

Encephlo — Security and Data-Flow Reference

Status: Internal source-of-truth document. Plain-English description of how the Encephlo Educator Pilot product handles institutional data. All other compliance artifacts (DPA, HECVAT-Lite, security one-pager) are derived from this document.

Last updated: September 18, 2026 **Owner:** Suraj Swaminathan, Founder — Suraj@encephlo.com **GCP project:** `relearn-learning-ai-tutor` (display name: "Encephlo") **Audience:** Customer security/IT teams, prospective Data Processing Agreement counterparties, internal engineering.

Honesty note: Encephlo is a pre-revenue early-stage company. We do not currently hold SOC 2, ISO 27001, or a formal HECVAT report. This document — and the derived HECVAT-Lite — are a good-faith self-assessment intended to be reviewed alongside source code and infrastructure access. Where a control is **not yet implemented**, this document says so explicitly under "Known Gaps."

1. Product Scope (what's in scope for an institutional pilot)

For a Fall 2026 pilot, the in-scope product is the **B2B Educator Pilot**:

Surface	In Scope (Pilot)?	Description
Educator Studio (web)	Yes	Faculty configure tutors, upload materials, sync Canvas, view Insights.
Student Web Tutor (web)	Yes	Students join via educator-issued link and chat with the configured tutor.
Canvas LMS sync	Yes	Read-only pull of course metadata, assignments, announcements, files.
Insights (educator)	Yes	AI-summarized class & per-student engagement, derived server-side.
Chrome Extension	No	Not deployed for institutional pilots.
macOS / iPadOS apps	No	Archived / paused; not in production.
D2C consumer product	No	Pilot users are not enrolled in the consumer subscription product.

The remainder of this document only describes the in-scope B2B surfaces.

2. Roles and Data Subjects

Party	Role under U.S. privacy law	Notes
The institution	Data controller	Owns the relationship with the data subjects (faculty, students).
Encephlo, Inc.	Data processor / Service provider	Processes data only on the institution's instructions to deliver the service.
Faculty user	Data subject (educator)	Creates a tutor, uploads materials, optionally connects Canvas.
Student user	Data subject (student)	Joins a tutor via educator link; chats with the AI tutor.
Encephlo subprocessors	Sub-processors	See §10 and <code>subprocessors.md</code> .

For FERPA purposes, Encephlo is positioned as a **"school official with a legitimate educational interest"** (FERPA 34 CFR § 99.31(a)(1)(i)(B)) acting on the institution's behalf, processing student education records solely to deliver the contracted service.

3. Data Categories

3.1 Educator data

- **Profile:** display name, title, first name, last name, email (Firebase Auth).
- **Tutor configurations:** subject, level, classroom rules, syllabus snippets, uploaded materials manifest.
- **Canvas tokens:** OAuth refresh token or Personal Access Token (PAT) — stored **server-side only** under `educators/{uid}/integrations/canvas` and not readable by any client (Firestore rules deny `read`, `write` from clients).
- **Course materials:** PDFs / text files the educator uploads or that arrive via Canvas sync. Stored in Cloud Storage at `educator-materials/{educatorId}/{tutorId}/...`.

3.2 Student data (the FERPA-sensitive set)

- **Profile:** display name, email (Firebase Auth).
- **Enrollment record:** which educator/tutor the student is enrolled in (`educators/{uid}/tutors/{tid}/enrollments/{studentUid}`).
- **Chat transcripts:** the student's questions and the tutor's replies, stored in `sessions/{sessionId}` (one session per chat thread). Includes timestamps and message counts.

- **AI-generated summaries:** per-student summaries (engagement, topics, needs-attention flag) cached at `educators/{uid}/tutors/{tid}/insights/current` . **Raw transcripts are never exposed to other students or written into the insights document — only derived natural-language summaries and counts.**

3.3 Operational data

- **Cloud Logging:** request logs, function execution logs, error stacks. May contain user IDs (Firebase UIDs) and high-level event names. Does **not** contain plaintext chat messages or material contents.
- **Vector embedding chunks:** numerical 768-dim vectors derived from educator- uploaded course materials. Stored at `educators/{uid}/tutors/{tid}/materials/{mid}/chunks/{chunkId}` . Not readable by any client (Firestore rules deny `read`, `write`).

