

ILLUSTRATIVE WORKED EXAMPLE

Campus generative AI policy review

A public, source-linked demonstration of the EX MULTIS review record

Reviewed September 3, 2026 | Public sample | Sources linked

BOTTOM LINE

Adopt a 90-day controlled pilot and an interim policy, but do not authorize regulated-data workflows until provider terms, institutional controls, and applicable law have been reviewed.

Conditions for approval

1. Define permitted, restricted, and prohibited uses by learning context rather than imposing one rule on every course.
2. Keep instructors responsible for course-level expectations, assessment design, and disclosure requirements.
3. Prohibit student records, confidential research, credentials, and other sensitive data in unapproved AI services.
4. Require human review of consequential outputs and independent verification of material citations.
5. Provide accessible alternatives, training, notice, and a process for students to question or appeal AI-assisted decisions.
6. Measure learning value, citation defects, privacy incidents, support demand, and unequal access before broader adoption.

01 - ASSIGNMENT AND ROUTE

The question and decision boundary

ORIGINAL QUESTION

Should North Valley Public University adopt a campus-wide generative AI policy for teaching and learning in the 2026-27 academic year?

North Valley Public University is fictional. This demonstration shows the structure and discipline of an EX MULTIS review; it is not a benchmark result, legal advice, or evidence that every run will reach the same conclusion.

Constraints supplied to the review

- Public university setting; federal baseline considered, but state law and local policy were not supplied.
- Teaching and learning only; admissions, employment, discipline, grading automation, and clinical uses are outside scope.
- No student personally identifiable information may enter an unapproved service during the pilot.
- Faculty retain academic authority, and final institutional approval remains human.

INTERSECT routing record

Step	Routing decision
Classify	Higher education policy, privacy, academic integrity, accessibility, and AI governance
Set risk	Elevated institutional review because the policy affects students, faculty, data handling, and assessment
Order authority	Applicable law and official guidance, then institutional policy and pedagogical evidence
Assign roles	Policy analyst, adversarial reviewer, source verifier, and human decision owner
Set the gate	No regulated-data use until privacy, security, procurement, and legal owners approve the provider

02 - INDEPENDENT WORK

Separate analyses before synthesis

Each role below formed its position before reconciliation. Separation reduces cross-anchoring; it does not guarantee statistical independence or eliminate shared error.

Policy analyst: Adopt with controls

A controlled policy can replace inconsistent informal use with common disclosure, verification, privacy, and accountability rules while leaving course-level judgment with instructors.

ASSUMPTION CARRIED INTO REVIEW

The university can provide training, approved tools, an accessible alternative, and meaningful pilot measurement.

Adversarial reviewer: Delay broad authorization

A campus-wide launch could normalize unsafe data handling, unequal access, unreliable citations, and unclear enforcement before provider contracts and support capacity are ready.

ASSUMPTION CARRIED INTO REVIEW

An interim policy can be adopted without implying that every AI use or provider has been institutionally approved.

Source verifier: Evidence supports governed experimentation, not blanket approval

The reviewed authorities support human-centered governance, documented risk management, privacy controls, and continuing evaluation. They do not prove that one campus policy will improve learning outcomes.

ASSUMPTION CARRIED INTO REVIEW

State requirements, collective-bargaining obligations, accessibility findings, and individual provider terms still require local review.

03 - CLAIM LEDGER AND FACT CHECK

Material claims remain attributable

Claim	Final status	Evidence record
FERPA categorically prohibits a university from using a third-party AI service with education records.	Contradicted	U.S. Department of Education materials describe circumstances in which a contractor or service provider may qualify under the school-official exception if the required conditions are met. Sources: S4, S5
The institution should define human responsibility and document AI governance decisions.	Supported	NIST describes governance, documented roles, transparency, human review, and ongoing risk management as central practices. Sources: S1
One identical classroom rule should apply across every discipline and assignment.	Disputed	The sources support coherent institutional governance while also emphasizing context, human agency, pedagogical suitability, and risk-based implementation. Sources: S1, S2, S3
Faculty approval alone is enough to permit student records to be entered into any AI tool.	Contradicted	FERPA rules and Department guidance impose conditions beyond an individual instructor's approval, and other laws or institutional requirements may also apply. Sources: S4, S5
A pilot should include measurement, monitoring, and a human decision before expansion.	Supported	NIST frames AI risk management as continuous and calls for measurement and governance across the lifecycle; education guidance emphasizes human-centered implementation and evaluation. Sources: S1, S2, S3

Correction caused by fact-checking

SHARED INITIAL ERROR

Three roles initially accepted the blanket statement. FERPA categorically prohibits a university from using a third-party AI service with education records.

