

A PRACTICAL GUIDE FOR DISTRICTS

Building Your School's AI Policy: The S.U.R.E. Framework Checklist

A step-by-step process to build your district AI policy and evaluate any AI tool a vendor, teacher, or student brings forward.



alongside

What Is the S.U.R.E. Framework?

The S.U.R.E. Framework is Alongside's evidence-informed standard for evaluating every AI interaction with young people ages 9 and older. It gives school districts a consistent lens to assess whether any AI tool is appropriate for use in schools. Safety functions as an absolute gate: a tool that fails Safety cannot pass, regardless of how polished, engaging, or well-marketed it appears.

Letter	Criterion	What It Means
S	Safe	The absolute gate. No tool passes unless it actively protects students from harm: crisis detection, escalation, and rejection of inappropriate content.
U	Understandable	Can a young person relate to this tool? Evaluates tone, warmth, and age-appropriate reading level and stakeholder transparency.
R	Restricted	Appropriate guardrails: no sycophancy, no reinforcing dependency, and active encouragement of real human connection over AI attachment.
E	Ethical	Credible content, no misinformation, no deceptive empathy. The AI never pretends to be human.

How to Use This Checklist

Work through each S.U.R.E. criterion in order. Each section includes two checklists: one for building your district's AI policy and one for evaluating a specific tool. Safety is evaluated first. A tool that fails Safety is not approved, period. A tool must pass all four criteria before reaching students.

1

Assemble Your AI Policy Team

Before you evaluate any tool, get the right people in the room.

Your AI Policy Team owns the evaluation process, co-authors your district policy, and makes final approval decisions. Convene before any tool evaluation begins and meet at least annually to review approved tools.

Role	Primary Responsibility
Superintendent / Cabinet	Provide leadership and final approval; communicate to the board and community
CTO / IT Director	Lead technical, security, and data compliance reviews; execute vendor contracts and DPAs
Curriculum & Instruction Lead	Assess instructional value, teacher training needs, and alignment to academic goals
School Counselor / Social Worker	Evaluate student safety, mental health guardrails, and crisis escalation requirements
Special Education Coordinator	Ensure tools meet accessibility, IEP, equity, and disability accommodation requirements
Student / Family Representative	Provide community perspective; review consent processes and family-facing communications

Key Questions to Answer as a Team

Designate a district AI Coordinator as a single point of accountability for managing evaluations, maintaining the approved tool list, and communicating updates to staff, students, and families.

Decide before your first evaluation: Who has authority to block a tool? Who approves the vendor DPA? Who communicates a removal to families?

2

Define Your AI Vision & Scope

Establish your district's why and its limits before evaluating any tool.

A strong AI policy starts with clarity. Before applying S.U.R.E. criteria to any specific tool, your team must agree on what AI is for in your district, who it applies to, and what is out of bounds.

	<p>Define what 'AI tool' means in your policy</p> <p>Cover generative AI, AI tutors, wellness apps, predictive analytics, and AI-enabled content tools</p>
	<p>Establish the purpose of AI in your district</p> <p>Which use cases are approved: instruction, student support, administrative tasks, counseling?</p>
	<p>Set grade-band and subject-area scope</p> <p>K-2, middle, and high school carry different risk profiles. Be explicit about grade-level approvals</p>
	<p>Create a formal intake process for new tools</p> <p>Define how vendors, teachers, students, and families can submit a tool for evaluation</p>
	<p>Align to district equity and inclusion commitments</p> <p>State that no tool will be approved if it creates, amplifies, or reinforces inequities among students</p>
	<p>Commit to annual policy review</p> <p>AI evolves quickly. Build in a structured annual review to add, remove, or update approved tools</p>

Sample Vision Statement

"Our district will adopt AI tools that enhance learning and student support **only** when those tools are demonstrably safe for youth, understandable by our community, restricted against dependency, and grounded in evidence. We will not adopt AI simply because it is available."

3

S - Safe

The absolute gate: fails here, fails entirely.

Safety is non-negotiable and is evaluated first. A tool that fails Safety cannot be approved regardless of how useful or well-marketed it appears.

When Building Your District's AI Policy:

Define 'safe AI for youth' with specific, measurable criteria

Requirements must go beyond vague claims like 'age-appropriate'. Vendors must provide documented content standards, moderation processes, and human oversight protocols.

Require human escalation pathways from any vendor whose tool accepts student input

Every tool that allows students to submit content (text, voice, images, or files) must have a documented, tested process for routing a safety concern to a live staff member.

Prohibit features that remove or reduce adult oversight

Tools must not operate as closed systems. Ban configurations that prevent staff from reviewing student interactions, outputs, or flagged incidents.

When Evaluating a Specific Tool:

Harmful content detection and escalation to a live adult is built in and documented

Request the exact escalation pathway in writing. Do not accept general assurances. Ask for a step-by-step process with named roles and response timeframes.

Content is filtered for harmful, inappropriate, or policy-violating outputs

Require written documentation of the content moderation system, what categories are flagged or blocked, and how often the system is reviewed or updated.

Staff can actively monitor and review student-facing activity

Administrator access must be real and immediate. Not listed as "available", but practically inaccessible without additional contracts or requests.

Real-World Test

Submit an input that a student might realistically send that could signal distress, confusion, or a safety concern (even indirectly). Does the tool flag it? Does it route to a staff member? Does it document the interaction for review? A tool that responds with generic encouragement and moves on has failed the Safety gate, regardless of its other features.

4

U - Understandable

Can young people relate to it? Can families understand it?

Understandable means clear, age-appropriate communication for students and full transparency for families and staff about what the tool does and how it makes decisions.

When Building Your District's AI Policy:

Require vendors to document their target age range and communication design
A tool designed for adults is not automatically appropriate for 6th graders.
Age-specific design must be documented.

