Where Does the Work Go?

The AI Workforce Transition Evidence & Measurement Toolkit

From AI productivity claims to evidence of changed work, real assignments, worker outcomes and business value.

THE QUESTION	THE OUTPUT
What did AI change?	Task-level work and capacity measurement
Where did capacity go?	Assignment, skills and workforce-flow records
What can we substantiate?	Evidence-to-decision worksheet and claim register

For CHROs · HR & L&D leaders · Transformation offices · CIOs · Operations and governance teams

Pilot edition v1.0 · 4 November 2026 · Companion Excel workbook included

SECTION 01 / AIGW PRACTITIONER TOOLKIT

The question is not only how much time AI saves.

The useful management question is what demonstrably happened after work changed.

A firm can shorten a task, extend capacity, train employees and report fewer staff—without any one of those facts proving the others. This toolkit treats each as a distinct, source-backed claim.

Use it to run a focused **3–5 workflow pilot**. Avoid using company-wide AI adoption as the unit of measurement.

Six questions to answer

CHECK	REQUIRED ANSWER
01	What specific unit of work changed?
02	Did total human effort change, including verification and rework?
03	Did quality, service and risk remain acceptable?
04	Where did the capacity go: assignments, absorbed demand, learning or something else?
05	What happened to people: capability, grade, pay, workload and progression?
06	What measurable value was achieved—and how confidently can it be attributed?

NOT AN AUDIT

This toolkit is not a validated maturity score, an AI layoff counter, an ROI certification or a causal evaluation. Missing evidence should remain explicitly unknown.

SECTION 02 / AIGW PRACTITIONER TOOLKIT

Eight domains. Eight different evidence questions.

AIGW Workforce Architecture distinguishes activity, work, workforce response and outcomes. [A1]

MEASUREMENT DOMAIN	CORE QUESTION
Technology	What is deployed, where and with what autonomy?
Work	Which accepted outputs, human minutes and handoffs changed?
Roles	Which responsibilities were removed, added or supervised?
Capability & learning	Who can demonstrate what skill on real work?
Authority	Who may approve, stop, override or reverse AI actions?
Workforce flows	Who was hired, transferred, retained, not replaced or exited?
Worker outcomes	What happened to pay, grade, progression, workload or choice?
Business outcomes	Did service, quality, risk, cost or value improve?

THE RULE

An enterprise can have extensive AI access but weak role redesign; large learning volumes but no verified assignments; and productivity improvement without known worker outcomes.

These domains are **not** a maturity staircase. Some pilots will have credible evidence in three domains and gaps in the others. The purpose is to identify those boundaries, not invent an aggregate score.

SECTION 03 / AIGW PRACTITIONER TOOLKIT

Select a pilot where claims can be tested.

Start with work units, not AI licences, departments or broad job titles.

An eligible pilot has: a named process owner; one comparable accepted output; observable human effort before and after; a stable observation period; a quality measure; and a way to link the result to a defined work unit.

PHASE	PRACTICAL ACTION	ACCOUNTABLE PERSON
Days 1–5	Select 3–5 workflows; define acceptance criteria and task boundaries; register baseline sample.	Operations + HR
Days 6–15	Observe AI-assisted work; capture time including checking/rework; measure quality and actual AI use.	Workflow owner
Days 16–22	Check assignments, human authority, capability and stock/flow records.	HRBP + governance
Days 23–30	Reconcile claims and alternative causes; select actions and schedule follow-up.	Transformation sponsor

Minimum evidence package

- Baseline sample definition, observation dates and an accepted-output denominator.
- Same measurement definitions after AI, including all human rework and exception handling.
- Observed AI-assisted volume, not tool licences or total team volume.
- A named quality gate, accountable owner and linkable source record.

STOP CONDITION	If the output mix or acceptance threshold changed materially, record the comparison as not established. Do not force an hour-saved figure from non-comparable work.

SECTION 04 / AIGW PRACTITIONER TOOLKIT

Measure human time, not theoretical FTE.

Use real comparable accepted units and include the human work that AI creates.

Signed monthly task-hours change = (baseline human minutes per accepted unit − AI-assisted human minutes per accepted unit) × eligible units per month × observed AI-assisted fraction ÷ 60.

The AI-assisted time must include prompting, validation, exception handling and human rework. Use representative observations to understand variability, not cherry-picked demonstrations. Do not sum overlapping tasks or observation periods: duplicate workflow coverage creates false hours.

INPUT	WHY IT IS REQUIRED
Baseline human minutes	What comparable accepted output previously required.
AI-assisted human minutes	What people actually do after AI, including review and correction.
Eligible units / month	A defensible volume denominator, not the whole enterprise.
Observed AI fraction	Share of comparable units that actually used the AI workflow.
Quality guardrail	Defect, service, safety and exception outcomes that may offset time gains.

