AI SAFETY & PRIVACY STANDARD FOR SCHOOLS

MEMORANDUM OF AGREEMENT

To Protect Students, Families and Educators

Highest Level of Protection | Ten Principles | Fully Enforceable

September 2026, Version 1.0

Preamble — Why This Agreement Exists

This Memorandum of Agreement exists for one reason: children come first.

Students are not products. Teachers are not beta testers. And schools are not a source for data collection or experiments. When artificial intelligence tools enter a classroom, they bring enormous potential, and equally enormous responsibility. This agreement sets one of the highest standards for how that responsibility must be met.

Likewise, the National Academy for AI Instruction is not a data company. The Academy is a national initiative dedicated to equipping educators with the knowledge, skills, and support to use artificial intelligence effectively and responsibly in teaching and learning.

This is a binding Memorandum of Agreement between the National Academy for AI Instruction (referred to throughout as "the Academy") and any AI Provider that the Academy engages to provide AI-powered services to educators or students. Every requirement in this agreement is mandatory. Vendors who cannot meet these requirements are not eligible to serve the Academy or to achieve the Standard. The Standard is a seal of excellence and a public commitment to the safety, privacy, and well-being of children, families, and educators. It signifies that an AI Provider has met the Academy's highest standards for responsible, trustworthy, and education-centered artificial intelligence.

Memorandum guiding principle: When in doubt, protect the students, families, and educators.

Beyond the Academy

The signatories to this Standard agree that the principles set out below do not belong only to the Academy but are available to any school district that desires the same principles to apply when purchasing AI Provider Educational Products. Each AI Provider that signs this agreement will create a clear, simple, and expeditious process through which any school district can request that the Provider meet or exceed the Standard in its own specific agreements for the AI Provider Educational Products (defined below). As a best practice, each AI Provider will encourage

school districts to adopt the Standard and will support districts seeking to incorporate its principles into their agreements. Each participating AI Provider will make the substantive protections of the Standard available to EDU Customers upon request within 90 days of the Effective Date, provided that such protections may be incorporated into the EDU Customer's existing data privacy agreement, licensing agreement, or addendum and need not be adopted in the identical form set forth in the Standard. The process should be easy to find and easy to use, and a district can make the request at any time. Once a district makes the request, the AI Provider will work diligently to incorporate the applicable terms into that district's agreement.

Key Definitions

As used in this Standard:

- "Student Data" means any information that is directly related to an identifiable student, including but not limited to: names, student ID numbers, grades, disciplinary records, prompts entered into AI Provider Educational Product, AI-generated outputs tied to a student, behavioral patterns, device identifiers, location data, memory files, audio or visual data, and metadata. This definition is intentionally broad. When in doubt about whether something is Student Data, treat it as Student Data.
- "Covered Data" means Student Data plus any data belonging to or generated by educators, administrators, or an EDU Customer (including the Academy) when using AI Provider Educational Products. Covered Data does not include Telemetry Data (as defined below).
- "Telemetry Data" means technical and operational data generated through the use or operation of the AI Provider Educational Products that has been de-identified so that it cannot reasonably be used to identify, relate to, describe, or be linked to an individual student, or educator. Telemetry Data may be associated with an EDU Customer solely for security, reliability, debugging, capacity planning, service health, aggregate performance monitoring, and to make product improvements. Telemetry Data may not be used, directly or indirectly, for generative-model training, student profiling, individual student user personalization of AI outputs, advertising, marketing, behavioral inference, or the creation or enhancement of profiles concerning individual users.
- "AI Provider" means the company providing the AI-powered service, including all subsidiaries, affiliates, contractors, and sub processors.
- "Training" means any use of data to develop, fine-tune, update, improve, or otherwise alter an AI model's weights, parameters, or behaviors.
- "AI Provider Educational Products" means services or features that use generative artificial intelligence technologies to generate outputs that are primarily designed and marketed for use by students, educators and administrators, with authenticated usage under an agreement between the educational entity and the AI Provider. For the avoidance of doubt, general-purpose productivity, collaboration, communication, search, cloud, development, or workplace-assistance products, including products that may be licensed or used by educational entities but are not primarily designed for educational purposes do not constitute an AI Provider Educational Product.
- "EDU Customer" means an educational entity located in the United States that licenses AI Provider Educational Products from AI Provider. For clarity, the Academy may be an EDU Customer when it licenses such AI Provider Educational Products.

- "Academy" means the signing institution, including its students, educators, staff, and administrators.
- "Breach" means any unauthorized access to, disclosure of, alteration of, or destruction of Covered Data.
- "Zero Data Retention (ZDR)" means that inputs (prompts) and outputs (responses) are not stored, logged, or retained beyond the immediate session required to deliver the service.
- "Synthetic Data" means data that is artificially generated to simulate patterns, scenarios, or edge cases. Under this agreement, Synthetic Data may be used only if it is not created from, seeded with, tuned on, or reasonably reversible to Covered Data.
- "Agentic or Action-Taking Feature" means any AI capability that can initiate, approve, transmit, modify, or execute actions on behalf of a user, online or offline, rather than only generating informational output.
- "Memory Files" means anything an AI Provider Educational Product saves about a user between sessions, stored facts, preferences, profiles, or conversation summaries, so it can "remember" them later. All Memory Files are Covered Data under this agreement.

PART I — THE TEN CORE PRINCIPLES

Each principle below is a binding commitment. The AI Provider must comply with all ten, across every AI Provider Educational Product it provides to an EDU Customer who has opted in to the principles. Complying with some principles but not others is not compliance.

These principles are effective from the date signed and remain in force for a term of two years (the “Effective Date”). To continue past that term, the agreement must be renewed in writing by both the Academy and the AI Provider every two years. The Academy may terminate the AI Provider's participation in the Standard immediately upon any material breach by the AI Provider that is not cured within the period set forth in Principle 5.

Compliance with Applicable Laws

All aspects of this Standard shall be implemented in accordance with all applicable federal, state, and local laws and regulations, including FERPA, COPPA, IDEA, and applicable student privacy, data protection, accessibility, civil rights, consumer protection, and education laws. Nothing in this Standard limits any right, protection, or obligation provided by law. Where this Standard provides stronger protections, those stronger protections shall apply to the extent permitted by law.

