



AIGW Policy Observatory
Global AI Governance & Workforce Transformation

Global AI Governance and Workforce Transformation Policy Observatory

AIGW EVIDENCE BRIEF

Before AI Replaces Jobs, It Redesigns Roles

What five enterprise cases reveal about task change, role architecture, capability transition and the employment effects of AI

3 October 2026

A decorative graphic in the bottom right corner consisting of several concentric, curved lines in shades of blue and purple, creating a sense of motion or a stylized 'C' shape.

CONTENTS

Contents

Executive Summary	03
-------------------	----

Key Findings	06
--------------	----

Role Architecture	10
-------------------	----

Seven Cross-Case Patterns	13
---------------------------	----

Five Cases, Five Different Workforce Outcomes	23
---	----

Workforce Transition Accounting	25
---------------------------------	----

External Benchmark	26
--------------------	----

Implications for Institutions	31
-------------------------------	----

AIGW Decision Framework	36
-------------------------	----

Notes and References	40
----------------------	----

EXECUTIVE SUMMARY

Executive Summary

The public debate about artificial intelligence and employment is often organized around a binary question:

- Will AI replace jobs?

That question matters, but it is frequently asked too early.

Across five enterprise cases examined by AIGW — Salesforce, DBS, PwC US, LONGi and JD.com — the first clearly observable changes do not usually appear as a clean one-for-one exchange between a technology and a job.

They appear inside the job.

Routine customer interactions move to AI agents. Data preparation and technical research become automated or compressed. Visual inspection moves toward machine vision. Customer-service knowledge moves behind the interface into AI training and configuration. Banking workflows are reorganized around human-AI collaboration. Junior auditors are expected to review, interpret and exercise professional skepticism earlier in their careers.

The role changes before the employment outcome becomes fully legible.

This Evidence Brief calls that intermediate layer role architecture.

A job is more than a title. It is a bundle of:

- tasks; • workflow position; • decision rights; • capabilities; • learning experiences; • career pathways; • and a staffing requirement.

AI can change each of these dimensions separately.

That helps explain why the same broad technological trend can produce very different workforce outcomes.

At Salesforce, Agentforce absorbed a large share of routine support interactions. Human support moved toward complex, consultative and proactive work; hundreds of employees were redeployed; reskilling accompanied the transition; and CEO Marc Benioff reported support headcount moving from roughly 9,000 to roughly 5,000 because fewer support staff were required. Later company reporting points to a mixture of redeployment, attrition and non-backfill rather than a documented 4,000-person AI layoff event.[1][2]

EXECUTIVE SUMMARY

At DBS, the bank embedded workforce design inside what it calls Operating Model Transformation. More than 11,000 employees in roles expected to be significantly changed by AI were identified for deeper upskilling or reskilling, 96% had commenced training at the reporting point, and employees moved into new career pathways including from customer service into relationship and wealth-planning roles. The bank later said it would avoid hiring new people into some jobs expected to be replaced by AI while retraining incumbents.[5][6][7]

At PwC US, AI is compressing research, data preparation, documentation and selected audit activities while increasing the expected importance of review, professional skepticism and judgment. PwC's own workforce thinking warns that removing too much junior work can damage the apprenticeship system that produces future experts. At the same time, reported internal planning points to a smaller US entry-level Audit and Tax intake by FY28, with technology, low attrition and Acceleration Center integration among the identified drivers.[9][10][11]

At LONGi, AI and machine vision substituted a substantial share of visual quality inspection. Human workers still double-check AI-flagged defects, while a documented former quality inspector, Zhang Chen, moved into an AI-trainer

role. LONGi also reports lower direct labour in parts of the quality process and lower manpower allocation at factory level. But the public evidence does not provide a worker-by-worker reconciliation that would translate those labour-efficiency gains into a layoff count.[14][15][16]

At JD.com, automated customer service has expanded from early chatbots into a large human-AI service system. JD says AI handled more than 4.2 billion customer inquiries during its 2025 11.11 promotion. A customer-service employee with nearly ten years of experience moved into an AI-trainer role where he diagnoses AI-response gaps, updates content, configures agents and helps tune the system. JD later said customer service had created multiple AI-related job categories and helped thousands of people complete job transformation or upgrading over roughly two years.[17][18][19]

These cases differ by country, sector, occupation and technology. They should not be treated as a representative sample of the global economy.

