



Personalized AI-Guided Education
Combined Policy & Compliance Documentation

Proficiency, Inc.

This document consolidates the following six PAIGE governance documents into a single reference:

- AI Ethics and Data Use Policy — Last Updated 2026-06-15
- Terms of Service — Effective June 22, 2026
- Privacy Notice — Effective June 22, 2026
- Subprocessor Disclosure — Last Updated 2026-06-15
- FERPA Statement — Last Updated 2026-06-15
- Accessibility Statement — Last Updated 2026-06-15

Table of Contents

AI Ethics and Data Use Policy	4
1. Purpose and Scope	4
2. Our Principles	4
3. How Paige Uses AI	5
3.1 Overview	5
3.2 What Data Is Sent to AI Models	5
3.3 Course-Level Data Isolation	6
3.4 AI Grounding and Guardrails	6
3.5 Confidence Handling and Failure Modes	6
3.6 Bias and Output Quality	7
4. AI Providers and Third-Party Data Processing	7
4.1 Provider Categories	8
4.2 Data Processing Commitments	8
4.3 Model Versioning and Stability	9
5. What Paige Is Not	9
5.1 Paige Is Not an Academic Integrity Enforcement Tool	9
5.2 Paige Is Not a Replacement for Instructors	10
5.3 Institutional Responsibility	10
6. Pedagogical Alignment	10
6.1 AI Assessment Scale (AIAS)	10
6.2 Institutional Policy Tiers	11
6.3 Alignment with Pedagogical Best Practices	11
7. Policy Review	11
8. Contact	12
PAIGE Terms of Service	13
1. Who May Use PAIGE	13
2. What PAIGE Is	13
3. Acceptable Use	13
4. Academic Integrity	14
5. Your Content and Data	14
6. AI-Generated Output	14
7. Service Availability and Changes	14
8. Disclaimers	14
9. Limitation of Liability	15
10. Termination	15
11. Governing Law	15
12. Changes to These Terms	15
13. Contact	15
PAIGE Privacy Notice	16
1. Data We Collect	16
2. How We Use Your Data	16
3. Who Can Access Your Data	17

4. Data Retention Policy.....	17
5. Your FERPA Rights.....	17
6. Third-Party Services.....	17
7. Data Security.....	18
8. Contact.....	19
Subprocessor Disclosure.....	20
1. Purpose.....	20
2. Data-Processing Posture.....	20
2.1 Training opt-out.....	20
2.2 Retention for abuse detection, and Zero Data Retention.....	20
2.3 Data Processing Agreements (DPAs).....	20
3. Subprocessors.....	21
3.1 OpenRouter — Language Model Routing.....	21
3.2 OpenAI — Embeddings and Voice.....	21
3.3 Cohere — Search Reranking (Optional).....	22
3.4 Langfuse — LLM Observability.....	22
3.5 Amazon Web Services (AWS) — Infrastructure and Configuration.....	23
4. Summary Table.....	23
5. Changes to This List.....	24
6. Contact.....	24
FERPA Statement.....	25
1. Purpose and Scope.....	25
2. Paige's Role Under FERPA.....	25
3. How Paige Protects Education Records.....	25
3.1 Access Is Limited by Role and Ownership.....	26
3.2 Consent Precedes Use.....	26
3.3 Students Can Access Their Data; Records Are Retained and Deleted Responsibly.....	26
3.4 The Audit Trail Cannot Be Altered.....	26
3.5 Education Records Are Encrypted.....	26
3.6 AI Providers Do Not Train on Student Data.....	27
4. Student FERPA Rights.....	27
5. Data Disclosure and Safety Escalation.....	27
6. Contact.....	27
Accessibility Statement.....	28
Our Commitment.....	28
Conformance Status.....	28
Feedback.....	28

AI Ethics and Data Use Policy

Proficiency, Inc. Product: Paige — AI Learning Assistant

Last Updated: 2026-06-15

1. Purpose and Scope

This document establishes Proficiency's principles, commitments, and practices governing the ethical use of artificial intelligence within the Paige platform.

Paige is an AI-powered learning assistant designed for use in higher education, primarily serving two-year and four-year colleges and universities. Paige integrates with course-specific textbook content to provide students with grounded, trustworthy academic support through multiple interaction modes.

This policy is intended for institutional administrators, faculty, department chairs, and other stakeholders evaluating or overseeing Paige's use within their institution. It describes how Paige's AI operates, what commitments Proficiency makes regarding transparency and responsible AI use, and how the tool fits into existing pedagogical frameworks.

For detailed information about data collection, retention, regulatory compliance, and student data rights, see Proficiency's Privacy & Compliance Overview. For the specific list of AI providers in the processing chain, see the Subprocessor Disclosure.

2. Our Principles

Proficiency's approach to AI in education is guided by the following principles. These are not aspirational statements — they are design constraints that shape how Paige is built and operated.

Student learning comes first. Paige exists to help students develop understanding, confidence, and effective learning skills. Every product decision is evaluated against this mission. Features that compromise learning outcomes in favor of engagement metrics or commercial interests are rejected.

AI supports human educators — it does not replace them. Paige provides academic support grounded in course materials. Educational decisions, grading, assessment, and academic policy remain the responsibility of faculty and institutions. Paige does not make autonomous decisions about students.

Transparency over reassurance. Proficiency discloses how Paige's AI works, what data it processes, and which third parties are involved. Where controls are still in progress, this document says so. Institutions deserve accurate information, not marketing language.

Data minimization. Paige collects and processes only the data necessary to deliver its educational function. Student data is never sold, never used for advertising, and never used to build commercial profiles.

