



PRIORITIZING STUDENTS WITH DISABILITIES IN AI POLICY



VERSION 2 (APRIL 2026)

TABLE OF CONTENTS

<u>EXECUTIVE SUMMARY</u>	3
<u>INTRODUCTION</u>	6
<u>GLOSSARY OF KEY TERMS</u>	10
<u>PILLAR 1:</u> <u>UPHOLDING CIVIL RIGHTS</u>	15
<u>PILLAR 2:</u> <u>DATA PRIVACY AND STUDENT PROTECTION</u>	21
<u>PILLAR 3:</u> <u>ACCESSIBILITY BY DESIGN</u>	24
<u>PILLAR 4:</u> <u>TRANSPARENCY, MONITORING,</u> <u>& ACCOUNTABILITY</u>	27
<u>FOUNDATIONS FOR IMPLEMENTATION</u>	31
<u>PROCUREMENT AND ADOPTION</u>	35
<u>CONCLUSION</u>	38
<u>ACTION STEPS FOR RESEARCHERS</u>	39
<u>GUIDANCE FOR TECHNOLOGY DEVELOPERS</u>	41
<u>APPENDIX</u>	44
<u>ACKNOWLEDGEMENTS</u>	49
<u>CONTRIBUTORS</u>	50

PRIORITIZING STUDENTS WITH DISABILITIES IN AI POLICY

EXECUTIVE SUMMARY

Artificial Intelligence (AI) is no longer a peripheral experiment in schools but a central component of instructional design for students with disabilities. As of the 2024–2025 school year, 57% of licensed special education teachers reported using AI to help develop IEPs or 504 plans—a massive 18-point increase from the previous year. This evolution is driven by a dual imperative:

the need to provide truly individualized learning experiences for neurodivergent and disabled students, and the urgent requirement to alleviate the administrative burden on a special education workforce facing unprecedented levels of burnout and attrition. Recognizing special education within AI policy is vital because these technologies serve as essential assistive "ramps" rather than mere conveniences, ensuring that students with disabilities maintain equitable access to the curriculum as mandated by law.

The rapid deployment of these technologies has outpaced the development of robust regulatory frameworks, raising profound concerns about data privacy, algorithmic bias, and the potential for overreliance on technology to dehumanize the educational experience. A 2025 systematic review of 11 major AI interventions for learning disabilities found that 0% were rated as "Low Risk" for bias; 70% were rated as "Moderate Risk"; and 30% were rated as "High/Serious Risk" due to a lack of randomized trials. Without immediate action, these tools risk deepening existing inequalities by misidentifying student capabilities through black box algorithms, exposing protected disability

73% of students with disabilities now use AI for school-related tasks, compared to 63% of their peers without documented disabilities, highlighting AI's role as a primary assistive tool (CDT, 2025).

Despite high adoption, only 22% of teachers have received formal training on AI risk management and data privacy (CDT, 2024).

information, and creating inaccessible interfaces that exclude the students they are meant to support.

The responsible integration of AI is an inclusion decision. This policy brief provides a roadmap to shift the narrative from restriction to access, positioning AI-integrated edtech as a lever for expanding educational opportunity. To manage this transition, the paper anchors all recommendations in the EDSAFE SAFE Framework (Safety, Accountability, Fairness, and Efficacy) through four operational pillars: Upholding Civil Rights, Data Privacy and Student Protection, Accessibility By Design, and Transparency, Monitoring, & Accountability. These pillars provide the actionable guidance necessary to protect students throughout the technology lifecycle.

Effective policy implementation requires leaders to anchor their work in four non-negotiable strategic foundations. This process begins with meaningful stakeholder engagement, shifting the institutional approach from mere consultation to authentic power-sharing with families and students. Decision-making must remain grounded in Learning Science and Universal Design for Learning (UDL), to ensure that all AI adoption is informed by a research-based understanding of how students learn. Furthermore, the development of comprehensive AI literacy is essential to building shared competency among educators, administrators, and families alike. Finally, a commitment to adaptive policy and practice facilitates the iterative feedback cycles and "Human-in-the-Loop" requirements necessary to maintain safe and effective oversight as technology evolves.

Recognizing that different stakeholders have distinct roles and responsibilities, this brief provides concrete, actionable guidance for several audiences, including but not limited to:

- **State Education Agencies (SEAs)** that set parameters, provide technical assistance, and, in many cases, procure statewide tools or oversee district-level procurement approvals

- **Local Education Agencies (LEAs)** that procure, implement, and monitor AI systems;
- **Educators** who integrate AI into practice while honoring student strengths and protecting rights
- **Students with Disabilities and Their Families** who use AI tools, evaluate their impact on learning, and advocate for accessible and effective systems that support their success, along with their families, who contribute essential knowledge about their children and hold systems accountable

To fully embrace inclusion, there's a symbiotic need to build capacity across the ecosystem, reaching beyond the stakeholders above. There are unanswered research questions about the potential impact, support, and harm AI can have on students with disabilities, as well as safeguards that need to be integrated into technology developers' work. To address this, the paper's addenda outline the necessary, actionable steps that should guide the work of researchers and technology developers in the field. These elements are part and parcel of the work ahead to guarantee safe, accessible, and inclusive design and deployment of these tools for students with disabilities.

AI is no longer a guest in classrooms; it is a permanent resident. The challenge is no longer whether to adopt these tools, but whether the ability exists to ensure they work for *everyone*. By centering students with disabilities in AI policy from the outset, the education system does more than fulfill a legal mandate. It unlocks a more robust, flexible, and intuitive educational landscape for every learner. Building for the margins strengthens the center. This roadmap offers a vision where technology honors human variation rather than erasing it, because designing for the few empowers the many.

INTRODUCTION

The recent widespread awareness of and emerging use cases for generative AI have sparked significant interest in its potential to support students with disabilities. The potential is already visible: students with motor disabilities who previously struggled with traditional art tools now use voice recognition to dictate commands into image-generating software, translating their words into digital artwork. Closed captioning technologies transform spoken words into real-time text, improving access to lectures and multimedia for students who are deaf or hard of hearing. With the promise of this technology so visible, it is also incumbent on all stakeholders to understand and manage the potential peril.

On May 7th, 2024, and again on December 3rd, 2025, the Educating All Learners Alliance (EALA) and New America convened educators, school administrators, parents, students, disability rights advocates, researchers, government, industry, and privacy experts to delve into the opportunities and challenges of deliberately designing and implementing AI to uphold civil rights protections for students with disabilities. Through presentations, group discussions, and breakout sessions, attendees shared their expertise to identify best practices and key issues in using AI in education to support students with disabilities. This brief reflects the priorities and updated guidance that emerged from these convenings.

This policy paper is designed to inform and guide policymakers, educators, and students with disabilities and their families on how to design, adopt, and govern AI educational policy to advance access, safeguard civil and individual rights, and improve educational outcomes for all students. Recognizing the transformative potential of AI in education, particularly for students with disabilities, this paper aims to equip stakeholders with a clear framework, guidance, and practical examples to leverage AI technologies effectively, transparently, and inclusively.

The SAFE Framework as the Anchor

While this brief is built on foundational federal frameworks, including the October 2023 Executive Order (14110) on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, the Blueprint for an AI Bill of Rights, and the Trump Administration's America's AI Action Plan, it centers the EDSAFE AI Alliance's SAFE Framework as the primary operational guide for responsible AI use in education:

- **Safety:** Prioritizing security, privacy, and the directive to "do no harm."
- **Accountability:** Defining clear responsibilities for all stakeholders.
- **Fairness:** Ensuring equity and actively mitigating algorithmic bias.
- **Efficacy:** Measuring and proving improved learning outcomes.




For all students, the EDSAFE framework ensures that AI systems are safe, effective, and intentionally designed to mitigate algorithmic discrimination. It mandates rigorous privacy protections and absolute transparency in how AI informs consequential decisions. Most importantly, it reinforces the necessity of a human in the loop, ensuring that every student has access to an educator who can intervene when challenges arise. By aligning AI policy with civil rights law and the SAFE Framework, we move beyond passive adoption to proactively harness technology to expand access and opportunity for all learners.

