Product Policy Document: Lexis+ AI / Protégé (platform): Conditional, Restr...

Product and vendor: Lexis+ AI and its assistant Protégé, LexisNexis. Arrives alongside the Lexis research subscription.

Way of working: Platform. Vendor-hosted legal AI; the vendor picks the models; documents go into the vendor's store (Vaults).

Status: Conditional, and the slowest of the group. Two gaps set it apart: LexisNexis publishes no provider-level zero data retention claim at all, and vendor-side personalization (Protégé "remembering" the user's work) is the flagship feature, on by design. Tier 1 use only, and only with personalization off, until a signed agreement supplies what the published layer does not.

Implements: Model Firm AI Use Policy (42/02). Exhibits: 04 (LexisNexis General Terms, retrieved 2026-08-05); 05 (CoCounsel and Protégé postures memo, 2026-08-05); 11 (Wave 1 exhibits pass memo).

Document date: 2026-08-07.

1. What this is

Lexis+ AI is LexisNexis's legal AI platform; Protégé is its personalized assistant layer. Research, drafting, and document work run against LexisNexis content and the firm's uploads, which live in platform "Vaults" (up to 50 vaults of 1 to 500 documents each; vault contents retained until the user deletes them; small uploads outside a vault purged at session end; vault results kept 90 days in "My Conversations") (Exhibit 05).

Two published facts require the restricted posture. First, **the missing claim:** every other vendor in this workstream publishes something about what its model providers may keep; LexisNexis publishes nothing about provider-level retention. Routing is described as multi-model on AWS infrastructure, OpenAI named, and whether LexisNexis holds ZDR terms with its providers is simply not stated anywhere public (Exhibit 05). Second, **the carve-out:** the no-training pledge ("we never use customer data to train AI models") sits next to "identifiable information is removed from AI interactions so that we can improve performance," which is a de-identified service-improvement use of firm interactions. Harvey and Legora bar in contract what this permits (Exhibit 05).

And the finding worth keeping whether or not the firm ever buys Protégé: **LexisNexis's own General Terms define the standard.** The GTC prohibits users from submitting Lexis content into third-party AI tools unless the tool is a "Compliant AI Tool": one with zero data retention enabled such that nothing is logged, stored, or retained beyond the active session, and an express no-training term by contract or published terms (Exhibit 04, GTC section 1.1). That is precisely the standard the firm should demand of every AI vendor, including LexisNexis, and on the published record Protégé would not obviously satisfy its own publisher's definition.

2. The Approved Tool List row (implements Section 3)

Not fillable yet. The template row, to complete from signed paper only:

Lexis+ AI / Protégé (platform). Tier / account type: [Firm subscription; order form reference] | **Training disabled:** [From signed terms; published pledge carries a de-identified improvement carve-out] | **Memory off / scoped: Personalization off org-wide is a precondition, not a preference** | **Retention terms:** [From signed terms: provider-level ZDR or stated retention; Vault and conversation retention; deletion at termination] | **Permitted input tier: Tier 1 until signature; permitted tier set at approval** | **Agentic capabilities:** [Enumerate at vetting] | **Approved date:** Pending | **Notes:** No published provider-ZDR claim as of 2026-08-07; hold the vendor to its own Compliant AI Tool definition

