Responsible AI Framework

AI & Governance | Applied adaptation | Portfolio version 1.0

A product-management interpretation of Govern, Map, Measure, and Manage for accountable AI-enabled workflows.

Problem it solves

Teams need a practical way to translate broad responsible-AI principles into product decisions, requirements, oversight, and ongoing operations.

When to use it

- AI affects users, decisions, access, safety, privacy, or trust.
- The workflow has meaningful failure consequences.
- Multiple functions share AI risk responsibilities.
- The product needs clear review, escalation, and incident processes.

When not to use it

- As a substitute for legal, regulatory, security, or compliance review.
- As a one-time launch checklist.
- When accountability and decision ownership are intentionally undefined.

Inputs

- Intended use and users
- Affected stakeholders
- Data and model documentation
- Potential benefits and harms
- Applicable policies and obligations
- Operational and incident context

Process

1. Govern - Define accountability, acceptable use, policies, roles, documentation, and decision rights.

2. Map - Understand context, users, impacts, data, dependencies, foreseeable harms, and misuse.

3. Measure - Evaluate quality, fairness, privacy, explainability, robustness, security, and human outcomes.

4. Manage - Prioritize and mitigate risks, monitor production, escalate incidents, communicate, and improve.

Decision points

- Who owns the product and risk decisions?
- Who can benefit or be harmed?
- What happens when the system is uncertain or wrong?
- Which risks are acceptable, mitigated, transferred, or avoided?
- What triggers escalation, suspension, or retirement?

Outputs

- AI use policy
- Impact map
- Risk register
- Evaluation and assurance plan
- Human-oversight requirements
- Incident and escalation plan
- Monitoring review

Success measures

- Risk coverage
- Evaluation coverage by segment
- Override and escalation rates
- Incident frequency and severity
- Time to detect and resolve
- User understanding
- Mitigation effectiveness

Portfolio application

I use this structure to turn responsible-AI concerns into practical product questions, acceptance criteria, review paths, monitoring requirements, and governance conversations.

Common pitfalls

- Treating responsible AI as a compliance document only.
- Using generic principles without workflow consequences.
- Ignoring affected non-users.
- Failing to monitor after launch.
- Confusing explainability with guaranteed correctness.

Skills demonstrated

- AI governance

- Risk framing
- Human oversight
- Evaluation
- Cross-functional leadership
- Incident readiness

Interview discussion prompts

- How do you operationalize responsible AI?
- How do you identify affected stakeholders?
- What happens when AI is uncertain or wrong?
- How do you decide whether a risk is acceptable?

References and attribution

NIST AI Risk Management Framework
The Govern, Map, Measure, and Manage functions informed this applied product-management interpretation.
https://www.nist.gov/itl/ai-risk-management-framework