They do, however, reveal a recurring transition architecture:

- Task change → workflow redesign → role redesign → capability shift → workforce reallocation

EXECUTIVE SUMMARY

The sequence is not deterministic and not always linear. Some organizations begin with workforce planning; others begin with a tool deployment. But the five layers help explain why a company can report major AI adoption without immediate layoffs, or reduce labour demand without a single discrete “AI layoff” event.

The synthesis produces seven central findings.

First, AI's earliest measurable employment effect is often task reallocation rather than job elimination.

Second, remaining human work tends to concentrate around exceptions, judgment, relationships, verification and system oversight.

Third, domain expertise often becomes transition capital. Product knowledge, customer knowledge, audit knowledge and production knowledge can be redeployed into AI-enabled roles.

Fourth, automation can remove the learning substrate of a profession. PwC makes this problem visible: when AI performs the routine work through which juniors traditionally learned, apprenticeship itself must be redesigned.

Fifth, workforce effects often appear through flows before they appear in stocks. Redeployment, reduced hiring, non-backfill, attrition and changed entry pathways can all alter labour demand without appearing immediately as a large layoff number.

Sixth, new AI roles are not necessarily permanent safe harbors. JD's AI-trainer case shows that the work created around AI can itself become a target for further automation.

Seventh, “human in the loop” is not a sufficient governance description. Escalation, verification, approval, veto, override and final authority are different forms of human control and should be measured separately.

The central implication is that leaders need a different unit of analysis.

Counting jobs is necessary.

But by the time the job count changes, many of the most consequential decisions have already been made inside the role.

KEY FINDINGS

Key Findings

01

1. Task change is the leading indicator

Across all five cases, the earliest strong evidence appears at task or workflow level: [1] [5] [9] [16] [17]

- customer questions delegated to Agentforce; • banking processes compressed or redesigned; • audit research and data preparation automated or accelerated; • visual inspection shifted to machine vision; • basic customer service shifted toward AI. [1] [5] [9] [16] [17]

This is why job counts alone are a lagging indicator of AI transition. [1] [5] [9] [16] [17]

02

2. Roles are being re-bundled around what remains human

The work that remains tends to require more: [2] [5] [9] [16] [17]

- exception handling; • judgment; • verification; • relationships; • business context; • system configuration; • AI supervision. [2] [5] [9] [16] [17]

The role often becomes narrower in routine execution and richer in higher-context work. [2] [5] [9] [16] [17]

03

3. Domain expertise is a major transition asset

Workers do not always need to abandon their existing expertise to remain relevant. [2] [5] [9] [16] [17]

Across the cases, existing knowledge was repurposed: [2] [5] [9] [16] [17]

- support expertise → Success, Renewals and Forward Deployed Engineering; • customer-service experience → relationship/wealth planning; • audit fundamentals → AI-supported review and skepticism; • inspection knowledge → AI training and quality-AI work; • customer-service knowledge → agent configuration and system tuning. [2] [5] [9] [16] [17]

The pattern is domain expertise + new AI capability, not “old expertise versus AI.” [2] [5] [9] [16] [17]

KEY FINDINGS

01

4. Training is necessary but insufficient

DBS provides large-scale training evidence. PwC reports substantial AI learning. Salesforce combines paid learning with internal mobility. LONGi and JD show worker pathways. [5] [10] [2] [16] [17]

The common lesson is that training becomes meaningful when it is connected to: [5] [10] [2] [16] [17]

- a redesigned role; • a destination job; • real work; • mobility; • progression. [5] [10] [2] [16] [17]

A training hour is not a transition outcome. [5] [10] [2] [16] [17]

03

6. AI can compress both work and learning

PwC makes the apprenticeship problem explicit, but the issue extends further. [9] [12] [16] [17] [2]