This paper provides a structured approach to AI governance through two integrated components: Four Strategic Foundations that serve as the prerequisite lens for all decision-making, and Four Operational Pillars that provide actionable guidance across the technology lifecycle.

Four Foundations: The Non-Negotiables

Before implementing specific policies, education leaders must anchor their work in these four non-negotiable foundations. These anchors ensure that AI use is not merely a technical update, but an opportunity to close historical access and opportunity gaps :

- **Meaningful Stakeholder Engagement:** Shifting from consultation to authentic power-sharing.
- **AI Literacy:** Building a shared competency across all stakeholder groups.
- **Learning Science Foundations:** Ensuring AI serves pedagogical goals rather than seeking them.
- **Adaptive Practice:** Committing to an iterative, responsive approach to governance.



Four Pillars: Actionable Guidance

Guided by these foundations, the following pillars operationalize the SAFE Framework to ensure that students with disabilities are protected and supported throughout the adoption process:

- **Pillar 1: Upholding Civil Rights** (*SAFE: Fairness*)
- **Pillar 2: Data Privacy and Student Protection** (*SAFE: Safety*)
- **Pillar 3: Accessibility by Design** (*SAFE: Efficacy*)
- **Pillar 4: Transparency, Monitoring, and Accountability** (*SAFE: Accountability*)



AI in education is evolving rapidly. New capabilities, use cases, and potential risks emerge with consistency, requiring ongoing attention and adaptation. By adopting the agile mindset modeled by leaders like the Santa Ana Unified School District (SAUSD), whose AI Compass suite mandates semi-annual reviews, we can ensure that technology remains a bridge to inclusion rather than a barrier to equity. The field of AI is not static, and neither can our policies and guidance be.

Metro Nashville Public Schools (MNPS) provides a leading example of how school districts can move beyond the "risk management" phase of AI to foster a culture of creative expression and student agency. In early 2026, MNPS launched the Lumi Story AI pilot, an initiative designed to enhance literacy by allowing students to use generative tools as "creative assistants" to develop characters and narratives. By framing AI as a tool for guidance rather than a replacement for critical thinking, the district ensures that students maintain ownership over their original ideas while exploring the possibilities of new technology.

The success of the Nashville model is rooted in its commitment to a student-led feedback loop. Rather than implementing policy in a vacuum, MNPS convened panels where students shared directly with administrators how they were utilizing AI for personalized learning — ranging from formatting scripts to generating targeted practice problems for ACT preparation. By prioritizing human connection and evaluating AI through the lens of student passion and purpose, MNPS offers a great example for districts to implement AI responsibly and effectively.

GLOSSARY OF KEY TERMS

Adoption: The process of integrating AI tools into educational practice through implementation planning, professional development, and systems integration. Successful adoption requires dedicated budgeting for ongoing support and infrastructure compatibility, as tools often fail due to under-resourced implementation rather than design flaws.

AI Literacy: The knowledge and skills that enable humans to critically understand, evaluate, and use AI systems and tools to safely and ethically participate in an increasingly digital world.

Americans with Disabilities Act (ADA): Federal civil rights law prohibiting discrimination against individuals with disabilities in all areas of public life, including schools. Title II of the ADA covers state and local government services, including public schools, and recently updated regulations require digital accessibility conforming to WCAG 2.1 AA standards.

Artificial Intelligence: AI is a technology that enables computer systems to learn, reason, and make decisions, performing complex tasks commonly thought to require human intelligence, such as recognizing speech, analyzing data, and solving problems.

Children's Online Privacy Protection Act (COPPA): COPPA imposes certain requirements on operators of websites or online services directed to children under 13 years of age, and on operators of other websites or online services that have actual knowledge that they are collecting personal information online from a child under 13 years of age.

Commercial Generative AI: Artificial Intelligence products and services that are designed for use in the public or business marketplace that can create new content—such as text, images, video, audio, or code—based on patterns and structures learned from vast datasets.

Decision Makers: Individuals and groups with authority to set policy, allocate resources, or approve technology adoption in educational settings. This includes school board members, district superintendents, principals, technology directors, curriculum coordinators, and nonprofit leaders who shape the use of AI in schools.

Education Technology (EdTech): A digital tool, platform, or ecosystem designed from its inception with the primary goal of addressing specific pedagogical needs, learning objectives, or educational administrative challenges. Unlike "general-use" technology (such as a standard word processor or generic video conferencing software) adapted for the classroom, purpose-built EdTech is grounded in the Learning Sciences and instructional design.

Educational Generative AI: refers to Artificial Intelligence systems—often based on Large Language Models (LLMs)—designed to create, simulate, or assist with educational content, such as generating text, images, code, or personalized, interactive learning materials.

Family Educational Rights and Privacy Act (FERPA): Federal law protecting the privacy of student education records. FERPA gives parents rights to access their child's records, request corrections, and control the disclosure of personally identifiable information. Schools must have written permission to release student records, with limited exceptions for school officials with legitimate educational interests.

Generative AI: A machine learning model that is programmed to take in data and generate the output of new data, including text, audio, images, and video. Examples include ChatGPT, Gemini, and other large language models. These tools can draft documents, answer questions, generate lesson plans, or create visual materials, but they require careful oversight to ensure accuracy and avoid bias.

Human in the Loop: A model of AI governance and design that requires active human oversight, intervention, and decision-making at critical stages of an AI system's lifecycle. In an educational context, it mandates that educators and students retain the final authority to interpret, validate, or override AI-generated outputs.

Individuals with Disabilities Education Act (IDEA): A federal law ensuring that children with disabilities receive a free appropriate public education (FAPE) in the least restrictive environment. IDEA mandates the creation of Individualized Education Programs (IEPs) and guarantees that parents are active participants in educational decision-making.

Learning Science: An interdisciplinary field that investigates how people learn, combining insights from neuroscience, psychology, education, and computer science to design more effective educational experiences.

Local Education Agency (LEA): A public board of education or other local authority legally recognized for its administrative control of public schools in a city, county, or school district. LEAs are responsible for the direct implementation of technology and disability services.

Privacy-Enhancing Technologies (PETs): A suite of digital tools, techniques, and methodologies designed to maximize the utility of data while minimizing the risks to individual privacy. Rather than simply "hiding" information, PETs use mathematical and architectural safeguards—such as differential privacy, synthetic data, and federated learning—to allow for complex data analysis and collaboration without exposing raw, personally identifiable information (PII)

Procurement: The strategic process of evaluating, selecting, and acquiring technology. For AI, this requires a cross-functional team (including special education directors and IT) to vet vendors for accessibility, data privacy, and alignment with educational quality frameworks before a contract is signed.

SAFE Framework: Developed by the EDSAFE AI Alliance, this framework guides the responsible use of AI in education. It stands for Safe, Accountable, Fair, and Effective, ensuring that tools are non-discriminatory and fulfill their intended educational goals.

Section 504 of the Rehabilitation Act: A civil rights law that prohibits disability-based discrimination in any program receiving federal funding. In schools, "504 Plans" are required to provide necessary accommodations and modifications to ensure students with disabilities have the same access to learning as their peers.

State Education Agency (SEA): The state-level department (e.g., a State Department of Education) responsible for supervising public education. SEAs distribute funding, set standards, and provide the technical assistance LEAs need to comply with civil rights and disability laws.

Technology Developers: Entities that design and license EdTech products. Developers and vendors are responsible for ensuring their AI tools are natively accessible, protect student privacy, and comply with civil rights laws like the ADA.

Title VI (1964 Civil Rights Act) and Title IX (1972 Education Amendments): Federal laws prohibiting discrimination in programs receiving federal funds. Title VI bans discrimination based on race, color, or national origin. Title IX prohibits sex-based discrimination in education programs.

Universal Design for Learning (UDL): A research-based educational framework to guide the design of learning environments that are accessible, inclusive, and challenging for every learner. Ultimately, the goal of UDL is to support learner agency, the capacity to actively participate in making choices in service of learning goals. It provides multiple means of engagement, representation, and action/expression so students with diverse abilities can access and demonstrate learning.

