

REPORT



The Enterprise Intelligence Readiness Report

A field report on how to assess whether the enterprise is ready to move intelligence from pilots into production decisions.

Mantra Labs

Mantra Labs Research / Published 2026-06-08 / mantra.ai



HOW TO USE THIS PUBLICATION

A decision aid, not a substitute for operating context.

A research-backed synthesis of external AI adoption and governance research, translated into a practical enterprise readiness model. It is not presented as original quantitative research by Mantra.

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 The gap
- 02 Where it breaks
- 03 What has to exist
- 04 Where to start
- 05 How to measure

01 THE READINESS GAP

Most enterprises are more model-ready than system-ready.

A model can be technically viable before the organization is ready to use it in production work.

Intelligence readiness depends on whether the workflow can receive a recommendation, the data can support it, the platform can route it, governance can review it, and users can adopt it.

The enterprise is not intelligence-ready until the decision path is ready.

Signal set

5 layers readiness model	6 signals diagnostic pattern	90 days first-move horizon
------------------------------------	--	--------------------------------------

The workflow is the first readiness test.

If the workflow is unclear, every AI use case becomes a sidecar. Users must translate output into action manually.

The right readiness question is: where will the recommendation land, and what will happen next?

Workflow readiness turns intelligence from an answer into an operating capability.

01 Decision moment	02 User owner
03 Action path	

03 FAILURE MODES

Six patterns usually explain why intelligence stalls.

The same constraints appear across industries: vague ownership, thin context, disconnected workflows, insufficient review, weak feedback loops, and low adoption readiness.

Ownership No one owns the decision that AI is meant to improve.	Context The model sees a prompt but not enough operating history.
Workflow The recommendation cannot trigger, route, or complete work.	Governance Review, permission, and evidence are added too late.
Feedback Outcomes are not captured after action.	Adoption The user does not trust or need the recommendation in daily work.

1 decision first production unit	4 signals minimum readiness test	0 sidecars target state
--	--	-----------------------------------

Assess readiness across five connected layers.

No single layer decides whether intelligence can scale. The constraint is usually the weakest connection between layers.

Workflow readiness Can the work receive and act on intelligence?	Data readiness Is the context accurate, timely, and decision-ready?
Platform readiness Can systems route the signal into the workflow?	Governance readiness Are review, permission, evidence, and audit paths built in?
Adoption readiness Will users trust and use the recommendation inside daily work?	

The first move should strengthen the lowest-readiness layer that blocks a high-value decision.

Readiness includes the right boundary for autonomy.

The enterprise needs to decide what should be automated, assisted, reviewed, or escalated before intelligence is scaled.

Governance is not only risk management. It is a design input that tells the system how far the recommendation can travel.

Autonomy should expand only where context, control, and outcome learning are already designed.

Permission Who can see, act, override, or approve.	Evidence What proof supports the recommendation.
Review Where human judgment remains required.	Auditability What the system records after the decision.

4 modes automate, assist, review, escalate	1 trail evidence path	Every decision needs feedback
--	---------------------------------	---

Evidence note. Treat governance as a design boundary for intelligent systems, not a final approval checkpoint.

Start where readiness and value intersect.

The right first use case is not always the most visible one. It is the one where the decision path is valuable, owned, and measurable enough to improve.

The first use case should prove the decision system, not the AI ambition.

Insurance servicing Use intelligence to route high-volume service requests when policy context and next actions are clear. Cycle time • containment • adoption	Financial review Use intelligence to assist fraud, credit, or compliance review where evidence and escalation are defined. Review time • exception rate • control adherence
--	---

- High decision volume
- Clear owner
- Available context
- Reviewable action
- Measurable outcome

Measure whether intelligence changes the work.

Model accuracy is not enough. Production value appears when work moves faster, decisions improve, users adopt, and governance remains intact.

Decision velocity

Cycle time	Queue time
Manual reconciliation	

Decision quality

Exception rate	Override rate
Outcome improvement	

Adoption

Usage by role	Accepted recommendations
Repeat use	

4 metrics minimum value set	1 loop outcome feedback	90 days evidence horizon
---------------------------------------	-----------------------------------	------------------------------------

The readiness architecture is a system, not a dashboard.

A production intelligence system connects signals, context, rules, recommendations, workflow actions, review boundaries, and feedback data.

01 Signals and source authority	02 Context model
03 Decision logic	04 Workflow surface
05 Review and audit trail	06 Outcome feedback

Evidence note. The architecture should make each recommendation usable, reviewable, and improvable inside the workflow.

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

Want to know where intelligence should move first?

Use the constraint advisory map to identify the lowest-readiness layer blocking the highest-value decision path.

mantra.ai