If AI takes the repetitive task, organizations must ask what the task previously taught the worker. [9] [12] [16] [17] [2]

Removing low-value execution can create a learning debt if the institution does not replace the experience embedded inside it. [9] [12] [16] [17] [2]

02

5. Workforce outcomes are plural

The five cases include evidence of: [2] [3] [7] [11] [15] [18]

- reskilling; • redeployment; • new role creation; • reduced staffing requirements; • lower labour input; • planned lower entry hiring; • non-hiring into some roles; • attrition and non-backfill. [2] [3] [7] [11] [15] [18]

These mechanisms can coexist in the same organization. [2] [3] [7] [11] [15] [18]

04

7. Better measurement is now a governance requirement

The most important missing data across the five cases are not technology metrics. [2] [5] [10] [16] [17]

They are worker-flow metrics: [2] [5] [10] [16] [17]

- who was affected; • who moved; • where they moved; • who left; • how pay changed; • how progression changed; • whether the new capability is durable. [2] [5] [10] [16] [17]

Enterprise AI disclosure is ahead on adoption and behind on transition quality. [2] [5] [10] [16] [17]

THE STARTING POINT

Why “Jobs Lost” Is the Wrong Starting Point

01

1. Why “Jobs Lost” Is the Wrong Starting Point

The jobs question is attractive because it appears measurable.

A job exists or does not.

An employee remains or leaves.

A company hires or reduces headcount.

But AI does not enter organizations at the level of the job title.

It enters through work.

A support agent answers a recurring question.

An auditor searches a technical standard.

A production worker checks a panel for defects.

A bank employee prepares a memo.

A customer-service representative explains a product rule.

These are tasks inside roles.

When AI begins performing or compressing those tasks, the organization has choices.

It can:

- leave the rest of the role unchanged;
- remove the task and give the worker different work;
- redesign the workflow;
- combine roles;
- create a new role;
- redeploy the employee;
- stop hiring into the role;
- allow attrition to reduce capacity;
- reduce headcount.

The employment outcome is therefore mediated by organizational design.

This is not an argument that AI does not destroy jobs.

Salesforce provides direct evidence of lower staffing requirements. LONGi provides direct evidence of lower labour requirements in parts of production.

PwC's hiring plans indicate entry-path contraction.

DBS has said it will avoid hiring into some AI-replaced jobs.

The argument is different:

- Job loss is one possible end state of a process that often begins earlier inside tasks, workflows and roles.

02

1. Why “Jobs Lost” Is the Wrong Starting Point

If researchers and executives start only with headcount, they miss the transition while it is being designed.

THE STARTING POINT

01

2. The Evidence Base: Five Cases, Not a Global Sample

This brief synthesizes five AIGW Deep Dives selected because they provide relatively strong public evidence connecting AI deployment to actual changes in work or workforce structure.

The cases span:

- the United States; • Singapore; • China; • service work; • regulated professional work; • banking; • advanced manufacturing.

That variation is useful for identifying mechanisms that recur across different settings.

But the sample has important limitations.

It is purposive, not representative.

The cases are unusually visible and unusually well documented.

They may therefore over-represent organizations with:

- strong digital infrastructure; • formal training capacity; • internal mobility; • resources to redesign jobs; • incentives to publicize successful worker transitions.

The evidence also relies heavily on first-party corporate reporting, supplemented by attributable journalism, institutional reports and regulatory sources.

The correct use of the five cases is therefore:

- mechanism discovery

—not:

- prevalence estimation.

This brief asks:

What forms can enterprise AI workforce transition take?

It does not claim:

How common is each form across the whole economy?

ROLE ARCHITECTURE

The Missing Middle: Role Architecture

The role is the organizational layer between a task-level technology change and an employment outcome.