Web Content Accessibility Guidelines (WCAG) 2.1 AA: International technical standards for making digital content accessible to people with disabilities. WCAG 2.1 Level AA is the baseline accessibility standard required under the ADA Title II final rule for state and local government services, including educational technology. These guidelines cover perceivability, operability, understandability, and robustness of digital content.

PILLAR 1: UPHOLDING CIVIL RIGHTS

To ensure AI systems are fair and designed to mitigate discriminatory practice, AI policies and practices must uphold the civil rights of students with disabilities under Section 504 of the Rehabilitation Act, the Individuals with Disabilities Education Act (IDEA), the Americans with Disabilities Act (ADA), and Titles VI and IX of the Civil Rights Act. AI-integrated tools must support individualized planning and decision-making while actively mitigating algorithmic bias, overreliance on automated recommendations, and assessments that fail to capture the full range of students' strengths and capabilities. Bias about disability is a well-documented phenomenon in AI models. Research from Penn State found that AI sentiment tools consistently flag sentences containing disability-related terms (like "blind" or "deaf") as toxic or negative, even when the context is neutral or positive. As such, it is incumbent on all stakeholders, from policy makers to family members, to actively monitor, evaluate, and intervene in the use of AI tools to prevent bias and protect the rights of students with disabilities.

Actionable Steps

State Education Agencies

- **Highlight and promote AI tools trained on representative datasets, which accurately reflect the diversity of the population they are intended to serve, including students with a full range of disabilities across all major IDEA categories.** The University of Illinois Speech Accessibility Project provides a prime example of a project that created a private, de-identified dataset for training machine learning models to understand a variety of speech patterns better. This initiative is a collaborative effort among the University of Illinois, Microsoft, Google, Amazon, and nonprofit organizations, aimed at creating representative training data to enable more accurate and inclusive AI speech recognition.

- **Develop and disseminate guidance clarifying how existing civil rights laws (IDEA, Section 504, ADA) apply to AI use in educational settings, including specific examples of how different AI applications may implicate these protections.** The Center for Democracy and Technology's 2022 report [Ableism and Disability Discrimination in New Surveillance Technologies](#) provides an outline of how technology can disproportionately harm people with disabilities in various settings, including in education. Resources could model and expand upon existing [Indicators and Matrices](#) developed under the Quality Indicators for Assistive Technology Services. Finally, SEAs should provide technical assistance and resources, including practical guidance, to help local education agencies develop and support internal capacity to procure, implement, and manage AI to support students with disabilities, such as [NCADEMI's procurement guide](#).
- **Require that final decisions about student placement, services, accommodations, or disciplinary actions affecting students with disabilities are made by humans, even when informed by AI systems, and include opportunities for parent participation and input.** The proposed [Artificial Intelligence \(AI\) Civil Rights Act](#) would prevent companies from using AI-powered algorithms to make critical decisions in areas that affect people's lives, such as housing, credit, employment, education, healthcare, and public benefits.

Algorithms can be designed to increase opportunity - The Michigan Economic Development Corporation (MEDC) launched the Michigan Career Portal, which uses AI to improve job placement by adopting a "skills-first" approach rather than simply matching resumes to keywords. Traditional manual placement processes are slow, often leading to missed opportunities and mismatches between students' skills and employers' needs. The platform analyzes a job seeker's experience, skills, projects, and educational background to identify relevant opportunities, going beyond basic keyword searches to pinpoint in-demand skills. This approach broadens the pool of qualified candidates for employers while providing students with more relevant, actionable job opportunities, resulting in higher interview call rates and better-tailored applications.

Local Education Agencies

- **Establish processes to evaluate and ensure compliance with applicable federal and state civil rights laws and regulations that may apply when adopting or using different types or uses of AI.** Particular attention should be given to student monitoring technologies, which may implicate multiple civil rights protections. Given the sensitivity and potential for disproportionate impact, agencies should require documented compliance before adoption and establish ongoing oversight to maintain alignment with civil rights protections. For additional guidance, see The Privacy and Equity Implications of Using Self-Harm Monitoring Technologies.
- **Select EdTech providers that actively strive to develop and implement AI models that mitigate harmful bias.** LEA procurement officers should demand "model cards" from vendors that disclose the model's limitations, intended use cases, and potential biases, enabling the organization to determine whether the AI is safe to use. LEAs can partner to investigate potential vendors and negotiate additional protective contractual provisions, as many LEAs have done under the auspices of the Student Data Privacy Consortium.
- **Mandate accessibility requirements when engaging vendors and select AI applications that are accessible to students with disabilities, including those who use assistive technology.** Verify WCAG 2.1 AA conformance and request Accessibility Conformance Reports (ACRs).

- **Require that AI adoption decisions are informed by input from special education directors, disability rights advocates, and families of students with disabilities, not just technology staff.** Agencies should prioritize engagement with individuals with lived experience and mandate representation from a wide range of disability communities, particularly those often underrepresented in technology design. A paper published by Stanford Accelerator for Learning, [AI + Learning Differences: A Future with No Boundaries](#), models and outlines how these processes can proactively identify and mitigate barriers in AI systems.

A strong example of this approach is the Franklin Township Community School Corporation in Indiana, which implemented AI-powered chatbots to support high school students who require special education services. Rather than a top-down approach to technology adoption, the initiative was driven by recognition of the need to support students with learning and thinking differences, enabling them to work more independently.

Key Aspects of the Design:

Inclusive Decision-Making: The initiative was not solely driven by IT staff. It involved input from educators who were struggling to provide individualized help, as well as stakeholders focused on the "promising" aspects of the technology, such as the district's technology integration supervisor, who is also a parent of a student with a disability.

Input from Families: The initiative addressed the real-world experiences of students like Makenzie Gilkison, who has dyslexia and felt that AI tools (such as customized chatbots and word prediction) were essential to her feeling less "stupid" and keeping up with her classmates.

Focus on Usability and Access: The process included identifying when features were not working (e.g., when a tool was unexpectedly turned off by the tech team) and highlighting the need to involve special education staff in maintaining access, rather than just technology staff.

Educators

- **Evaluate how AI tools may interact with students' IEPs and 504 plans before adoption.** If an AI tool is being used to support a student with a disability, consider whether it aligns with or undermines their documented accommodations and goals. Before using AI tools to draft supporting content for a student, review guidance, like CDT's guidance on AI use in IEP development, to understand privacy risks, bias concerns, and requirements for maintaining individualized, human-centered planning. AI should support, not replace, the individualized planning process. Qualified professionals and IEP teams should review AI-generated recommendations about student placement, services, or interventions.

Resource: As part of this process, use tools like The National Education Association's Accessibility Tree to evaluate whether AI tools demonstrate strong indicators of accessibility and universal design.

Students with Disabilities and Their Families

- **Students should be aware of their rights under IDEA, Section 504, and the ADA, and ensure that accommodations provided by schools meet their individual needs rather than relying solely on AI recommendations.** To effectively educate students on disability rights and self-advocacy, schools can integrate the Changing Perspectives curriculum for comprehensive lessons and reflection, alongside legal resources from the ADA National Network and the U.S. Department of Justice to clarify specific rights. This foundation is strengthened through NCLD-recommended workshops on disclosure and accommodation requests, as well as the use of digital tools like Read&Write and Microsoft's Immersive Reader to demonstrate available support technologies. Finally, organizing alumni panels and peer support groups bridges the gap between theory and practice, allowing students to learn through the lived experiences and mentorship of others.

- **Parents or guardians should request information from their child's school on which AI tools are being used and how those tools respect their child's IEP or 504 plan accommodations.** This includes understanding whether AI is being used to make or inform decisions about student placement, services, or educational programs. Families should raise questions during IEP meetings and request documentation of the school's review alignment with accommodations and privacy protections. For additional guidance and strategies for effective advocacy, see AI Policy Lab's [A Parent's Guide to AI Advocacy in Education](#).