PRINCIPLE 1 | No Student or Educator Data for Training — Ever, With One Narrow Safety Exception

The AI Provider absolutely may not use any Covered Data, including prompts, responses, uploaded files, behavioral signals, metadata, or any derivative thereof, to train, fine-tune, update, benchmark, or otherwise improve any AI model, at any time, for any purpose, except as expressly permitted under the Narrow Safety and Security Exception below.

This prohibition:

- Applies to all data types without exception: text, audio, images, video, biometrics, metadata, interaction logs.
- Applies to de-identified, aggregated, transformed, or other Covered-Data-derived datasets. De-identification is not a loophole, and Covered Data may not be used as source material for synthetic data generation.
    - De-identified data may be used for Telemetry Data (as outlined in Key Definitions).
- Does not prohibit the AI Provider from using wholly Synthetic Data that was created without using Covered Data as source material and is not reasonably linkable or reversible to any student, educator, or EDU Customer record.
- Applies retroactively: the AI Provider must not use any previously collected Covered Data for training.
- Survives the termination of this agreement indefinitely.
- Applies to all subsidiaries, sub processors, and third-party partners of the AI Provider (excluding any third-party models implemented by the user and not provided by the AI Provider directly).

Narrow Safety and Security Exception

Covered Data may be used for safety and security purposes only under all the following conditions:

- **Safety and Security Functionality.** The exception applies only to processing that is reasonably necessary to detect, prevent, investigate, mitigate, or remediate harm to users or threats to the security or integrity of the AI Provider Educational Products, including self-harm risk, child sexual abuse material, grooming, threats of violence, bullying, malicious activity, security incidents, vulnerabilities, or unauthorized access. Safety or security functionality may be technically integrated into a broader model, system, or product that performs other functions, provided that the processing of Covered Data under this exception remains strictly limited to the safety or security purposes described in this paragraph.
- **Purpose-Limited Use.** Covered Data processed under this exception may be used solely for the applicable safety or security purpose and may not be used for general model training or improvement, user profiling, personalization, advertising, behavioral inference, market research, development of unrelated products or functionality, or other commercial purposes.
- **Minimum Necessary.** The AI Provider may process only the minimum amount of Covered Data reasonably necessary to accomplish the applicable safety or security purpose and must use reasonably available technical and organizational measures to minimize the Covered Data processed.
- **Retention Limits.** Covered Data processed under this exception may be retained only for as long as reasonably necessary to accomplish the applicable safety or security

purpose or to satisfy applicable legal obligations, after which it must be deleted or de-identified in accordance with this Standard.

Auditability. Upon request, the AI Provider must maintain appropriate records sufficient to demonstrate that processing under this exception is limited to the permitted safety or security purposes, including the nature and purpose of the processing, applicable retention periods, and relevant safeguards. Such records must be reasonably available for review and audit, subject to reasonable confidentiality, security, and privilege considerations.

- **No General Model Improvement.** Nothing in this exception permits Covered Data to be used to train, retrain, fine-tune, evaluate, or otherwise improve a general-purpose or generative AI model, except to the limited extent that a specific safety or security classifier, safeguard, or functionality is being evaluated or improved for the safety or security purposes expressly permitted above.
- **Interpreted Narrowly.** This exception will be interpreted narrowly. The burden is on the AI Provider to demonstrate that any processing of Covered Data under this exception satisfies each of the foregoing requirements.

Annual Confirmation Required:

Upon request, the AI Provider must deliver a written annual statement, in plain language, signed by a senior officer to the Academy, confirming that (1) no Covered Data was used for training outside the Narrow Safety Exception in the preceding twelve months, and (2) no data, learnings, or model improvements derived from Covered Data crossed from safety systems into any other model or product.

PRINCIPLE 2 | Data Minimization by Design

The AI Provider may only collect, process, or store data that is strictly necessary to deliver the specific AI Provider Educational Product for which an EDU Customer has contracted. Nothing more.

For clarity, strict necessity may include limited contextual or operational data elements, such as school, district, class, course, assigned teacher, device, or coarse location information, only when they are genuinely required to deliver the AI Provider Educational Product.

Prohibited practices (no exceptions without explicit written approval from the EDU Customer):

- Precise geolocation tracking or continuous location monitoring, except where a documented safety or operational need is noted in these principles and/or Addendum A.
- Behavioral tracking, keystroke logging, or continuous passive attention monitoring.
- Long-term profiling of students or educators.
- Biometric data collection of any kind (facial recognition, voice prints, gaze tracking, emotional inference).

- Collection of data beyond what is necessary for the immediate session, administrative controls, or oversight.
- Where memory or personalization is offered, the AI Provider must give the EDU Customer admin controls to disable it or disable those features upon request, and the EDU Customer must disable any persistent memory or personalization features that are not governed by the EDU Customer's configured retention and deletion policies or by equivalent retention and deletion controls that provide the same or greater protection.
- Collection of data from students under 13 without verifiable parental consent compliant with COPPA, where such consent is required by applicable law and AI Provider is a controller. In scenarios where the AI Provider is a processor, the EDU Customer must obtain such consent or, in the absence of such consent, select which AI Provider Educational Products to disable. The AI Provider will provide mechanisms to enable EDU Customers to record which students under 13 have or have not obtained such consent.
- **Privacy Impact Assessment (PIA) Required, Including New Features:**

Before any material new feature, capability, or data collection practice is activated for an AI Provider Educational Product, the AI Provider must complete a Privacy Impact Assessment (PIA) in accordance with its standard practices. For purposes of this requirement, a change is "material" if it is reasonably likely to increase student privacy or safety risk, including through the collection or use of new categories of data or new purposes for existing data, biometrics, profiling, agentic authority, significant changes to data retention, new or materially changed safety functionality, or other changes reasonably likely to increase such risk. Upon request, and subject to confidentiality terms, the AI Provider will share relevant PIA documentation. See Principle 9 for more details.

PRINCIPLE 3 | The Customer Owns Its Data — Completely

All Covered Data, including every prompt, response, file, and AI-generated output, belongs exclusively to the EDU Customer and its students. The AI Provider has no ownership interest in any of it, ever.

The AI Provider must:

- Never sell, license, share, or transfer Covered Data to any third party for any purpose, other than to provide the services requested by the EDU Customer, comply with the law, or ensure the safety of minors.
- Never use Covered Data for product development, research, or any purpose other than delivering the contracted service. No Covered Data may ever be used for advertising.
- Provide the EDU Customer with the ability to export all Customer Content and other Covered Data reasonably capable of export at any time in standard, non-proprietary, machine-readable formats (such as JSON, CSV, or PDF). Any category of Covered Data excluded from export must be identified and reasonably justified, such as ephemeral security artifacts or Telemetry Data that cannot reasonably be associated with an EDU Customer user or account.

