

Seven Questions on AI Dependency

For boards and risk committees overseeing consequential AI systems.

If the answer to question four requires a special investigation, the operational answer is no.

1. Which models in consequential workflows are supplied rather than fully controlled, and what is the contractual and practical notice period for each?
2. For each, what is the measured elapsed time from a deprecation notice to a validated replacement in production?
3. Where the institution runs its own models, what is the supported life of the inference stack - runtime, libraries, toolchain and hardware - and who tracks it?
4. Can the institution produce, for any decision in the last twelve months, the specific model and version that produced it?
5. What is the defined behavior when a model component is unavailable or materially changed?
6. Which controls protecting customers or patients are enforced outside the model?
7. Who owns the replacement path for each dependency, and has the path been exercised?