

10-Minute Self-Check

A fast self-check for identifying the biggest production risks in an AI system. Nine areas, one question each. Score honestly — the value is in what you cannot answer.

■ Green — acceptable for current stage ■ Yellow — risk exists, needs improvement ■ Red — serious production risk

AREA	QUESTION	G	Y	R
Architecture	Can we explain the full AI workflow end-to-end, including what happens on failure?			
Reliability	Do we have timeouts, retries, and fallback behaviour for every LLM call?			
Observability	Can we trace one user request through the whole system, including the LLM call?			
Evals	Do we know if output quality gets worse after prompt, model, or retrieval changes?			
RAG & Context	Do we know why specific context was selected and included?			
Cost Control	Do we know the cost per request, workflow, or user?			
Security & Privacy	Can users only access data they are allowed to see?			
Deployment	Can we roll back application, prompt, model, and configuration changes?			
Governance	Is there a named owner and a review process for AI behaviour changes?			

Interpretation

- Mostly green** The system has a healthy production foundation. Continue with the full checklist to find the non-obvious risks in each area.
- Several yellow** The system may work, but production risk exists. Prioritise the yellow areas before usage grows or access widens.
- Any red** Treat this as a serious risk. Red in security, observability, reliability, or deployment should be addressed before scaling the system.

Scored it? Record the full review on the next page — one line per area, so the result survives the meeting it was discussed in.

Production Readiness Scorecard

One page to summarise a full review. Use it after working through the checklist areas, not instead of them.

AREA	SCORE	KEY FINDING	PRIORITY
Architecture			
Reliability & Failure Handling			
Observability			
Evals & Quality			
RAG & Context Design			
Cost Control			
Security & Privacy			
Deployment & Release			
Governance & Risk			

Summary

OVERALL READINESS — LOW / MEDIUM / HIGH

THREE BIGGEST BLOCKERS

RECOMMENDED NEXT STEP

This is the working version of a larger open framework. The full checklist — nine detailed areas, discovery questions, and an audit report template — is public at github.com/manuelblinkert/ai-production-readiness-checklist. Use it, fork it, adapt it to your stack.

Written by Manuel Blinkert, AI Systems Architect. If the self-check surfaced more red than you expected and you want the same review run properly on your system, that is what I do: manuelblinkert.com.