PILLAR 2: DATA PRIVACY AND STUDENT PROTECTION

To ensure AI systems are safe and do not harm students, they must be built upon comprehensive privacy and security practices that protect sensitive disability-related information. These systems must align with the Family Education and Privacy Act (FERPA), state privacy laws, and fundamental disability rights protections. To uphold these standards, districts and providers must rigorously apply the principles of data minimization, purpose limitation, and timely deletion. Implementing these strong security measures is essential when collecting, using, or sharing student information, particularly disability data used to evaluate AI effectiveness or train algorithms.

Actionable Steps

State Education Agencies

- **Establish comprehensive data privacy and security standards** specifically designed to protect sensitive disability-related student information. Develop and disseminate model privacy policies and templates, such as data-sharing agreements and vendor contracts. The Maryland State Department of Education's 2026 AI Policy provides a suite of resources to support the responsible and safe implementation of AI, ensuring it strengthens teaching and learning while protecting student privacy.
- **Clarify legal protections for AI integration** by providing guidance that explains how existing federal and state laws, including FERPA and the Children's Online Privacy Protection Act (COPPA), apply to AI use. This clarification should focus on upholding the rights of students with disabilities and ensuring EdTech providers meet high-quality safety and legal standards. The Public Interest Privacy Center outlines the laws that educators and administrators should be considering before integrating AI-powered tools.

- **Provide clear guidance regarding the risks associated with “free” or “freemium” AI tools.** Emphasize the importance of using approved platforms that provide adequate protection for student data, particularly for students with IEPs or 504 plans, where sensitive information is most vulnerable. CESA 6 (Cooperative Educational Service Agency) offers resources on protecting student data from "hidden AI risks" and provides "Dos and Don'ts" guidelines.

Local Education Agencies

- **Implement robust privacy safeguards to ensure all contracts with AI providers include specific provisions to protect sensitive student information, particularly disability status and IEP and 504 data.** Establish clear requirements for data ownership, strict use limitations, and mandatory deletion timelines. Privacy and Security Conditions from New America provides support in identifying critical privacy considerations and essential questions for vendors.
- **Mandate manual configuration of privacy settings to prohibit the use of default privacy settings for any AI tools.** Before deployment, LEAs should manually review and configure settings to maximize student protections, ensuring that sensitive disability-related data is not used for model training or unauthorized data sharing. This includes establishing guidelines for educators' data input. Develop clear policies that limit the specific types of student data that may be entered into AI platforms. These guidelines should provide educators with practical steps for handling identifiable data related to specialized instruction and student support plans.
- **Disseminate evaluation toolkits to help school leaders evaluate the potential impact of different AI tools based on their risk level.** These toolkits should guide the formulation of local policies that account for the use of AI across educational contexts. The Responsible AI and Tech Justice Guide for K-12 Education report has examples of strategies for using technology responsibly and ethically.

Educators

- **Distinguish between using AI as a general productivity tool and inappropriately entering protected student information - such as names or specific disability details - into unauthorized AI tools.** Before using any AI application, verify that the tool is on the district's approved list and that it's in compliance with local privacy policies. The National Education Association has outlined key recommendations for educators to reduce the risk of privacy issues. This includes being transparent about the use of AI in a student's education experience, through disclosure and discussion with the student and their family. Commit to conducting a thorough, manual review of all AI-generated outputs for accuracy and appropriateness. New America has published a list of questions for educators to consider regarding responsible AI use.

Students with Disabilities and Their Families

- **Students should proactively learn about their rights under federal privacy laws to ensure their personal information remains protected while using digital learning tools.**
- **Parents/Guardians are encouraged to advocate for their children by asking school leaders which AI tools are in use and how student privacy is protected, particularly for sensitive information about your child's disability, IEP, or 504 plan.** Families can build their foundational knowledge and prepare for these conversations by using the AI Literacy Family Toolkit.

The US Department of Education's guide on Protecting Student Privacy outlines how data should be handled by schools and technology providers.

PILLAR 3: ACCESSIBILITY BY DESIGN

To ensure AI systems are effective and improve learning outcomes, AI systems must be designed using Universal Design for Learning (UDL) principles and meet WCAG 2.1 AA accessibility standards at a minimum. Tools must be compatible with assistive technologies, offer multiple means of engagement and expression, and provide accessible alternatives when AI-generated content or interfaces create barriers for students with disabilities. When accessibility by design is the default, tools can evolve to meet the needs of students with disabilities. Recent AI projects, like the Speech Accessibility Project, have seen AI misunderstanding rates for students with speech disabilities drop from 20% down to 12% due to specialized training data. Technical access, however, should be paired with operational transparency, requiring policymakers and decision-makers to provide adaptable communication to explain AI integration to students and families.

Actionable Steps

State Education Agencies

- Provide model communication templates and resources that LEAs can adapt to explain AI use to families, including multiple-language and plain-language explanations of various AI applications and their implications for students with disabilities. A leading example of this practice comes from the Oklahoma State Department of Education's Office of AI and Digital Learning, which provides districts with ready-to-use family guides in multiple languages and plain language consent forms that clarify how AI can and should be used.
- Establish advisory structures that include students with disabilities, families, and disability rights advocates in state-level AI policy development and evaluation. Accessibility by design requires shifting from 'designing for' to 'designing with' the disability community. Inclusive oversight will clarify how emerging AI tools interact with established legal protections and prevent discrimination, as highlighted in the Public Interest Privacy Center report, From Data Privacy to

Discrimination: Examining the Legal Ramifications of AI in Schools (PIPC).

Local Education Agencies

- **Communicate openly about the use of AI in education, actively working with the local community to tailor how AI is implemented.** Use the Student Privacy Communications Toolkit: For Schools & Districts (FPF) as a starting point for conversations about adopting or using any education technology.
- **Promote AI literacy, empowering educators and students with disabilities and their families to engage in informed discussions about protecting privacy and critically assessing the use and appropriateness of AI in education.** This can be incorporated into existing digital literacy curricula, such as these materials from Teaching Privacy, or developed independently, as seen in this curriculum from MIT Media Lab. A host of resources is available by grade level and time span to promote AI literacy through National AI Literacy Day. AI literacy can also be developed through online games that teach about AI risks, such as predictive yearbook awards, training self-driving cars, and simulated hiring with AI.
- **Refine policies in a participatory manner to detail how AI will be used to support students with disabilities, ensuring these policies are clearly communicated to all stakeholders.** To be effective, these discussions should engage a cross-functional team of leaders – such as superintendents, curriculum and instruction chiefs, EdTech and AI directors, and disability service coordinators – as well as those in implementation and enablement roles who manage the day-to-day integration of new technologies. Because district structures vary by size and resources, the focus should remain on including those who hold the primary responsibility for academic strategy, technical infrastructure, and student support services. This work should include creating accessible feedback mechanisms where students with disabilities, families, and educators can report when AI tools create barriers or function differently than expected.

Educators

- **Invite feedback from students with disabilities about whether AI tools are helping or creating barriers, and share that feedback with administrators so tools can be improved or discontinued if necessary.** Rather than treating AI integration as a static decision, educators should formalize student-led audits to determine if a tool is aiding a student's IEP goals or introducing new hurdles. By formalizing the feedback loop, educators can provide administrators with the qualitative evidence needed to refine classroom technology or discontinue tools that fail to meet students' needs.

For Students with Disabilities and Their Families

- **Students should provide developers and school leaders with feedback on the features that facilitate their learning.** By documenting what works in practice, students provide the essential "human-in-the-loop" data that developers and LEAs require to move beyond theoretical compliance and toward practical accessibility. Families should also provide feedback on the school's implementation and communication about the use of AI, including whether such tools are accessible and effective for students with disabilities, and if additional AI tools should be procured for use as productivity tools to aid in the development of IEPs or 504s, or other disability early warning screening.