- Give the EDU Customer direct control over deletion and retention of Covered Data through native administrative tools, so the EDU Customer can set its own retention policies and initiate deletion itself. Where the AI Provider acts solely as a data processor, it must delete or retain Covered Data according to the EDU Customer's configured policies and documented instructions.
- Honor deletion fully and on a contractually agreed timeline. Once the EDU Customer initiates deletion (whether through administrative tools or a verified written request), the AI Provider must permanently delete it from active systems within 180 days, subject to EDU Customer-configured retention policies, documented legal holds, and retention required by law. Backup copies will be deleted or rendered commercially unrecoverable according to documented service-specific purge cycles, not to exceed 180 days, except where retention is required by law or a documented legal hold. Upon request, the AI Provider will provide written confirmation when the deletion is complete.
- Maintain Covered Data solely within the European Union or North America and allow the EDU customer to identify the geographic location or region in which that EDU Customer's Covered Data is stored or processed either through administrative tools or upon request.
- Covered Data may be retained or excluded from deletion workflows only to the extent and for so long as reasonably necessary for evidence-preservation, applicable legal compliance purposes, or as needed for purposes consistent with the Narrow Safety and Security Exception above (such as authorized safety, security, abuse-prevention, investigation, and incident-response purposes).

NOTE ON "ZERO DATA RETENTION" SERVICES: If the AI Provider operates under a Zero Data Retention model (such as OpenAI's ZDR API), prompts and outputs will not be stored by the Provider at all. In such cases, the Academy is solely responsible for maintaining its own records of interactions if records are needed. This arrangement must be explicitly acknowledged by the Academy in writing and reflected in Addendum A.

User Control of Memory Files:

Upon request, EDU Customers can control whether an AI Provider Educational Product "remembers" Student Data. If an AI Provider Educational Product is configured to retain any Student Data about an individual, that person, where applicable, must be able to:

- **See it.** View the Student Data the tool has saved about them, in plain language, at any time.
- **Erase it.** Delete Memory Files at any time, with clear and simple actions. Deletion must be real and permanent, not just hidden from view.
- **Turn it off.** Use the tool with memory switched off entirely, without losing access to the service.

The AI Provider may never require a Memory File as a condition of using the service and may never use a Memory File for any purpose other than serving that individual user. Where

persistent memory is permitted, upon request, the EDU Customers can control whether an AI Provider Educational Product "remembers" Student Data. If an AI Provider Educational Product is configured to retain any Student Data about an individual, that person (or, for younger students, their parent or guardian working through the EDU Customer), where applicable, must be able to see it, erase it, and turn it off as provided below.

Administrative Access Required:

The AI Provider must provide the EDU Customer with a secure administrative means to obtain information about the Provider's handling of Covered Data. Compliance may be provided through an administrative portal, report, API, or documented request process. The information provided must be sufficient for the EDU Customer to understand the categories of Covered Data processed, applicable retention periods, data storage locations, subprocessors, and access controls.

PRINCIPLE 4 | Human Oversight and Explainability

AI tools support educators and students. They do not replace human judgment. No AI system covered by this Standard may make decisions for the EDU Customer's educators and students, without allowing for meaningful human review or prior approval.

The disclosure, testing, and risk-control requirements in this Principle are required transparency artifacts and deployment conditions, not merely aspirational statements.

Required transparency practices:

- The AI Provider must publish, in plain language, a document explaining how its AI Provider Educational Products generate outputs, a sufficiently detailed summary of the data used for training AI models, with a reference to any relevant card(s) or similar artifacts that provide additional information on the training program pertaining to the model, and what its known limitations and failure modes are. This document must be updated whenever the product changes substantially (see Principle 9).
- The AI Provider must provide clear in-product indicators whenever students or educators are interacting with AI-powered products. An initial indicator that the user is interacting with an AI technology at the beginning of the user session is sufficient where the user remains in a clearly identified AI-powered experience for that session.
- The AI Provider must conduct regular fairness testing results across diverse education populations.
- The AI Provider must provide the EDU Customer with the ability to enable or disable AI Provider Educational Products, including to account for grade level, age, and use case. This shall include the ability to apply more restrictive settings for students under the age of 13. The EDU Customer may maintain the relevant grade-level and age information and determine which available AI Provider Educational Products to enable or disable.

AI-Specific Safety and Risk Controls:

- The AI Provider may not intentionally design, deploy or maintain an AI system that generates content the AI provider knows is materially false or misleading. This requirement is not violated solely because an AI-generated response contains inaccurate information, provided the AI Provider maintains reasonable safeguards and takes appropriate corrective action when it becomes aware of a significant problem.
- The AI Provider may not provide any EDU Customer user with companion-style or relationship-oriented features that are designed to foster emotional attachment or dependency, to keep students engaged for longer than the learning task requires, simulate friendship, or encourage disclosure of sensitive information.
- The AI Provider must evaluate risk mitigations for content related to self-harm, suicide, sexual content, abuse, violence, and other high-risk topics likely to arise in school use, including crisis-level protocols for content relating to self-harm, suicide, or violence.
- Any agentic or action-taking capability, including sending messages, submitting work, changing settings, accessing third-party systems, or making purchases or commitments on behalf of a user, must be able to be disabled for student users. For any deployment authorized for student use, externally consequential agentic features that take action on behalf of students outside of the boundary of the AI Provider's systems must be disabled by default at the tenant or policy level unless an authorized EDU Customer administrator affirmatively enables them. However, safety or security functionality designed to identify, detect, prevent, or mitigate harm to students, including bullying, self-harm, abuse, or threats of violence, need not be disabled by default so long as it does not independently take an externally consequential action without appropriate human review or EDU Customer-approved controls. AI systems must not act beyond their defined scope, permissions, or authorized limits. Any such capability may be enabled only with explicit EDU Customer approval, scoped permissions, appropriate human oversight, and complete audit logging.
- The product must provide clear, age-appropriate notices that AI outputs may be incorrect.

