

8 Questions Your L&D AI Vendor Is Not Expecting

Use these before your next demo, contract renewal, or vendor selection.

A real answer is specific, verifiable, and references documentation. A dodge is vague, redirects to marketing language, or puts the burden of governance entirely on you.

A REAL ANSWER SOUNDS LIKE

A DODGE SOUNDS LIKE

1

What training data was used to build the model that generates or recommends our learning content?

“Our content model was trained on [named corpus]. Here is our data card with provenance, date range, and known limitations.”

“We use the latest AI technology and our proprietary algorithms ensure quality.”

2

Can you show us a bias audit for how the system recommends learning paths across different employee demographics?

“Here is our most recent audit report, conducted [date] by [auditor]. It tested for disparate impact across [dimensions].”

“Our platform treats all learners equally. We do not discriminate.”

3

If the system scores an assessment incorrectly or recommends the wrong learning path, what is the override process?

“L&D admins can override any recommendation or score in [specific workflow]. Here is the override log and escalation path.”

“Our accuracy rate is over 95%, so overrides are rarely needed.”

4

What happens to our employee learning data after the contract ends? How long is it retained and who can access it?

“Data is returned or deleted within [X days] of contract termination. Here is our data processing agreement and retention policy.”

“We comply with all applicable data protection regulations.”

Questions 5-8

Use the answers to separate evidence from reassurance.

A REAL ANSWER SOUNDS LIKE

A DODGE SOUNDS LIKE

5

Does the system make any decisions the EU AI Act would classify as high-risk? How are you preparing for compliance?

“Our skills assessment module falls under Annex III, Category 4. Here is our conformity assessment timeline and our designated point of contact.”

“We are monitoring regulatory developments and will update our platform accordingly.”

6

Can a learner see why the system recommended a specific course or flagged a skill gap? Is this explanation auditable?

“Learners see a plain-language explanation on every recommendation. Admins can access the full decision log with weighted factors.”

“Our AI is transparent and we are committed to explainability.”

7

Who is liable if AI-generated compliance training content contains an error and an employee acts on it?

“Section [X] of our agreement addresses content liability. AI-generated compliance content is flagged for human review before publication.”

“Content is generated by AI but the customer is responsible for reviewing all materials.”

8

How do you test for hallucinated content in AI-generated courses, and how often do false claims make it to production?

“We run [specific validation process]. Our false-positive rate for factual claims is [X%], measured quarterly. Here is the latest report.”

“Our models are highly accurate and continuously improving.”

THE PATTERN TO WATCH FOR

A real answer names specifics: dates, documents, processes, limits. A dodge uses words like “committed to,” “continuously improving,” or “industry-leading” without pointing to anything you can verify.