ROLE ARCHITECTURE

Role Architecture: Six Dimensions of a Job

FIGURE 1.1

analytical framework

Role architecture: six dimensions of a job

A role is often treated as a title. That is too coarse for AI transition. A useful role architecture has at least six components:

Role dimension	Question and examples
Task portfolio	What does the worker actually do? Examples: answer customer questions; inspect products; prepare data; review evidence. [2] [5] [10] [16] [17]
Workflow position	Where does the worker sit in the process? Examples: first-line executor; exception handler; reviewer; escalation point; system operator. [2] [5] [10] [16] [17]
Decision rights	What can the worker: • approve; • reject; • override; • escalate; • change? [2] [5] [10] [16] [17]
Capability bundle	What knowledge and skills are required? Examples: domain expertise; AI fluency; judgment; data literacy; relationship capability; system configuration. [2] [5] [10] [16] [17]
Learning pathway	How does a novice become competent? Examples: repetition; apprenticeship; simulation; certification; mentoring; supervised practice. [2] [5] [10] [16] [17]
Staffing envelope	How many people are required to perform the role at a given volume of work? [2] [5] [10] [16] [17]

Source: AIGW synthesis across five purposively selected cases [\[2\]](#) [\[5\]](#) [\[10\]](#) [\[16\]](#) [\[17\]](#)

Note: AIGW framework: domain knowledge and decision authority can persist even as task execution changes. Case outcomes differ by operating model and staffing needs.

TRANSITION STACK


The AIGW Transition Stack

FIGURE 1.2

five analytical layers

The AIGW Transition Stack

Five analytical layers — not a mechanical or universally linear sequence.



Source: AIGW synthesis across five purposively selected cases [2] [5] [10] [16] [17]

Note: These are analytical layers, not a fixed or deterministic sequence; organizations can move across them in different orders.

CROSS-CASE PATTERNS

Seven Cross-Case Patterns

Across sectors and countries, the cases show recurring mechanisms without implying one universal path.

CROSS-CASE PATTERNS

Patterns One to Three: Work, Human Tasks and Transition Capital

01

Why the stack matters

This stack prevents a common reasoning error:

- AI can do Task X • therefore Job Y disappears.

A more defensible chain is:

- AI changes Task X • which changes Workflow A • which may change Role Y • which changes Capability Z • after which the organization decides how to allocate people and staffing.

That final step is partly technological.

It is also managerial.

02

5. Pattern One: Work Changes Before Employment Becomes Legible

Across the five cases, task and workflow evidence is generally stronger and earlier than employment evidence.

This is not surprising.

Organizations can measure:

- conversations handled; • cycle time; • defect detection; • training participation; • AI-agent deployment;

before they can cleanly explain:

- which workers moved; • which positions disappeared; • which vacancies were not backfilled; • how much attrition was voluntary; • what happened to pay and progression.

03

Salesforce

The work transition became visible through Agentforce handling customer-support interactions, changes in resolution and case volume, and human escalation.

Only later did the workforce picture become clearer: hundreds redeployed, non-backfill described, and a smaller reported support organization.[1][2][3]

04

DBS

AI use cases and Operating Model Transformations were measured alongside role analysis and training.

But the public record is more precise about the number of employees identified for learning than about the number who ultimately changed occupations.[5][6]

05

PwC

The technology and role direction are visible:

- AI-supported research; • data preparation; • workflow automation; • earlier review and judgment.

The long-term employment outcome is less settled. Reported hiring targets indicate fewer US entry seats, but the plan is forward-looking and multi-causal.[9][10][11]

06

LONGi

The strongest process metrics include:

- inspection efficiency; • labour intensity; • defect rate; • yield; • automation.

The worker-transition evidence is concrete but narrower: named individuals moved into AI-related roles, while the full workforce distribution remains unknown.[15][16]

CROSS-CASE PATTERNS

01

JD

AI inquiry volume is enormous, and JD reports thousands of job transformations or upgrades.

Yet the current customer-service headcount and transition denominator are not public.[18][19][20]

Cross-case inference — The first observable AI labor-market effect may therefore be:

- a change in the amount, type and allocation of work before it becomes:
- a clean change in the number of jobs.

This matters for policy.

If institutions wait for a large layoff event before recognizing transition, they may intervene after the role architecture has already changed.