As part of the Academy's instruction, the Academy will instruct that educators cannot use AI for the following prohibited uses, no exceptions:

- Automated grading without human review.
- Automated disciplinary decisions or recommendations based on AI analysis of educator or student behavior.
- Automated placement decisions (course selection, ability grouping, special education referrals)
- Emotional or psychological assessment of educators or students
- Surveillance of educators or students for behavioral prediction purposes

Model Specification / Governance Document:

The AI Provider must maintain and publish documentation (sometimes called a "model spec," "model card," "system card," or "application card") that describes the capabilities and limitations of the AI Provider Educational Products and available oversight mechanisms. This

documentation must be publicly accessible and kept current. The enforceable commitments on data use are defined in this Standard, not in that governance document alone. The AI Provider may also make available, either publicly or to customers, information about AI governance policies or programs.

PRINCIPLE 5 | Real Accountability — Enforceable, Not Just Promised

A privacy agreement without teeth is not a privacy agreement. This agreement creates genuine, enforceable obligations with real consequences.

Third-Party Audits:

The AI Provider must pursue or maintain current, independent third-party certifications appropriate to the sensitivity of educational data. Required certifications include those identified below. If the AI Provider does not maintain a required certification at the time of signing, the AI Provider must be actively pursuing that certification and a target date by which the certification will be obtained must be specified in **Addendum C**.

- SOC 2 Type II (annual renewal required)
- ISO 27001 (Information Security Management)
- ISO 27701 (Privacy Information Management)
- ISO 42001 (AI Management Systems) or equivalent
- FedRAMP Moderate authorization or an equivalent federal authorization, where required for the Educational Product's EDU Customer, deployment, or use.

All current audit reports and certification documentation must be provided to the Academy upon request.

Breach Notification:

In the event of any actual or suspected Breach affecting an EDU Customer's use of the AI Provider Educational Product, the AI Provider must:

- Notify the EDU Customer without unreasonable delay and, in any event, no later than 72 hours after becoming aware of a confirmed or reasonably suspected Breach affecting Covered Data. Initial notice may be based on the facts then known but must clearly identify what remains under investigation.
- Provide material updates on a rolling basis.
- Cooperate fully with any regulatory investigation or legal proceeding.

Contractual Provisions for EDU Customers:

- If requested, the AI Provider agrees to make the substantive protection of the Standard available to EDU Customers within 90 days of the Effective Date through a clear, simple, and expeditious process. Such protections may be incorporated into an EDU Customer's existing data privacy agreement, licensing agreement, or addendum and need not be adopted in the identical form set forth in the Standard, provided that the resulting terms meet or exceed the substantive protections of the Standard. As a best practice, the AI

Provider will encourage EDU Customers to adopt the Standard and will support them in incorporating its protections into their agreements.

- Once the substantive protections of the Standard are incorporated into an EDU Customer's agreement, those protections will be contractually enforceable by the EDU Customer. In the event of a breach of those incorporated protections, the EDU Customer will have the contractual remedies available under its agreement with the AI Provider, including, without limitation, the right to terminate for an uncured material breach and to seek damages or other remedies available under its agreement or applicable law.
- For clarity, making the Standard's substantive protections available to an EDU Customer reflects commitments undertaken by the AI Provider only. Nothing in this Standard or in the AI Provider's process for making these protections available creates, expands, modifies, or imposes any obligation, responsibility, representation, warranty, or liability on the EDU Customer or requires the EDU Customer to change its existing practices, configurations, or use of the AI Provider Educational Products.

Contractual Remedies for the Academy:

- To the extent permitted by applicable law, the Academy may enforce the commitments set forth in the Standard.
- Upon an alleged breach, the parties will engage in good faith to identify appropriate corrective or remedial measures consistent with this Agreement. A remediable breach must be cured within a period that is reasonable under the circumstances, taking into account the nature and complexity of the issue, the risk presented, and the time reasonably required to implement and validate remediation. If the AI Provider has commenced remediation and is diligently pursuing it, the Academy and/or AFT may not revoke the AI Provider's participation solely because remediation is not yet complete. If the parties disagree as to whether a breach occurred, the appropriate time to cure the breach, or whether the remediation is adequate, the matter must be escalated to senior representatives for good-faith consultation before revocation. Failure to provide a required transparency artifact or evaluation item is a material breach and grounds for the Academy to pause or deny deployment, subject to this cure process.

- If a material breach or repeated breach is not cured or otherwise resolved in accordance with the notice, remediation, and escalation process above, including through corrective measures that provide substantially equivalent or greater protection, the Academy and/or AFT may revoke the AI Provider's participation in the Standard and pursue any remedies available under this Agreement or applicable law. Upon the effective date of revocation, the AI Provider must immediately cease representing that it participates in or is currently recognized under the Standard. The Academy and/or AFT may publicly announce that the AI Provider is no longer a participant in the Standards, including the effective date of the revocation and, where appropriate, a general description of the basis for the revocation, provided that any such public statement is accurate, does not disclose confidential information protected under this Agreement, and does not characterize a disputed matter as finally determined. Except where immediate disclosure is reasonably necessary to prevent or mitigate ongoing material harm, the Academy and AFT will provide the AI Provider with at least five business days to review and comment on any proposed public statement identifying the AI Provider. Revocation does not

relieve the AI Provider of any Standard obligation incorporated into an agreement with an EDU Customer or any obligation that expressly survives termination.

- Nothing in the notice, consultation, or cure provisions of this Section limits the right of the Academy, AFT, or an EDU Customer, as applicable, to seek immediate injunctive or other equitable relief where reasonably necessary to prevent or mitigate ongoing material harm, unlawful use or disclosure of Covered Data, prohibited AI training, or a material threat to student safety or security.
- The AI Provider will defend the Academy and AFT against any third-party claim brought by an EDU Customer or its end user arising directly from the use of an AI Provider Educational Product and will pay any damages awarded in a final judgment, and any settlement approved by the AI Provider, resulting from that claim. These obligations apply only if the Academy or AFT: (a) promptly notifies the AI Provider of the claim in writing; (b) gives the AI Provider control of the defense and settlement; and (c) provides the assistance, information, and authority reasonably requested by the AI Provider, and the AI Provider shall reimburse the Academy or AFT for reasonable out-of-pocket expenses incurred in providing such assistance.
- Nothing in this Section limits the right of the Academy, AFT, or an EDU Customer, as applicable, to seek actual damages, injunctive relief, or other remedies available under the applicable agreement or law. If any remedy or enforcement provision in this Section is determined to be unenforceable, the remaining provisions of this Agreement will remain in full force and effect.

