

Tenet — Objection Handling Battlecard

FIELD	VALUE
STATUS	V1, working document
LAST REFRESHED	2026-05-24
OWNER	Caleb
AUDIENCE	Anyone running Tenet sales conversations — founder, future SDR/AE hires, partners pitching on our behalf
LENGTH	~12-minute read
COMPANION DOCS	Conversion Playbook · Pro Value One-Pager · Pricing Guide · Securly Reverse Engineering · Briefing 04 — Safety & Compliance

0. TL;DR

A practical battlecard for the objections we actually hear, in rough frequency order. **The single most common objection in our market is "we already have Securly transparency" — it gets \$2 to itself.**

Format for each: the objection in their words, the real concern under it, what to say, what to leave behind, what to NOT do.

Index

#	Objection	Frequency
1	"Why do we need this on top of Securly / GoGuardian / Lightspeed?"	Very high
2	"Securly already does AI transparency"	Very high — front and center
3	"What about GoGuardian Beacon?"	High
4	"We're a Microsoft / Edge district"	High
5	"\$8/student is too expensive"	Medium
6	"We want to wait to see how the AI policy landscape settles"	Medium
7	"Our teachers don't want another tool"	Medium
8	"What if your classifiers are wrong?" (false positives / negatives)	Medium
9	"How do you compare to BrainFreeze?"	Low-medium

#	Objection	Frequency
10	"We don't have budget this year"	Medium
11	"We want to build this ourselves"	Low
12	"What happens to our data if you get acquired or shut down?"	Low (but critical when it comes up)
13	"We need an RFP"	Low (but high-effort when it happens)

1. "Why do we need this on top of Securly / GoGuardian / Lightspeed?"

Their words

"We already pay for Securly. They cover this."

The real concern

The CTO is being asked to add another vendor on top of an existing safety stack. Their internal political capital and procurement-headcount-hours are limited. They need to understand whether Tenet replaces something or adds to something.

What to say

"Securly, GoGuardian, and Lightspeed are web filters — they block URLs and monitor browsing. They were architected before AI chat surfaces existed, and their AI features today are extensions to that core. Tenet is the opposite — built for the inside of the AI chat: what gets typed, what gets uploaded, what gets returned. The two are complementary, not redundant. Almost every Tenet customer keeps their existing filter; we just add the AI-specific layer."

Leave-behind

[Competitor Comparison Sheet](#) — feature-by-feature matrix.

What NOT to do

- Don't trash-talk the incumbent. Almost every prospect has bought one of these products in the last 3 years; criticizing the purchase criticizes the buyer.
- Don't claim Tenet replaces them. It doesn't, and overclaiming kills credibility.

2. "Securly already does AI transparency" — the headline objection

This is the most common objection and the most dangerous one to mishandle. It gets a full treatment.

Their words

Variations heard:

- "We already pay Securly for transparency. They monitor student prompts to ChatGPT."
- "Securly is rolling out AI features in their next release — why wouldn't we just wait for that?"
- "Our filter already sees what students are typing into AI tools."
- "Isn't your value-add just the redaction overlay?"

The real concern

Two related concerns sit underneath this:

1. **The CTO believes Securly already solves the problem.** They have heard a Securly pitch, seen a demo, or read the press release. They are pattern-matching Tenet to "another version of the same thing."
2. **The CTO has not actually examined the data flow.** Most CTOs have not read Securly's technical architecture; they trust the brand and the filter's existing posture.

Why the objection is wrong (the technical reality)

Per the repo's reverse engineering of Securly's AI chat product (see [securly-reverse-engineering.md](#), based on the production Securly extension shipped on managed Chromebooks):

Securly's AI transparency works as follows:

1. Securly's content script (`content_ai.min.js`) hooks the AI platform's input box.
2. On send, it **intercepts the prompt** and posts the text via a service-worker channel to Securly's backend (`POST chat.securly.com/crextn/sendmessage`).
3. The POST body includes `{site, sender:"USER", message, uuid, uid}` — meaning **the raw student prompt** (any PII, any roster data, any free text the student pasted in) and the student's school email (`uid`) are transmitted to Securly's infrastructure.
4. Securly's server-side classification decides allow / block / rewrite. The decision comes back to the content script.
5. AI responses are similarly POSTed (`sender:"AI"`).

This is the architectural reality of "transparency" — every prompt and every response leaves the student's device and enters Securly's cloud.

