The Work Assignment Matrix

Map the work. Then decide who does it.

A two-part method for HR and business leaders: build the work taxonomy, then place every task with one of four answers.

From ***AI Will Scale Good Leadership and Bad Leadership:*** *The CHRO's Method for Redesigning Work Before AI Scales What's Broken* by Liz Corey · Toolbox No. 6

AI doesn't start with the tool. It starts with the work. Most AI conversations ask one question: can AI do this? Almost everything is partly automatable now, so the answer is always yes, somewhat. The better question is whether it should, and under what supervision. You can't answer that until you can see the work task by task.

A single job holds twenty, fifty, sometimes hundreds of tasks. AI will take some, speed up others, support a decision in a few, and have no business in the rest. The unit of analysis is not the job. It's the task.

PART ONE

Build the work taxonomy

Don't start with jobs. Don't start with people. Start with the work, and build it in four levels.

LEVEL	WHAT IT IS	RECRUITING EXAMPLE
L1 Function	The business capability	*Human Resources*
L2 Process	The value stream inside it	*Recruiting*
L3 Activity	A group of related work	*Screen candidates*
L4 Task	One unit of work you can assign	*Review a resume. Conduct a phone screen. Enter a candidate disposition.*

Capture for every task

- How often it happens, and how much volume
- How long one occurrence takes, and how much that varies
- Who does it today
- What capability it truly requires
- What value it creates
- Whether it's legally required
- Whether it needs to happen at all

How to collect it

- **Start with the people closest to the work.** Ask what they actually did this week, not what the job description says.
- **Interview, don't hand out a form.** A form asks people to summarize. A conversation lets them remember.
- **Verify with system data.** Self-reported time is directionally useful and precisely wrong.
- **Keep the room on the task, not the name.** People come back into the picture once the work is clear.

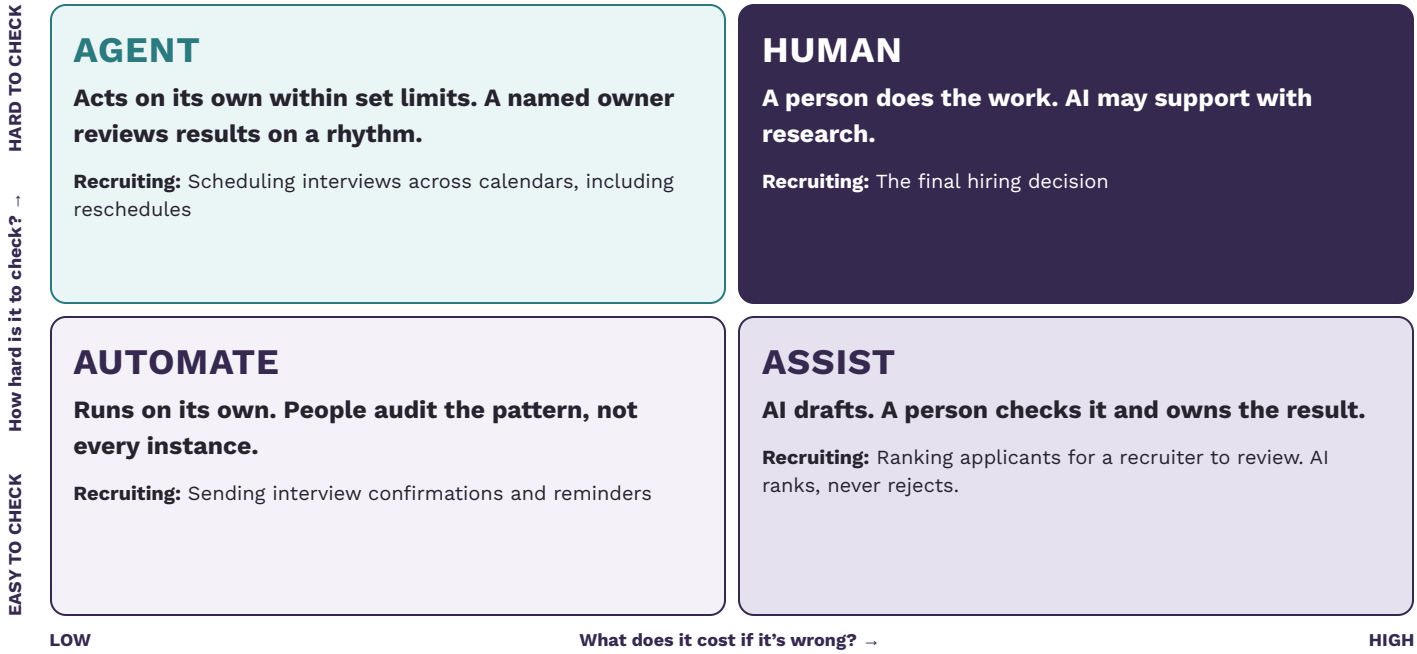
The taxonomy tells you what the work is. The matrix tells you who should do it. Ask the question most teams skip first: should this task exist at all? Cutting a task beats automating it.

© 2026 Elizabeth A. Corey. All rights reserved. GROWTH Execution Model™.

