

Seven Decisions Before Scaling AI

An executive guide to responsible, valuable adoption

Move from promising pilots to governed AI capability with clear ownership, evidence, human oversight, and operating discipline.

01

Executive summary

Responsible AI is not a review at the end. It is a set of operating decisions made before scale increases the consequences.

● Value

Name the business outcome and the evidence required to justify continued investment.

● Ownership

Assign accountability for the product, the decision, the data, and the residual risk.

● Boundaries

Define unacceptable uses, required human judgment, and release conditions.

● Learning

Monitor value, performance, drift, incidents, and adoption as one system.

Use this guide before procurement, before pilot expansion, and whenever a use case changes data, users, workflow, or consequence.

The seven decisions

Leaders should be able to answer each decision in plain language before an AI use case reaches consequential scale.

-
- | | | |
|----|---------|---|
| 01 | Outcome | What business or stakeholder result justifies using AI? |
|----|---------|---|
-
- | | | |
|----|----------------|---|
| 02 | Accountability | Who owns the product, decision, and effect? |
|----|----------------|---|
-
- | | | |
|----|------|--|
| 03 | Data | What data may be used, and under which conditions? |
|----|------|--|
-
- | | | |
|----|------|--|
| 04 | Risk | Which consequences are unacceptable regardless of value? |
|----|------|--|
-
- | | | |
|----|-----------------|---|
| 05 | Human oversight | Where must people review, intervene, or remain responsible? |
|----|-----------------|---|
-
- | | | |
|----|-----------|---|
| 06 | Operation | How will the use case be supported, monitored, and changed? |
|----|-----------|---|
-
- | | | |
|----|----------|--|
| 07 | Evidence | What proof is required before expanding access or scope? |
|----|----------|--|
-

Govern the complete use case

Risk follows the workflow and consequence, not only the model. Review the entire path from input to action and feedback.

Purpose	Decision supported, expected benefit, affected stakeholder
Input	Source, permission, quality, sensitivity, representativeness
Model	Capability, limitation, testing, change control, supplier dependency
Workflow	User, context, automation level, exception handling
Action	Decision or output, reversibility, consequence, communication
Feedback	Monitoring, appeal, incidents, drift, learning, retirement

A familiar model can create high risk in a consequential workflow. A novel model can be manageable in a bounded, reversible use case.

A proportionate control model

Match the strength of controls to consequence, reversibility, scale, and uncertainty.

EXPLORE

- Synthetic or low-sensitivity data
- Small expert user group
- No consequential automation
- Explicit learning questions

PROVE

- Named product owner
- Documented testing and limitations
- Human review and fallback
- Value and risk instrumentation

SCALE

- Formal release evidence
- Operational monitoring and support
- Incident and change management
- Periodic assurance and retirement criteria

Stages should have evidence gates. Time spent in a pilot is not evidence that the use case is ready to scale.

Leadership workshop questions

Use these questions to reveal unresolved operating decisions before technical progress creates pressure to launch.

01

Why is AI the right mechanism for this outcome?

02

Who accepts the residual risk if the use case performs as designed but causes harm?

03

Which users or stakeholders could be disadvantaged?

04

What data use would violate expectation, permission, or policy?

05

Where can a person meaningfully intervene?

06

How will performance and drift be detected in the real workflow?

07

What event should pause or withdraw the use case?

08

What evidence would justify broader access or automation?

A 30-60-90 day readiness plan

Establish the operating foundations while focusing on a small set of bounded, valuable use cases.

FIRST 30 DAYS

- Define ownership and principles
- Inventory active use cases
- Select a risk classification
- Identify data and supplier dependencies

DAYS 31-60

- Apply the seven decisions
- Create stage and evidence gates
- Instrument value and risk
- Design incident and change paths

DAYS 61-90

- Review pilot evidence
- Approve, reshape, or stop use cases
- Publish operating responsibilities
- Set the next learning agenda

Responsible scale is the result of repeatable decisions and evidence, not a one-time policy approval.