Tenet by TrueMade AI

THE FOUNDATION

On-Device Architecture

The sensitive work happens in the student's browser. Prompts never reach Tenet's servers.

Most K-12 AI governance products route student prompts through a vendor-operated cloud filter. That creates a new place your students' words are stored, a new data agreement, and a new breach surface. Tenet is built the other way. Detection, data-loss prevention, the safety classifiers, and rule enforcement all run **inside the Chrome extension on the student's device**. Tenet's backend only sends configuration and receives sanitized analytics. It is a privacy guarantee that is built into the architecture, not promised in a policy.

WHAT CROSSES THE DEVICE BOUNDARY

- **Leaves the device:** the DLP-scrubbed prompt to your own AI vendor, and sanitized, aggregate analytics to a storage bucket you choose.
- **Comes back:** the AI vendor's response, your district configuration and roster, and model updates.
- **Never reaches Tenet:** raw prompts, raw AI responses, real student names, uploaded file contents, or conversation transcripts.

WHY IT MATTERS TO YOUR TEAM

Privacy officer
There is no Tenet-side store of student prompts to disclose, audit, or breach.

CTO / IT director
A compromise of Tenet's backend cannot leak prompts that were never in it. No proxy, no slowdown.

FERPA lead
Tenet holds no conversation transcripts. The only records that exist are sanitized and live in your storage.

Procurement
Removes much of the standard vendor security review: data residency, retention, and deletion of raw content.

THE HONEST SCOPE

- Detection runs in milliseconds with no cloud round-trip added to traffic, and Tenet pairs alongside your existing web filter rather than replacing it.
- Coverage is managed Chrome today, on Chromebooks and managed Windows. Other browsers and native mobile apps are not covered yet.