Human oversight is preserved. Institutions retain the ability to review, question, and govern how Paige is used within their courses. Proficiency provides the information and access controls necessary for institutions to fulfill their oversight responsibilities.

3. How Paige Uses AI

3.1 Overview

Paige provides three modes of AI-powered academic support. Each mode is grounded in course-specific content — Paige does not generate responses from the open web or from general-purpose training data.

Mode	Function	AI Role
RAG (Retrieval-Augmented Generation)	Students ask questions about their course material. Paige retrieves relevant passages from the course textbook and generates a response grounded in that content.	Retrieval + generation from courseware
Socratic	Paige engages students in guided questioning to test and deepen understanding, rather than providing direct answers.	Guided dialogue grounded in courseware
Study Tools	Paige generates flashcards, practice questions, and study guides drawn from the student's course content.	Content generation from structured course data

3.2 What Data Is Sent to AI Models

Paige uses third-party AI models for response generation. Proficiencycee does not own, operate, or host these models. When a student interacts with Paige, the following data is sent to a third-party AI provider:

- The student's question or prompt
- Retrieved textbook passages relevant to the question (RAG and Socratic modes)
- Recent conversation history within the current chat session (for conversational continuity)
- A system prompt defining Paige's behavior, tone, and constraints

When using voice features, Paige additionally sends:

- Audio recordings of student voice input to a third-party speech-to-text provider for transcription
- Generated text responses to a third-party text-to-speech provider for audio playback

Audio data is processed in real time and is not stored by Proficiencycee beyond the immediate request. Third-party provider retention and processing terms are governed by each provider's published data-handling terms and by Proficiencycee's account configuration with that provider (see Section 4.2).

Paige does not programmatically include any student identity data (name, email, student ID, enrollment status) in what it sends to AI models. System prompts, retrieval queries, and structured data payloads are constructed without student identity fields.

Because conversational AI requires sending the student's message to generate a response, Proficiency cannot filter personal information a student voluntarily includes in their own messages. This is why Proficiency's training opt-out at every AI provider (Section 4.2) matters: by configuring each provider account to exclude Paige's data from model training, Proficiency ensures providers do not train on any data received through the platform, regardless of what that data contains.

The following data is not sent to AI models by Paige's architecture:

- Student identity fields (name, email, student ID)
- Grades, enrollment status, or LMS data
- Data from other students or other course sections
- Historical data beyond the current chat session's conversation context

3.3 Course-Level Data Isolation

Each student's AI interactions are scoped to their enrolled course. Paige enforces course-level isolation through metadata filtering on every retrieval query. A student enrolled in Introduction to Nutrition will only receive responses grounded in that course's textbook content — never content from another course or another institution.

This isolation is enforced at multiple layers — database, service, and retriever — and cannot be bypassed by user input. The student's course association is determined server-side from their authenticated account, not from any client-provided parameter.

Students within the same course share access to that course's content. Finer-grained isolation at the section level (e.g., distinguishing between sections with different textbook editions) is on Proficiency's development roadmap and will be implemented as multi-section deployment requires it.

3.4 AI Grounding and Guardrails

Paige's AI responses are grounded in course-specific content through its Retrieval-Augmented Generation architecture. This means:

- Responses are anchored to textbook passages and course materials approved by the institution or publisher
- Paige does not draw from the open web, general knowledge, or unverified sources
- Paige is instructed via system prompts to stay within the scope of the provided course content
- Temperature and other model parameters are configured to prioritize adherence to source material over creative generation

Additionally, Paige includes a moderation layer that filters inappropriate content in student messages. This layer is designed for student safety, not for academic integrity enforcement (see Section 7).

3.5 Confidence Handling and Failure Modes

Paige does not attempt to answer every question with equal confidence. When a student's question falls outside the scope of their course content or when retrieval confidence is low, Paige adjusts its behavior rather than guessing.

RAG and Socratic modes use document relevancy scores to assess confidence across three zones:

- **High confidence:** Paige answers directly, grounded in retrieved textbook content.
- **Moderate confidence:** Paige may still respond, but is prompted more strongly to stay within what the textbook actually says. It may ask clarifying questions or guide the student toward related content that is well-covered in the courseware.
- **Low confidence:** Paige declines to answer and communicates that the question appears to fall outside the scope of the course material. Socratic mode allows slightly more flexibility in this zone to support exploratory dialogue, but still refuses when content relevance is insufficient.

Study Tools uses a separate confidence mechanism. When the AI extracts parameters from a student's request (e.g., which chapter, which topic), it assigns a confidence score to those parameters. If confidence is low, Paige asks the student to clarify rather than generating materials from an uncertain interpretation.

This approach reflects Proficiency's position that an AI tool that says "I don't know" when appropriate is more trustworthy than one that always produces an answer.

3.6 Bias and Output Quality

Because Paige uses Retrieval-Augmented Generation, the content of its responses is drawn primarily from the course textbook rather than from the AI model's general training data. This significantly limits the surface area for model-level bias, since the source material is curated and approved by the institution or publisher.

However, the generation layer can still influence how textbook content is presented — for example, in emphasis, framing, or the selection of which retrieved passages to highlight. Proficiency addresses this through several measures:

- **Quality metrics:** Paige is evaluated against automated metrics for faithfulness (does the response accurately reflect the source material?) and answer relevancy (does the response address what the student asked?).
- **Model selection:** Different tasks use different models selected for their suitability — smaller, more constrained models for structured tasks like parameter extraction, and larger models for generation where nuance is required.
- **Query reformulation and reranking:** Before generating a response, Paige reformulates the student's question for better retrieval and re-ranks retrieved passages for relevance, reducing the likelihood of irrelevant or misleading content reaching the generation step.
- **Content moderation:** A dedicated moderation model screens both student input and AI output for safety and security concerns.

