

Data Privacy at Acumen

What we do to protect your students' data, and how Acumen fits within FERPA

Acumen is an AI-native platform for developing professional judgment in accounting and finance, built by a CPA and former FASB staff member on faculty at BYU. Because Acumen works directly with student coursework and grades, protecting that data is a first-order design constraint, not an afterthought. This overview explains what data Acumen touches, how it's used, and how the platform fits within your institution's FERPA obligations — in plain language, for anyone evaluating Acumen before the lawyers get involved.

The short version

- Student work is used only to grade, coach, and support your course — never sold, never used for advertising, never used to train a third party's general-purpose AI model.
- AI grading runs through Anthropic's Claude API, which does not train its models on submitted work by default.
- Data is encrypted in transit and at rest, and access is scoped so students and professors only see what they're authorized to see.
- Acumen operates as a FERPA “school official” — your institution stays in control of the data; we're a service provider to it, not an owner of it.
- A full Data Privacy Addendum is available for your legal or IT team to review and execute alongside your agreement — this document is the plain-English version.

What data does Acumen touch?

Four kinds. Roster and account information — name, institutional email, class enrollment. Submitted coursework — memos, case analyses, and Discovery Case interview transcripts. Grading and feedback data — AI-generated and professor-reviewed scores and comments. And basic usage data — login and submission timestamps. That's the full list; Acumen doesn't collect anything beyond what it needs to run the class.

How does the AI grading actually work?

When a student submits work, Acumen sends the relevant text to Anthropic's Claude API, which generates a grade and feedback for that submission and returns it. Nothing more happens with that data on Anthropic's side: by default, Anthropic does not use API inputs or outputs to train its models. One nuance worth knowing, current as of August 2026: Anthropic now retains prompts and responses sent to its most capable models for up to 30 days, purely for safety and abuse monitoring, with access restricted to a small, logged set of reviewers. That's an industry-wide shift, not something specific to Acumen, and it doesn't change the no-training policy — but we'd rather you hear it from us than discover it later.

Is student data ever sold or used for advertising?

No. Acumen uses student data for one purpose: to provide, support, and improve the Service for your institution. It is never sold, never rented, never used for behavioral advertising, and never used to train a general-purpose AI model beyond the per-request use described above. The only other use is de-identified, aggregated analysis for product improvement — the kind of data that can't be traced back to an individual student.

How is the data secured?

Encryption in transit and at rest. Access controls that scope every student's data to their own enrollment, and every professor's access to their own class. A centralized authorization layer that governs every grading, rubric, and roster change. Rate limiting on the endpoints that call the AI models. These aren't aspirational — they're what's running in production today, most recently verified in an internal security review in August 2026.

Who else touches the data?

A small set of infrastructure providers, each under contract and each restricted to using the data solely to provide their service to Acumen — none of them can use it for their own purposes.

Service	What it does	What it sees
Anthropic (Claude)	Powers AI grading, feedback, and the Judgment Coach	Submitted work, one request at a time, to generate that student's output
Supabase	Database, authentication, file storage	Everything Acumen stores — accounts, submissions, grades
Vercel	Hosts the application	Requests passing through; nothing persisted independently
Resend	Sends notification emails	Name, email address, notification content
Sentry	Alerts us to bugs and errors	Technical error metadata only — not student content

What happens when a course ends, or an institution stops using Acumen?

Data is retained only as long as it's needed to run the course and preserve grade history. On request, or when an agreement ends, Acumen will export an institution's data and then delete it from production systems within 30 days, aside from routine backup rotation and any short-term retention required by law or described above.

Want something more formal?

This overview is meant to answer the questions that come up before a contract is on the table. When you're ready to move forward, Acumen has a full Data Privacy Addendum — the FERPA “school official” designation, permitted-use limits, breach notification terms, subprocessor list, and retention/deletion commitments, in contract language — for your legal or IT team to review and execute alongside your agreement.

Questions? Reach out to Aaron Jacob, Acumen.