

AI Risk Audit — 30-Minute Agenda

Exactly what happens on the call (no fluff, decision-ready output).

1) Pick one production workflow (5 min)

- What the AI does and where it runs (product, ops, support, finance, etc.)
- Who is affected and what “failure” means in dollars or risk

2) Define pass/fail (8 min)

- Success criteria (accuracy, latency, safety, compliance)
- Current metrics (or what you wish you had)

3) Identify top failure modes (10 min)

- Hallucinations, policy drift, tool misuse, escalation hotspots
- Where logs/evals are missing (blind spots)

4) Map the 10-day Reliability Sprint (7 min)

- Eval rubric + test suite
- Failure-mode heatmap
- Prioritized remediation plan + go/no-go checklist

Outcome

You leave with a clear Sprint plan, a measurement framework, and a realistic path to shipping AI with confidence.