Proficiency does not claim that Paige's outputs are free of bias. Where bias can enter — through the textbook content itself, through model generation tendencies, or through the framing of

retrieved passages — Proficiency monitors output quality and refines its pipeline as issues are identified.

4. AI Providers and Third-Party Data Processing

Proficiency uses third-party AI services to power Paige's core functionality. This section discloses the categories of providers involved and the data each receives.

4.1 Provider Categories

Category	Function	Data Received
LLM Inference	Generates responses to student queries (RAG, Socratic modes), generates study tools content, performs content moderation, and reformulates queries	Student question, retrieved textbook passages, conversation context, system prompt
Embeddings	Converts textbook content into vector representations for retrieval	Textbook content chunks (no student data during normal operation); student question text (for query embedding at retrieval time)
Reranking	Re-ranks retrieved textbook passages by relevance to the student's question	Student question text, retrieved textbook passages
Voice	Converts student voice input to text (speech-to-text) and generated responses to audio (text-to-speech)	Audio recordings of student voice input; generated response text
LLM Observability	Traces and monitors AI performance for reliability and quality	LLM prompts and responses (operational telemetry)

The specific providers in each category are named in Proficiency's Subprocessor Disclosure, which is available to institutional partners.

4.2 Data Processing Commitments

Proficiency is committed to ensuring that no third-party AI provider uses student data for model training. To enforce this commitment:

- **Data-use-for-training opt-out is in effect at every AI provider that receives student data.** Proficiency has configured each provider account to exclude Paige's inputs and outputs from model training, fine-tuning, and product improvement. This relies on the providers' published no-training-on-API-data positions, which Proficiency enforces through its provider account settings.
- **Providers retain data transiently for abuse detection.** Excluding data from training does not mean the provider holds no copy. As a standard part of their API terms, providers retain inputs and outputs briefly for abuse and misuse detection. Eliminating this requires Zero Data Retention (ZDR) mode, which Proficiency has deliberately not enabled because it

materially restricts model choice and availability — a considered tradeoff. See the [Subprocessor Disclosure, Section 2.2](#).

- **Proficiency publishes a Subprocessor Disclosure** identifying all third-party services that process student data, the purpose each serves, the category of data each receives, its DPA status, and a link to each provider's public data-handling documentation. It is available to institutional partners and is updated when providers change.
- Provider changes that materially alter data processing will be communicated to institutional partners in advance.

Pilot posture on Data Processing Agreements: Proficiency obtains a DPA from each AI provider where one is available and meaningful, and relies on the training opt-out above where a DPA is not available. A DPA is signed with OpenAI; a DPA with Langfuse (observability) has been requested and is pending signature. Cohere and OpenRouter do not offer a separate DPA for this usage. A DPA is not a legal prerequisite for the pilot, but is obtained where it adds protection. Per-provider status is recorded in the [Subprocessor Disclosure](#).

4.3 Model Versioning and Stability

Paige selects AI models based on task suitability, and maintains fallback models for reliability. Proficiency commits to the following:

- **No material changes to AI providers or model behavior mid-semester without institutional notice.** If a change to primary models or providers is required during an active semester (e.g., for reliability or safety reasons), institutional partners will be notified before the change takes effect.
- Routine updates to model parameters, prompt engineering, and quality improvements may occur without notice, provided they do not materially alter the student experience or data processing practices.
- Fallback models may activate automatically for reliability purposes. Any change to the fallback model roster will be communicated to institutional partners.

5. What Paige Is Not

This section addresses the most common concerns faculty raise when evaluating AI-powered educational tools.

5.1 Paige Is Not an Academic Integrity Enforcement Tool

Proficiency takes academic integrity seriously and recognizes that it is a significant concern for faculty adopting AI-powered tools.

Paige is designed to promote learning, not to provide shortcuts. Several design choices reflect this:

- **Socratic mode** guides students through questioning rather than providing direct answers, encouraging active reasoning and self-assessment.
- **Grounded responses** are drawn from course-specific textbook content, not the open web. Students cannot use Paige to retrieve arbitrary information outside their course scope.
- **Content moderation** filters inappropriate messages for student safety.

- **Confidence-based refusal** means Paige declines to answer when content relevance is insufficient, rather than generating speculative responses.

These are pedagogical design choices, not enforcement mechanisms. They are intended to create an environment where learning is the natural path of least resistance.

Proficiency does not guarantee that students cannot use Paige — or any AI tool — to circumvent learning. Overly restrictive controls create poor learning experiences, generate false positives that penalize honest students, and ultimately drive students to less transparent and less safe alternatives where no guardrails exist at all.

Proficiency's position is that the most effective approach to academic integrity in the context of AI is a combination of thoughtful tool design, clear institutional policies, and ongoing dialogue between technology providers and educators. We are committed to:

- Continuing to refine Paige's design to encourage genuine learning
- Listening to faculty feedback on how students interact with the tool
- Evaluating new approaches to supporting academic integrity as the field evolves
- Being transparent about what Paige can and cannot control

5.2 Paige Is Not a Replacement for Instructors

Paige does not grade, assess, or make educational decisions. It does not determine course outcomes, modify curricula, or communicate with students on behalf of faculty. Paige is a study tool — it helps students engage with course material between classes and outside office hours, complementing the instruction faculty provide.

Faculty retain full control over how Paige is used within their courses, including the option to not use it at all.

5.3 Institutional Responsibility

Academic integrity policies, honor codes, and acceptable use definitions for AI tools remain the responsibility of each institution. Proficiency recommends that institutions establish clear AI use policies for their students and communicate expectations for how tools like Paige should be used within their courses.