Tenet does the **opposite**: PII detection, classifier inference, and pseudonymization all happen **inside the student's browser**. The prompt is examined, redacted if needed, and then sent to the AI vendor. **The Tenet backend never sees the prompt content**, only sanitized categorical events (`message_sent` , `blocked` , classifier severity tier).

What to say (the two-sentence version)

"Securly's transparency mode works by sending every student prompt — raw, with whatever PII the student pasted — to Securly's servers for classification. Tenet does that work on the device, so the prompt never leaves the browser. If your district counsel is comfortable with the data flow Securly creates, they will love what Tenet does instead; if they're not yet comfortable, that's the reason to evaluate Tenet."

What to say (the longer version, for an architecturally-curious CTO)

"Both products are trying to solve a real problem — make it possible for districts to govern AI without going dark on what students are doing. We just made the opposite architectural choice. Securly puts the classifier on their servers; Tenet puts it on the device. The Securly choice means they can update classifiers without pushing new code; it also means every prompt creates a separate data flow to their infrastructure, which most district DPAs need to be expanded to cover. The Tenet choice means we ship classifier updates with extension updates (Chrome auto-updates handle that), but no prompt ever leaves the browser, which sidesteps the FERPA data-flow question entirely. Which choice is better for your district depends on what your counsel and superintendent are comfortable with. We are happy to do a side-by-side technical walkthrough."

Leave-behind

- The full breakdown above, ideally rendered into a short technical blog post (TODO; on the [GTM strategy §5.4 content list](#))
- [Briefing 04 — Safety & Compliance](#)
- [Security & Compliance Datasheet](#) — has the data-flow diagrams

What NOT to do

- **Never accuse Securly of being non-compliant.** They are likely compliant under their own DPA. The point is architectural difference, not legal violation.
- **Never make this binary.** Most districts will continue running Securly for web filtering. Tenet is the AI-layer addition.
- **Never demand the CTO take a side in the meeting.** This is a technical decision their counsel needs to weigh. Plant the seed; let them research.
- **Never lead with the technical reverse engineering in the first meeting.** It's heavy. Save the deep dive for the architectural call, or for written follow-up.

What to expect after the objection lands well

- A 1-2 week pause while the CTO talks to district counsel or the data privacy officer

- Either silence (lost; they decided not to disturb the existing setup) or a follow-up call where counsel is on the line (won; the conversation is now substantive)
-

3. "What about GoGuardian Beacon?"

Their words

"GoGuardian Beacon already does self-harm detection. We have that."

The real concern

GoGuardian Beacon is a real, well-known product. The CTO is asking whether Tenet's safety-classifier story has anything Beacon doesn't.

What to say

"GoGuardian Beacon scans browsing activity for self-harm signals across the web — broadly across all sites, not specific to AI chat. The signal it catches is mostly search queries, documents, and visited pages. Tenet's self-harm classifier runs on the actual content of AI conversations — the prompts the student is typing into ChatGPT or Claude or Gemini, and the responses coming back. Those are different signals; Beacon misses most of the AI-specific ones. Best case, they're complementary."

Leave-behind

Briefing 04 — [Safety & Compliance](#), specifically the four-path self-harm architecture description.

What NOT to do

- Don't claim Beacon doesn't work. It does, for what it covers.
 - Don't oversell Tenet's classifier — be honest about false-positive rate (see §8).
-

4. "We're a Microsoft / Edge district"

Their words

"We don't run Chromebooks. We're a Surface district / Edge district / Entra ID."

The real concern

They want to know if Tenet will work for them at all.

What to say

The honest answer:

"Tenet today is Chrome-first and Google Workspace-first — about 90% of K-12 nationally is Chromebook-dominant, and we built for that majority first. Edge support is on our roadmap; we revisit when we get the first serious Microsoft-only district lead, and you may be that trigger. Want me to be transparent: today, we cannot deploy to you the same way we deploy to a Google district. If this is a deal-breaker, I'd rather tell you now than waste your time. If you're open to a 6-12 month conversation while we build Edge support — driven in part by your needs — let's have that conversation."

Leave-behind

None. This is a "let's not pretend" conversation.

What NOT to do

- Don't say "we'll build it for you." We won't without a serious commitment.
- Don't pretend our Chrome posture transfers to Edge. The DOM adapters are platform-specific; the architecture concepts transfer, but the code doesn't.
- Don't lose them entirely — keep their name in the pipeline as a "Microsoft trigger" — we revisit when we have 3+ identified Microsoft district leads.