PART TWO

Place every task

Every task lands in one of four boxes. Two questions decide which one: **what does it cost if it's wrong**, and **how hard is it to check?**



Cost if it's wrong

ANY YES = HIGH

1. **Financial:** Could a mistake cost money above what a leader can approve without escalating?
2. **Legal:** Could a mistake create legal or compliance exposure?
3. **People:** Could a mistake affect someone's job, pay, or career?
4. **Reversibility:** Would a mistake be hard or impossible to undo?
5. **Trust:** Would a mistake damage trust with candidates, employees, or customers?

How hard it is to check

ANY YES = HARD

1. **Time:** Would checking it take almost as long as doing it yourself?
2. **Visibility:** Could a mistake go unnoticed?
3. **Steps:** Would you have to check every step, not just the final result?
4. **Standard:** Is there no clear standard to check it against?

Checking is done by the person who owns the work.

The rules

- **Unsure?** Start in Assist. A task moves up only after its error pattern is known.
- **Always Human:** termination, discipline, medical information, and legal exposure, whatever the scores say.
- **Landed in Human?** Ask what would make it easier to check. Redesign the task, then score it again.
- **No Agent** without a job description and a named owner.
- **Every task** gets an owner and a review date.
- **Revisit placements** as checking gets easier. The matrix is a snapshot. Governance is the film.

© 2026 Elizabeth A. Corey. All rights reserved. GROWTH Execution Model™.

WORKED EXAMPLE

Shortlisting applicants

Same task, two designs. Redesigning how the work is done moved it to a different box.

THE WAY MOST COMPANIES RUN IT

AI screens applicants and rejects the ones it scores low. No one ever sees who was screened out.

Cost if wrong: **High** (Legal, People, Trust)
Hard to check: **Yes** (Visibility)

Box: HUMAN

REDESIGNED

AI ranks every applicant and never rejects. The recruiter sees the full list. Candidates are scored against a written performance profile.

Cost if wrong: **High** (unchanged)
Hard to check: **No** (all four answers are now no)

Box: ASSIST

The cost of being wrong didn't change. How easy it is to check did. That's the redesign work: you don't just sort tasks into boxes. You change tasks so they can be checked.

After placement: what each box requires

BOX	WHAT TO PUT IN PLACE
Automate	Set an audit rhythm. Someone reviews the pattern of results on a schedule, not every instance.
Assist	Name the person who checks each output and owns the result. The checkpoint is built into the flow.
Agent	Write the agent job description first: purpose, scope, decision rights, escalation path, a named owner, and performance standards. No document, no agent.
Human	Protect the time. AI supports the preparation, and the person does the work.

Then decide where the hours go

Freed capacity doesn't redirect itself. It gets absorbed by whatever is loudest. Senior leaders set the direction for the time AI frees, and every leader carries it into their operating rhythm. Redesign the human roles around what's left before anyone touches headcount.

Why the second question matters. Refining and checking AI responses accounts for roughly 60 percent of an agentic task's cost. Place tasks by what checking really takes, not by what the demo can do. (McKinsey, July 2026)

Want it done faster? The Work Assignment Classifier places your task list for your team to review. Early access. Call for pricing.
See what your leaders can carry first: take the free AI Leadership Load Assessment at onthegrow.ai

© 2026 Elizabeth A. Corey. All rights reserved. GROWTH Execution Model™.

Worksheet: Place One Task

Answer for the task as it runs today. One worksheet per task.

Task: _____

0 Should this task exist at all?

If not, stop. Cutting a task beats automating it. ☐ Keep ☐ Cut

1 What does it cost if it's wrong?

ANY YES = HIGH

Financial: Could a mistake cost money above what a leader can approve without escalating?	☐ Yes ☐ No
Legal: Could a mistake create legal or compliance exposure?	☐ Yes ☐ No
People: Could a mistake affect someone's job, pay, or career?	☐ Yes ☐ No
Reversibility: Would a mistake be hard or impossible to undo?	☐ Yes ☐ No
Trust: Would a mistake damage trust with candidates, employees, or customers?	☐ Yes ☐ No

2 How hard is it to check?

ANY YES = HARD

Time: Would checking it take almost as long as doing it yourself?	☐ Yes ☐ No
Visibility: Could a mistake go unnoticed?	☐ Yes ☐ No
Steps: Would you have to check every step, not just the final result?	☐ Yes ☐ No
Standard: Is there no clear standard to check it against?	☐ Yes ☐ No

3 Read the box.

Low cost + Easy **AUTOMATE**	High cost + Easy **ASSIST**
Low cost + Hard **AGENT**	High cost + Hard **HUMAN**

Cost: ☐ Low ☐ High

Check: ☐ Easy ☐ Hard

Box: _____

4 Apply the rules.

- **Unsure?** Start in Assist. A task moves up only after its error pattern is known.
- **Always Human:** termination, discipline, medical information, and legal exposure, whatever the scores say.
- **Landed in Human?** Ask what would make it easier to check. Redesign the task, then score it again.
- **No Agent** without a job description and a named owner.
- **Every task** gets an owner and a review date.

Owner: _____ **Review date:** _____

© 2026 Elizabeth A. Corey. All rights reserved. GROWTH Execution Model™.