

The AI Policy Brief for School Leaders

July 2026 edition · The cited research, a decision framework, and a vendor checklist — written to be forwarded to a board, a faculty meeting, or a parents' association as-is.

PART ONE · THE EVIDENCE

The research brief.

The cited research in one place — and the single shape it forms: students are using AI to skip the thinking, and the cognitive baseline was already falling before it arrived.

95%

Student-AI activity is overwhelmingly work-completion, not learning.

Across 1.2 million monitored student-AI sessions in 1,300+ districts, 95% of interactions were students attempting to have the AI complete assigned work entirely. *Education Week / Securly, March 2026.*

84%

Faculty closest to the work see originality declining.

84% of college faculty report AI is directly reducing student originality and deep engagement — near-universal agreement among the people who grade the work. *College Board, February 2026.*

67%

Students themselves now report the cognitive cost.

67% of students say using AI for schoolwork is eroding their critical thinking — up from 54% earlier in the same school year. *RAND Corporation, March 2026.*

48→62%

Adoption crossed into the majority in eight months.

The share of students using AI for homework rose from 48% to 62% between May and December 2025, driven by middle and high schoolers — the age range where habits of effortful learning form. *RAND, March 2026.*

−7 pts

The U.S. cognitive baseline was already weakening.

U.S. nine-year-olds posted their first-ever NAEP math decline (−7 points, 2020–22) and the largest reading drop since 1990. AI arrived during a learning crisis, and made shortcuts trivial. *NCES, NAEP 2022.*

−15 pts

The decline is global — and predates the pandemic explanation.

PISA 2022 recorded the steepest math decline in the program's history across 81 economies; OECD analysis found the drop began before COVID and continued after reopenings. *OECD, December 2023.*

Full citations and context: orchen.ai/the-problem

Three honest options, and what each one costs.

Option A — Ban AI on school networks and devices.

What it buys: a clear public position and short-term board comfort. **What it costs:** the activity moves to personal devices within weeks. The school loses all visibility, the policy becomes unenforceable theater, and students learn that the institution's rules and reality diverge. Roughly 40% of schools have chosen this; monitoring data shows usage continues regardless.

Option B — Permit consumer AI with guidelines.

What it buys: honesty about reality and low cost. **What it costs:** the school is endorsing tools designed for task completion, with no pedagogy, no school visibility, no crisis routing, and student data on consumer terms. Guidelines without an environment to apply them in are advice, not policy — and consumer terms of service were never written to satisfy FERPA or COPPA obligations, which quietly shifts data liability to the school.

Option C — Provide a governed environment, and set policy around it.

What it buys: an AI the school configured to teach rather than complete work, structured visibility for faculty and parents, crisis detection that routes to a counselor, and a defensible data posture. **What it costs:** budget, an implementation effort, and understanding that a school governs its own environment rather than the open internet — which is exactly why that environment needs to be one students genuinely prefer to use. This is the only option where the school is making a decision rather than reacting to one.

A note on who's telling you this. This brief is published by Orchen, which sells a governed AI environment — so Option C is, transparently, our business. We've tried to state all three options at their strongest. Whichever direction your school takes, the checklist below applies to any vendor you evaluate, including us.

PART THREE · THE POLICY CHECKLIST

Twelve questions to settle before any deployment.

Use this in a leadership meeting or send it to any vendor — including us. The quality of the answers tells you most of what you need to know.

	Does the AI complete student work, or does it require student thinking before it proceeds? <i>Ask for a live demonstration of a student asking for the answer directly.</i>
	Is the pedagogical behavior architectural, or a mode students can switch off?
	What do teachers see, and is it synthesis or surveillance? <i>Raw transcript access at scale is both a privacy problem and a workload problem.</i>
	What do parents see, and who controls that — the school or the vendor?

	How long are student conversations retained, and is the answer specific? <i>"Per our privacy policy" is not a number. Get a number.</i>
	Is student data used to train AI models — the vendor's or their providers'?
	What happens when a student types something concerning? <i>Who at the school is notified, how fast, and is there a tracked lifecycle?</i>
	Is staff access to student data logged, and can the school review the log?
	How does under-13 consent actually work — and what happens without it?
	Can the school export everything and exit cleanly? <i>Ask for the export format and the post-termination deletion window in writing.</i>
	What certifications exist today — and what does the vendor volunteer about its roadmap? <i>A vendor that names its gaps unprompted is showing you how it behaves under pressure.</i>
	Can a pilot be scoped small enough that failure costs weeks, not budget?

Our answers to all twelve questions are public — at **orchen.ai/trust** and **orchen.ai/platform**. To see the environment live, book a walkthrough at **orchen.ai** or write to **emil@orchen.ai**.