

AI Feature Release Readiness Scorecard

A 10-question check for teams changing a prompt, model, retrieval setup, or tool workflow. It is a conversation starter, not a safety certification or legal advice.

How to use it

For the next customer-facing AI release, score one point for every **Yes**. Mark an owner or note beside every No. A useful outcome is not a perfect score; it is one specific decision that becomes reviewable before release.

#	Question	Yes	No / owner / note
1	Is there one named owner for the next release decision?		
2	Can the team name the user task that the AI feature must complete?		
3	Do you have examples from normal traffic and known failures?		
4	Are unacceptable or critical failures explicitly defined?		
5	Are model, prompt, tool, and retrieval versions recorded?		
6	Can at least some behavior be checked deterministically or by execution?		
7	Does the team track latency and cost alongside task quality?		
8	Is the expected human-review volume acceptable and owned?		
9	Can the feature be limited, disabled, or rolled back?		
10	Can a new incident become a redacted regression case?		

Score and choose one next move

Score	Recommended next move
0-3	Define one narrow release decision, a named owner, and unacceptable failures before adding more tooling.
4-6	Turn useful ingredients into a versioned test set, documented versions, and explicit thresholds.
7-8	Build a repeatable gate and make the exception, escalation, and rollback path explicit.
9-10	Keep calibrating with incident learning, reviewer feedback, and controlled change management.

Use this responsibly

Use synthetic, approved, or properly redacted cases. Keep secrets and unnecessary personal data out of test fixtures. This scorecard does not guarantee safety, compliance, or product outcomes; engineering, security, legal, and operational decisions remain with your team.