Proficiency is available to collaborate with institutions on developing these policies and to provide information about Paige's capabilities and limitations to support that process.

6. Pedagogical Alignment

Proficiency designed Paige with existing pedagogical frameworks and institutional AI policies in mind. This section describes how Paige maps to the frameworks faculty are already using to evaluate AI tools.

6.1 AI Assessment Scale (AIAS)

The AI Assessment Scale, adopted by hundreds of institutions worldwide, defines five levels of AI use in academic work. Paige operates primarily at Level 2 (AI Planning) and Level 3 (AI Collaboration), depending on the mode:

Mode	AIAS Level	Rationale
RAG	Level 2–3	Supports research and understanding; does not produce assignment-ready output
Socratic	Level 2	AI helps the student think; the student does the thinking
Study Tools	Level 2	Generates study materials (flashcards, practice questions), not final academic work

Paige does not operate at Level 4 (Full AI) or Level 5 (AI Exploration). It is not designed to produce papers, complete assignments, or generate work that a student would submit as their own.

6.2 Institutional Policy Tiers

Most institutions are adopting tiered AI policies — typically Allowed, Conditionally Allowed, and Prohibited. Paige is designed to fit the Conditionally Allowed tier: a tool with clear boundaries, pedagogical justification, and transparency.

Faculty who permit Paige in their courses may find the following sample syllabus language useful as a starting point:

"Students in this course have access to Paige, a textbook-grounded AI study assistant. Paige may be used for reviewing course concepts, generating practice materials, and self-testing. Paige is not a substitute for completing assignments independently and does not provide assignment answers."

Proficiency can work with faculty to develop course-specific language that reflects their policies and learning objectives.

6.3 Alignment with Pedagogical Best Practices

Paige's design reflects principles from current research on AI in education:

- **Process over product:** Socratic mode emphasizes the reasoning process, not the final answer — aligning with the pedagogical shift toward valuing drafts, reflection, and metacognition.
- **Active recall:** Study tools generate practice questions and flashcards that require students to retrieve information, a technique with strong evidence for long-term retention.
- **Bounded context:** By grounding responses in the course textbook, Paige provides the same bounded learning environment that faculty create when they specify which resources students should use for an assignment.

7. Policy Review

This policy is reviewed and updated:

- At least annually
- Before any material change to AI providers, data processing practices, or the scope of student data collected

- In response to regulatory changes affecting student data privacy in higher education

Changes to this policy will be communicated to institutional partners. The "Last Updated" date at the top of this document reflects the most recent revision.

8. Contact

For questions about this policy, to obtain Proficiency's Subprocessor Disclosure, or to discuss Paige's data practices with your institution:

Email: paige@perceivant.com

Proficiency, Inc.

This document is provided for informational purposes and does not constitute legal advice. Proficiency's specific obligations to each institution are defined in the service agreement executed between the parties. In the event of a conflict between this document and that agreement, the agreement governs.

PAIGE Terms of Service

Effective Date: June 22, 2026

These Terms of Service ("Terms") govern your access to and use of PAIGE (Personalized AI-Guided Education) ("PAIGE," "the Service"), an AI-powered learning assistant provided by Proficiency, Inc. ("Proficiency," "we," "us"). By using PAIGE, you agree to these Terms. If you do not agree, do not use the Service.

PAIGE is provided to you through your educational institution. Your use of PAIGE is also subject to your institution's policies, including its academic integrity and acceptable-use policies.

1. Who May Use PAIGE

PAIGE is intended for use by enrolled students at participating colleges and universities, and by authorized faculty and staff. PAIGE is designed for users aged 18 and older. You must access PAIGE through an account provisioned in connection with your institution.

You are responsible for keeping your account credentials secure and for all activity that occurs under your account.

2. What PAIGE Is

PAIGE is a study tool. It provides academic support grounded in your course materials through three modes:

- **RAG:** Answers questions using passages retrieved from your course content.
- **Socratic:** Guides you through questioning rather than giving direct answers.
- **Study Tools:** Generates study materials such as flashcards, practice questions, and study guides from your course content.

PAIGE is intended to support your learning. It is not a substitute for your own work, for your instructors, or for your institution's resources.

3. Acceptable Use

When using PAIGE, you agree that you will not:

- Use PAIGE to complete graded work in a way that violates your institution's academic integrity policy.
- Attempt to access another user's account, data, or course content.
- Attempt to bypass, disable, or interfere with PAIGE's safety, moderation, or access controls.
- Use PAIGE to generate, store, or transmit content that is unlawful, harassing, or harmful to others.
- Attempt to reverse engineer, scrape, overload, or disrupt the Service or the systems that support it.

- Use PAIGE for any purpose other than the educational purpose for which it is provided.

We may suspend or terminate access that violates these Terms or your institution's policies.

Safety monitoring and escalation. PAIGE automatically screens messages for serious safety concerns and blocks content it identifies as harmful (such as content indicating violence, self-harm, or child sexual abuse material). If we believe a message reflects a real and present threat to you, to others, or to your institution, we reserve the right to share relevant information with your professor or your institution so that someone can respond. This is a best-effort safety measure made using human judgment; it is not an automatic action and not a guarantee that any particular content will be detected.

4. Academic Integrity

PAIGE is designed to encourage genuine learning, not to provide shortcuts, and it is not an academic-integrity enforcement tool. You remain responsible for complying with your institution's and your instructors' policies on the use of AI tools. When in doubt about whether a particular use is permitted in a course, ask your instructor.

5. Your Content and Data

You retain ownership of the questions, messages, and other content you submit to PAIGE ("Your Content"). You grant Proficiencyce a limited license to process Your Content solely to operate and provide the Service.