3. Required setup (implements Section 3): the ask list

1. **Open with their own words.** Ask LexisNexis to confirm, in the agreement, that Lexis+ AI and Protégé meet the GTC's own "Compliant AI Tool" definition (Exhibit 04): zero data retention beyond the active session at the model-provider level, and an express no-training commitment. If the vendor's paper cannot satisfy the vendor's published standard, that answer is itself the vetting result.
2. **Provider-level retention, in writing:** which model providers serve the firm's requests, what each may retain, and whether ZDR terms exist for every route. This is the gap; do not let the platform-side retention schedule answer a provider-side question.
3. **Close the improvement carve-out for firm content:** no use of the firm's interactions, de-identified or otherwise, to improve the service, or a written statement of exactly what "improvement" processing occurs.
4. **Personalization off, org-wide, as a configuration the firm controls:** Protégé's memory is the policy's Section 4 cross-matter vector as a flagship feature. The agreement or admin documentation must show an organization-level off switch, not a per-user preference. If personalization is ever wanted, it needs matter-scoped configuration recorded on the row, per the policy.
5. **Retention schedule in the agreement:** Vault retention, the 90-day conversation retention, deletion mechanics, deletion at termination with certification.
6. **The DPA and subprocessor list** (third-party reporting lists Anthropic among subprocessors, unverified; if true, the covered-model question applies and must be answered) (Exhibit 05).
7. **Pricing-model questions** per policy Section 3, vetting item 6.
8. **File everything, complete the row, then move the status.**

4. What may go in (implements Section 4)

Until signature: Tier 1 only, with personalization off for any firm user who touches the product at all. No client documents into Vaults. No matter facts beyond the Placeholder Card's residue.

This product's specific hazard is memory sold as the feature. Everywhere else in the package, vendor memory is a setting to disable; here it is the product's selling point, on by design and opt-out (Exhibit 05). A tool whose value proposition is remembering your work across sessions is, for a multi-client practice, a tool whose value proposition is the cross-matter contamination vector. It may still be adoptable with the personalization architecture understood, scoped, and controlled; it cannot be adopted by treating the memory as a chatbot nicety.

5. Verification (implements Section 5)

Unchanged: the ladder, the statute check, the no-delegation rule. Lexis-grounded output gets the same four steps; Shepard's living inside the product is a convenience for step 4, not a substitute for doing it.

6. The record (implements Sections 8 and 10)

Platform architecture, with disclosed granularity: Vaults until deleted, session uploads purged at session end, conversations 90 days (Exhibit 05). The matter-file export habit applies. Note the interaction with the personalization question: whatever Protégé "remembers" is also a record the firm does not hold; part of Section 3, step 4 is learning where that memory lives and how it is purged.

Billing: subscription overhead; pass-through rules if the order form meters anything.

7. Vendor facts and evidence grades

"We never use customer data to train AI models," with de-identified improvement carve-out. Source: Exhibit 05 (product page, retrieved 2026-08-05) | **Grade:** Published (marketing) | **Date:** 2026-08-05

Vault and retention specifics (50 vaults; until-deleted; session purge; 90-day conversations). Source: Exhibit 05 | **Grade:** Published | **Date:** Same

Personalization flagship, on by design, opt-out. Source: Exhibit 05 | **Grade:** Published | **Date:** Same

No provider-level ZDR claim published. Source: Exhibits 05, 11 | **Grade:** Absence verified on the published record | **Date:** Confirmed 2026-08-05

GTC "Compliant AI Tool" definition (ZDR plus no-training as conditions for third-party AI use of Lexis content). Source: Exhibit 04, GTC section 1.1 | **Grade:** Contract text (LexisNexis's own published terms) | **Date:** Retrieved 2026-08-05

Anthropic among subprocessors. Source: Third-party reporting only | **Grade:** Unverified; ask at vetting | **Date:** Noted 2026-08-05

Not established: provider-level retention terms (the central gap); the personalization control architecture; the subprocessor list. All are Section 3 asks.

8. Re-vetting triggers (implements Section 12)

1. **Signature**, or the vendor's answer to the Compliant AI Tool question (either way, the answer updates this document).
2. **A provider-ZDR claim appears** on the published layer (check at each annual review; its arrival would move Protégé's rank).
3. **Personalization defaults or controls change.**
4. **The GTC's Compliant AI Tool definition changes** (Exhibit 04's source; it is the firm's benchmark clause and its dilution would be worth knowing about in its own right).
5. **Confirmation of Anthropic routing** (raises the covered-model question).
6. **Pricing model changes.**