

— REFERENCE ARCHITECTURE · EDUCATION

Education AI Teacher Workflow

The agentic workflow behind a teacher's day — planning, authoring, grading, insights, differentiation, and family comms. Grounded, FERPA-safe, and teacher-approved.

Architecture · agents · grounding · fairness · approval · specs →



Draft freely, approve carefully

Six specialized agents do the teacher's heavy lifting — but everything is grounded in your curriculum, FERPA-checked, fairness-reviewed, and signed off by the teacher before it reaches a student or family.

Six agents, one teacher

- Lesson & Unit Planner — standards-aligned plans, pacing, and activities from your curriculum.
- Assessment Author — quizzes, assignments, and rubrics tagged to the standards you're teaching.
- Grading & Feedback — rubric-aligned scores and specific feedback drafts, for you to approve.
- Class Insights — mastery trends, coverage gaps, and which students are at risk.
- Differentiation & IEP — reading-leveled, accommodation-aware variants from one base plan.
- Family & Student Comms — progress notes and messages home, in your voice.

Safe by design

- Every plan, item, and message draws only on the approved curriculum + standards — cite-or-refuse.
- Student PII is redacted before any prompt leaves the boundary (FERPA / COPPA-aware).
- Roster, gradebook, and accommodations are read to ground the work — never leaked into output.
- Every draft and decision is logged to an audit trail for teachers and admins.

Checked, then teacher-approved

- A fairness & bias review checks grades and feedback for rubric adherence and consistency.
- Nothing is posted autonomously — the teacher reviews, edits, and signs every output.
- Risk-tiered review focuses your time on flagged, high-stakes items; routine work is auto-routed.
- You can disable any agent — the human always has the final say.

Intake & orchestration

— Teacher / Request

Plain-language intent

EX A 7th-grade teacher types: 'Plan next week's unit on linear equations, and tell me who's still behind on fractions.'

TECH Web / LMS dashboard, LTI 1.3 launch; OAuth2/OIDC SSO; per-teacher + per-class scope; Request envelope carries class + standard codes

CTRL Teacher identity + class scope bound at entry; Consent + data-use policy applied; Rate-limit + abuse filter before any model runs

FIN A cheap intent classifier routes simple asks to templates so they never reach a premium model.

FALL An ambiguous request is clarified in one prompt rather than guessed at.

— Planner / Orchestrator

Plans + routes the work

EX It splits the request into 'plan the unit' and 'find students behind on fractions', routing each to a different agent.

TECH Orchestrator (Semantic Kernel / LangGraph); Deterministic, versioned routing policy; Per-task step + token budget, circuit breakers

CTRL Tool-RBAC allow-lists per agent; Class-scope enforcement on every agent; No agent acts outside the teacher's roster

FIN Per-task budget caps spend; a cheap model routes, premium tier only on genuinely hard generation.

FALL Budget or breaker trips -> return the work done so far and hand the rest to the teacher; state preserved.

Planning & authoring

— Lesson & Unit Planner

Standards-aligned plans

EX Drafts a 5-day linear-equations unit with objectives, a warm-up, guided practice, and an exit ticket per standard.

TECH Retrieval-grounded generation over curriculum; Standards-framework tagging; Editable plan templates

CTRL Cite-or-refuse to curriculum + standards; Read-only curriculum access; Teacher edits before use

FIN Reuses cached unit templates; the model fills gaps rather than writing every plan from scratch.

FALL If the curriculum lacks a standard's material, it flags the gap instead of inventing content.

— Assessment Author

Quizzes + rubrics

EX Generates a 10-item quiz and a 4-level rubric on solving two-step equations, tagged to the standard.

TECH Item + rubric generation grounded in standards; Difficulty calibration; Answer-key + distractor rationale

CTRL Standards-tagged items only; Bias / fairness checks on items; No high-stakes use without teacher review

FIN Draws from an item bank first; new items are generated only to fill coverage gaps.

FALL Low-confidence items are marked 'draft - review' rather than added to a graded assessment.

Grading & insights

— Grading & Feedback

Rubric scores + feedback

EX Scores 28 short-answer submissions against the rubric and drafts two-sentence feedback for each - for the teacher to approve.

TECH Rubric-aligned scoring; Evidence-linked feedback generation; Submission ingestion from the LMS

CTRL Teacher signs every grade; Rubric adherence enforced; Consistency + fairness check before review

FIN Batches submissions and reuses the rubric context, cutting per-item token cost.

FALL Uncertain or borderline scores are escalated to the teacher rather than recorded.