03

Exceptions

AI handles routine or predictable cases.

Humans handle:

- ambiguity; • edge cases; • unusual defects; • unresolved customer issues; • complex audit questions.

05

Relationships

As information retrieval becomes cheaper, human value can shift toward:

- trust; • persuasion; • empathy; • client context; • relationship management.

Salesforce and DBS both show this movement.

02

6. Pattern Two: Human Work Moves Toward Exceptions, Judgment and System Oversight

A recurring narrative around AI is that humans will simply “do higher-value work.”

The five cases give that phrase more precision.

The remaining human work often falls into several categories.

04

Judgment

Humans interpret rather than merely retrieve.

Examples include:

- deciding whether an AI-supported audit conclusion is defensible; • deciding how to respond to a complex customer case; • deciding whether an AI-flagged defect is valid.

06

System oversight

Humans increasingly work on the AI system rather than only through it.

Examples: AI trainers; agent configuration; quality-AI engineering; workflow design; model verification.

The pattern is strong.

But it contains a risk.

If “higher-value work” requires more judgment, organizations must ensure workers actually have the experience and support required to perform it.

That leads directly to the learning problem.

CROSS-CASE PATTERNS

01

7. Pattern Three: Domain Expertise Becomes Transition Capital

The most consistent worker-level mechanism across the five cases is not replacement of domain knowledge.

It is repositioning of domain knowledge.

03

DBS: customer-service experience moves toward relationship work

DBS reports movement from customer-service roles into relationship and wealth-planning pathways.[5]

The transferable asset is not only bank product knowledge.

It is the ability to understand customers and operate inside the bank's service system.

05

LONGi: inspection experience becomes model-training and verification knowledge

Zhang Chen's expertise in identifying defects becomes useful for AI training.[16]

The value shifts from:

- seeing every defect personally
- toward:
- knowing what the system should recognize and when it is wrong.

02

Salesforce: support expertise becomes customer and implementation capability

Salesforce's named transitions show support engineers moving into Success, Renewals and Forward Deployed Engineering.[2][4]

The transferable assets include:

- product knowledge; • troubleshooting; • customer communication; • ambiguity management.

04

PwC: audit knowledge remains the basis for supervising AI

PwC continues to emphasize:

- accounting fundamentals; • professional skepticism; • review; • CPA qualification; • audit methodology.

AI increases the need to challenge output; it does not remove the need to understand what the output should mean.[9][10]

06

JD: customer-service expertise becomes AI operations

Yuchen's customer-service experience helps him diagnose AI gaps, update knowledge and configure agents. [17] This is transition capital: product and client knowledge, process knowledge, pattern recognition, troubleshooting, professional judgment and institutional memory. [17] Ignoring it can treat knowledgeable workers as obsolete or lead employers to hire externally while losing business context the system still needs. [17]

CROSS-CASE PATTERNS

Pattern Four: Automation Can Create Learning Debt

01

8. Pattern Four: Automation Can Create Learning Debt

The strongest example appears in PwC, but the mechanism is broader.

03

LONGi and JD: learning changes when execution moves to the system

A new AI trainer must understand the original work deeply enough to recognize errors.

If future workers enter after AI already performs the task, where will that original domain understanding come from?

This is not fully answered in the public evidence.

It becomes a future workforce-design problem.

02

PwC: the apprenticeship paradox

PwC explicitly warns that shrinking the junior base too aggressively can weaken the pipeline that produces future experts.[9]

If AI performs more of the foundational work, junior auditors may be expected to:

- review earlier; • exercise skepticism earlier; • communicate at a higher level earlier.

But the organization must recreate the experience that traditionally produced those capabilities.

PwC's response includes:

- in-person training; • small-group reviewer development; • coaching; • simulations; • structured development; • human + AI skills.[10][12]

04

Salesforce: redeployment depends on prior expertise

Salesforce's support-engineer example moves existing product and customer knowledge into Forward Deployed Engineer or customer-success work. [4] If support entry-level work shrinks, employers need a new route to develop that expertise.

