COIEP Security & Privacy

Trustworthy AI for K-12 special education — grounded, tenant-isolated, FERPA-aligned.

COIEP is an AI assistant for IEP teams, built by the University of Wyoming with privacy and grounding as the starting design constraints, not as features added later. This page summarizes how we handle student data, the legal framework we operate under, and where we differ from general-purpose chatbots.

What COIEP does not do

> *A note for readers comparing tools after the recent Denver Public Schools decision to restrict OpenAI: the issue there was a consumer privacy policy that permits training on user inputs. COIEP is built differently.*

- **OpenAI does not train on your students' data.** Requests are sent under our enterprise agreement with zero data retention.
- **We do not share student records with advertising, marketing, or analytics vendors.** There is no third-party tracking on authenticated pages, and student data is not exported to any product analytics service.
- **We do not vectorize student PII.** Our retrieval index (Pinecone) contains only public content: academic standards and descriptions of evidence-based practices. Student identifiers, IEP narrative text, and session transcripts stay in the primary database.
- **We do not claim HIPAA compliance, because HIPAA is the wrong framework for school records.** IEP data is governed by FERPA. We say so plainly because vendors that conflate the two often misunderstand both.

What COIEP does do

- **Schema-level tenant isolation.** Every IEP, student, session, conversation, and document is scoped to an `organizationId` at the database schema. The foreign key is required. One district's data cannot appear in another district's session by query design.
- **Per-request audit logging.** Every API call is logged with user, endpoint, and timestamp, retained for review.
- **Encryption in transit and at rest.** Database connections require TLS. The managed Postgres instance encrypts data at rest with AES-256.
- **Organization-scoped API credentials.** API keys are issued per organization, not globally. Revoking access for a district revokes it cleanly.
- **Grounded generation.** AI responses are grounded in academic standards and evidence-based practice descriptions, retrieved at request time. The model is not free-associating from internet pretraining when drafting goals or PLAAFP statements.

Compliance roadmap

- **In flight this quarter:** parent-consent workflow surfaces, per-record access audit views for district admins, and TLS termination at the ingress layer of our self-hosted deployment.
- **Planned, not yet started:** formal SOC 2 Type I scoping. We will not claim certifications we have not earned.
- **Legal framework:** FERPA. We operate as a school official under direct control of the district, with data used only for the educational purpose specified.

Compared to a general-purpose chatbot

Dimension	ChatGPT / Claude.ai (consumer)	COIEP
Training on your inputs	Possible under consumer terms	Disabled via zero data retention
Data scoping	Single user account	Per-organization, enforced in schema
Grounding source	Model pretraining	Standards + evidence-based practice library
Audit trail	None visible to the district	Per-request log, retained
Tenant isolation	Not applicable	Required foreign key on every record
Legal framework for K-12	Not designed for it	FERPA, stated explicitly

Who to ask

- **Technical and security questions:** Brad Anderson, Fruition — brad.anderson@fruition.net
- **Licensing and legal questions:** University of Wyoming Office of Technology Transfer
- **Product questions:** Tiffany Hunt and Ling Zhang, College of Education, University of Wyoming