— Class Insights

Mastery + risk insights

EX Flags that 6 students are below mastery on fractions and that the class is weak on word problems overall.

TECH Knowledge-tracing / mastery model; Gradebook + assessment aggregation; Standards-coverage analytics

CTRL Insights scoped to the teacher's class; No profiling beyond educational purpose; Accommodations respected in flags

FIN Runs on incremental gradebook events, not full recomputation, keeping compute cheap.

FALL Sparse data -> it reports low confidence rather than over-claiming a trend.

Differentiation & comms

— Differentiation & IEP

IEP / 504 / ELL variants

EX *Produces a reading-leveled version of the worksheet and an extended-time variant for two students with IEPs.*

TECH Accommodation-aware rewriting; Reading-level + translation support; Variant generation from a base asset

CTRL Accommodations pulled from the record, applied faithfully; FERPA-scoped access; Teacher approves each variant

FIN Generates variants from the approved base asset rather than authoring each from scratch.

FALL If an accommodation is unclear, it defaults to the most supportive option and flags for the teacher.

— Family & Student Comms

Messages home, your voice

EX *Drafts a warm note home: 'Maya improved on equations this week; here's one way to support fractions at home.'*

TECH Template + tone-matched generation; Record-grounded personalization; Multi-language drafting

CTRL FERPA-safe: only shareable data included; No message sent without teacher approval; PII handling audited

FIN Template-driven drafts keep most messages near-free; the model personalizes only the specifics.

FALL If shareable data is insufficient, it asks the teacher rather than guessing details.

Grounding — curriculum + records

— Curriculum + Standards

The only content source

EX *Retrieves the district's linear-equations unit and the matching standards so plans and items stay aligned.*

TECH Vector store + hybrid search + rerank; Standards-framework index; Versioning + effective-date awareness

CTRL Only approved, signed materials indexed; Access scoped by course / grade; Stale-edition guard

FIN Index once, query cheaply; rerank only top-k; stable units cached to skip retrieval.

FALL No grounded source -> the agent declines and points the teacher to the gap rather than guessing.

— Roster / LMS / Gradebook

Roster / gradebook (read)

EX *Reads the gradebook, roster, and accommodations to ground grading, insights, and messages in real data.*

TECH LMS / SIS read via LTI / OneRoster; Field-level access scoping; Read-only credentials

CTRL FERPA-aligned access controls; PII masking in prompts; Audit of every record read

FIN Reads cached, scoped class data rather than hot SIS calls per task.

FALL If the record is unavailable, it proceeds on session context only and marks outputs provisional.

Safeguards

— Grounding + Safety + FERPA

Grounded + FERPA-safe

EX *Blocks a feedback draft that accidentally referenced another student's score, returning a clean version.*

TECH Grounding verifier (cite-or-refuse); Safety + age-band classifiers; PII / leak detection

CTRL Unsafe or leaking output blocked, not shown; Citations required on claims; Escalation to teacher on flags

FIN Deterministic checks are near-zero cost and run before review - far cheaper than a FERPA incident.

FALL Any check fails -> the draft is withheld and routed back with the reason logged.

— Fairness & Bias Review

Bias + consistency gate

EX *Catches that the grading agent scored two equivalent answers differently and flags both for the teacher.*

TECH Consistency scoring across submissions; Bias / tone screening; Rubric-conformance checks

CTRL Must pass before reaching the teacher; Disagreements escalate, not auto-resolve; Decision recorded for audit

FIN A short review pass is far cheaper than an unfair grade or a complaint later.

FALL If fairness can't be confirmed, the item goes to the teacher unresolved rather than through.

Teacher approval

— Teacher Approval (HITL)

Teacher signs everything

EX The teacher skims the unit plan, tweaks two questions, approves 26 grades, edits one message, and signs off.

TECH Review queue + dashboard; One-click edit + approve / reject; Override feeds back into the agents

CTRL Human sign-off mandatory for anything student-facing; Full audit trail of AI + human actions; Teacher can disable any agent

FIN Risk-tiered review focuses teacher time on flagged + high-stakes items, auto-routing the routine.

FALL Anything uncertain lands here - the default path is always a human.

YOUR TURN

Bringing AI into teaching safely?

I help schools and edtech teams put agentic AI into the teacher workflow — grounded in your curriculum, FERPA-safe, fairness-checked, and teacher-led. On Azure.

DM me “teacher”

Atiqul Islam

Fractional Cloud & AI Architect for SMBs · atiqulislam.dev