How PAIGE collects, uses, and protects your data is described in the [Privacy Notice](#). The third-party providers that process data on PAIGE's behalf are listed in Proficiencyce's Subprocessor Disclosure. Your data is not used to train third-party AI models.

6. AI-Generated Output

PAIGE uses artificial intelligence to generate responses. While PAIGE is designed to ground its responses in your course materials and to reduce inaccuracies, it remains an AI system and can still produce output that is incomplete, inaccurate, or misleading, even when it appears confident. You should independently verify important information and should not rely on PAIGE as the sole source for academic, factual, or other consequential decisions. PAIGE's responses do not represent the views of Proficiencyce or your institution.

7. Service Availability and Changes

We may modify, suspend, or discontinue any part of the Service at any time. We may update PAIGE's features, models, and underlying providers to maintain reliability, safety, and quality. Material changes that affect data processing during an active academic term are communicated to institutions in advance, as described in the AI Ethics and Data Use Policy.

8. Disclaimers

The Service is provided "as is" and "as available," without warranties of any kind, whether express or implied, including warranties of merchantability, fitness for a particular purpose, accuracy, and non-infringement. Proficiencyce does not warrant that the Service will be uninterrupted, error-free, or that AI output will be accurate or suitable for any particular purpose.

9. Limitation of Liability

To the maximum extent permitted by law, Proficiencyce will not be liable for any indirect, incidental, special, consequential, or punitive damages, or for any loss of data, arising out of or related to your use of the Service. Nothing in these Terms limits liability that cannot be limited under applicable law.

10. Termination

You may stop using PAIGE at any time. Your account is provisioned and managed through your institution, and the retention and deletion of your account data are described in the [Privacy Notice](#). We may suspend or terminate your access if you violate these Terms or your institution's policies, or if your institution's agreement with Proficiencyce ends.

11. Governing Law

These Terms are governed by the laws of the State of Indiana, without regard to its conflict-of-laws principles. Proficiencyce, Inc. is incorporated in the State of Indiana. Any dispute arising out of or relating to these Terms or your use of the Service will be resolved in the state or federal courts located in Indiana, and you consent to the jurisdiction of those courts.

12. Changes to These Terms

We may update these Terms from time to time. The "Effective Date" above reflects the most recent revision. The current version is always accessible from within the application.

13. Contact

For questions about these Terms:

Email: paige@perceivant.com

Proficiencyce, Inc.

This document is provided for informational purposes and does not constitute legal advice.

PAIGE Privacy Notice

Effective Date: June 22, 2026

PAIGE (Personalized AI-Guided Education) is a learning assistant provided by Proficiency. This notice explains how we collect, use, and protect your information in compliance with the Family Educational Rights and Privacy Act (FERPA) and applicable privacy laws.

1. Data We Collect

When you use PAIGE, we collect the following information:

- **Account information:** Your email address, first and last name, and your association with your institution (through your enrolled course).
- **Conversation data:** Chat messages you send and responses generated by PAIGE, including the learning mode used (RAG, Socratic, Study Tools).
- **Usage analytics:** Session activity, feature usage patterns, message counts, and overall usage volume for system monitoring and improvement.
- **Feedback and survey responses:** Voluntary feedback you submit through in-app surveys or message feedback controls.
- **Audit events:** Timestamped records of account actions (login, consent, profile updates, account deletion) for FERPA compliance.

2. How We Use Your Data

Your data is used solely for educational and operational purposes:

- **Personalized learning:** Providing contextually relevant responses based on your course content and conversation history.
- **System improvement:** Aggregated, anonymized usage patterns help us improve response quality and reliability.
- **Safety and moderation:** Detecting inappropriate content, abuse patterns, and potential safety concerns.
- **Compliance and auditing:** Maintaining FERPA-required audit logs of data access and account activity.

We do not sell your data, use it for advertising, or share it with third parties for non-educational purposes.

Safety escalation. PAIGE automatically screens messages for serious safety concerns (such as content indicating violence, self-harm, or child sexual abuse material). When such content is detected, the message is blocked. In addition, if we believe a message reflects a real and present threat to you, to others, or to your institution, we may share relevant information with your professor or your institution so that someone can respond and help. This is a best-effort safety measure made using human judgment — not an automatic action, a guarantee, or a legal determination.

3. Who Can Access Your Data

Access to your data is strictly limited:

- **You:** You can view and update your profile, and delete individual conversations, at any time. See Data Retention below for how account data is handled when you are no longer enrolled.
- **Professors:** Professors can view aggregate usage statistics and content analytics for their course only. They cannot read your individual chat messages.
- **Administrators:** System administrators may access individual records for technical support, security investigation, or FERPA compliance purposes only.
- **No other parties** have access to your personal data without your explicit consent or a legally required disclosure.

4. Data Retention Policy

- **Active accounts:** Your data is retained for as long as you are enrolled and being served by PAIGE — typically the duration of the academic term — so you can reference earlier conversations as the term progresses.
- **End of enrollment:** When you are no longer being served by PAIGE (for example, at the end of the term), our policy is to soft-delete the associated account and then permanently delete it after a 30-day grace period, which allows recovery from an accidental or erroneous deletion. Account deletion is managed by your institution; PAIGE does not offer student self-service account deletion.
- **Conversation history:** Chat messages are retained to support continuity of learning. You can delete individual conversations at any time from the chat interface, and a deleted conversation is permanently removed. At the end of each academic term, chat data is cleared.

