

AI VENDOR DILIGENCE. MEMO

# Lexora AI Triage

---

Prepared for partner@example-legal.com.au

12 June 2026 · Confidential

RECOMMENDATION

MAYBE

*Could be the right answer; the pitch did not surface enough to be sure.*

# Strip the marketing.

## What it actually does

Strip the marketing and Lexora is a SaaS inbox sitting alongside your existing email. An LLM (the pitch does not say which model; the customer success deck shows it is GPT-4o) reads incoming messages, classifies them into categories you define, drafts a suggested reply, and surfaces a queue your team works through. The classification model is not proprietary; the value is in the inbox UI, the audit trail, and the human-in-the-loop workflow.

## Usefulness for you

For your stated use case (cutting an hour of daily triage on around fifty support enquiries) Lexora is in the right shape. The classification and draft-reply pattern is well suited to the kind of repetitive routing you described, and the human-in-the-loop step preserves the partner sign-off your practice needs. The risk is overpaying for the wrapper, not that the wrapper does not work.

## Red flags

- The pitch claims 'proprietary AI'. They are routing your data through GPT-4o with a system prompt. That is not a disqualifier but the pricing should reflect it.
- Quoted price (\$24,000 a year, twelve-month minimum) is roughly five times the raw API cost for fifty messages a day. You are paying for the UI and the audit trail.
- Pitch does not name where data is stored or whether it is used to train the underlying model. Your concern about the dodged data-storage question is well placed.
- The demo used dummy data. Insist on running your own representative messages through a paid trial before signing the annual commitment.

## RECOMMENDATION

# Maybe

Could be the right answer; the pitch did not surface enough to be sure.

## RATIONALE

*Maybe is the call. The product shape fits your use case, the time saving is real, and the human-in-the-loop pattern preserves the partner sign-off you need. The reason this is not a Go is the gap between the pitch ('proprietary AI', dodged data-storage question, dummy-data demo) and what the tool actually is. That gap matters because it shapes the price negotiation and the data risk profile. A two-month paid pilot, at one twelfth of the annual price, with your real messages, will turn this into a clear Go or a clear Stop. Refuse to sign the annual commitment until that pilot has run and the data-storage question is answered in writing.*

# Before you sign anything.

---

## Questions to put to the vendor in writing

Send these as one email. Watch how fast they answer, and which ones they answer with specifics.

- 01 Which underlying model does Lexora use, is it OpenAI or Anthropic, and what is your contractual position with that provider on data retention and training.

---

- 02 Where is our data physically stored, in which jurisdiction, and is it ever processed outside Australia.

---

- 03 What happens to our data, queues, and audit trail when our subscription ends.

---

- 04 Show me two customers running roughly fifty messages a day in a regulated professional services context, who have been live more than six months.

---

- 05 What is the actual classification accuracy on first pass on a customer dataset similar to ours, and what does the human-in-the-loop edit rate look like at month one versus month six.

---

- 06 Will you offer a paid one-month pilot with our real messages, at a price that is roughly one twelfth of the annual quote, before we commit.

---

- 07 If we cancel after the pilot, what specifically happens to the model, to our prompts, and to our audit trail.

---

---

## Worth considering instead, or alongside

### A. Two-month paid pilot with Lexora, then revisit

Lowest-friction path. You get the real read on classification accuracy at one twelfth the cost. The trade-off is you have signalled buying intent, which weakens your negotiating position at month two if you do proceed.

### B. Build a thin internal tool on top of the OpenAI or Anthropic API

Two to four weeks of a junior engineer's time and roughly \$2,000 a year in API costs. You own the model choice and the audit trail completely. The trade-off is you also own the UI and the on-call when it breaks at 4pm on a Friday.

### C. Stay with the current process for one more quarter

Cost is the hour a day you are already losing. Benefit is that the AI tooling market in this niche will be meaningfully different in three months. The trade-off is the hour a day compounds.

## IMPLEMENTATION

If you want a second opinion before you commit, or help running the vendor through the questions above, email [dpexconsulting@gmail.com](mailto:dpexconsulting@gmail.com) with the subject "Vendor diligence" and a copy of this memo attached. I read every message myself.

*Daniel Xuereb*

DPEXCONSULTING · SYDNEY