5. "\$8/student is too expensive"

Their words

"\$8 per student is more than we budgeted." / "Securely is only \$4." / "Our budget is tight."

The real concern

Usually one of three:

1. **Real budget constraint.** Their AI line item is small or nonexistent and they need to find money.
2. **Anchoring.** They've compared to per-K-12-student pricing of \$3-4/year and not yet understood our AI-active basis.
3. **Value not yet proven.** They are saying "expensive" but mean "I don't yet believe it's worth \$8/year."

What to say

"Worth understanding two things before we talk dollars. First, our \$8 is per AI-active student — typically grades 7-12 — not full K-12. If you have 10,000 students and 4,000 in grades 7-12, you pay for 4,000, which usually puts the all-in cost at or below comparable per-K-12 platforms. Second, if value isn't yet clear, that's a different conversation — let's set up a 30-day Pro trial at no cost so you can see what Pro does on your actual data before signing anything."

If the objection persists:

"If price is genuinely the issue, our standing 2-year is 8% off and 3-year is 12% off. We don't discount the standing 1-year price. What's the gap between \$8 and what would make this work?"

Leave-behind

[Pricing Presentation Guide](#). [Tier Comparison Chart](#).

What NOT to do

- **Never discount below \$4/student/year** without explicit founder approval. The floor protects unit economics and future-pricing power.
- **Never volunteer a discount.** Make them ask.
- **Never claim we're "the cheapest."** We are not. We are the most economical per-K-12-student because of the AI-active basis, which is a different statement.

6. "We want to wait to see how the AI policy landscape settles"

Their words

"We're going to wait a year. The policy landscape is too uncertain right now."

The real concern

This is usually one of:

1. **Genuine paralysis.** The CTO doesn't know what to recommend to the superintendent.
2. **Hidden political situation.** A board member or parent has strong opinions and the CTO is avoiding the conversation.
3. **Soft no.** They've decided not to evaluate but won't say it directly.

What to say

"Totally fair — and that's exactly why Basic is free and deployable in 30 minutes. Waiting doesn't mean doing nothing. The fastest way to be ready when the landscape settles is to have AI governance running in the background so you have data, not speculation, when your board asks. Want to install Basic this month and revisit Pro in 90 days? That commits you to nothing and gives you actual ground truth on what's happening with AI in your buildings."

This re-frames the wait as: deploy Basic now, wait on Pro. The CTO often accepts.

What NOT to do

- Don't pressure for a Pro decision today. The "wait" objection is real.
- Don't disappear after the no. Add to the warm-watch list; check back at a sensible interval (90 days minimum).

7. "Our teachers don't want another tool"

Their words

"Our teachers are exhausted. They don't want one more thing to learn."

The real concern

Real change-fatigue. Often combined with bad experiences from previous EdTech rollouts.

What to say

"Tenet doesn't require teacher training to install — it runs in the background as a managed extension and the basic experience for teachers is exactly the same as if it weren't installed. Teachers only encounter Tenet when (a) they want to author class-specific rules in Pro, in which case the rule editor is a 5-minute tutorial, or (b) the DLP warning popup fires when they paste a roster into ChatGPT — which is the moment the value lands. For Basic specifically, teachers don't need to do anything. There's nothing to learn."

For Pro:

"Our teacher pilot training is a 30-minute video and a one-page quick reference. Districts that have done it report teacher adoption of class rules in week 2."

Leave-behind

[Teacher Onboarding One-Pager](#). [Teacher Pilot Training Guide](#).

What NOT to do

- Don't promise zero teacher impact. For Pro, classroom rules ARE a teacher touchpoint. Be honest.
- Don't oversell the simplicity. The DLP popup IS an interruption — that's the value, but acknowledge the cost.

8. "What if your classifiers are wrong?"

Their words

"What's your false-positive rate?" / "What happens if a kid says something innocuous and the self-harm overlay pops up?" / "What if you miss a real crisis?"

The real concern

Legitimate question from a Director of Student Services. They want to understand the failure modes before deploying.

What to say

"Every classifier has FP and FN rates; pretending otherwise loses your trust. Here's where we are: the self-harm classifier has been retrained to suppress academic-context false positives — for example, students writing about Romeo and Juliet, or research papers on suicide prevention. The local crisis-resource overlay is designed to be non-judgmental — even a false positive surfaces a helpful resource, not an accusation. On the false-negative side, we run a four-path architecture — prompt-side classifier, immediate-threat hard-stop classifier, response-side monitoring, and a vendor-only safety net for cases the prompt-side missed. The full architecture is in our ML system doc; happy to walk through it with your team."