5. Your FERPA Rights

FERPA gives eligible students certain rights over their education records. Those rights are held against your educational institution, which is responsible for honoring them; PAIGE supports your institution in meeting these obligations. Under FERPA, you have:

- **The right to inspect and review** your education records.
- **The right to request amendment** of records you believe are inaccurate or misleading, and to a hearing if your institution declines.
- **The right to consent** to disclosures of personally identifiable information from your records, except where FERPA permits disclosure without consent.
- **The right to file a complaint** with the U.S. Department of Education if you believe your FERPA rights have been violated.

To exercise these rights, contact your institution. For records maintained specifically within PAIGE, your institution's PAIGE administrator can assist.

6. Third-Party Services

PAIGE uses third-party AI service providers to generate responses and operate safely:

- **Language model provider** (OpenRouter): Generates responses to your questions, reformulates queries for better search, moderates content for safety, and produces study materials. Receives your messages, recent conversation context, and relevant passages from your course material.
- **Embeddings and voice provider** (OpenAI): Converts course material and your questions into a searchable form, and — if you choose to use voice features — transcribes your speech and synthesizes spoken responses. Receives course-material text, your question text, and (for voice) your audio. Voice features are optional and only used when you choose to use them; when you do, your audio recording is sent to OpenAI to provide that feature.
- **Search reranking provider** (Cohere, when enabled): Re-orders search results so the most relevant course passages reach the model. Receives your question text and retrieved course passages.
- **Observability provider** (Langfuse): Records traces of model requests and responses so we can monitor reliability and quality. May receive prompts and responses, which can include text you enter.
- **Infrastructure provider** (Amazon Web Services): Hosts the application, database, and configuration. Does not receive your conversations for any AI purpose.

A note on voice. Voice input and output are optional features you turn on yourself — your microphone is never used unless you start a voice recording, and your audio is sent only when you do. Because a voice recording can be considered sensitive personal information, we want to be clear: choosing to use voice means choosing to send that audio to our speech provider (OpenAI) for the sole purpose of transcribing or generating speech. Your audio is not stored by PAIGE, is not shared with anyone beyond that provider, and is not used to train any model.

We have configured each AI provider so that your data is not used to train or improve their models. As a standard part of their service, these providers may keep your messages for a short time to detect abuse or misuse, after which it is discarded — but they do not use it to train their models. These providers process data under their own published privacy and data-handling terms. A fuller account of each provider, the data it receives, and links to its data-handling documentation is available in our Subprocessor Disclosure.

7. Data Security

We implement technical and organizational safeguards to protect your data, including:

- Encrypted data transmission (HTTPS/TLS)
- Secure account authentication appropriate to your role: passwords are stored only in hashed form and never in plain text, and one-time email login codes expire after a short period and cannot be reused
- Role-based access controls limiting data access to authorized personnel
- Audit logging of data access events. To minimize the personal data held in these permanent records, IP addresses are logged only for privileged (faculty and administrator) accounts; student activity is recorded without an IP address.

8. Contact

If you have questions about this privacy notice or your data, contact your institution's PAIGE administrator, or reach the Proficiency team at paige@perceivant.com.

This notice may be updated to reflect changes in our practices or applicable law. The current version is always accessible from your Settings page.

Subprocessor Disclosure

Proficiency, Inc. Product: Paige — AI Learning Assistant

Last Updated: 2026-06-15

1. Purpose

This disclosure lists the third-party services ("subprocessors") that process data on Paige's behalf. For each subprocessor, it states the purpose it serves, the category of data it receives, and a link to that provider's public data-handling documentation.

Proficiency provides this disclosure so that institutional partners can evaluate Paige's data-processing chain directly, rather than relying on a confidential agreement. It complements the [Privacy & Compliance Overview](#) and the [AI Ethics and Data Use Policy](#).

2. Data-Processing Posture

2.1 Training opt-out

Proficiency has configured each AI subprocessor account so that Paige's inputs and outputs are not used to train, fine-tune, or improve the provider's models. This relies on each provider's published no-training-on-API-data position, enforced through Proficiency's account settings with that provider.

2.2 Retention for abuse detection, and Zero Data Retention

Excluding data from training is not the same as the provider holding no copy of it. As a standard part of their API terms, model providers retain inputs and outputs transiently — typically for a limited window — for abuse and misuse detection (for example, to detect attempts to generate prohibited content). This is true across the major providers, not specific to any one of them.

The only way to eliminate this transient retention is the provider's Zero Data Retention (ZDR) mode. Proficiency has deliberately not enabled ZDR, because doing so materially restricts the available models and provider routing that Paige depends on. This is a considered tradeoff: Paige accepts standard abuse-detection retention (data not trained on, held briefly, then discarded) in exchange for access to the models that make the product work.

2.3 Data Processing Agreements (DPAs)

Proficiency obtains a DPA from each subprocessor where one is available and meaningful, and relies on the training opt-out (Section 2.1) as the safeguard where a DPA is not available. The per-provider status is recorded in each provider's entry in Section 3 and summarized in Section 4: a DPA is signed with OpenAI; a DPA with Langfuse has been requested and is pending signature; Cohere and OpenRouter do not offer a separate DPA, and the training opt-out is in effect for both.

See the [Privacy & Compliance Overview, Section 8.5](#) for how this posture maps to Proficiency's FERPA position.