INTERPRETATION	
	A positive hour difference describes measured human effort—not cash savings, positions removed, workers redeployed or additional customer value. Negative differences remain visible. The workbook shows positive quality-cleared capacity separately.

SECTION 05 / AIGW PRACTITIONER TOOLKIT

Stocks are not flows.

Do not read a headcount snapshot as a history of how people moved.

Expected ending headcount = starting headcount + external hires + transfers in − transfers out − voluntary exits − involuntary exits + other net adjustments.

Compare this calculated ending stock with the actual ending stock for the same population, unit and period. A discrepancy calls for reconciliation—it does not measure AI displacement.

FLOW OR MEMO	HOW TO RECORD IT
External hire	People who entered the defined workforce population.
Transfers in / out	Actual dated movement across the population boundary.
Voluntary / involuntary exits	Separate departure reasons; do not treat attrition as layoffs.
Other net adjustment	Known scope or reporting changes, explained in notes.
Non-backfill	Memo: vacancy remains unfilled after an exit; not a second exit.
Reduced hiring vs plan	Memo: actual intake below a dated approved counterfactual plan.

EXAMPLE LIMIT

A net decrease of 6,716 employees, as documented for one Wipro quarter, is not a count of AI-caused layoffs. The corresponding hires, exits, internal moves and attribution are needed before an employment claim. [A2]

SECTION 06 / AIGW PRACTITIONER TOOLKIT

Document where capacity went.

The destination of released time is a separate empirical question.

POSSIBLE DESTINATION	MINIMUM EVIDENCE
More work delivered	Additional accepted output, demand and quality over the same period.
Reassignment	Named destination function/project, actual dated assignment, retention at follow-up.
Capability investment	Training plus demonstrated competence and application on real work.
Demand / vacancy absorption	Staffing plan, vacancies, backfills, workload and explicit decision evidence.
Worker or business benefit	Workload, pay, quality, revenue, cost or risk evidence, with attribution limits.

Capacity destinations can overlap

A single employee might spend some released hours on training, move to another project and contribute to higher throughput. Do not add these activities as if they were mutually exclusive people or jobs. A destination should be described with units, time period and evidence.

The decisive questions are: what specific assignment changed, who authorized it, how long did it last, and what result became observable? A plan for redeployment is not a completed move; a completed move is not automatically a better career.

SOURCE TEST

Wipro's 20,000-worker-equivalent capacity statement is an attributed executive claim. The public reporting does not reveal the calculation method, distinct persons, destination assignments or long-term worker outcomes. [A2][A3]

SECTION 07 / AIGW PRACTITIONER TOOLKIT

Follow capability into real work.

Learning is an input; the pathway to demonstrated responsibility is the outcome question.

For a cohort whose work is changing, track each stage independently. Missing downstream evidence cannot be supplied by upstream training totals.

EVIDENCE STAGE	MEASURE	COMMON FALSE INFERENCE
Access / enrolment	Eligible and enrolled people.	Every learner is AI fluent.
Completion	Completed modules, by cohort.	Every completer is proficient.
Proficiency	Assessed work sample or observed performance.	A badge is productive ability.
Assignment	Paid project/role and destination date.	A recommendation is a move.
Career quality	Duration, grade, pay, choice, retention and progression.	Any move is a successful transition.

Add a worker-quality check

- Has work intensity increased because of AI monitoring, checking or exception handling?
- Who has access to meaningful training and destination roles? Is the opportunity equitable across cohorts?
- Are people moving into durable capability-building work—or narrow roles likely to be automated again?

CASE EVIDENCE

TCS reports nearly half of internal allocations occur through AI-enabled Talent Marketplace recommendations. This supports an institutional matching process; it does not by itself establish the quality or permanence of resulting careers. [A4][A5]

SECTION 08 / AIGW PRACTITIONER TOOLKIT

Authority and evidence are part of productivity.

A faster workflow can still be weakly governed or poorly evidenced.

CONTROL POINT	ASK THE PROCESS OWNER
Permission	What may the AI read, write, execute or trigger?
Review / approval	Which outputs require human authorization?
Override / rollback	Who can stop or reverse an action; is it tested?
Escalation	Which failures or risks force a human handoff?
Final accountability	Who owns results, audits and adverse incidents?

Claim-to-source discipline

Every important statement belongs in the workbook **CLAIM REGISTER** with an original URL, evidence date, publisher, precise support statement and unresolved alternative explanation. Classify it as **Supported / Qualified / Unsupported / Not yet evidenced**.

Those labels are practical claim statuses, not an automatic substitution for AIGW Tracker E1–E4 evidence grades. Public corporate accounts can verify that a company reported a result without independently validating the causal interpretation.

DECISION GATE

Where approval or override rights are not documented for consequential AI actions, record a governance gap. Do not describe the process as “human-controlled” solely because an operator appears somewhere in the workflow.