If they push for numbers:

"Held-out evaluation: 72.5% pass rate on a 40-item golden set for self-harm v1, 62.5% for v2 (which trades recall for multiclass severity). We publish these in our ML release notes. The exact production rate depends on the prompt distribution at your district — happy to do a 30-day silent evaluation against your Basic install before turning on Pro alerting."

Leave-behind

[ML Classifier System](#) (deep technical reference; share only if they're technical enough to want it). [Briefing 04 — Safety & Compliance](#).

What NOT to do

- Don't claim a low FP rate without numbers. Vague claims sound dishonest.
- Don't be defensive. This is the right question to be asked.
- Don't promise to never miss a crisis. We will. The right framing is "we catch what humans wouldn't, and we miss some that humans would; our job is to maximize the union of those two."

9. "How do you compare to BrainFreeze?"

Their words

"We've been looking at BrainFreeze. They have a similar story."

The real concern

BrainFreeze is a real, growing competitor with overlapping positioning. The CTO wants a fair comparison.

What to say

"BrainFreeze and Tenet have overlapping problem statements — AI in K-12 governance — but different architectures and pricing. BrainFreeze is district-flat or per-teacher pricing (\$0.50/user/year district or \$99/year individual teacher). Tenet is per-AI-active-student. For a 5,000-student district with 2,000 AI-active students and 200 teachers: BrainFreeze district plan is \$2,500/year; Tenet Pro is \$10,000/year. We're more expensive. We do more — classroom rule profiles, roster-aware DLP, on-device pseudonymization, counselor alert dispatch — but if your district doesn't need those, BrainFreeze is a perfectly good answer. Want me to do a side-by-side feature comparison for your district size?"

Leave-behind

[Competitor Comparison Sheet](#).

What NOT to do

- Don't claim Tenet is "better." For some districts, it isn't. Be honest about overlap and pricing trade-offs.
- Don't take a swing at BrainFreeze. They're a credible competitor doing real work in the space.

10. "We don't have budget this year"

Their words

"Budget's already set; we couldn't take this to the board until next cycle."

The real concern

True budget constraint, OR a polite way of saying "we're not interested but I don't want to say no."

What to say

"Hear you. Two paths: (1) Basic is free now, and we put Pro on the agenda for your next budget cycle — you'll have 6-12 months of Basic data to make the case with. (2) If there's a specific use case Pro would address — like counselor alerting — we can sometimes structure a Q4 deploy with a payment cycle starting your next FY. Which fits better?"

Leave-behind

Basic GTM Strategy §1.3 — Basic install as the no-budget-required entry point.

What NOT to do

- Don't push for an out-of-budget Pro deal. Districts who buy outside cycle often regret it and don't renew.
- Don't lose the relationship over the "no this year." Schedule the next-cycle conversation.

11. "We want to build this ourselves"

Their words

"We have an in-house dev team; we could build this."

The real concern

Usually a face-saving response to a price objection, or a CTO with engineering background who genuinely thinks it's tractable.

What to say

"You could. We've open-sourced parts of the approach in our public materials. Realistically: the rule engine, classroom rules, basic DLP regex — your team could build in 2-3 months. The on-device ML classifiers (self-harm, jailbreak, illicit) plus the four-path safety architecture is more like 12-18 months and requires ongoing model training. The seven adapter integrations (ChatGPT, Claude, Gemini, etc.) plus keeping them working against weekly DOM changes is another 6-9 months of one engineer's full attention. If your team has 24 months of slack and you want to own the codebase, build it. If you want to spend that time on things only your district can do — curriculum, instructional design, your own ML — buy Tenet."

Leave-behind

None. This is a judgment-call conversation.

What NOT to do

- Don't dismiss the build. K-12 IT teams have real engineering capability.
- Don't lock in. If they decide to build, leave the door open for a future "this is harder than we thought" conversation in 12 months.

12. "What happens to our data if you get acquired or shut down?"

Their words

"You're an early-stage company. What's the continuity plan?"

The real concern

Real procurement question, especially from district counsel. They've been burned before.