Public Trust Page:

The AI Provider must maintain a publicly accessible trust page listing all current certifications, recent audit summaries, and any known past security incidents affecting education customers, updated at least quarterly. In addition, if an AI Provider is terminated from the Standards pursuant to this Section, it shall note that it has been terminated on the trust page, and include the public statements provided by the Academy and/or AFT.

PRINCIPLE 6 | Security That Matches the Stakes

Student data is among the most sensitive data that exists. Children cannot "reset" their identities after a breach. Security must reflect that reality.

Minimum required security practices:

- Covered Data is protected at rest and in transit using controls appropriate to the applicable security use case and technical architecture, including current industry-standard cryptographic protocols where encryption is applicable.
- Role-based access controls ensuring that only authorized personnel can access Covered Data, with the principle of least privilege applied throughout.

- Multi-factor authentication is required for all AI Provider personnel with access to any system containing Covered Data.
- Annual penetration testing by an independent third party, with results made available to the Academy upon request, subject to confidentiality restrictions.
- A documented, tested incident response plan that the Academy can review upon request, subject to nondisclosure terms.
- Vulnerability disclosure program that allows external researchers to report security issues, with a commitment to investigate and remediate within defined timelines.

Security Roadmap:

The AI Provider must share with the Academy a security maturity assessment and forward-looking security roadmap showing planned improvements, subject to reasonable confidentiality and nondisclosure restrictions. If disclosure of the roadmap would create a material security risk, the AI Provider may instead provide an equivalent independent assurance artifact sufficient to demonstrate its current security maturity and planned security improvements.

Subprocessor Security:

Any third-party subprocessor used by the AI Provider must meet the same security standards required of the AI Provider under this agreement. The AI Provider is responsible for ensuring subprocessor compliance and for any breach caused by a subprocessor.

PRINCIPLE 7 | Meaningful Consent and Transparency for Families

Families deserve to know when and how AI is being used in their children's education. Plain language, not legal language. Accessible to all, not just the tech-savvy.

Required transparency materials:

- Within ninety (90) days following the Effective Date, the AI Provider shall develop and make available educational transparency materials regarding its AI Provider Educational Products, including: (a) a plain-language guide describing the material features and capabilities of such products, including, as applicable, what such features and capabilities do, whether they are enabled or disabled by default, and any relevant administrative controls; (b) a frequently asked questions (FAQ) resource addressing common topics relating to AI product functionality, data privacy, safety, security, and responsible use; and (c) such other explanatory materials as the AI Provider reasonably determines are appropriate to support educator and family understanding of AI Provider Educational Products.
- During such ninety (90)-day period, the AI Provider will meet with the Academy to review the proposed materials and obtain the Academy's feedback regarding the level of detail, format, and accessibility of the materials for the Academy's audience. The AI Provider will consider such feedback in good faith in preparing the final materials.
- Following completion of the materials, the AI Provider will make them broadly available to its United States-based education customers, including through direct customer communications and publication on a publicly available website maintained by the AI

Provider. The AI Provider will also offer informational webinars relating to the materials and will work collaboratively with the Academy to review and update the materials on an annual basis.

- The parties acknowledge that the materials described in this Section are intended solely to provide general educational and informational content regarding AI Provider Educational Products. Such materials are not product warranties, service commitments, or representations made to any EDU Customer or other third party, and no third party shall have any right to rely upon or enforce this Section. Any claim by the Academy that the AI Provider has failed to satisfy its obligations under this Section shall be addressed exclusively through the notice, discussion, and cure procedures set forth in Principle 5 (Contractual Remedies for the Academy), and the remedies provided therein shall constitute the Academy's sole and exclusive remedies with respect to this Section.

Family Inquiry Process:

- Where the AI Provider is a data controller, it must provide a clear, accessible process for families to ask questions about how AI is used, request information about data collected about their child, and raise concerns. As part of their EDU Customer agreements, school districts will engage their relevant stakeholders to develop and communicate a process through which families may submit such requests directly to such AI Provider. Where the AI Provider acts solely as a data processor, it will reasonably assist the EDU Customer in using available functionality to respond to such requests (such as administrator portals or data subject request forms).

PRINCIPLE 8 | Equity, Accessibility, and Inclusion

AI tools must work fairly for every student, regardless of race, language, disability, or economic background. Equity is not optional.

Because equity and accessibility are validated through evidence, the requirements in this principle function both as contractual disclosure duties and as ongoing evaluation conditions for deployment. The Academy may withhold approval or suspend use until the required evidence is provided and any material gaps are addressed.

Required commitments:

- The AI Provider must build accessibility into the core product, not as an afterthought. Required accessibility features include, screen reader compatibility, keyboard navigation, captioning for audio/video content, support for text-to-speech and dictation, and translation support for English Language Learners.
- The AI Provider must test its tools across diverse student populations and must not deploy tools that perform materially worse for students of color, students with disabilities, or English Language Learners without disclosing known limitations and presenting a remediation plan.
- The AI Provider Educational Products shall offer the same features to all students equitably, without the ability for one student to purchase different features within the

same subscription. If the EDU Customer has paid for a service, all features of that service must be available to all educators and students equally.

- The AI Provider must maintain a clear, accessible process for reporting bias, accessibility failures, or harm, including channels accessible to students, educators, and families.
- Within 12 months, the AI Provider will fund an independent, academically peer-reviewed equity and accessibility analysis describing how AI products used in schools perform across relevant student populations. The researchers shall be selected by the Academy or an independent academic institution and shall have complete control over the research questions, methodology, analysis, findings, conclusions, peer review, and publication, with advisory support from AI Provider's data scientists and/or researchers. The AI Provider shall have no role in selecting, directing, supervising, approving, or removing the independent researchers and may not influence, limit, delay, edit, suppress, or prevent the research or its publication. The evaluation may rely on appropriately licensed or public datasets, properly authorized research studies, privacy-protected aggregate findings, or wholly Synthetic Data permitted under Principle 1. Nothing in this requirement obligates or authorizes the AI Provider to collect, infer, retain, or use Covered Data or sensitive student characteristics, or to use Covered Data for training, benchmarking, model improvement, or any other prohibited purpose. For clarity, this requirement is a commitment solely between the AI Provider and the Academy and no EDU Customer or other third party may rely upon or enforce it. Any claim by the Academy or AFT that the AI Provider has failed to satisfy this requirement shall be addressed exclusively through the notice, discussion, cure, and escalation procedures set forth in Principle 5 (Contractual Remedies for the Academy), and the remedies provided there shall be the Academy's and AFT's sole and exclusive recourse with respect to this requirement.