When AI removes a developmental task, institutions need a substitute learning mechanism: simulation, shadowing, supervised exception handling, rotations, scenario libraries, deliberate practice, oral defense of AI-supported work and earlier client exposure. Preserve capability formation without keeping inefficient manual work indefinitely. [9] [10]

Document preparation, data reconciliation, inspection and customer answering create outputs while teaching workers normal patterns, error modes, evidence, customer response and exceptions. If AI removes that practice faster than organizations replace it, learning debt accumulates.

CROSS-CASE PATTERNS

Workforce Effects Appear Through Flows Before Stocks

FIGURE 2.1

measurement framework

Workforce effects appear through flows before stocks

Stocks: snapshots	Flows: transitions and mechanisms
Stocks	Total headcount; number of jobs; number unemployed. [2] [5] [10] [16] [17]
Flows	Internal redeployment; external hiring; hiring avoided; attrition; non-backfill; entry-level intake; role creation; retirement; involuntary exits. [2] [5] [10] [16] [17]

Source: AIGW synthesis across the five cases [\[2\]](#) [\[5\]](#) [\[10\]](#) [\[16\]](#) [\[17\]](#)

Note: A lower staffing requirement can emerge through several flows, with different consequences for workers. Do not infer one mechanism from a changed stock.

CROSS-CASE PATTERNS

Pattern Five: Workforce Effects Appear Through Flows

01

9. Pattern Five: Workforce Effects Appear Through Flows Before Stocks

Employment analysis often focuses on stocks:

- total headcount; • number of jobs; • number unemployed.

The five cases show why flows can be more informative during AI transition.

Relevant flows include:

- internal redeployment; • external hiring; • hiring avoided; • attrition; • non-backfill; • entry-level intake; • role creation; • retirement; • involuntary exits.

03

DBS: future hiring changes before incumbent displacement

CEO Tan Su Shan said DBS would not hire new people into some jobs expected to be replaced by AI while retraining people already in affected roles.[7]

This is an employment effect even if no incumbent is dismissed.

A future vacancy disappears.

05

LONGi: lower labour input without a worker reconciliation

LONGi reports lower direct labour and lower manpower allocation in parts of the operation.[14][15]

The labor requirement changes.

The worker-flow decomposition remains incomplete.

02

Salesforce: headcount change with mixed mechanisms

Support headcount reportedly moved from roughly 9,000 to roughly 5,000.[3]

The public evidence also describes:

- redeployment; • attrition; • non-backfill.

The exact contribution of each remains unresolved. [2]

The stock changed.

But understanding the transition requires the flows.

04

PwC: entry-path contraction

Reported internal planning points to fewer annual US Audit and Tax associate hires by FY28.[11]

Again, the effect can appear as:

- fewer people entering before it appears as:
- existing employees leaving.

06

JD: transition counts without current stock reconciliation

JD reported thousands of transitions but has not published a complete source/destination flow matrix. [18] AIGW workforce-flow accounting separates headcount stocks from hires, internal moves, hiring avoided, non-backfill, attrition, involuntary exits, new-role entry and training completion. These mechanisms can change labour demand without one dramatic layoff event. [2] [5] [10] [16] [17]

CROSS-CASE PATTERNS

Pattern Six: New AI Roles Can Be Moving Safe Harbors

01

10. Pattern Six: New AI Roles Can Be Moving Safe Harbors

One of the easiest policy mistakes is to identify a new AI occupation and assume it is a durable destination.

JD shows why that assumption is risky.

Yuchen moved from customer service into AI training and system configuration.[17]

But JD's current JoyAI product already advertises automation of:

- conversation analysis; • document analysis; • FAQ mining; • knowledge enhancement;

specifically to reduce the cost of AI-trainer knowledge-base configuration.[21]

The destination role is itself being redesigned.

02

First-order automation

AI automates the original task.

Example: answer routine customer questions.

03

First transition

Worker moves into: training/configuring the AI.

04

Second-order automation

AI begins automating: knowledge extraction; FAQ generation; configuration work.

The human role moves again toward:

- evaluation; • business-context diagnosis; • workflow architecture; • exception design; • governance.