PILLAR 4: TRANSPARENCY, MONITORING, & ACCOUNTABILITY

To ensure clear accountability and stakeholder oversight, schools must disclose when and how AI is used in educational decisions that affect students with disabilities, particularly regarding placement, services, accommodations, and assessment. This pillar put in place the mechanisms to move beyond initial implementation toward long-term accountability. AI systems require ongoing evaluation through stakeholder feedback, bias audits, and effectiveness monitoring, with clear accountability structures in place for when tools cause harm, perpetuate discrimination, or fail to meet promised outcomes. By institutionalizing these transparent reporting and evaluation protocols, education leaders can ensure that AI remains a reliable and supportive tool for every learner's success.

Actionable Steps

State Education Agencies

- **Establish statewide evaluation protocols** that require LEAs to conduct semiannual audits of AI tools. These evaluations should specifically track accessibility compliance, privacy protections, and effectiveness, disaggregated by data by disability category. The Center for Democracy and Technology offers an overview of the legislative trends across legislatures, providing guardrails and guidance to districts and schools.
- **Modernize state reporting and vendor accountability** by publishing annual reports detailing statewide AI use, including technical provider performance, privacy incidents, and common accessibility highlights. SEAs should require vendors to meet high standards of transparency, safety, and efficacy to support students with disabilities.
- **Provide technical assistance to LEAs on balancing student privacy** protections with the collection of necessary disaggregated data. This support should help LEAs to evaluate if AI tools are effectively supporting students across disability categories and intersecting identities, without compromising individual confidentiality.

- **Foster ethical governance and review** through sustainable forums for feedback and collaboration, such as the AI Advisory Group, to ensure that the implementation of AI in schools reflects ethical standards and the real-world needs of students with disabilities. For high-risk systems – such as those that make recommendations about placement, services, or disciplinary actions – SEAs should consider conducting annual third-party audits. Refer to UNESCO’s Recommendation on the Ethics of Artificial Intelligence for a global framework on establishing Fact-Based and ethical AI oversight.

Local Education Agencies

- **Establish accountability and reporting structures** to ensure AI performance remains aligned with educational values. This involves conducting quarterly reviews of AI tools to evaluate effectiveness for students with disabilities, including analysis of disaggregated outcome data, accessibility compliance, privacy incidents, and reported barriers. Refer to WestEd’s Four Key Assessment and Accountability Trends for guidance on modernizing leadership oversight.
- **Conduct rigorous quarterly performance audits** to evaluate the functional effectiveness of AI tools for students with disabilities. These reviews should analyze disaggregated outcome data and accessibility compliance, and be supported by clear, predefined criteria for discontinuing any tools that fail to meet safety or instructional standards. Montana Digital Academy published an AI Integration Framework for School Districts to invest in, integrate, and innovate their practices, with a focus on privacy policies.
- **Allocate dedicated funding for ongoing evaluation**, monitoring, and adaptation of AI systems, beyond the initial procurement. Budgeting should account for the ongoing costs of accessibility audits, specialized data analysis, and the continuous collection of stakeholder feedback to ensure tools stay relevant and safe. The White House released a list of organizations that have committed to providing resources for AI in education.

- **Implement formal systems for community feedback and public verification** of AI policies to ensure long-term accountability. The Arizona Institute for Education and Collaboration recommends that these mechanisms support ongoing collaboration between school leaders and specialized service providers. Stakeholders can include speech-language pathologists, occupational therapists, school psychologists, counselors, and teachers of students with visual impairments. By creating formal pathways for the community to identify AI use cases and voice concerns, LEAs can proactively identify and address potential unintended consequences or technology misuse.
- **Adopt an age-appropriate AI literacy and data rights curriculum that empowers students with disabilities to navigate these new technologies.** This instruction should help students understand what data AI tools collect about them and how it is used, while also teaching them how to identify when a tool creates a barrier to their learning. Ensure that students know their rights to access, question, and challenge AI-informed decisions, and who to contact when the system is not working as intended. Education Week outlines a developmental framework for introducing AI to K-12 students, emphasizing that educators should shift from protecting younger children from AI's personification to teaching older students to critically evaluate and use AI ethically. By categorizing AI engagement into four stages—early elementary, upper elementary, middle school, and high school—the resource provides specific strategies for building AI literacy while ensuring that technology complements, rather than replaces, essential problem-solving skills.

Educators

- **Implement whole-child monitoring practices** that go beyond basic academic data to ensure AI tools support students' overall well-being. This includes tracking indicators such as student engagement, sense of belonging, and self-advocacy to ensure that technology supports growth for the whole child, not solely on increasing test scores. The National Education Association recommends applying the Web Content Accessibility Guidelines and Universal Design for Learning to evaluate and select AI tools that are effective for every learner.
- **Document and share both successes and challenges** to inform an adaptive mindset of AI implementation for students with disabilities. Maintaining a record of bright spots and hurdles in the classroom provides the essential evidence needed to identify which AI interventions truly support specialized instruction and which need adjustment. North Carolina's EVERY Framework serves as a primary model for this practice, emphasizing that educator-led evaluation is mandatory at every step to ensure AI remains a supportive scaffold.

Students with Disabilities and Their Families

- **Participate in student advisory groups** that evaluate AI tools and recommend whether to adopt or discontinue them. Student involvement in these formal review processes ensures that the tools used truly meet the functional needs and preferences of the learners themselves. The Consortium for School Networking provides a resource on engaging students in AI evaluation through surveys and feedback forms to identify areas for improvement.
- **Use the IEP process to proactively inquire** about how AI tools are being monitored for effectiveness for your child, and request support when the tools are not working as intended. A resource from the Center for Innovation, Design, and Digital Learning outlines the essential principles for using AI in IEP development and family engagement.

FOUNDATIONS FOR IMPLEMENTATION

Leveraging AI to support students with disabilities effectively and responsibly rests on four essential foundations: meaningful stakeholder engagement, comprehensive AI literacy, grounding in learning science foundations, and a commitment to adaptive practice. Together, these elements ensure that technology serves as a bridge to accessibility rather than a barrier to inclusion.

Meaningful Stakeholder Engagement

Effective AI policy is built upon the authentic participation of those most impacted by these systems, specifically students with disabilities and their families. Meaningful engagement should move beyond mere consultation, evolving into a model of shared power in which stakeholders are involved in decision-making from the earliest stages of procurement and policy design. To achieve this, districts should ensure all participation opportunities are fully accessible and clearly communicate how stakeholder input directly shapes final decisions.

This commitment to accountability ensures that feedback results in visible, documented changes to AI governance. The guidance throughout this brief reflects this "all-in" approach, providing concrete, actionable steps for state agencies, local districts, educators, and families to work in tandem. By centering the lived experiences of students with disabilities, educational leaders can ensure that AI implementation remains grounded in equity and responsive to the real-world needs of the community.

AI Literacy

AI literacy is a non-negotiable requirement for safe, effective, and ethical use of AI in education. Educators, administrators, students with disabilities, and their families all need an understanding of how AI works, its capabilities and limitations, and its implications for learning, privacy, and civil rights. Professional development in AI should be sustained to build staff capacity to understand core concepts, critically evaluate tools, and implement them responsibly.

Resources such as the U.S. Department of Labor’s [AI & Disability-Inclusive Hiring Framework](#) outline key competencies, including understanding how machine learning works, what training data means, and how algorithms make decisions.

Effective communication about AI requires that all stakeholders, from school boards to teachers to families, share a common vocabulary.

Schools and districts should establish and communicate

clear definitions of terms such as AI, machine learning, generative AI, algorithmic bias, and training data. This shared language enables productive conversations about policy, procurement, and practice. The glossary included in this document provides a starting point for building this common understanding.

Professional development should enable educators and administrators to recognize and mitigate AI bias, particularly bias that may disproportionately harm students with disabilities or other marginalized groups. Leaders need to evaluate AI tools critically, asking questions about accessibility, privacy protections, evidence of effectiveness, and alignment with educational goals. They can make informed decisions about when AI is and is not appropriate for supporting students, and understand how to protect student privacy when using AI tools, including understanding the risks of free or “freemium” versions that may use student data for model training. Professional development should not be a one-time workshop but an ongoing learning process that evolves as AI technology and use cases develop.