SECTION 09 / AIGW PRACTITIONER TOOLKIT

Three cases. Three different evidence limits.

Use public cases as reasoning examples, not as plug-in estimates for your organization.

EVIDENCE EXAMPLE	WHAT IT DOES — AND DOES NOT — SHOW
Wipro Original Reuters report · AIGW analysis	**Reported:** 20,000-worker-equivalent AI capacity; leader describes redeployment. **Unknown:** the equivalence methodology, distinct worker moves and their outcomes. **Lesson:** keep capacity separate from individual employment claims.
TCS Official human-capital report · AIGW analysis	**Reported:** nearly half of internal allocations use AI-recommendation-enabled matching. **Unknown:** percentage of recommendations accepted, enduring career outcomes. **Lesson:** distinguish recommendations, assignments and career quality.
Salesforce Company account · AIGW case study	**Reported:** hundreds of support-engineer redeployments and non-backfill; support headcount lower. **Unknown:** full decomposition of headcount change across exits and movements. **Lesson:** document dated workforce flows, not just net stock.

What this means for AIGW users

A useful measurement system does not assume that the company with the largest training programme or capacity claim has the best transition. It asks what each source actually establishes, what it omits, and what an operational pilot could measure next.

CASE SELECTION

These examples are purposively chosen to illustrate different measurement boundaries. They are not a representative sample of companies or a ranking.

SECTION 10 / AIGW PRACTITIONER TOOLKIT

A worked example with deliberately limited claims.

All numbers on this page are fictional and are also editable in the DEMO ONLY workbook tab.

FICTIONAL WORKFLOW	BEFORE / AFTER	VOLUME & AI SHARE	RESULT / DECISION
FAQ tagging	12 min → 7 min	1,000 units; 70%	+58.3 hours/month; quality pass; continue measuring
Complex escalation	25 min → 29 min	80 units; 50%	−2.7 hours/month; quality fail; redesign
Onboarding summary	18 min → 10 min	300 units; 60%	+24.0 hours/month; quality unknown; hold

For the first task:

(12 − 7) × 1,000 × 0.70 ÷ 60 = **58.3 measured monthly task-hours**. All human review and rework are assumed to be included in the seven-minute figure. The result is a modeled calculation on the stated fictional inputs, not audited savings.

The second task must not disappear from reporting because the result is negative. The third task must not be counted as quality-cleared just because the speed estimate looks attractive.

NOT INFERRED

No employee was “saved,” laid off, reassigned or promoted in this example. No cash saving, customer value or increased revenue follows from the task calculation alone.

After the pilot

Ask the workflow owner to evidence the actual use of released time. The HR partner must independently document role and assignment changes; the finance and service owners must validate any business-value claim.

SECTION 11 / AIGW PRACTITIONER TOOLKIT

Turn evidence into one defensible decision.

Use the 90-minute review agenda, record the next proof request, and keep the original-source trail.

TIME	DECISION MEETING BLOCK
15 min	Confirm accepted work units, before/after comparability and sampling.
20 min	Review measured human effort, quality, exceptions and adoption.
15 min	Inspect capacity destinations and actual assignment evidence.
15 min	Reconcile workforce flows, cohort skills and worker outcomes.
15 min	Review authority, source strength and alternative explanations.
10 min	Decide: continue, redesign, invest, add controls, monitor or pause.

Companion workbook: START HERE · WORK MEASUREMENT · WORKFORCE FLOWS · CAPABILITY CAREER · AUTHORITY MAP · CLAIM REGISTER · DECISION LOG · DASHBOARD · DEMO ONLY · DATA DICTIONARY. The synthetic demo is excluded from live totals.

Original sources & further evidence

[A1] **AIGW Policy Observatory** — Workforce Architecture for the AI Enterprise (2026-10-12).

[A2] **AIGW Policy Observatory** — When AI Frees Capacity, Where Does the Work Go? (Wipro) (2026-11-01).

[A3] **Reuters** — Wipro's AI push frees capacity equivalent to 20,000 workers, CTO says (2026-09-10).

[A4] **Tata Consultancy Services** — TCS FY2026 Human Capital reporting (FY2026).

[A5] **AIGW Policy Observatory** — From Training to Assignment (TCS Deep Dive 13) (2026-10-22).

[A6] **Salesforce** — How Salesforce Is Reshaping Its Workforce in the Age of AI (2026-04-29).

[A7] **AIGW Policy Observatory** — Salesforce and Agentforce - structured case study (2026).

Data protection: Collect only appropriately authorized cohort/role-level workforce information. Do not insert personal employee identifiers into a public or broadly shared workbook.

Method: AIGW synthesis for practical evidence organization. Not a validated scoring model, employment impact evaluation, financial attestation or legal opinion.