3. Subprocessors

3.1 OpenRouter — Language Model Routing

Field	Detail
Purpose	Routes all language-model requests: response generation (RAG and Socratic modes), study-tools content generation, content moderation, and query reformulation. OpenRouter forwards these requests to the underlying model providers.
Models used	Meta Llama family — separate models within the family for generation, content moderation, and lightweight tasks such as query reformulation. Exact model versions change as the family evolves and are available on request.
Data received	Student questions, recent conversation context, retrieved passages from course material, and system prompts. May contain personal information a student voluntarily types into a message.
Training opt-out	In effect. OpenRouter provides account-level controls to exclude inputs and outputs from training and to restrict routing to providers that do not train on customer data.
DPA status	No separate DPA. Training is disabled at the routing layer; transient abuse-detection retention applies per Section 2.2.
Public documentation	OpenRouter Privacy Policy , Provider Data Retention , Zero Data Retention controls

Note on downstream routing. OpenRouter is a routing layer; the underlying model providers process the data sent to them under their own terms. Paige currently pins Meta Llama models, so the effective downstream provider is Meta's model infrastructure as served through OpenRouter's selected endpoints.

3.2 OpenAI — Embeddings and Voice

Field	Detail
Purpose	Generates vector embeddings of course material and student questions for retrieval. When voice features are enabled, also transcribes student speech and synthesizes spoken responses, using OpenAI's speech-to-text and text-to-speech models. Exact model versions are available on request.
Data received	Course-material text chunks and student question text (for embeddings). When voice is used: audio recordings of the student's voice, and response text to be spoken aloud.
Training opt-out	In effect. OpenAI does not train on API data by default and offers

Field	Detail
	enterprise data-retention controls.
DPA status	DPA signed. Transient abuse-detection retention applies per Section 2.2.
Public documentation	Enterprise privacy at OpenAI , How your data is used to improve model performance

3.3 Cohere — Search Reranking (Optional)

Field	Detail
Purpose	Re-ranks passages retrieved from course material by relevance to the student's question, using Cohere's rerank model, so the most relevant passages reach the model. Exact model version is available on request.
Status	Optional. Reranking runs only when a Cohere API key is configured; if it is absent, Paige falls back to vector-search ordering and sends no data to Cohere.
Data received	Student question text and the retrieved course-material passages.
Training opt-out	In effect. Cohere offers an enterprise opt-out from training on customer prompts and generations.
DPA status	No separate DPA offered for this usage. Training opt-out is in effect; transient abuse-detection retention applies per Section 2.2.
Public documentation	Cohere Privacy Policy , Cohere Enterprise Data Commitments

3.4 Langfuse — LLM Observability

Field	Detail
Purpose	Records traces of model requests and responses so Proficiency can monitor reliability, latency, and output quality. Integrated as a tracing wrapper around model calls.
Data received	Model prompts and responses, which may include retrieved course passages and text a student typed into a message.
Deployment	Langfuse Cloud. The pilot uses Langfuse's hosted service; trace data is sent to Langfuse Cloud. Langfuse is required for the pilot (reliability, latency, and output-quality monitoring).
Training opt-out	Langfuse is an observability platform and does not train models on customer trace data.
DPA status	DPA requested; pending signature.

Field	Detail
Public documentation	Langfuse Security & Privacy , Langfuse Data Region and self-hosting options

3.5 Amazon Web Services (AWS) — Infrastructure and Configuration

Field	Detail
Purpose	Hosts the application, the PostgreSQL database, and the Redis cache. AWS Systems Manager Parameter Store holds cost-governance configuration (token caps, model kill switch).
Data received	As the hosting provider, AWS stores all application data at rest (accounts, chat history, usage analytics). It does not receive student conversations for any AI purpose; Parameter Store receives only configuration values, not student data.
Training opt-out	AWS does not access or use customer content for training, marketing, or advertising.
Public documentation	AWS Data Privacy FAQ , Data Protection and Privacy at AWS

4. Summary Table

Subprocessor	Role	Data Received	DPA	Optional?
OpenRouter (→ Meta Llama)	Language-model generation, moderation, reformulation	Student questions, context, course passages, system prompts	None (training off)	No
OpenAI	Embeddings; voice transcription and synthesis	Course/question text; student audio (voice features)	Signed	Voice is optional
Cohere	Search reranking	Student question text, retrieved passages	None (training off)	Yes (key-gated)
Langfuse	LLM observability / tracing	Prompts and responses (may contain student text)	Requested; pending	Yes (credential-gated)
Amazon Web Services	Hosting, database, configuration	All application data at rest; configuration values	Hosting terms	No

All providers exclude Paige's data from model training (training opt-out in effect); all retain data transiently for abuse detection per Section 2.2 (ZDR not enabled).

5. Changes to This List

Proficiency updates this disclosure when a subprocessor is added or removed, or when a subprocessor materially changes how it handles data. Provider changes that materially alter data processing during an active semester are communicated to institutional partners in advance, consistent with the [AI Ethics and Data Use Policy, Section 4.3](#).

6. Contact

For questions about this disclosure or to discuss Paige's data-processing chain with your institution:

Email: paige@perceivant.com

Proficiency, Inc.

Links and provider terms are maintained by the providers and may change; this disclosure is updated when a subprocessor is added, removed, or materially changes how it handles data. This document is provided for informational purposes and does not constitute legal advice.

FERPA Statement

Proficiency, Inc. Product: Paige — AI Learning Assistant

Last Updated: 2026-06-15

1. Purpose and Scope

This statement describes how Paige supports institutions in meeting their obligations under the Family Educational Rights and Privacy Act (FERPA). It is written for institutional administrators, registrars, legal counsel, and compliance officers evaluating Paige for use with their students.

This statement is part of Proficiency's institutional trust documentation. It is the companion to the [Privacy & Compliance Overview](#), which describes Proficiency's broader data, security, and regulatory posture, and the [Subprocessor Disclosure](#), which lists the third-party services that process data on Paige's behalf. It is distinct from the student-facing Privacy Notice that students review and consent to within the application.

2. Paige's Role Under FERPA

Paige operates under FERPA as a school official — a third-party service provider granted access to education records to perform an institutional function the institution would otherwise perform itself, under the institution's direct control.