3.4 What we do NOT collect

- No credit-card data from students or faculty (Stripe is for the D2C product only and will not be invoked in the pilot).
- No first-party tracking cookies, advertising pixels, or cross-site tracking.
- No screen recordings or persistent screenshots. (Chrome Extension's "what's on my screen" feature is **not in pilot scope**.)
- No biometric data.
- No precise geolocation.
- No data sold to third parties — ever.

4. Where the Data Lives (residency)

Data	Service	Region	Notes
Firestore (all app data)	Cloud Firestore	<code>nam5</code> US multi-region (Iowa + South Carolina)	US-only at rest.
Cloud Storage (uploads)	Firebase Storage bucket <code>relearn-learning-ai-tutor.firebaseio.com</code>	US multi-region	US-only at rest.
Cloud Functions (server)	Firebase Functions v2	<code>us-central1</code> (Iowa) — the v2 default region	All callable / scheduled / streaming endpoints.
Web app compute	Firebase App Hosting (Cloud Run backed)	<code>us-central1</code>	Serves Next.js SSR.

Auth	Firebase Auth	Google global	Account records (email, displayName, UID).
AI inference	Vertex AI Gemini API	global endpoint	See §5 — disclosed honestly.
Logs	Cloud Logging	us-central1 (default)	30-day retention by default.

Disclosure (Vertex AI region): We use Vertex AI's `global` endpoint rather than a single regional endpoint. Both models on our inference path (`gemini-3.8-flash` primary, `gemini-3.5-flash` fallback) are also offered in Google's `us` multi-region; moving the path there is scheduled (see §13, gap 3) and this disclosure will be updated when it lands. Per Google's [Cloud Data Processing Addendum \(CDPA\)](#) and the [Vertex AI service-specific terms](#), Google does not use Customer Data submitted to Vertex AI to train its foundation models, and the CDPA's confidentiality, security, and sub-processing commitments apply globally. We will pin Vertex AI to a US regional endpoint as soon as the production model becomes available there.

5. AI Inference Architecture (Vertex AI)

We migrated from AI Studio (consumer Gemini API key) to **Vertex AI** in April 2026 specifically to inherit Google's enterprise data-processing commitments.

Path	Provider	Auth	Model	Notes
Tutor chat (B2B student)	Vertex AI	Application Default Credentials (ADC) on Cloud Functions service account	<code>gemini-3.8-flash</code> (global); automatic one-time retry on <code>gemini-3.5-flash</code> if the primary fails before responding	No customer training, per CDPA. Same provider and terms for both models.
Tutor configurator (educator)	Vertex AI	ADC	<code>gemini-3.8-flash</code> (global), same retry	Same.
Insights summarization	Vertex AI	ADC	<code>gemini-3.8-flash</code> (global), same retry	Same.
Vector embeddings (RAG)	AI Studio Gemini API	API key (<code>GEMINI_API_KEY</code>)	<code>gemini-embedding-2</code> (text mode, 768 dims)	Transitional — see §13 Known Gaps.

Tool calling and retrieval: the B2B student chat path may invoke a server-side retrieval tool (`retrieve_course_materials`) that queries our vector index to fetch relevant chunks of educator-uploaded materials. **The B2B chat path does NOT call out to Google Custom Search or YouTube Data API** — those tools exist only on the D2C consumer path, which is out of scope for the pilot.

5.1 Answer keys and solution material (academic-integrity control)

Educators may mark an uploaded material as an **answer key**. Two distinct controls apply, and they are deliberately different in strength. We describe both plainly because the difference matters to an academic-integrity review.

Control 1 — students cannot retrieve an answer key. (Hard control.) `answer_key` is a member of `STUDENT_BLOCKED_MATERIAL_KINDS` (`functions/src/b2b/material-retrieval.ts`). The student-facing manifest filter and `fetchMaterialExcerpt` both reject that kind, so a student cannot open, list, or fetch an answer key even if they learn or guess its material ID. This is enforced in the retrieval code, not by an instruction to the model.

