

# Learning-AI Build-vs-Buy Worksheet

Run this gate for each AI feature before you commit. The drivers, not the model leaderboard, pick the lane. Convert every option to cost per learner per month before deciding.

---

## A. The feature and the option

- We named the AI feature and what learner data it touches
- We considered all three options: buy a product, call a managed API, self-host an open model
- For a new or unproven feature, we start on a managed API to validate before optimizing

## B. The four drivers

- Usage volume: learners, minutes, or messages per month is estimated
- Accuracy bar: how wrong the feature is allowed to be is defined
- Data sensitivity: whether the learner data can legally leave our walls is answered
- Lock-in tolerance: how painful switching providers later would be is assessed

## C. Cost per learner per month

- We converted seats, minutes, credits, and tokens to one unit: cost per learner per month
- We compared that figure to the engineering cost of self-hosting (GPU + ML/infra time)
- Self-hosting is only chosen past a high, steady-volume break-even — not to 'save fees'

## D. Privacy, lock-in, and tracking

- A GDPR Article 28 data processing agreement is in place; the no-training clause is explicit
- Sub-processors, retention, and (for US records) FERPA constraints are checked
- If the feature evaluates learners, EU AI Act high-risk duties are accounted for
- A thin abstraction layer is in place; the feature emits xAPI statements to our analytics