This designation carries specific commitments. As a school official, Proficiency:

- Uses student education records only for the legitimate educational purpose for which access was granted.
- Does not sell, share, or commercially exploit student education records.
- Does not disclose education records to third parties except as the institution directs or as permitted under FERPA.
- Supports the institution in honoring student rights to inspect and amend their education records, and deletes those records responsibly under its retention policy when the student is no longer being served.
- Maintains audit records sufficient to show who accessed student data and when.

Proficiency's specific obligations to each institution are set out in the service agreement between the parties. That agreement, together with the institution's direct control over the education records, is the basis of Proficiency's FERPA position.

There is no such thing as "FERPA certification." Compliance is demonstrated through verifiable practices and contractual commitments — described below — rather than a credential.

3. How Paige Protects Education Records

The commitments above are backed by controls built into the product and its infrastructure.

3.1 Access Is Limited by Role and Ownership

Access to student data is governed by role-based access control:

Role	Access
Student	Own data only — own chats, own account, own usage.
Professor	Aggregated, anonymized usage data for their own course sections. No access to any individual student's conversations or identifiable data.
Administrator (Proficiencyce)	Narrow, audited access to individual student data only in response to a formal complaint or a system-flagged safety concern.

Ownership is enforced at the data layer: a request for another student's data is refused, and the data is not returned in any form, including in error messages.

3.2 Consent Precedes Use

Students must accept the Privacy Notice before they can use Paige's core features. No student data is sent to a third-party AI service until the student has been informed of Paige's data practices and has consented. The consent and its timestamp are recorded.

3.3 Students Can Access Their Data; Records Are Retained and Deleted Responsibly

Students can view their own data within Paige and can delete individual conversations, which are removed immediately. Paige does not offer student self-service deletion of the account itself or bulk data export; these are deliberate choices, and FERPA does not require either. Requests to inspect or amend records are supported through the institution, consistent with FERPA.

Education records are retained while the student is being served by Paige — typically the academic term. Once the student is no longer being served, Proficiencyce's policy is to soft-delete the associated account and then permanently delete it after a 30-day grace period, which allows recovery from an accidental or erroneous deletion. This retention-and-deletion cycle, rather than deletion on student demand, is how Proficiencyce ensures student data is not held longer than the educational purpose requires.

3.4 The Audit Trail Cannot Be Altered

Paige records access to student data and administrative actions in an audit trail that is enforced as immutable at the database level: audit records cannot be deleted or modified by any application path, administrative session, or direct database connection. When a student account is deleted, audit entries are preserved but the student's identity is disassociated, so the record of access remains intact without retaining the deleted student's identity.

3.5 Education Records Are Encrypted

Student data is encrypted in transit and at rest. The [Privacy & Compliance Overview](#) describes Paige's encryption and infrastructure controls in full.

3.6 AI Providers Do Not Train on Student Data

Paige relies on third-party AI services to deliver its core functionality. Proficiencyce has configured each provider so that student inputs and outputs are not used to train or improve the providers' models. The [Subprocessor Disclosure](#) names each provider, the data it receives, and its data-handling commitments.

4. Student FERPA Rights

Students retain their FERPA rights when using Paige. These include the right to:

- **Inspect** their education records maintained in Paige.
- **Request correction** of records they believe are inaccurate or misleading.
- **Consent** to disclosures of their education records, except where FERPA permits disclosure without consent.
- **File a complaint** with the U.S. Department of Education regarding an alleged FERPA violation.

Because the institution is the owner of the education records, students exercise these rights through their institution. Proficiencyce supports the institution in fulfilling them.

5. Data Disclosure and Safety Escalation

Proficiencyce does not disclose student education records to third parties except as the institution directs or as FERPA permits.

Paige screens messages for serious safety concerns. Where a message reflects what appears to be a real and present threat to a student, to others, or to the institution, Proficiencyce may share relevant information with the institution so that someone can respond. This is a best-effort safety measure based on human judgment, consistent with FERPA's provisions for disclosure in connection with a health or safety emergency. It is not an automated determination or a guarantee.

6. Contact

For questions about this statement or to discuss Paige's FERPA posture with your institution:

Email: paige@perceivant.com

Proficiencyce, Inc.

This statement is provided for informational purposes and does not constitute legal advice. Proficiencyce's specific obligations to each institution are defined in the service agreement executed between the parties. In the event of a conflict between this statement and that agreement, the agreement governs.

Accessibility Statement

Proficiencyce, Inc. — Paige (AI Learning Assistant)

Last Updated: 2026-06-15

Our Commitment

Proficiencyce is committed to making Paige accessible to everyone, including people who use assistive technology. We design and build Paige to meet the Web Content Accessibility Guidelines (WCAG) 2.1, Level AA, and we treat accessibility as an ongoing part of how we develop the product.

We give particular attention to the experiences students use most — signing in, the chat interface, and the interactive learning tools — so that the core of Paige is accessible to as many students as possible.

Conformance Status

Accessibility is an ongoing effort, and we continue to identify and address areas for improvement. Paige's conformance is documented in our Accessibility Conformance Report (VPAT® 2.5, WCAG Edition), which details conformance with WCAG 2.1 Level A and Level AA; the report is available to institutions on request. A separate Accessibility Conformance Report covering the broader Proficiencyce courseware platform (WCAG 2.0–2.2 and Section 508) is also available on request. We continue to identify and address accessibility gaps as new versions of Paige are released.

Feedback

We welcome your feedback on the accessibility of Paige. If you encounter a barrier, please let us know.

Email: paige@perceivant.com

This statement reflects Paige's current accessibility posture and is updated as that posture changes.