Transparency Documentation:

The AI Provider will maintain publicly available documentation, regarding its responsible AI governance and transparency practices for the AI Provider Educational Products. Upon written request from the Academy, AI Provider will provide any current transparency documentation, including responsible AI reports relevant to the AI Provider Educational Products. This documentation must be written in plain language, include aggregated performance data, and shared with the Academy within 30 days of their request.

PRINCIPLE 9 | No Feature Creep and No Vendor Lock-In

Educators, not vendors, decide how AI Provider Educational Products evolve in classrooms. The EDU Customer must never become trapped in a vendor relationship, where it cannot exit.

Feature change controls:

- No new features that involve new data collection practices, new data processing activities, new autonomous or agentic actions, or material changes to how the AI operates may be activated in the EDU Customer's environment without notice and ability to manage settings. Such features may be enabled or disabled by the EDU Customer's administrator(s).

The AI Provider will provide notices concerning changes to the AI Provider Educational Products, applicable terms, and data practices in accordance with the EDU Customer's licensing agreement, product terms, and data protection agreement, as applicable. Such notices may be provided electronically, including by email, through an applicable administrative portal, or through an identified website. The data protection terms applicable to an EDU Customer's subscription will remain unchanged during the subscription term, subject to the exceptions and conditions provided in the applicable data protection agreement.

Data portability and exit rights:

- Covered Data must be exportable at any time by the EDU Customer, at no cost, in standard, non-proprietary, reasonably usable formats.
- Upon termination of this Standard for any reason, the AI Provider must provide the EDU Customer with the ability to export all exportable Covered Data for at least 30 days, at no charge, before deleting such data.
- The AI Provider may not impose technical, contractual, or other barriers that materially impede or increase the difficulty or cost of transitioning Covered Data or services to another provider.
- The AI Provider's breach and security-incident notification obligations survive termination for as long as the AI Provider or any subprocessor retains, possesses, or can access the EDU Customer's Covered Data.
- Any prohibition on the sale of the EDU Customer's Covered Data or its use for training, developing, or improving AI or machine-learning models survives termination indefinitely.

PRINCIPLE 10 | Long-Term Data Responsibility

The AI Provider's obligations to protect student data do not end when the contract ends. Some commitments are permanent.

Data retention limits:

- Covered Data may not be retained longer than is necessary to deliver the service, subject to a maximum retention period as described in Principle 3. Session data, including prompts and outputs, will be retained and deleted in accordance with the retention setting selected by the EDU Customer.

Permanent obligations that survive contract termination:

- The prohibition on using Covered Data for training (Principle 1), permanent.
- The prohibition on selling or sharing Covered Data (Principle 3), permanent.
- The obligation to delete Covered Data at the direction of the EDU Customer survives for 2 years post-termination or until the AI Provider no longer holds any Covered Data, whichever occurs later.
- The obligation to report any post-termination breach survives for 2 years post-termination or until the AI Provider no longer holds any Covered Data, whichever occurs first.

Parties to This Memorandum of Agreement

THE ACADEMY	Legal name: National Academy for AI Instruction
AI PROVIDER	Legal name: Microsoft Corporation
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____
AI PROVIDER	Legal name: _____

PART II — GENERAL TERMS OF AGREEMENT

1. Applicable Law

This agreement is governed by the laws of the state in which the Academy is located. The AI Provider agrees to comply with all applicable federal, state, and local privacy laws in its provision of AI Provider Educational Products.

1A. Relationship to Existing DPAs and Online Terms

EDU Customers purchasing AI Provider Educational Products may opt to include these principles in their specific agreements with AI Providers. Nothing in this Standard reduces, limits, or supersedes any greater protection provided by applicable law or an EDU Customer's

agreement with an AI Provider. The EDU Customer's agreement should specify which document controls in the event of a conflict; however, participation in or compliance with the Standard may not be used to reduce, waive, or otherwise diminish any protection or right that would otherwise apply to the EDU Customer, students, or Covered Data.

2. Sub processors

The AI Provider must make publicly available and maintain a current list of all subprocessors that have access to or process Covered Data. The AI Provider must provide EDU Customers with a mechanism to obtain reasonable advance notice of any new subprocessor that will process Covered Data. EDU Customers must have a reasonable opportunity to object based on legitimate privacy, security, or compliance concerns, and the AI Provider must work in good faith to address or remediate those concerns. If the EDU Customer ultimately does not approve of a new subprocessor, then the EDU Customer may terminate any subscription for the affected product without penalty or termination fee by providing, before the end of the relevant notice period, written notice of termination. The AI Provider remains fully responsible and liable for its subprocessors' compliance with all applicable obligations relating to Covered Data.

3. Term and Termination

This agreement is effective from the date signed and continues for the term specified in the applicable service agreement. AI Provider may terminate its participation in this agreement with 60 days' written notice.

4. Amendments

Any amendment to this agreement must be in writing and signed by authorized representatives of all parties.

5. Severability

If any provision of this agreement is found to be unenforceable, the remaining provisions continue in full force. The parties agree to replace any invalid provision with a valid provision that achieves, as closely as possible, the same purpose.

6. Notices

All notices under this agreement must be in writing and delivered by email with delivery confirmation or by certified mail to the designated contact persons listed in Addendum A.

Signatures

By signing below, each party agrees to all principles in this agreement and confirms they have the authority to bind their organization.

FOR THE ACADEMY

Randi Weingarten

Randi Weingarten (Sep 7, 2026 15:58:36 EDT)

Signature

Randi Weingarten, President, American Federation of Teachers

Printed Name and Title

09/07/2026

Date

FOR THE AI PROVIDER

Brad Smith (CELA) (Sep 7, 2026 07:41:17 PDT)

Signature

Brad Smith, Vice Chair and President, Microsoft Corporation

Printed Name and Title

09/07/2026

Date

FOR THE AI PROVIDER

Signature

Printed Name and Title

Date

FOR THE AI PROVIDER

Signature

Printed Name and Title

Date

FOR THE AI PROVIDER

Signature

Printed Name and Title

Date

ADDENDUM A — Academy-Specific Configuration

This addendum is completed separately for each AI Provider engagement and becomes part of this agreement. This Addendum is to be completed by both parties within 90 days of Effective Date. Addendum A applies solely to the engagement between the Academy and the applicable AI Provider and will be completed and signed only by those parties.

