vendor will provide documentation of model adaptation logic and material changes.
- ☐ Institutional requests for decision explanations will receive defined response timelines.
- ☐ Audit logs and relevant reporting access will be retained and made available upon request.

Documentation Prompts

- ☐ Who at the institution owns this principle operationally?
- ☐ What review cadence keeps this question active over time?
- ☐ What threshold or event would trigger a formal review?

Nondiscrimination

Why it matters

Bias in training data can produce uneven outcomes across learner groups unless institutions monitor for divergence.

Institutional Self-Audit Questions

- Do we know how the system performs across demographic groups?
- Have we defined thresholds for outcome divergence?
- Who reviews disaggregated outcome data and on what cadence?

Vendor Questions

- Can you provide disaggregated outcomes reports?
- Have independent bias audits been conducted?
- Can the model be recalibrated on institution-specific data?

Contract & SLA Language to Require

- ☐ Vendor will provide documentation of model adaptation logic and material changes.
- ☐ Institutional requests for decision explanations will receive defined response timelines.
- ☐ Audit logs and relevant reporting access will be retained and made available upon request.

Documentation Prompts

- ☐ Who at the institution owns this principle operationally?
- ☐ What review cadence keeps this question active over time?
- ☐ What threshold or event would trigger a formal review?

Respect for Autonomy

Why it matters

Learners should understand, question, and override system recommendations where appropriate.

Institutional Self-Audit Questions

- Do learners know when AI is shaping their learning path?
- Can learners override recommendations without penalty?
- Is there a documented appeals process for contested decisions?

Vendor Questions

- What learner-facing transparency does the system provide?
- What override mechanisms exist?
- Are there downstream penalties for overriding recommendations?

Contract & SLA Language to Require

- ☐ Vendor will provide documentation of model adaptation logic and material changes.
- ☐ Institutional requests for decision explanations will receive defined response timelines.
- ☐ Audit logs and relevant reporting access will be retained and made available upon request.

Documentation Prompts

- ☐ Who at the institution owns this principle operationally?
- ☐ What review cadence keeps this question active over time?
- ☐ What threshold or event would trigger a formal review?

Accountability

Why it matters

Institutions need clear ownership and authority when adaptive systems produce harmful or divergent outcomes.

Institutional Self-Audit Questions

- Who owns the institutional response when harm is identified?
- What authority exists to suspend or modify deployment?
- How are accountability decisions documented and revisited?

Vendor Questions

- What is the vendor response timeline when concerns are raised?
- Can the institution restrict or terminate use without vendor concurrence?
- How are liability and indemnification handled?

Contract & SLA Language to Require

- ☐ Vendor will provide documentation of model adaptation logic and material changes.
- ☐ Institutional requests for decision explanations will receive defined response timelines.
- ☐ Audit logs and relevant reporting access will be retained and made available upon request.

Documentation Prompts

- ☐ Who at the institution owns this principle operationally?
- ☐ What review cadence keeps this question active over time?
- ☐ What threshold or event would trigger a formal review?

About LX Studio

LX Studio is the learning design and learning technology consulting practice of the University of Central Oklahoma. We help institutions evaluate adaptive learning platforms, run institutional AI audits, and design learning environments that perform in the age of AI.

References

Georgieva, M., & Stuart, J. (2025). Ethics is the edge: The future of AI in higher education. EDUCAUSE Review.

Georgieva, M., Webb, J., Stuart, J., Bell, J., Crawford, S., & Ritter-Guth, B. (2025). AI ethical guidelines. EDUCAUSE Library.

Gándara, D., et al. (2024). Detecting and mitigating algorithmic bias across racialized groups in college student-success prediction.