Speakable

WHITEPAPER · TRUST & SAFETY

Responsible AI Use in K-12 Schools

How Speakable delivers AI-powered feedback and grading while protecting student data, keeping teachers in control, and setting firm boundaries for student use.

Audience: District and school leaders, IT and data-privacy officers, procurement teams
Published by: Speakable Technologies, Inc.
Version: 1.0

1 Executive Summary

Artificial intelligence has moved from the periphery of education technology to the center of everyday instruction. For district leaders, this shift raises an unavoidable question: how do you capture the benefits of AI, such as faster feedback, more consistent grading, and more practice for every student, without exposing students to the risks that have made many institutions hesitant to adopt it?

Speakable is built to answer that question. It gives teachers AI-powered support for feedback and grading without using student data to train AI models and without giving students open-ended AI chat. Safety and compliance are not features layered on after launch; they are architectural decisions made from the first line of code.

This whitepaper explains the three principles that govern how AI operates inside Speakable, details the compliance framework the platform is designed to meet, and answers the questions districts most often raise during evaluation and procurement.

Contents

2 The AI Adoption Dilemma in K-12

Schools face real pressure to adopt AI. Teachers are stretched thin, feedback cycles are slow, and the demand for individualized practice outpaces the hours available in any school day. AI can help close these gaps. Yet the same capabilities that make general-purpose AI powerful also make it difficult to deploy responsibly with students.

Three concerns come up repeatedly when districts evaluate AI tools. The first is data: where does student information go, who can access it, and is it being used to train models that live outside the district's control? The second is oversight: if an algorithm assigns a grade, who is accountable for that grade, and can an educator intervene? The third is exposure: open-ended chatbots can produce unpredictable, inappropriate, or off-task content, and they invite students to stray far from the lesson at hand.

A responsible AI tool for schools has to resolve all three concerns at once. Addressing only data privacy while leaving students in front of an open chatbot, or keeping teachers in the loop while quietly training models on student work, leaves a district exposed. Speakable's design treats these not as trade-offs to balance but as constraints to satisfy simultaneously.

3 Speakable's Design Principles for Safe AI

Three principles define how AI behaves inside Speakable. Each maps directly to one of the concerns above, and each is enforced by the product's architecture rather than by policy alone.

PRINCIPLE 1: DATA PRIVACY

Student PII is not sent to the AI model, and your data never trains it

Speakable is designed so that student personally identifiable information (PII) is not sent to the AI model provider. When a student completes an activity, only the student's work itself, such as a recording or a written response, is passed to the model for evaluation. Identifying details like the student's name and email address are not included. The model returns a result (feedback and a suggested score), and Speakable associates that result with the correct student inside its own system. Identity and work are kept separate at the point of AI processing.

Beyond that separation, student data is never used to train AI models, is never sold to third parties, and is not retained beyond what is needed to deliver Speakable's service. Because student work does not feed model training and is never paired with identifying information at the model, sensitive information cannot resurface in another context or leave the boundary of the service the district has authorized.

PRINCIPLE 2: HUMAN OVERSIGHT

Teachers approve every AI score

Speakable evaluates student work against the teacher's rubric automatically, but those scores are suggestions, not verdicts. The teacher reviews each score, adjusts it where their judgment differs from the model's, and finalizes the grade before anything is recorded. The division of labor is deliberate: the AI accelerates the work of assessment, and the teacher remains the accountable decision-maker. This keeps professional judgment, and responsibility for a student's grade, with the educator.

PRINCIPLE 3: BOUNDED STUDENT ACCESS

Students never have unsupervised AI access

Speakable includes no chatbot and no free-form prompting for students. Every interaction a student has with the platform is scoped to an activity a teacher designed. Students practice within the guardrails their teacher has set, and cannot direct the AI outside that scope. This eliminates the most common source of risk in student-facing AI, the open-ended prompt, while preserving the practice and feedback that make the tool valuable.

> **The AI assists. The teacher decides. The student stays within the boundaries the school has set.**

4 Compliance and Security Framework

Beyond its design principles, Speakable is built to meet the security and compliance requirements K-12 institutions need in order to use AI with students confidently. The platform processes student data only as directed by the institution and never for commercial purposes such as advertising or profiling. The table below summarizes the standards Speakable is designed to meet.

Standard	What it covers and how Speakable meets it
FERPA	The Family Educational Rights and Privacy Act protects student education records. Speakable processes student data only as directed by the institution and never for commercial use.
COPPA	The Children's Online Privacy Protection Act governs the online collection of data from children under 13. Speakable is compliant and never uses student data for advertising or profiling.
SOPPA	The Student Online Personal Protection Act and similar state-level student data privacy laws impose additional obligations on operators of educational services. Speakable is built to meet these requirements.
Data Processing Agreement	Districts can request a signed Data Processing Agreement (DPA) for procurement and compliance review, documenting how student data is handled contractually.

Taken together, these commitments mean that student data stays within the service the district has authorized, is governed by contract, and is never repurposed for commercial ends.

5 Frequently Asked Questions

Q Does student data ever train Speakable's AI?
No. Student data, including recordings, scores, and written responses, is never used to train AI models. It is processed only to deliver Speakable's service to your school.

Q What does Speakable actually send to the AI model?
Only the student work itself, such as a recording or written response. Student personally identifiable information, including name and email address, is not sent to the AI model provider. When the model returns feedback and a suggested score, Speakable associates that result with the correct student inside its own system.

Q Who controls final grades: the AI or the teacher?
The teacher. AI scores are a starting point. Teachers review every score, adjust as needed, and approve grades before they are finalized.

Q Can students have open chats with Speakable's AI?
No. Speakable does not include an open chatbot for students. Every student interaction is tied to a specific activity a teacher built, and students cannot go outside that scope.

Q What compliance certifications does Speakable hold?
Speakable is FERPA, COPPA, and SOPPA compliant. Districts can request a signed Data Processing Agreement.

Q Does Speakable share student data with third parties?
No. Student data is never sold or shared with third parties. It is processed only to deliver Speakable's service to your institution.

Q How can districts request a Data Processing Agreement?
Contact the Speakable team using the details in the next section, and we will provide a DPA for your procurement and compliance review.

6 Conclusion & Next Steps

AI can make a meaningful difference in how quickly students receive feedback and how much targeted practice they get, but only if it is deployed in a way that schools can trust. Speakable's position is that safety is not a constraint on that value; it is the precondition for it. By ensuring that student data never trains its models, that teachers approve every score, and that students never encounter an open-ended AI, Speakable lets districts adopt AI with students confidently rather than cautiously.

AI built for schools. Safe to use with students.

Every decision in Speakable is designed to protect student data, keep teachers in control, and keep students within the boundaries you set. To move forward, districts can:

→ Request a signed Data Processing Agreement (DPA)
→ Review Speakable's privacy and security documentation
→ Schedule a walkthrough for the district evaluation team

Get in touch: speakable.io/contact · speakable.io/resources/ai-safety