A1. Contact Persons

Academy Privacy Officer	Name: _____Email: _____ Phone: _____
AI Provider Privacy Contact	Name: _____ Email: _____ Phone: _____
Breach Notification Contact (Academy)*	Name: _____ Email: _____ Phone: _____
Breach Notification Contact (Provider)	Name: _____ Email: _____ Phone: _____

*Note: Academy must include Breach Notification contact in its administrator portal as well as above.

A2. Service Description

AI Provider Educational Product Name & Version	_____
Intended Users	Educators only Students (grades: _____) Staff _____
Use Cases Authorized	_____
Use Cases Prohibited	Automated grading without review | Disciplinary decisions | Companion-style student interactions | Autonomous action-taking without approval | Others: _____

Data Retention Setting	Session only (ZDR) Maximum _____days Custom (see attached schedule)

A2A. Approved Operational Data Elements (if any)

List any contextual or operational data elements the AI Provider may process because they are strictly necessary for the authorized use case (for example: school, district, class, course, assigned teacher, device identifier, or coarse location/state information):

In providing Online Services to our customers (here, the Academy), Microsoft acts as a processor. As a processor, Microsoft processes personal data only on behalf of and in accordance with the customer's documented instructions. These documented instructions include the terms of the applicable Data Protection Addendum (DPA), Product terms, Service Specific terms as well as the customer's configuration, administration, and use of the services within its environment to support its specific business needs.

The data processed by Microsoft therefore depends on the services licensed and deployed by the customer and how those services are used and configured. For a particular customer, Microsoft's records as processor typically include:

1. **The purposes of processing** described in the applicable DPA and related contractual documentation.
2. **The Online Services licensed and deployed by the customer**, such as services within the Azure, Microsoft 365, or Dynamics 365 families. Information about the functionality and processing activities associated with these services is publicly available in Microsoft Learn documentation:
    - Azure: Azure documentation | Microsoft Learn
    - Microsoft 365: Microsoft 365 documentation | Microsoft Learn
    - Dynamics 365: Dynamics 365 documentation - Dynamics 365 | Microsoft Learn
3. **The processing activities and service outcomes associated with those services**, as documented in product documentation and reflected in the product experience. These activities vary depending on the functionality used by the customer.
4. **The customer's actual use and configuration of the services**, including actions taken by users and administrators that instruct Microsoft to perform processing operations.
5. **Audit logs** that reflect processing activity carried out on the customer's behalf. These records provide a factual account of compute and service operations initiated through customer use of the Online Services. For example, where a user stores a document in OneDrive, Microsoft's processing includes receiving, storing, and making that document available through the service. Service logs and related records may provide evidence of such processing activities and often offer a level of detail that goes beyond high-level "categories" of processing by creating an audit trail of specific service functions performed.

Accordingly, the data Microsoft processes as a processor is determined by the customer's chosen services, configurations, and usage of those services. We can assist in identifying the relevant services and processing activities for a particular customer deployment.

A3. ZDR Acknowledgment (if applicable)

If the AI Provider operates under Zero Data Retention, check all that apply and sign below:

- The AI Provider does NOT retain prompts or outputs beyond the session
- The Academy acknowledges it is solely responsible for maintaining its own records of AI interactions
- The Academy has implemented its own secure logging system to preserve records where required
- Both parties have reviewed and agreed to the data ownership implications of ZDR

Academy Representative Initials: _____ AI Provider Representative Initials: _____ Date: _____

A4. Approved Sub Processors

Sub Processor Name	Purpose	Data Location
See below		

List of AI Provider Subprocessors available at:
https://servicetrust.microsoft.com/DocumentPage/dbd39f46-4f02-4653-8f66-2d7a3b75937c

The Academy may subscribe to receive updates to the Subprocessor List

A5. Session Data

The Academy may set retention periods for session data (prompts and outputs). Academy intends for session data to be deleted at the end of each user session unless the Academy has expressly selected a longer retention period for a specific educational purpose.

ADDENDUM B — Technical Requirements and Standards

This addendum sets out the technical implementation requirements that AI Providers must meet. These are not aspirational — they are minimum requirements for eligibility.

B1. Encryption Standards

The following are the minimum encryption requirements. Providers may exceed these requirements but not fall below them.

Requirement	Minimum Standard
Data at rest	AES-256 encryption or equivalent
Data in transit	TLS 1.2+ or higher
Key management	Keys must be rotated at least annually; keys must not be stored alongside the data they protect
Backup data	Same encryption standards as primary data; backups must be inventoried and included in deletion processes
API communications	All API calls must use HTTPS with certificate pinning recommended

B2. Access Controls

- Role-based access control (RBAC) must be implemented such that AI Provider personnel can access Covered Data only to the extent their specific role requires
- Access logs must be maintained for all access to Covered Data by AI Provider personnel.
- Privileged access (e.g., database administrator access) must require additional authentication factors beyond standard MFA
- Departing employees must have their access revoked within 24 hours of termination
- All access must be logged with timestamps, user identifiers, and actions taken

B3. Identity and Authentication

- Multi-factor authentication (MFA) is required for all AI Provider personnel with any access to production systems
- Students and educators accessing the AI Provider Educational Product must authenticate through the Academy's existing identity provider (SSO via SAML 2.0, OAuth 2.0, or equivalent) wherever technically feasible
- Session tokens must expire after a period of inactivity.
- Password policies must be able to meet or exceed NIST SP 800-63B Digital Identity Guidelines when configured by the Academy.

B4. Audit Logging

- All access to Covered Data must be logged with sufficient detail to reconstruct a complete audit trail
- Logs must be retained for a minimum of 180 days.
- Log integrity must be protected (logs must not be modifiable by standard users)
- The Academy must be able to query audit logs through an administrative interface or upon request

B5. Penetration Testing

- Annual external penetration testing is required at minimum; semi-annual testing is strongly preferred for systems serving under-13 students
- Testing must cover both the application layer and the underlying infrastructure
- Critical and high-severity findings must be remediated within 30 days of identification
- Medium-severity findings must be remediated within 90 days
- A summary of findings and remediation status must be shared with the Academy within 30 days of each test

B6. Required Security and Compliance Certifications and Assurances — Verification Checklist

At the time of signing and annually thereafter, the AI Provider must provide documentation confirming the status of each applicable certification, assurance, or other item listed below.