Learning Science Foundations

Decisions regarding the use and integration of AI to support students with disabilities must be anchored in learning science - the research-based understanding of how people learn. AI is not a solution in search of a problem; it should be adopted only when evidence demonstrates it can effectively support specific, well-defined learning goals for students with disabilities.

Resource: The [Blueprint for Action: Comprehensive AI Literacy for All](#) offers a strategic framework and actionable steps for districts to build a robust foundation of AI knowledge and skills across their entire educational community.

Whether an AI tool is a dedicated educational product or a general-purpose application, its utility should be grounded in the principles of Universal Design for Learning (UDL). By centering multiple means of engagement, representation, and action, UDL ensures that learning environments are intentionally flexible from the outset, rather than retrofitted after barriers to access have already surfaced.

When evaluating or developing AI tools, educational stakeholders should ask:

- Does this tool provide multiple pathways for students to engage with content that account for different interests, motivations, and engagement needs?
- Does it present information through multiple modalities, such as text, audio, visual, and interactive formats, so students can access content in ways that work for them?
- Does it allow students to demonstrate understanding and express what they know through varied formats?
- Is the tool fully compatible with assistive technologies students already use, such as screen readers, alternative input devices, or communication systems?

When making decisions about AI adoption, stakeholders should:

- Prioritize AI tools with demonstrated evidence of effectiveness for students with disabilities, ideally using the Every Student Succeeds Act (ESSA) evidence tiers as a framework.
- Require tech providers to provide transparent information on the research basis for their claims to support students with disabilities.
- Look for evidence that AI tools have been tested with diverse disability subgroups, not just general student populations.
- Continuously evaluate whether AI tools are achieving intended learning outcomes for students with disabilities in your individual context.
- Be particularly cautious about AI tools making unsupported claims about addressing learning disabilities, behavioral challenges, or developmental differences without robust evidence.

Adaptive Practice

An adaptive mindset is a key component for ensuring that AI policies and practices for students with disabilities are integrated safely, effectively, and responsibly. The insights in this paper underscore the necessity of a systematic commitment to reassessing whether these emerging tools are serving their intended purpose of supporting learners' diverse needs. Because the speed of this technology is unprecedented, any policy that remains static risks immediate obsolescence or inadvertent infringement on student civil rights. Student feedback cycles are an integral part of this approach, as the lived experiences of learners with disabilities provide critical evidence for determining the efficacy of AI tools. This requires all stakeholders, from policymakers to those in the classroom, to commit to iterative strategies, acknowledging that as the landscape shifts, so too must the pedagogical applications.

PROCUREMENT AND ADOPTION

Successful AI implementation requires a clear separation between procurement (selecting and purchasing tools) and adoption (implementing and sustaining use). Many AI tools fail not because of poor design, but because districts under-resource implementation—lacking budget for training, technical assistance, integration with existing systems, and evaluation to determine whether tools work as intended. Districts should assess readiness before procurement, engage stakeholders in the selection process, negotiate comprehensive contracts that protect student privacy and ensure accessibility, and commit to discontinuing tools that fail to deliver promised outcomes.

Readiness Assessment

Successful AI integration requires a clear distinction between procurement (the selection and purchase of tools) and adoption (the long-term implementation and sustained use). Many AI initiatives fail not due to poor software design, but because districts under-resource the adoption phase by lacking the necessary budget for professional development, technical integration, and efficacy evaluation. Before any purchase is made, districts must conduct a formal readiness assessment to identify infrastructure gaps and ensure that the digital ecosystem is prepared to support students with disabilities.

Resource: Use frameworks like [aiEDU's AI Readiness Rubric](#) to evaluate infrastructure, professional development capacity, data governance, and stakeholder engagement readiness. This assessment helps districts identify gaps that should be addressed before successful implementation.

Procurement: Actionable Strategies

Procurement is the primary gatekeeping mechanism for ensuring AI tools meet rigorous standards for accessibility, privacy, and instructional fairness.

- **Districts should use cross-functional vetting teams and standardized processes**, such as the [EDSAFE AI and Project Unicorn Procurement Checklist](#), to ensure that Superintendents, IT leaders, and Special Education directors collaboratively evaluate tools prior to approval.
- **State and local agencies should mandate transparency and accessibility conformance** by requiring providers to deliver [Accessibility Conformance Reports](#) (ACRs) using the [Voluntary Product Accessibility Template \(VPAT®\)](#), and provide standardized evidence of efficacy for students with disabilities, specifically aligned with [ESSA evidence tiers](#).
- **Educational leaders should negotiate rigorous data protections and contracts** that include specific language, such as the models found in [Alabama's AI Policy Template](#), prohibiting the training of AI on student data without explicit permission and mandating regular compliance audits.
- **Decision-makers should prioritize assessment integrity and fairness** when procuring AI-powered measurement tools. Refer to the [Duolingo English Test Case Study](#) for frameworks on maintaining validity and efficacy in AI-driven assessments.
- **Local agencies should maintain a transparent, public-facing inventory of approved AI tools** to ensure that only systems meeting [GADRRS](#) data-governance standards are used in classrooms and shared with the community.

More Procurement Resources: [NIST AI Risk Management Framework](#), [ACLU Algorithmic Equity Toolkit](#), and [Student Privacy Compass: Vetting Generative AI](#).

Adoption: Actionable Strategies

Successful adoption is a non-linear, ongoing cycle of improvement that requires an intentional commitment to the human side of technology integration. Districts should recognize that the technical, instructional, and social support required for effective AI use constitutes the "hidden cost" of technology, often exceeding the initial procurement price. To navigate this, leaders should utilize the Digital Promise AI-Enabled Edtech Adoption Cycle to guide the rollout of technology through evidence-based, manageable stages rather than a single event.

- **Districts should budget for the full lifecycle of AI implementation** by treating technical assistance, integration, and sustained professional learning as essential, non-negotiable financial components.
- **Administrators should protect educators' capacity by ensuring instructional staff have dedicated, compensated time to learn AI tools and adapt them** to meet specific IEP and 504 goals at a manageable pace.
- **State and local education agencies should provide localized implementation resources**, such as the five foundational recommendations outlined in Integrating AI-Enabled EdTech in PK16 Education, to support the functional efficacy of AI for students with diverse learner profiles.
- **School systems should institutionalize communities of practice** that allow educators to collaborate both synchronously and asynchronously to share successful strategies, troubleshoot barriers, and document "bright spots" in AI-enabled instruction.
- **Leadership should maintain a commitment to discontinuing any AI tool** that fails to deliver promised outcomes or creates new accessibility barriers, ensuring that accountability persists even if a contract is currently active.

CONCLUSION

AI holds the potential to be the ultimate equalizer, but only if we treat disability rights as a primary architectural requirement rather than a secondary feature. The path forward requires a relentless focus on lived experience and rigorous evidence. As we apply the Four Strategic Foundations and Operational Pillars, we are doing more than just implementing software; we are upholding a civil rights mandate. The true measure of AI's success in education will be found in its ability to honor the variability of every learner and ensure that "access" is a reality for all, not just a promise for some.

ACTION STEPS FOR RESEARCHERS

The responsible use of AI tools for students with disabilities relies on researchers to investigate their efficacy, identify potential algorithmic biases, and provide the evidence base necessary for informed policy and practice. While these technologies offer transformative potential for personalized learning, they also introduce significant risks that must be proactively mitigated. By investigating the efficacy of AI tools and building the evidence base for effective practice through the lenses of civil rights and UDL, researchers can establish the necessary foundation to inform the field..

Pillar 1: Upholding Civil Rights

- **Conduct studies on how AI can be used to support students with disabilities**, including identifying opportunities for AI tools to be used to support students with disabilities, the effectiveness of integrating AI into Individualized Education Programs (IEPs) planning structures and processes, and how AI can be made more accessible for students with disabilities. Specific intervention studies are needed in these areas, as multiple literature reviews have explored the historical dimensions of AI applications for students with disabilities. More rigorous research on ensuring accessibility and promoting opportunities for students with disabilities to actively participate and have control over their learning experiences is necessary
- **Conduct studies on how the potential harms of using AI in education for students with disabilities can be mitigated**, including bias and discrimination, physical and emotional safety, increased surveillance, data re-identification, anthropomorphization of AI, and AI's impact on brain development.

