

DLP & Pseudonymization

Private student data is scrubbed from prompts and files before they ever leave the browser.

Students paste a lot into AI tools: a classmate's name, an email, a paragraph from a record. Tenet's data-loss prevention pipeline catches that private information **on the device, before the prompt reaches the AI vendor**, and it covers both what a student types and the files they upload, including PDFs, Office documents, and images by way of text recognition.

A FOUR-PHASE PIPELINE

1

Pattern PII

Recognizes common personal data such as emails, phone numbers, and IDs by pattern. **Included in Basic.**

2

Roster-aware

Uses your roster to catch real student names and identifiers that no generic pattern would. **Pro.**

3

Pseudonymization

Swaps real names for consistent stand-ins so the prompt still makes sense to the AI. **Pro.**

4

Context classifier

Judges sensitive context that simple matching misses. **Pro.**

WHY PSEUDONYMIZATION IS THE STANDOUT

- Plain redaction breaks the prompt: "help [REDACTED] and [REDACTED] resolve a conflict" is useless to the AI.
- Tenet instead substitutes stable stand-in names, so the AI gives a genuinely useful answer while the real identities never leave the device.

WHAT IS BASIC, WHAT IS PRO

TENET BASIC

Pattern-based PII redaction across chat and every upload format. Free.

TENET PRO

Adds roster-aware detection, name pseudonymization, and the context classifier.