

The AI Transformation *Loop Workbook*

A six-page companion to the article
"Agentic AI Transformation Is Not a Program"



Four steps. One observability layer. Repeat forever.

HOW TO USE THIS WORKBOOK

- Print it. Fill it in. Not in one sitting.
- Page 2 today. Page 3 with your team this week.
- Pages 4–5 over the next 90 days.
- Page 6 when you come back for the next loop.

Assess

Where are you actually?

1. Run the Canvas.

15 questions. 8 to 10 minutes.

[askajay.ai/canvas](#)

2. Record your pillar scores.

From the Canvas result, write each pillar score (1–5) and identify the lowest.

Strategic Alignment	Is AI connected to your business strategy?	/ 5
Data Infrastructure	Can your data teams access what AI needs?	/ 5
Operational Processes	Do you know every AI system running right now?	/ 5
Talent & Culture	Do the right 5 people have the AI capability?	/ 5
Ethics & Governance	Has your governance stopped a deployment?	/ 5

3. YOUR CONSTRAINT

My weakest pillar is: _____

*This is where your first loop starts: the lowest score, even if another pillar feels more urgent.
The constraint is where transformation begins, and constraints shift as you move them.*

One thing to remember:

*With agents in production, watch Data Infrastructure
and Ethics & Governance most closely.*

Agents act on whatever data they can reach, and only governance can stop them. If two pillars tie for lowest, start with one of these.

4. Commit.

Date I confirmed this with my team: _____ Accountable name: _____

Focus

The weakest pillar is the only pillar for 90 days.

The declaration.

Read this aloud to your team. Then write it on a slide.

*For the next 90 days, the only pillar that matters
to our AI transformation is:*

_____ write the pillar name here

*The other four pillars are on hold.
They will be reassessed on:*

_____ date, 90 days from today

Signed: _____

Why focus matters.

A steering committee sees five pillars, authorizes dozens of initiatives, distributes budget evenly, and produces symmetric mediocrity.

Every pillar moves a little. No pillar moves enough to matter.

***You are picking one. The board will be uncomfortable.
Good. That is the signal you have focus.***

REMEMBER

- A program allocates. A practice delivers.
- The program holds the ceiling: capital, board oversight, multi-year bets.
- The Tuesday practice moves the floor: how people actually work.

Practice

One practice. Every Tuesday. Protected like a board meeting.

Find the row for your weakest pillar. That is your practice for the next 90 days.

STRATEGIC ALIGNMENT WEAKEST

The Weekly Alignment Check

30 minutes. Top 3 AI roadmap items re-confirmed against top 3 business problems.
Catches strategy drift, because agents drift silently.

DATA INFRASTRUCTURE WEAKEST

The Data Honesty Meeting

One hour. What data do we claim to have? What can we actually access without a 6-week cleanup?
Surfaces the gap before your agents exploit it.

OPERATIONAL PROCESSES WEAKEST

The Weekly Inventory

Every AI system running, written down, owned by a name. Includes agents bought without IT knowing.
An inventory that matches IT's list exactly has probably missed some.

TALENT & CULTURE WEAKEST

The Uncomfortable Upskilling

5 people whose decisions matter, trained deeply, before the 5,000.
The 5 are the ones agents will argue back with. Pick them accordingly.

ETHICS & GOVERNANCE WEAKEST

The Monthly Stop

Once a month, decide whether a live system should be paused. At least one real pause a quarter.
If governance never stops anything, it isn't governing. It's documenting. See MVG.

MY PRACTICE

My weekly practice is: _____

Scheduled: every _____ day at _____ time

Measure

Three questions every 90 days. Plus the Observability Layer check.

The three questions.

1

Did the weakest pillar score improve on the Canvas?

Yes / No. If no, the practice is not moving the number.

Yes

No

2

Did the practice actually happen every week?

Not "mostly." Every week. Rescheduling counts as not happening.

Yes

No

3

What did the practice surface we didn't know before?

Write it here. If the answer is "nothing," the practice has become a ritual. Redesign it.

If any answer is "no" or "nothing," the practice is theater.
Kill it. Design a sharper one. Score moves last. Awareness moves first.

The Observability Layer check.

If you have agents in production, answer these too. Any "no" is a red flag.

☐

Instrumentation

Every agent decision logged, queryable and reviewable in real time. Weekly reports are too slow.

☐

Thresholds

Three pre-written stop conditions per agent. Written BEFORE deployment. Auto-stop, no meeting.

☐

Named Escalation

One human, 24/7, documented authority to pull the plug. One name, on call.

“The agents still running in two years
will be the most closely watched ones.”

Come back for the *next loop.*

The organizations that get through this will have three things in common:

*small practices, run every week,
clear sight of every agent in production,
and a named person watching each one.*

Swap the usual question:
"What is our agentic AI transformation strategy?"

for a narrower one:

**What are we doing on Tuesday,
and who is watching the agent
we shipped on Friday?**

WHAT TO DO NEXT

- 1 Take the Canvas.**
15 questions. 8 to 10 minutes. → askajay.ai/canvas
- 2 Subscribe to The Brief.**
Every other Tuesday. Original thinking on AI governance and transformation. → askajay.ai
- 3 Read the full argument and its sources.**
Updated October 2026. → askajay.ai/thinking