Pillar 2: Data Privacy and Student Protection

- Research and explore how utilizing privacy-enhancing technologies in AI may help avoid and mitigate potential privacy harms for students with disabilities.

To learn more, see the [Privacy-Enhancing Technologies](#) work of Georgetown University's Massive Data Institute. Refer to New America's guidance on [privacy-enhancing technologies \(PETs\)](#), such as differential privacy, federated learning, and synthetic data generation, to evaluate AI effectiveness while minimizing the exposure of individual student data, particularly sensitive disability-related information.

Pillar 3: Accessibility by Design

- **Investigate how AI literacy initiatives can be designed to be accessible and relevant for students with disabilities**, families, and educators across diverse communities, examine the impacts of AI tools on different student populations, and explore how developers can apply Universal Design for Learning and the Learning Sciences to create inclusive, effective, and pedagogically aligned AI tools. A recent initiative from the [Stanford Accelerator for Learning](#) highlights how AI can support learners with disabilities when tools are co-designed with individuals with learning differences and grounded in principles of accessibility and student agency.

Pillar 4: Transparency, Monitoring, & Accountability

- **Conduct longitudinal studies examining how the efficacy and effectiveness of AI tools for students with disabilities change over time**, including whether benefits persist, diminish, or improve with sustained use. Additional research is needed to determine the efficacy of AI-enhanced curricular materials for students with unique abilities.
- **Identify and utilize accurate measures and indicators of the effectiveness of AI when used for its intended purpose** or other foreseeable purposes, the level of risk associated with such uses, and how to identify potential harms for students with disabilities, such as bias and discriminatory outputs. Look at [Towards an Educator-Centered Method for Measuring Bias in Large Language Model-Based Chatbot Tutors](#) for a developing method on measuring and mitigating potential harms from chatbot tutors. Also see [Using Demographic Data as Predictor Variables: a Questionable Choice](#).

GUIDANCE FOR TECHNOLOGY DEVELOPERS

To move beyond minimum compliance with accessibility, privacy, and disability laws, technology developers should embrace **responsible stewardship** as the foundational architecture for AI tools. By adopting benchmarks that reflect promising practices for inclusive design, such as Universal Design for Learning and the ADA Title II final rules, the education technology community can prioritize innovation that does not leave students with disabilities behind. This design imperative requires a direct partnership between developers and the disability community to build a digital ecosystem in which capacity is not just supported but actively expanded for every learner. Centering the lived experiences of students and the expertise of families and researchers shifts the focus from reactive repair toward proactive, inclusive innovation.

Pillar 1: Upholding Civil Rights

- Design AI systems with accessibility and flexibility built in from the beginning, not retrofitted after development. Apply Universal Design for Learning (UDL) principles to ensure such tools are accessible for students with diverse needs. When evaluating or developing algorithms, actively seek approaches that reduce discriminatory impacts and bias, as discussed in "Less Discriminatory Algorithms." See CAST's AI and UDL guidance to work within the minimum conformance of WCAG 2.1 AA, consistent with the recent ADA Title II final rule.

Pillar 2: Data Privacy and Student Protection

- **Expand existing comprehensive data governance policies** to explicitly detail the lifecycle of student information within AI systems. This includes providing clear declarations on how data is collected, protected, and destroyed, with specific disclosures regarding whether student information and data are used to train or refine underlying AI models.

- **Design service agreements and contracts** to address the multiple components of student data privacy. These should include clear provisions for access controls, de-identification practices, and specific retention limits, using models such as the University of Florida's GADRRS framework as a guide for comprehensive contractual protection.
- **Apply core privacy principles** while tracking the metrics necessary to evaluate AI efficacy, including data minimization and purpose limitation. It is essential to ensure that training datasets and interventions are representative of students with disabilities to accurately measure effectiveness. Fair Now has created an AI Governance Framework to guide AI developers in designing and deploying AI systems that mitigate bias and privacy violations.
- **Incorporate customizable privacy controls directly into AI interfaces** to support LEAs oversight of student rights. For example, developers can integrate "just-in-time" warnings that alert educators and administrators when the system detects that student personal information has been entered.

Pillar 3: Accessibility by Design

- **Provide clear, plain-language disclosures about the purpose, capabilities, limitations, and safety policies of AI.** The Center for Democracy and Design & Technology's Plain Language Hub presents a model and explanation of this communication strategy.
- **Work with local education agencies to ensure that AI meets educational standards, safety, and accessibility requirements.** Refer to the new final rule on accessibility under Title II of the Americans with Disabilities Act (ADA) and the Office of Elementary and Secondary Education's information on standards and assessments.

- **Co-design tools with students with disabilities, advocates, researchers, educators, and families that are aligned with the Learning Sciences and UDL** to ensure that the tools are pedagogically effective, age-appropriate, and accessible to all learners. The Universal Design for Learning Product Certification, a collaboration between CAST and Digital Promise funded with support from the Educating All Learners Alliance, helps inform educators who are evaluating the accessibility and usability of edtech products. To earn this certification, a company must prove they used Universal Design for Learning principles and conducted User Testing with people with disabilities during the development of their tool.

Pillar 4: Transparency, Monitoring, & Accountability

- **Regularly assess AI for potential harms** to students with disabilities by conducting data analysis, algorithmic impact assessments, and privacy impact assessments. Partners should be contractually obligated to comply with established protections for student data and to regularly review and audit partners' data management practices. For more information on evaluating AI tools, see the NIST AI RMF Playbook.
- **Develop an accessibility roadmap** with goals, priorities, and key activities for continuously improving conformance with WCAG 2.1 AA and beyond. Embed accessibility experts in the design and development workflow of AI products to ensure conformance throughout the production process. The National Education Association provides guidance on applying key accessibility principles.
- **Identify the purpose AI is designed to achieve** by providing standardized statements using the Every Student Succeeds Act (ESSA) Tiers of Evidence to indicate the tools' efficacy for such purpose.
- **Encourage and facilitate independent third-party research** on the efficacy or benchmarking of AI for its intended purpose and how it supports students with disabilities. Publish progress reports accessible to current and prospective customers. See the Software and Information Industry Association's guidelines for the responsible development of AI tools for education.

APPENDIX

Additional Resources

Accessibility & Universal Design

ADA Title II Final Rule on Accessibility of Web Content and Mobile Apps (U.S. Department of Justice), <https://www.ada.gov/resources/2024-03-08-web-rule/>

The Impact of AI in Advancing Accessibility for Learners with Disabilities (EDUCAUSE Review), <https://er.educause.edu/articles/2024/9/the-impact-of-ai-in-advancing-accessibility-for-learners-with-disabilities>

Quality Indicators for Assistive Technology Services (QIAT Consortium), <https://www.qiat.org/>

Three Strategies to Enhance Accessibility Using AI (CoSN), <https://www.cosn.org/three-strategies-to-enhance-accessibility-using-ai/>

Universal Design for Learning Guidelines (CAST), <https://udlguidelines.cast.org/>

Voluntary Product Accessibility Template (VPAT®) (ITI), <https://www.itic.org/policy/accessibility/vpat>

Web Content Accessibility Guidelines (WCAG) 2.1 (W3C), <https://www.w3.org/TR/WCAG21/>

AI Literacy & Professional Development

AI4K12 Initiative (AAAI/CSTA), <https://ai4k12.org/>

AI Literacy Curriculum (MIT Media Lab), <https://www.media.mit.edu/projects/ai-literacy/overview/>

The Moral Machine (MIT), <https://www.moralmachine.net/>

The Most Likely Machine (Interactive AI Learning Tool), <https://www.aaai.org/teach-ai-in-k-12/the-most-likely-machine/>