Control 2 — the tutor is grounded in the key, and the answers are stripped from its output. (Hard control on output; the grounding itself is intentional.) Since August 2026 the tutor's own context *does* include the answer key, so it explains a problem using the professor's method rather than inventing its own. The security boundary therefore sits on the **output**, not the prompt: `functions/src/b2b/answer-key-grounding.ts` computes a protected set defined as

```
tokens(answer key) - tokens(everything the student can already see)
```

and scans every outgoing response server-side, blocking any text containing a protected value before it reaches the student. What survives that subtraction is, by construction, material that exists only in the key: final results and answer-bearing proper nouns. The scan runs in code the model never sees, so it cannot be prompt-injected around. This design replaced an earlier prompt-instruction approach after internal red-teaming defeated the prompt version 16 times out of 16; that finding is why the boundary was moved.

What this does not do. The scan closes the **verbatim** channel. It does not stop a paraphrase, and it cannot stop a student inferring an answer from a correct explanation. We state this explicitly rather than claim a guarantee we cannot support. Educators control each material's visibility at any time, and an educator who wants the stricter posture can simply not upload the key.

6. Authentication and Identity

- **Firestore Authentication** is the identity provider for both educators and students. Supported methods today: email/password, Google sign-in, Microsoft sign-in.
- **Educator onboarding** captures title, first name, last name; these compose the public tutor display name (e.g., "Prof. Van Horn's Microeconomics tutor").

- **Student onboarding** for the pilot is via educator-issued join link (`/join/{code}`); the student authenticates with email/password or Google.
- **Session tokens** are Firebase ID tokens (short-lived JWTs). Refresh tokens are stored in browser local storage by the Firebase SDK, scoped to the Encephlo origin.
- **Custom auth domain:** `login.encephlo.com`.

There is **no SSO via SAML or OIDC against the institution's IdP today**. We can add SSO during the pilot if requested; it is not a blocker for the pilot to launch.

7. Authorization (data isolation)

Authorization is enforced by **Firestore Security Rules** and **Cloud Storage Security Rules** (see `firestore.rules` and `storage.rules` in the repository). Highlights:

- **Educator data** (`educators/{uid}/...`) is readable/writable only by the authenticated educator with matching UID.
- **Canvas tokens** (`educators/{uid}/integrations/...`) are denied to all clients — Cloud Functions read them via the Admin SDK.
- **Roster rows** (`enrollments/{studentUid}`) are readable by the educator; client writes are denied (Cloud Functions write them on join).
- **Materials & vector chunks** (`materials/{mid}/chunks/...`) are denied to all clients; the chunks subcollection is touched only by the ingest trigger and the retrieval Cloud Function via Admin SDK.
- **Sessions** (`sessions/{sessionId}`) are readable/writable only by the user whose `userId` field matches their UID.

A student cannot read another student's transcripts. An educator cannot read transcripts directly from Firestore — the educator only sees AI-derived summaries from the insights document. A student cannot read other students' enrollments or summaries.

Cloud Functions that perform privileged work (insights generation, Canvas sync, material ingest) authenticate the caller via the Firebase ID token and verify ownership of the target resource before acting.

8. Encryption

- **In transit:** all client ↔ server and server ↔ Google traffic is TLS 1.2+. Firebase / Google Cloud do not accept plaintext HTTP for any production endpoint.
- **At rest:** Google Cloud encrypts all stored data with AES-256 by default, using Google-managed keys. We do not currently use customer-managed encryption keys (CMEK); this is available on Google Cloud and can be discussed for a future enterprise tier.

9. Logging, Monitoring, and Retention

- **Cloud Logging** captures function execution logs and request logs. We log structured events (e.g., `[findRelevantChunks] tutor=... candidates=16 returned=6`) but **never log the plaintext student chat text or full material contents**. We may log message counts, similarity scores, and user IDs to debug issues.
- **Default log retention**: 30 days (Cloud Logging's `_Default` bucket).
- **Firestore data retention**: retained until the educator deletes the tutor (cascade delete via `deleteEducatorTutor`) or the institution requests deletion under §11.
- **Insights cache**: regenerated nightly by `scheduledEducatorInsights` and on demand via the Refresh button. The previous version is overwritten — there is no history.