What to say

"Good question to ask. Two pieces. First, our architecture means there's very little 'your data' on our servers — district configuration and sanitized categorical events. No raw prompts, no AI responses, no PII. So even in a worst-case shutdown, your students' content isn't trapped anywhere. Second, our DPA includes a data-export clause and a 90-day post-termination data retention requirement before purge. If we were acquired, our DPA terms transfer to the acquirer — that's contractually binding. If we shut down, you get a 90-day window to export your config and roll off."

Leave-behind

[DPA Strategy & Tenet Exhibit C](#), specifically the continuity and termination clauses.

What NOT to do

- Don't pretend we're a Fortune 500. We are an early-stage company; counsel can see that on SoS filings.
 - Don't promise we'll never be acquired. Many early-stage SaaS companies are.
-

13. "We need an RFP"

Their words

"We require an RFP for all software contracts over \$X."

The real concern

Procurement-mandated process. Not really negotiable.

What to say

"Understood — we respond to RFPs. Heads up before you draft: about 40 of the questions in a typical K-12 EdTech RFP are already answered in our [Security & Compliance Datasheet](#), and you can paste those answers directly. If your RFP has a draft template, send it to us and we can flag where our existing documentation already addresses each section. Saves us both time. Turnaround on the full RFP response is typically 5-10 business days from receipt."

Leave-behind

[Security & Compliance Datasheet](#). [DPA + Exhibit C](#). [Privacy Policy](#). [Pro Value One-Pager](#).

What NOT to do

- Don't dodge the RFP. Some districts can only buy via RFP; refusing kills the deal.
 - Don't write a custom response for every section. Use the datasheet as your library; only write custom answers for district-specific items.
-

14. Cross-cutting principles for handling any objection

1. **Listen before answering.** Restate the concern in your own words. "So what you're saying is X — am I hearing you right?" Half of objections evaporate once the CTO feels heard.

2. **Acknowledge the legitimate part.** Most objections contain a real concern; pretending otherwise loses trust.
3. **Differentiate stated reason from real reason.** "We don't have budget" often means "I'm not sold yet." The right response addresses the real reason, not the stated one.
4. **Don't be defensive.** Objections are gifts — they tell you what's blocking the deal. The CTO who has no objections is the CTO who won't sign.
5. **Have a fact, not just an opinion.** "On-device > cloud" is a slogan; the Securly POST endpoint to `chat.securly.com/crxtxn/sendmessage` is a fact. Facts win.
6. **Know what you don't know.** "Let me get back to you within 2 business days" is a perfectly acceptable answer for anything technical you're not sure about. Beats improvising and being wrong.

15. What we explicitly do NOT do in objection handling

- **No FUD (Fear, Uncertainty, Doubt) about competitors.** State the architectural facts; let the prospect decide. FUD damages our credibility and is wrong on principle.
- **No "limited time" pricing offers** as a closing tactic. We are not a B2C SaaS. False urgency reads as desperation.
- **No bait-and-switch on Pro features.** Whatever the one-pager says Pro includes, Pro includes. Don't promise extras to close.
- **No name-dropping prospects to other prospects** without explicit permission from the named district. Roger's AR district has consented; future references require the same explicit consent.
- **No promising features we don't have on the roadmap.** "We're going to build that" is a future-looking statement subject to regulatory framing; don't commit to features informally.

16. Triggers to revisit this battlecard

Trigger	What we'd revisit
Same objection comes up 3+ times	Add to this doc; refine the response
First lost deal where the objection wasn't covered here	Add the new objection; analyze why we lost
Securly releases a new AI feature	§2 needs technical re-evaluation; re-do the reverse engineering
GoGuardian releases major Beacon update	§3 needs refresh
BrainFreeze pricing changes	§9 needs refresh
New competitor enters the market	Add to the comparison table and add an entry here

Trigger	What we'd revisit
First RFP response written	§13 gets actual content from the RFP responses
Annual review (2027-05-24 default)	Full refresh

17. Companion reading

- [Basic→Pro Conversion Playbook](#) — when and how to deploy each objection response
- [Pro Value One-Pager](#) — the positive pitch
- [Pricing Presentation Guide](#) — deeper pricing detail
- [Outreach Sequences](#) — email templates that pre-empt some of these objections
- [Securly Reverse Engineering](#) — source for §2
- [ML Classifier System](#) — source for §8
- [Briefing 04 — Safety & Compliance](#) — for §2 and §8 walkthroughs
- [Competitor Comparison Sheet](#) — for §1, §3, §9
- [Security & Compliance Datasheet](#) — for §13 RFP responses
- [DPA + Exhibit C](#) — for §12 continuity questions