ExcelinEd

Guardrails for Student-Facing Digital Tools in K-12 Schools

Frequently Asked Questions | 2026

What Are AI-Powered Educational Tools?

AI-powered educational tools are digital programs that deliver instruction, feedback or tutoring directly to students in real time. Unlike general AI tools like ChatGPT or AI companions, AI-powered educational tools are designed to personalize instruction and provide timely feedback. When designed and used responsibly, these tools can improve student outcomes.

Learn more about AI companions here.

Why Do AI-Powered Educational Tools Need Guardrails?

A recent study reveals that currently, the risks of AI in K-12 education outweigh the benefits. AI-powered educational tools are being adopted quickly in the classroom while oversight and safety standards lag behind. Guardrails help protect student safety and privacy; prevent manipulative or inappropriate interactions; enable teachers to monitor and review student AI conversations; and ensure operators are transparent and responsive to school districts. Common sense guardrails will help districts deploy AI to support educators and students alike.

What Guardrails Are Needed?

To protect students, AI-powered educational tools used in classrooms must follow strict privacy, security, transparency and design expectations.

Data Privacy and Security

- Comply with all federal and state student privacy and security laws
- Maintain evidence of strong data-protection through a third-party audit
- Limit unnecessary data collection
- Prohibit operators from commercially using student data
- Prohibit operators from using personally identifiable information or biometric identifiers to train, improve or develop AI models
- Require operators to permanently destroy student data when the contract ends
- Require strict limits on use of biometric identifiers
- Require any subcontractor with access to student data to follow the same privacy and security rules as the primary operator

Transparency

- Keep records of all student inputs and AI outputs for 12 months
- Explain each tool's purpose, expected benefits and evidence base to families
- Monitor AI outputs for accuracy, bias and error rates
- Disclose the implementation conditions needed for positive student outcomes
- Report safety or privacy incidents in a timely manner

Software Design

- Prohibit addictive design features
- Limit human-like behaviors designed to engage in personal or emotional relationships
- Prevent automated completion of assignments
- Redirect off-task behavior
- Block access to inappropriate or non-educational content

Why Must Operators Keep Student Records for 12 Months?

A 12-month retention period allows states and districts to look back on and investigate any material safety or privacy incidents from the past school year. This period continues after a contract ends, and operators must permanently destroy the records once it expires. Retaining student inputs and AI outputs any longer would leave operators and their tools more vulnerable to data breaches that could expose sensitive student information.

How Do We Know Which Tools Actually Improve Learning?

To show educational efficacy, AI-powered educational tools must demonstrate measurable, verifiable impact on student outcomes through one of three pathways:

- Show alignment with Tier I or II evidence standards under the Every Student Succeeds Act (ESSA).
- Meet the Tier IV evidence standard under ESSA, which requires a tool to demonstrate a rationale based on current research, and participate in an Outcomes-Based Contract.
- Pilot the tool through a state-administered sandbox program.

What Are Outcomes-Based Contracts (OBC)?

An OBC is a performance-driven payment mechanism that ties a portion of an operator's financial compensation to the achievement of agreed-upon, measurable student outcome goals. The Center for Outcomes Based Contracting recommends that at least 40 percent of the overall payment be contingent on meeting these outcomes. A high-quality OBC includes several essential components:

- Identifying the student population targeted by the tool.
- Establishing measurable student outcome goals and evaluation metrics.
- Defining dosage thresholds and other implementation benchmarks needed to achieve target student outcomes.
- Making payments contingent on student outcomes.
- Clearly outlining implementation responsibilities for the operator and the district.
- Maintaining a continuous data collection, reporting and monitoring plan.

This approach balances rigorous evidence with flexibility for innovation by giving districts a mechanism to try new products with less financial risk, ensuring districts only fully pay for tools that successfully deliver results. Furthermore, recent research from WestEd indicates that participating in OBCs can help products build causal evidence of their impact.

What Is a State-Administered Sandbox Program?

A sandbox program offers a safe, structured pathway to pilot emerging educational technologies while building empirical evidence of their impact. To participate, operators must pass rigorous security audits and adversarial safety testing (i.e., red teaming). Parents are provided with an opt-out mechanism from the pilot with no academic penalty, and an independent evaluator holistically assesses the tool's performance before it can be recommended for state approval.

How Will Districts Be Supported in Evaluating AI-Powered Educational Tools?

The state will establish and maintain an evaluation framework that addresses safety, privacy, transparency, pedagogical alignment, educational efficacy, accessibility and interoperability. The state will continuously review and publish a publicly available list of AI-powered educational tools that meet these framework criteria. By providing model policies, procurement guidance, sample OBCs and continuous monitoring checkpoints, the state removes the heavy lifting from local districts while ensuring adopted tools are effective.

Should AI-Powered Educational Tools Be Banned?

No. When designed and used responsibly, AI-powered educational tools can personalize learning and improve student outcomes. Research shows that mastery-based AI platforms increase student math scores, and students who receive AI-supported instruction demonstrate higher reading comprehension than those who only receive traditional face-to-face instruction. Strong guardrails ensure AI-powered educational tools achieve their intended benefits while keeping students safe.

How Will Parents Be Informed About AI Use in the Classroom?

Districts must provide families with a plain-language notice at the beginning of each school year and shortly after deploying a new AI-powered educational tool. As AI becomes more integrated into educational technology products, districts should pause to conduct due diligence before adopting and deploying a new tool. Requiring districts to identify each tool, explain its purpose, summarize the supporting evidence, disclose any biometric data collection and describe their policy for AI-powered educational tools adopted by teachers without formal district approval helps ensure they have carefully reviewed a tool before using it with students.

Does This Model Policy Apply to Tools Used by Teachers?

No. This model policy does not apply to tools used exclusively for lesson planning or preparation. It also does not apply to learning and classroom management systems (ex: Google Classroom), student monitoring systems (ex: GoGuardian) or assessment delivery platforms. Educators should be encouraged to use AI tools that enhance their professional expertise and instructional capacity, in alignment with district policies.

Does This Model Policy Conflict with the Federal AI Executive Order?

No. The executive order explicitly preserves state authority over child safety protections, as well as state government procurement and use of AI.

Appendix

- AI Red Teaming: A structured testing effort to find flaws and vulnerabilities in an AI system, often in a controlled environment and in collaboration with developers.
- Algorithmic Impact Assessment: A documented evaluation of an AI system's purpose, data sources, risks, accuracy and potential bias.
- Children's Internet Protection Act (CIPA): Federal law that requires schools and libraries receiving E-rate federal funds to implement internet safety policies and technology protections to block or filter harmful online content.
- Children's Online Privacy Protection Act (COPPA): Federal law that requires websites to receive parental consent before collecting personal information from children under 13.
- Family Educational Rights and Privacy Act (FERPA): Federal law that protects the privacy of student education records and gives parents rights to access and control disclosure of those records.
- ISO/IEC 42001: An international management system standard that sets requirements for how organizations govern, assess and continuously improve AI systems.
- SOC 2: An independent audit framework, typically performed by a Certified Public Accountant (CPA) firm, that verifies a company has strong data security and privacy controls in place.

See Excel*in*Ed's Guardrails for Student-Facing Digital Tools in K-12 Schools model policy for more.