We do not have a SIEM today; alerting is via Google Cloud's built-in operational alerts and email notifications to the founder.

10. Subprocessors

A complete, dated list lives in `subprocessors.md`. In summary, the pilot subprocessors are:

1. **Google LLC** (Firebase, Cloud Firestore, Cloud Functions, Cloud Storage, Cloud Logging, Vertex AI Gemini, Firebase App Hosting, Firebase Auth)
2. **Instructure Inc.** (Canvas LMS — only as the customer's authorized integration source; data flows from the institution's Canvas instance to Encephlo at the educator's request)
3. **Microsoft Corporation** (Microsoft Identity Platform — only if a user chooses Microsoft sign-in)

Not used in the pilot: Stripe (D2C billing only), Resend (internal-only founder digest), ElevenLabs (deprecated and removed), Google Custom Search API and YouTube Data API (D2C web tools only — not invoked on the B2B path).

11. Data Subject Rights and Deletion

- **Access**: an educator can view all of their tutor configs, materials, and insights through the Educator Studio. A student can view their own chat history through the Student Web Tutor.
- **Correction**: users can edit their `displayName` via account settings.
- **Deletion (educator)**: the `deleteEducatorTutor` callable cascades a delete across the tutor doc, enrollments, materials, chunks, and the insights doc. Educator account deletion is available on request to suraj@encephlo.com (we will provide a self-service path before GA).
- **Deletion (student)**: student account deletion on request to suraj@encephlo.com — turnaround target ≤ 30 days.

- **Institutional deletion:** at the end of the pilot, the institution may request full deletion of all institution-affiliated data; turnaround target ≤ 30 days.

12. Incident Response

- Founder (Suraj Swaminathan) is the security incident commander.
- Reachable at suraj@encephlo.com.
- **Breach notification commitment for the pilot:** Encephlo will notify the institution of any confirmed unauthorized access to its data **without undue delay and in any case within 72 hours** of confirmation, per the pilot DPA.
- We will share the incident scope, affected data subjects, root cause, and remediation plan as facts become available.

13. Known Gaps and Roadmap

Honest disclosure of what is **not yet** in place. Each item has an owner and a target.

#	Gap	Mitigation in flight	Target
1	<code>apphosting.yaml</code> contains plaintext API keys for Gemini AI Studio and Google Custom Search.	Rotate keys; migrate to Secret Manager (<code>secret: SECRET_NAME</code> references).	Before the pilot starts (target: pre-Fall 2026).
2	Vector-embedding path still uses AI Studio Gemini API (not Vertex AI).	Migrate <code>gemini-embedding-2</code> calls to Vertex AI. Behavior is currently preserved by registering both <code>googleAI()</code> and <code>vertexAI()</code> Genkit plugins.	Q3 2026.
3	Vertex AI uses the <code>global</code> endpoint, not a US-only regional endpoint.	Both models on the path (<code>gemini-3.8-flash</code> , <code>gemini-3.5-flash</code>) are now offered in Google's <code>us</code> multi-region. Pin the inference path to <code>us</code> and re-verify latency; non-global pricing is about 10% higher.	Q4 2026.

4	No SOC 2 / ISO 27001 / formal HECVAT report.	This HECVAT-Lite is our self-assessment for pre-pilot. SOC 2 Type 1 is a 2027 goal.	2027.
5	No SAML/OIDC SSO against the institution's IdP.	Microsoft sign-in is available today; SAML is a roadmap item if the pilot requires it.	On request.
6	No customer-managed encryption keys (CMEK).	Available on Google Cloud; discuss for enterprise tier.	On request.
7	No formal pen test or third-party security audit.	Internal review only.	Before broad GA.

14. Document Map

- `security-and-data-flow.md` — this document (source of truth).
- `subprocessors.md` — current subprocessor list.
- `dpa-template.md` — Encephlo Pilot Data Processing Agreement template.
- `hecvat-lite.md` — Higher Education Community Vendor Assessment Toolkit (Lite) self-assessment.

This document is a living artifact. Material changes — new subprocessors, new regions, security-relevant code changes — must be reflected here within 30 days.