U.S. Department of Education materials describe circumstances in which a contractor or service provider may qualify under the school-official exception if the required conditions are met.

CORRECTED DECISION RULE

The final report replaces the categorical prohibition with a conditional rule: do not disclose protected records unless counsel and privacy owners confirm a valid basis, direct control, authorized use, redisclosure limits, and all other applicable requirements.

04 - DISSENT AND HUMAN JUDGMENT

Material disagreement stays visible

Policy timing

Majority view: Adopt an interim policy alongside a controlled pilot.

Minority view: Delay any adoption until approved-provider and accessibility reviews are complete.

HUMAN RESOLUTION

Adopt only minimum safeguards now; defer broad tool authorization and protected-data use.

Uniformity across courses

Majority view: Use institution-wide disclosure and privacy baselines.

Minority view: A single disclosure rule may be too rigid for different disciplines and assignments.

HUMAN RESOLUTION

Set common minimum safeguards while instructors publish course- and assignment-specific expectations.

Recorded human decision

Illustrative decision owner: Provost with CIO, privacy, accessibility, faculty, and student representatives

DECISION

Approve a 90-day pilot and interim policy; do not approve regulated-data workflows.

Return for a stop, revise, or expand decision after the pilot evidence is reviewed.

Pilot measures

- Material citation-defect and correction rate
- Privacy, security, and academic-integrity incidents
- Student and faculty usefulness, accessibility, and support demand
- Differences in access or outcomes across courses and student groups
- Frequency and rationale for human overrides

05 - LIMITS AND SOURCE MAP

What remains unresolved

- No state law, collective-bargaining agreement, accreditation rule, disability accommodation record, or existing university policy was supplied.
- No specific AI provider contract, retention schedule, training-use term, security assessment, or data-residency commitment was reviewed.
- The cited sources support governance principles and legal boundaries; they do not establish that this pilot will improve learning outcomes.
- Source status and product capabilities can change. Material conclusions must be rechecked at the time of an actual decision.

Source map

S1 - NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0)

National Institute of Standards and Technology | January 2023

Used for: Governance, documentation, human roles, measurement, monitoring, and lifecycle risk management

<https://doi.org/10.6028/NIST.AI.100-1>

S2 - Artificial Intelligence and the Future of Teaching and Learning

U.S. Department of Education, Office of Educational Technology | May 2023

Used for: Human-centered educational implementation, policy development, opportunities, and risks

<https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>

S3 - Guidance for generative AI in education and research

UNESCO | September 2023; page updated January 2026

Used for: Privacy, human agency, inclusion, validation, policy development, and pedagogical suitability

<https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>

S4 - FERPA regulations, including 34 CFR 99.31

U.S. Department of Education, Student Privacy Policy Office | Current public regulation text reviewed September 2026

Used for: Conditions under which prior consent is not required and requirements applying to outsourced service providers

<https://studentprivacy.ed.gov/ferpa>

S5 - Responsibilities of Third-Party Service Providers under FERPA

U.S. Department of Education, Privacy Technical Assistance Center | Department guidance

Used for: School-official exception, direct control, authorized purposes, redisclosure limits, and contracting practices

https://studentprivacy.ed.gov/sites/default/files/resource_document/file/Vendor%20FAQ.pdf

DEMONSTRATION BOUNDARY

This worked example demonstrates the report structure and correction trail. It does not establish product accuracy, eliminate the need for source verification, or replace institutional legal, privacy, security, accessibility, procurement, or professional review.