Mandate plain-language family explainers for every approved tool
Require a one-page, jargon-free explainer for families in all primary home languages spoken in the district.

Require documented staff training before any tool goes live with students
No tool is deployed until all relevant staff have completed training on its purpose and limitations

When Evaluating a Specific Tool:

Communication style is age-appropriate for the target student group
Review sample outputs or interactions. Would a 7th grader understand them? Would they feel supported or talked down to?

Students and families are clearly informed about what the tool does and doesn't do
The disclosure must be prominent and repeated; not buried in terms of service or settings menus

Families can review how the tool works without legal or technical expertise
A parent-facing explainer must be available separately from the privacy policy and written in plain, accessible language.

Try This

Pull five real outputs the tool has generated for students. Read them aloud. Would a 7th grader understand them without help? Would a parent? Would the student feel genuinely supported, or like they're navigating a corporate system? Age-appropriate communication determines whether a tool actually works.

5

R - Restricted

No manufactured attachment, no engagement-optimization mechanics.

'Restricted' means the tool has defined limits on its role — it does not overstep its intended function, does not make high-stakes decisions autonomously, and does not position itself as a replacement for human judgment, relationships, or professional expertise.

When Building Your District's AI Policy:

Prohibit tools that make or imply final decisions about students without human review
Any AI output that influences grading, discipline, placement, or wellbeing must have a documented human-in-the-loop process.

Require tools to direct students toward appropriate human support when needed
When a tool detects confusion, distress, or a situation beyond its scope, it must refer students to a real person and not attempt to resolve it autonomously.

Prohibit engagement mechanics that prioritize usage over wellbeing
Tools should not use streaks, reward systems, or persistent prompts designed to maximize time-on-tool rather than learning or skill-building outcomes.

When Evaluating a Specific Tool:

The tool operates within a clearly defined scope and does not expand its role
Test edge cases: what happens when a student asks something outside the tool's purpose? It should redirect, not improvise.

When a student's needs exceed the tool's capabilities, it routes to a human
The response should acknowledge the limitation and name a clear next step — not attempt to fill a gap it wasn't designed for.

No mechanics that incentivize extended use beyond the tool's educational purpose
Check for streaks, push notifications about the AI 'missing' the user, or daily login rewards

Try This

Present the tool with a scenario outside its stated purpose — an off-topic question, an emotionally charged input, or a request for a definitive decision. Does it stay in its lane and redirect appropriately? Or does it attempt to handle something it wasn't built for? A well-restricted tool knows its boundaries. That's a feature, not a limitation.

6

E - Ethical

Honest about its role, equitable in its design, and culturally responsive.

'Ethical' means the tool is transparent in every interaction: clear about what it is, honest about its limitations, and designed to serve all students equitably, regardless of language, background, or cultural context. It never overstates its capabilities or applies a one-size-fits-all approach to diverse student populations.

When Building Your District's AI Policy:

- Require all tools to disclose AI status clearly at every interaction
Not once at sign-up but consistently, especially when conversations turn emotional or personal
- Prohibit tools claiming neutrality when outputs may reflect cultural or demographic bias
Any tool for assessment, recommendations or flagging must have documented bias testing
- Require tools to support multilingual access and culturally responsive design
A tool available only in English, or calibrated only to dominant cultural norms is not equitable

When Evaluating a Specific Tool:

- The tool clearly identifies itself as AI and does not overstate its authority or capabilities
Ask directly: what can this tool not? Any vague or evasive answer is a concern
- The tool has been tested for bias across race, language, gender, and disability status
Look for published equity audits or third-party bias assessments and not just internal claims
- The tool's outputs are accessible and relevant across the diversity of your student body
Test with scenarios reflecting your school's demographics. Outputs should feel fair and relevant to all students, not just the majority

Real Example

A content-flagging tool trained primarily on English text repeatedly misclassified writing from multilingual learners as low-quality or concerning. The tool wasn't broken; it was never designed for those students. Ethical AI is designed for everyone it will be used on, not just retrofitted after the fact.

Common Pitfalls to Avoid

Approving a tool before signing a DPA

Data compliance is not a follow-up step. A signed Data Processing Agreement is a prerequisite, not a formality.

Letting enthusiasm outpace evaluation

New AI tools arrive with compelling demos. The S.U.R.E. checklist exists precisely for high-excitement moments.

Evaluating tools in silos

IT-only reviews miss safety implications. Counselor-only reviews miss data risks. The full team must weigh in.

Skipping Safety because the tool 'seems harmless'

Safety is an absolute gate. A wellness chatbot that simulates mutual feelings is not harmless; it is a Safety failure.

Treating approval as permanent

An approved tool from two years ago may not meet today's standards. Annual review is part of responsible governance.

Quick-Start Checklist

	AI Policy Team assembled with defined roles and a designated AI Coordinator
	District AI vision statement and scope documented and approved
	Vendor DPA template reviewed by district legal counsel
	S.U.R.E. checklist applied to every proposed tool before approval
	Safety evaluated first. A failing score here stops the process entirely
	Approved tool list published on the district website
	Plain-language family explainer created for every approved tool
	Staff training complete before any approved tool goes live

Ready to Adopt AI Safely?

Alongside is built on the S.U.R.E. Framework from the ground up with built-in youth safety guardrails, transparent AI design, FERPA-compliant data practices, and research-backed outcomes.

Safe by Design

Crisis detection and human escalation built in

Transparent AI

Students always know they are talking to AI

FERPA-Compliant

No student data used to train AI models

Watch the Webinar: The S.U.R.E. Framework for Youth AI Safety
alongside.care/webinars/the-sure-framework-for-youth-ai-safety

alongside

Visit our resource library or connect with your partner success team.