Product Policy Card: Embedded AI Features

Scope: Every product, cross-cutting. This card covers the AI features inside software that is not on the Approved Tool List as an AI tool: videoconference assistants, PDF tools, scheduling apps, e-signature platforms, browsers, document management systems, and email clients.

Way of working: Not applicable. An embedded feature is whatever the host product makes it; that is the problem.

Status: The rule is live now. This card implements policy text that is already in force (42/02, Section 3): an embedded AI feature is a tool, and it stays off for client work until vetted.

Implements: Model Firm AI Use Policy (42/02), Section 3. No vendor exhibits; this card states firm rules, not vendor facts.

Document date: 2026-08-07.

Why this card exists

AI features now arrive inside software the firm already trusts for something else. The Zoom account the firm bought for meetings now offers a meeting summarizer. The PDF tool offers a document chat. Each of these features sends content to a vendor's AI service under terms nobody at the firm has read, because nobody bought it as an AI tool. The trust the firm extends to the host product does not extend to the AI feature; the feature is a new tool with a new data path, and the policy treats it as one.

This card is written to be used at the moment of temptation: the button is right there, the meeting is ending, and the summary would be convenient.

The rules

1. **Do not turn on an embedded AI feature for client work.** If the feature is not a row on the Approved Tool List, it is off. This applies to features that are on by default: turn them off.
2. **Treat a meeting assistant as a participant.** Do not let an AI notetaker or summarizer join a client call, a deposition, a mediation, or any privileged conversation. It records; where the recording goes is governed by terms the firm has not vetted. This includes assistants that other parties bring: if an unknown notetaker joins, ask whose it is and object on the record if it stays.
3. **Do not chat with a client document inside a PDF tool or browser.** The chat sends the document to the vendor's AI service. That is a Tier 2 input into an unvetted tool, which the policy's Tier 3 rule prohibits.
4. **Ask before you assume a feature is safe because the host is.** The firm's document management system, email client, and research platforms all now bundle AI features. The host being approved for its original purpose does not approve its AI feature.
5. **Propose, do not improvise.** Anyone may propose an embedded feature for vetting: send the Administrator the product name, the feature, and why it would help. The Administrator runs the Section 3 vetting (terms, training, retention, account binding, security), and an approved feature becomes a row on the Approved Tool List like any other tool.
6. **When in doubt, it is off.** The cost of asking is one email. The cost of guessing wrong is a client's information in an unvetted vendor's AI pipeline.

What "off" looks like

The Administrator maintains, as part of onboarding and the annual review, a short checklist of the firm's installed software with known AI features and where each off switch lives. New software joins the checklist at purchase. This keeps rule 1 executable: nobody can turn a feature off if nobody knows it shipped.

Re-vetting triggers

Host products add AI features in ordinary updates, without asking. At each annual policy review, and whenever a host product announces an AI feature, the Administrator re-checks the checklist above. A feature that was off can come back on after an update; the per-machine check is part of the review.