

Responsible AI Readiness Checklist

Technology creates potential. People create value.

Adopting AI responsibly is a transformation challenge, not only a technical one. It succeeds through the same governance, ownership, and human readiness that any meaningful change requires. This checklist is where that conversation starts.

BEFORE YOU BEGIN

This is not a compliance audit or a pass-or-fail test. It is a practical way to see where your organization already has a foundation, and where the real work of governance still needs to happen. Click each box that is already true, or print this page and check it by hand. The boxes you cannot check are not a failure — they show you exactly where to focus next.

1 Transparency

People should understand when AI is involved in a decision that affects them, and how it works, in plain language.

- ☐ We disclose to users when AI is involved in a decision that affects them.
- ☐ We can explain, in plain language, how our AI systems reach their outputs.
- ☐ A human review option exists for high-stakes AI decisions.
- ☐ Our teams apply their own judgment instead of defaulting to what the AI says.

2 Accountability

Every AI system needs a named human owner. Responsibility cannot be delegated to the algorithm.

- ☐ Every AI system we use has a named human owner.
- ☐ That owner has the authority to pause, modify, or retire the system when needed.
- ☐ AI-related incidents are escalated and resolved within a defined timeframe.
- ☐ Ownership is written down somewhere, not just known informally.

3 Fairness and Non-Discrimination

AI systems should not produce outcomes that disadvantage people based on protected characteristics.

- ☐ AI systems used in decisions about people are tested for bias before launch.
- ☐ Bias monitoring continues after deployment, not only at launch.
- ☐ We have defined fairness metrics for systems that affect people.
- ☐ A disparate impact finding would trigger a real review, not just a note in a file.

4 Privacy and Data Stewardship

AI systems should collect only the data they need, with clear consent and real protection.

- ☐ We know what personal data feeds into each AI system we use.
- ☐ A privacy impact assessment has been completed for systems using personal data.
- ☐ Data retention limits are defined and actually enforced.
- ☐ Consent for AI use of personal data is documented where required.

5 Human Oversight and Control

People should stay in meaningful control of consequential decisions. AI supports human judgment, it does not replace it.

- ☐ High-stakes AI-assisted decisions require a human review step before action is taken.
- ☐ Reviewers have the time, information, and authority to meaningfully evaluate AI output.
- ☐ Override mechanisms exist and are actually used when something looks wrong.
- ☐ We track how often humans override the AI, and review why.

Reading your results: A mostly-checked list means you already have a real foundation to build on. A mostly-unchecked list is common, and simply shows you where to start. Either way, the next step looks the same: turn these checkboxes into documented governance, clear ownership, and a readiness plan that carries through implementation and ongoing oversight.

TOP THREE PRIORITIES

Before you close this page, name the three items that matter most for your organization right now.

1. _____
2. _____
3. _____

Ready to see how your organization compares?

Nancy Constable helps organizations turn checklists like this into documented governance, clear ownership, and a readiness plan for responsible AI adoption. Reach out at peakmomentumadvisory.com to start the conversation.