FROM THE APPENDICES OF DEFENSIBLE AI

The 12 Supplier Questions

Due-diligence questions to put to any AI vendor before, and after, you sign. If a supplier cannot answer one, that is where your Liquid Evidence will come from.

---

01 Who is the named human owner of this system's decisions, on your side and on ours?

Why it matters: accountability that lives in a RACI nobody reads is not ownership. A name, with authority.

02 What evidence is captured at the moment of decision: inputs, model version, logic path, and context?

Why it matters: this is the Evidence Loom. If it isn't captured at decision time, it cannot be reconstructed later.

03 Can you replay any individual decision from six months ago, end to end, today?

Why it matters: the reconstruction test. Ask for a live demonstration, not a slide.

04 How are model updates versioned, and how do you preserve the ability to defend decisions made by previous versions?

Why it matters: a model update that orphans past decisions converts your entire history into Liquid Evidence.

05 How do you detect, measure, and report drift, and what thresholds trigger action?

Why it matters: this is the Intelligent Loop. Decay caught late is a failure already in progress.

06 Is there a kill-switch, who holds it, how fast does it act, and when was it last tested?

Why it matters: the Liability Kill-Switch. The authority and the mechanism to halt a decision before it becomes irreversible.

07 What are this system's V, H, and C scores by your own assessment, and what is the evidence for the C?

Why it matters: a supplier who cannot score their own system has not thought about defensibility. Validate the claimed C yourself.

08 What audit rights do we have: logs, model documentation, evaluation results, and incident history?

Why it matters: contestability you cannot inspect is a promise, not a control.

09 Where does liability sit when the system is wrong, and what does your contract actually say about algorithmic harm?

Why it matters: indemnities written for software outages rarely survive contact with an AI liability claim.

10 What data was this model trained on, what are the known limitations and failure modes, and where is that documented?

Why it matters: undocumented limitations become your operational surprises and your regulator's findings.

11 How is human oversight designed into the workflow, and what prevents it being quietly optimised away?

Why it matters: John's million-pound transfer went through because a review step was removed for "frictionless" completion.

12 If we exit, what happens to the evidence: decision records, logs, and the ability to defend historic decisions?

Why it matters: defensibility must survive the contract. Exit without the evidence is exit into liability.

Scoring tip: treat each unanswered question as a deduction from the supplier's claimed Contestability. Then run R = (V × H) / C before you sign.

From **Defensible AI: How to Survive the Era of Liquid Evidence** by Aziz Ahmed · Amazon UK & US · Live calculator at www.defensibleai.net