


EXECUTIVE GUIDE



A CXO Guide to Sequencing AI Use Cases

A practical guide for choosing the first production AI use case by readiness, decision value, and operating fit.

Mantra Labs

Research and advisory team / Published 2026-06-04 / mantra.ai



HOW TO USE THIS PUBLICATION

A decision aid, not a substitute for operating context.

A practical decision framework informed by current AI adoption, workflow redesign, and risk-management research. Validate scoring thresholds against your own economics and governance requirements.

Use the models to structure a leadership conversation, identify the first workflow or system boundary worth examining, and define what evidence would be needed before committing to implementation.

Contents

- 01 What it means
- 02 Warning signals
- 03 Diagnostic model
- 04 Why it matters
- 05 Diagnostic map
- 06 Cost of debt
- 07 How to reduce it
- 08 How Mantra helps

01 PRINCIPLE

Start with the decision, not the model.

Most AI prioritization starts with technology possibility. Production value starts with a decision that is valuable, frequent, owned, and measurable.

The first use case should teach the organization how intelligence enters the workflow, not just prove that a model can generate an answer.

If the decision path is weak, the AI use case will inherit the weakness.

02

WARNING SIGNALS

Do not start with a use case that lacks a workflow home.

No clear owner No role is accountable for accepting, rejecting, or acting on the recommendation.	No action path The output cannot create a task, trigger a workflow, or update a system.
No review model The team has not defined what should be automated, assisted, or escalated.	

03 SCORING MODEL

Score use cases across five readiness dimensions.

Decision value The decision has enough volume, margin, risk, or service impact to matter. What changes if this decision improves?	Context readiness The system can assemble the data, rules, history, and evidence required. What does the user need to trust the output?
Workflow fit The recommendation has a clear surface and action path inside daily work. Where will the recommendation land?	Governance boundary The organization knows what can be automated, assisted, reviewed, or escalated. What should never happen without review?
Feedback capture The system can measure whether the recommendation improved the outcome. How will the system learn after action?	

04 WHY IT MATTERS

The wrong first AI use case creates organizational drag.

A visible use case can attract attention but still fail if the workflow, data, and review model are not ready.

A well-sequenced use case creates a reusable production pattern: context, action, governance, feedback.

The first AI use case should create a system pattern the enterprise can reuse.

Use these signals to choose the first use case.

Decision value High volume, high cost, high risk, or high customer impact. Outcome metric, frequency, owner, and decision latency.	Context readiness Required data exists but needs better assembly or source authority. Systems, fields, rules, evidence, and freshness.
Workflow fit The user already makes the decision inside a defined work path. Surface, role, task, action, and exception path.	Governance Human review and audit trail can be clearly defined. Permissions, review rights, escalation, and evidence.
Feedback Outcome can be captured within weeks, not quarters. Accepted recommendations, overrides, outcome quality, and cycle time.	

Avoid use cases that look impressive but cannot prove operating value.

Demo-first use case Looks strong in a workshop but has no daily workflow owner.	Data-starved use case Requires context the enterprise cannot reliably assemble yet.
Control-heavy use case Needs governance that has not been designed into the workflow.	Low-frequency use case Does not generate enough usage or feedback to learn quickly.
No-measure use case Cannot show whether the decision improved.	

Sequence the first production use case in four moves.

01 Map one decision path	02 Design context and action
03 Build review and feedback	04 Measure before scaling

What the operating response usually requires.

Product & Experience Engineering Design the surfaces where customers, teams, and leaders can complete the work.	Data & Intelligence Engineering Turn fragmented signals into decision context, recommendations, and feedback loops.
Core Platform Modernization Modernize the system paths that let work move across products, teams, and channels.	Compliance Orchestration Build permission, evidence, control, and review into the way systems operate.

SOURCES & EVIDENCE

The evidence behind the publication.

External claims are grounded in the sources below. Mantra frameworks and recommendations are identified as applied interpretation rather than original quantitative research.

01 The state of AI: How organizations are rewiring to capture value

McKinsey & Company

Survey evidence linking workflow redesign with reported enterprise value from generative AI.

02 2026 Work Trend Index: Agents, human agency, and the opportunity for every organization

Microsoft WorkLab

Research on organizational readiness, workflow redesign, governance, and human-agent operating models.

03 Artificial Intelligence Risk Management Framework (AI RMF 1.0)

National Institute of Standards and Technology

A governance framework for mapping, measuring, managing, and governing AI risk.

Continue the work

Need to choose the right first AI use case?

Use constraint diagnostic to find the decision path where value, readiness, and adoption are strongest.

mantra.ai