AI · TEMPLATE

Adding AI Features Without *Breaking* Your Compliance Posture

KAANSYSTEMS.COM/LIBRARY/AI-FEATURES-COMPLIANCE-POSTURE · JULY 9, 2026

ABOUT THIS TEMPLATE

Bolting an LLM onto a regulated product opens four new ways for PHI to leave your compliance boundary. The pre-flight checklist that catches them before you ship.

THE TEMPLATE

Run this before every AI feature that could touch regulated data. Each item is yes / no / not-applicable. Any `no` on a required item blocks the ship until it's resolved.

Legal + vendor

- [ ] BAA signed with the model vendor (or PHI is de-identified before it leaves our systems)
- [ ] Vendor added to the public sub-processor list
- [ ] Vendor data-use terms reviewed: confirmed no training on our inputs
- [ ] Vendor's own retention of our data is documented and acceptable

Data flow

- [ ] Traced every field that goes into the prompt; confirmed the minimum necessary
- [ ] Confirmed no PHI in URL parameters, feature flags, or analytics events for this feature
- [ ] Response path checked: PHI in the output isn't written anywhere out of scope

Logging + retention

- [ ] Prompts + responses are NOT logged in plaintext to general application logs
- [ ] Any retained prompts/responses live in a scoped, encrypted, in-boundary store
- [ ] Retention period for prompt/response data is set deliberately, not left at default

- [ ] Observability traces for this feature strip or exclude prompt/response content

Output safety

- [ ] Consequential outputs have a human in the decision, not just a notification
- [ ] Users can tell an AI generated the content
- [ ] There's a documented plan for when the model is confidently wrong

Ongoing

- [ ] This feature is in scope for the next access review + evidence pipeline
- [ ] A re-review triggers when the vendor changes terms or we change the data flow

The checklist takes an hour to run for a new feature. The alternative — finding an undisclosed sub-processor and an out-of-scope PHI log during a customer's security review — takes a lot longer and costs the deal.

HOW TO USE IT

Run the pre-flight per feature, not per company. The BAA you signed for the summarization feature doesn't cover the new vendor you're using for embeddings. Each distinct model vendor and each distinct data flow gets its own pass.

Watch the vendor terms. AI providers update their data-use and retention terms far more often than your database host does. Put a quarterly reminder on the calendar to re-read the terms for every model vendor on your sub-processor list. A default that flips from "no training" to "training unless you opt out" is the kind of change that doesn't send you an email.

The teams that get this wrong aren't reckless. They're moving fast on a feature that looks like every other feature and isn't. The pre-flight is cheap precisely because it turns an invisible compliance change into a visible, answerable checklist — before the feature ships, not after the audit finds it.