Certification, Authorization, or Assurance	Status	Note
SOC 2 Type II (annual)	**Current** Pending N/A	Microsoft 365: Renewed 2025-09-30 (bridge letter through August 2026)
ISO 27001	**Current** Pending N/A	Expiration: Microsoft 365: 2027-07-28
ISO 27701 (Privacy)	**Current** Pending N/A	Expiration: Microsoft 365: 2027-07-28
ISO 42001 (AI Management)	Current **Pending** N/A	Target receiving by December 31, 2027
CSA STAR Level 1 or 2	**Current** Pending N/A	Microsoft 365: Created or renewed on October 14, 2025 (no expiration) (see: STAR Registry Listing for Microsoft 365 | CSA).
FedRAMP Moderate (U.S. districts)	Current Pending **N/A**	Not available for Microsoft Education Products.
COPPA Assurance	**Current** Pending N/A	Microsoft's education services support customers' compliance with COPPA. Where Microsoft processes children's personal information on behalf of an educational institution, Microsoft acts as a school-authorized service provider, and the educational institution is responsible for providing any required notice and obtaining any required parental consent. See the applicable Microsoft Product Terms and Data Protection Addendum.
FERPA Compliance Assurance	**Current** Pending N/A	Compliance acknowledged in DPA: Licensing Documents
IDEA Compliance Attestation	**Current** Pending N/A	Microsoft 365 offers various accessibility tools to ensure everyone can use its features effectively. Microsoft 365 supports IDEA Compliance by its school customers. Please see conformance reports: Accessibility Conformance Reports | Microsoft Accessibility and Microsoft 365 accessibility features: Create an inclusive

		classroom with Microsoft accessibility tools - M365 Education | Microsoft Learn
ADA / Section 508 Accessibility Audit	**Current** Pending N/A	Please see conformance reports: Accessibility Conformance Reports | Microsoft Accessibility

ADDENDUM C — Notes on Current AI Provider Offerings

This addendum documents special considerations for each AI Provider who signs this agreement.

Microsoft Corporation:

- Microsoft has not completed an independent assessment under ISO 42001 (AI Management Systems) for its AI Provider Educational Products. Microsoft is targeting completion of the certification by December 31, 2027.
- Principle 2's Prohibited practices do not apply to Speaker Coach or Speaker Progress, products designed to help students with public speaking. During a user-initiated speaking or presentation activity, these features process only the data needed to provide coaching, feedback, or progress insights - not for monitoring, profiling, or other secondary purposes.

ADDENDUM D — Academy Implementation Checklist

Before any pilot with live Covered Data, production deployment, or activation of a materially new feature, the Academy's designated Privacy Officer must confirm that each item below has been completed. This checklist functions as the agreement's assessment-to-adoption workflow and becomes part of the agreement record.

Pre-Deployment (Complete Before Any Rollout)

#	Action Required	Done / Date
1	Sign the agreement with the AI Provider and complete all Addenda.	_____
2	List the approved product, intended users, authorized uses, prohibited uses, and retention setting in Addendum A.	_____
3	Verify that required certifications are current. If one is still being pursued, record the target completion date in Addendum C.	_____
4	Review the Provider's data map. Confirm that it collects only what is needed and list any approved operational data in Addendum A.	_____
5	Confirm that Covered Data will be stored and processed only in North America or the European Union, and record how the Academy can identify the region used.	_____
6	Confirm the Provider's current subprocessors, their purposes, and the process for notice and objections when a new subprocessor is added.	_____
7	Set retention and deletion rules, including session-data settings, deletion from active systems, backup purge cycles, and any legal-hold or safety exceptions.	_____
8	If using Zero Data Retention, complete the acknowledgment in Addendum A and decide how the Academy will keep any records it needs.	_____
9	Turn off memory and personalization unless they are covered by approved retention and deletion controls. If enabled, confirm users can see, erase, and turn off memory.	_____
10	Determine what consent is required by law, document it, and configure access for students who do not have required consent. Confirm the Provider can record consent status for students under 13 when needed.	_____
11	Configure the product by grade, age, and use case, including more restrictive settings for students under 13 where appropriate.	_____
12	Review evidence of encryption, access controls, MFA, audit logging, penetration testing, incident response, vulnerability reporting, and the security roadmap or equivalent assurance.	_____
13	Test the Provider's administrative tools or request process for data categories, retention, storage location, subprocessors, access controls, logs, deletion, and export.	_____

14	Confirm that the Provider completed its privacy review for the deployment and any material new feature, and obtain relevant documentation upon request.	_____
15	Review the Provider's plain-language explanation of how the product works, its training-data summary, known limits, failure modes, and current model or system documentation.	_____
16	Review content-safety controls and crisis-response protocols for self-harm, suicide, sexual content, abuse, violence, and other high-risk topics.	_____
17	Confirm that companion-style or relationship-oriented features are not available to Academy users.	_____
18	Confirm that student action-taking features can be disabled. Keep externally consequential actions off by default unless an authorized administrator approves scoped permissions, human oversight, and audit logging.	_____
19	Review and approve the AI Use Statement, family FAQ, age-appropriate AI notices, uncertainty notices, and the process for family questions or concerns.	_____
20	Complete the initial feature-and-control review with the Provider, document who can turn each control on or off, and create the plain-language guide. Confirm the process for required 30-day notices.	_____
21	Review accessibility and fairness evidence, equal access to paid features, known gaps and remediation plans, and the process for reporting bias, accessibility failures, or harm.	_____
22	Train educators on approved uses, prohibited high-risk uses, privacy and safety expectations, human review, product limits, and how to report concerns.	_____
23	Give families the required notices and materials before launch and complete any consent process required by law.	_____
24	Test data export and confirm that the Academy can leave the service without technical or contractual barriers	_____
25	Document breach contacts and procedures and establish the Academy's process for immediate triage and coordination during the first 72 hours.	_____

Annual Review (Complete Each Year)

- Receive and review AI Provider's annual training non-use confirmation statement
- Verify that all certifications remain current and obtain updated documentation
- Review and update the data map if any practices have changed
- Review of AI Provider's bias testing results and equity report
- Review any changes to content-safety controls, crisis-response protocols, companion-style features, agentic capabilities, and student-facing uncertainty notices before continuing or expanding deployment

- Confirm deletion of any data that has reached the end of its agreed retention period
- Re-train educators on any new features or changed practices
- Update the Academy's AI Use Statement and family FAQ if anything has changed
- Conduct an Academy-side privacy risk review of the deployment
- Confirm all sub processors remain approved and their certifications are current