National AI Literacy Day Resources, <https://www.nationalailiteracyday.org/>

Survival of the Best Fit (AI Hiring Simulation), <https://www.survivalofthebestfit.com/>

Teaching Privacy: Digital Literacy Curriculum, <https://teachingprivacy.org/>

Civil Rights, Equity & Algorithmic Bias

Ableism and Disability Discrimination in New Surveillance Technologies

(Center for Democracy & Technology), <https://cdt.org/wp-content/uploads/2022/05/2022-05-23-CDT-Ableism-and-Disability-Discrimination-in-New-Surveillance-Technologies-report-final-redu.pdf>

Algorithmic Bias: The State of the Situation and Policy Recommendations

(Ryan S. Baker, Aaron Hawni, and Seiyon Lee, OECD Digital Education Outlook), https://www.oecd.org/en/publications/oecd-digital-education-outlook-2023_c74f03de-en/full-report/algorithmic-bias-the-state-of-the-situation-and-policy-recommendations_a0b7cec1.html

Avoiding the Discriminatory Use of Artificial Intelligence (U.S. Department of Education, Office for Civil Rights),

<https://www.ed.gov/media/document/avoiding-discriminatory-use-of-artificial-intelligence-112274.pdf>

The Fallacy of AI Functionality (Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew D. Selbst),

<https://dl.acm.org/doi/10.1145/3531146.3533158>

Late Applications: Protecting Students' Civil Rights in the Digital Age (Center for Democracy & Technology),

<https://cdt.org/insights/report-late-applications-protecting-students-civil-rights-in-the-digital-age/>

Less Discriminatory Algorithms (Upturn), <https://www.upturn.org/work/less-discriminatory-algorithms/>

Evidence & Research

Exploring the Impact of AI in K-12 Education (Artefact),

<https://www.artefactgroup.com/case-studies/exploring-the-impact-of-ai-in-k-12-education/>

Towards an Educator-Centered Method for Measuring Bias in Large

Language Model-Based Chatbot Tutors, (Emma Harvey, Allison Koenecke, and Rene F. Kizilcec), <https://openreview.net/forum?id=REidmDpL9r>

Using Demographic Data as Predictor Variables: A Questionable Choice

(2023, Ryan Shaun Baker, Lief Esbenshade, Jon Vitale, and Shamyia Karumbaiah), <https://doi.org/10.5281/zenodo.7702628>

Using Evidence to Support EdTech Adoption in Schools (U.S. Department of Education, Office of Educational Technology), <https://tech.ed.gov/evidence/>

General AI in Education Frameworks & Guidance

AI in PK-12 Education (2025, Digital Promise),
<https://doi.org/10.51388/20.500.12265/269>

Bringing AI to School: Tips for School Leaders (2023, ASCD, ISTE, NASSP, NAESP, and AASA), <https://www.aasa.org/resources/resource/bringing-ai-to-school-tips-for-school-leaders>

EdSAFE AI Alliance, <https://www.edsafeai.org/>

Historical Context

Teaching Machines: The History of Personalized Learning (2021, Audrey Watters), <https://teachingmachines.es>

Responsible Procurement & Adoption

AI + Learning Differences: Designing a Future with No Boundaries (Stanford Accelerator for Learning),

https://acceleratelearning.stanford.edu/app/uploads/2025/07/AI-Learning-Differences-Designing-a-Future-with-No-Boundaries_Final.pdf

AI Tools: Artificial Intelligence Tools for SPS Students and Staff (Seattle Public Schools), <https://www.seattleschools.org/student-portal/technology-supports-for-families/digital-tools/ai-tools/>

Alabama AI Policy Template for LEAs (AI for Education),
<https://www.aiforeducation.io/ai-resources/state-ai-guidance>

Algorithmic Equity Toolkit (ACLU of Washington), <https://www.aclu-wa.org/AEKit>

Legal Considerations: AI and Business Contracts (Public Interest Privacy Center and F3 Law), <https://f3law.com/>

NIST AI Risk Management Framework (AI RMF) (National Institute of Standards and Technology), <https://www.nist.gov/itl/ai-risk-management-framework>

NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, Version 1.0 (National Institute of Standards and Technology), <https://www.nist.gov/privacy-framework>
Responsible AI and Tech Justice Guide for K-12 Education (Kapor Foundation), <https://kaporfoundation.org/wp-content/uploads/2024/01/Responsible-AI-Guide-Kapor-Foundation.pdf>
Student Data Privacy Consortium Resource Registry, <https://sdpc.a4l.org>

Special Education & IEP Resources

Artificial Intelligence in Special Education: Resources and Recommendations (CAST), <https://www.cast.org/what-we-do/artificial-intelligence/>
Heightened AI use in special education brings elevated risks (K-12 Dive), <https://www.k12dive.com/news/artificial-intelligence-special-education-Section-504-benefits-risks-privacy-IDEA-IEP/804535/>
NEA AI and Accessibility Decision Tree, <https://www.nea.org/professional-excellence/student-engagement/tools-tips/ai-and-accessibility>

State & District Policy Examples

Boston Public Schools Artificial Intelligence Guidance, <https://bostonpublicschools.helpdocs.io/article/7nvt595hpb-guidance-on-the-use-of-artificial-intelligence-in-bps>
Seattle Public Schools AI Tools List, <https://www.seattleschools.org/student-portal/technology-supports-for-families/digital-tools/ai-tools/>
Washington State AI Guidance for K-12 Schools (Office of Superintendent of Public Instruction), <https://www.k12.wa.us/student-success/resources-subject-area/digital-citizenship-internet-safety/artificial-intelligence-guidance>

Student Privacy & Data Protection

K-12 Privacy Policy Guide: How to Quickly Spot Red Flags (Public Interest Privacy Center), <https://publicinterestprivacy.org/privacy-policy-red-flags/>
Privacy-Enhancing Technologies (Georgetown University Massive Data Institute), <https://mdi.georgetown.edu/privacy-enhancing-technologies/>

The Privacy and Equity Implications of Using Self-Harm Monitoring Technologies (Future of Privacy Forum),

<https://studentprivacy.org/resources/privacy-equity-implications-self-harm-monitoring-technologies/>

Privacy and Security Considerations (New America),

<https://www.newamerica.org/oti/blog/artificial-intelligence-in-schools-privacy-and-security-considerations/>

Student Data Privacy Toolkit Parts 1, 2 & 3 (CoSN),

<https://www.cosn.org/tools-and-resources/resource/student-data-privacy-toolkit-part-1-2-3/>

Student Privacy Communications Toolkit: For Schools & Districts (Future of Privacy Forum), <https://studentprivacy.org/resource/student-privacy-communications-toolkit-for-schools-districts/>

Vetting Generative AI Tools for Use in Schools (Future of Privacy Forum), https://fpf.org/wp-content/uploads/2024/10/Ed_AI_legal_compliance.pdf_Flnal_OCT24.pdf

ACKNOWLEDGEMENTS

Supported by InnovateEDU's Educating All Learners Alliance and New America, this policy paper was made possible by the diverse perspectives and expertise shared at the 2024 and 2025 Prioritizing Students with Disabilities in AI Policy convenings. We would like to acknowledge the invaluable contributions of the educators, school administrators, parents, disability rights advocates, government, industry, and privacy experts who participated in this convening and provided subsequent revisions. Please note that the views expressed in this paper are not endorsements made by organizations. This work was funded by the EDSAFE AI Alliance in partnership with Educating All Learners.

CONTRIBUTORS

Barbara Pape

Benetech

Brett Roer

Bruce Alter, PT, AT
Consultant

CAST (Center for Applied
Special Technology)

Center for Democracy &
Technology (CDT)

Chad Sussex

DC Special Education
Cooperative

Metro Nashville Public
Schools

National Center on
Accessible Digital
Educational Materials &
Instruction (NCADEMI)

National Center for
Learning Disabilities
(NCLD)

National Digital Inclusion
Alliance

National Education
Association (NEA)

NewSchools Venture
Fund

Rachel S. Phillips

The Study Group