AIGW concept: the moving safe harbor — A new job created by AI is not automatically protected from AI.

A durable reskilling strategy should ask:

- Which capability becomes more valuable as the technology improves?

not only:

- Which new title exists today?

This suggests that transition programs should emphasize:

- judgment; • system reasoning; • domain expertise; • error analysis; • coordination; • accountability.

Those capabilities may be more durable than routine interaction with a particular AI interface.

CROSS-CASE PATTERNS

Pattern Seven: “Human in the Loop” Is Not One Thing

01

Escalation, verification and review

The categories are not interchangeable. Escalation routes a case to a person when defined conditions arise; verification checks a system result; review evaluates quality, sufficiency or reasoning. At LONGi, workers double-check AI-flagged defects—evidence of verification, but not a public description of the complete authority structure. Salesforce and JD report escalation or handoff in customer service, yet do not disclose a full decision-right matrix.

02

Approval, override and accountability

Approval is human authorization before action; veto or override can block or reverse system output; final authority identifies who retains formal accountability. DBS describes AI-supported credit-memo preparation followed by credit-manager review and a human final credit decision. PwC audit retains professional judgment and accountability. These examples support distinct forms of involvement, not one common architecture across employers.

03

AIGW governance rule

Future disclosures should identify the trigger, stage, human task, information available, authority to approve or override, and accountable decision-maker. A handoff alone does not establish approval; checking an output does not establish independent validation. A person may interact with a system without having authority to change its result, while responsibility may sit with someone acting at a later step. Disclosures should say when staff can withhold approval, reverse an output or escalate incomplete evidence. This is a recommendation from the synthesis, not a finding that every case implemented each control.

WORKFORCE ACCOUNTING

From Evidence to Workforce Accounting

What organizations choose to measure and redesign shapes how task-level change reaches workers.

CROSS-CASE MATRIX

Five Cases, Five Different Workforce Outcomes

FIGURE 2.2

Purposive cases; mechanisms differ

Five cases, different workforce outcomes

Case and sector	Task and human-role change	Workforce response and visible signal
Salesforce — Enterprise software / support	Routine support moves to AI; people handle complex and proactive cases.	Support staffing: ~9,000→~5,000. Company describes redeployment, attrition and non-backfill; not 4,000 AI layoffs. [1] [2] [3] [4]
DBS — Banking	AI enters OMT; service work shifts toward relationship and wealth planning.	11,000+ identified; 96% commenced. Some hiring restraint; no quantified net AI reduction. [5] [6] [7]
PwC US — Professional services / audit	Routine research, data preparation and documentation compress; review, skepticism and judgment gain weight.	Smaller US Audit/Tax entry-level intake planned through FY28; forward-looking and multi-causal. [9] [10] [11]
LONGi — Solar manufacturing	Machine vision inspects; people verify flagged defects and train systems.	Named inspector became an AI trainer. Labour metrics fell in some quality/factory work; no worker-level reconciliation. [14] [15] [16]
JD.com — E-commerce / customer service	Basic service automates; people handle complex cases and train/configure AI.	Thousands reportedly transitioned/upgraded. Current service headcount and comparable denominator undisclosed; net effect unknown. [17] [18] [19] [20] [21]

Source: AIGW synthesis; case-specific sources appear in the table [\[2\]](#) [\[5\]](#) [\[10\]](#) [\[16\]](#) [\[17\]](#)

Note: Selected cases are not representative or directly comparable; no workforce rate is inferred.

WORKFORCE ACCOUNTING

Why Workforce Outcomes Differ

01

12. Five Cases, Five Different Workforce Outcomes

The value of the synthesis is not that the cases converge on one employment outcome.

They do not.

This is why statements such as:

- “AI is creating jobs”

or:

- “AI is destroying jobs”

are often analytically under-specified.

Both can be true inside the same transition.

A company can:

- create AI-trainer roles; • redeploy employees; • reduce hiring; • reduce staffing in an old function; all at once.

The relevant question is:

- What is the complete workforce flow?

02

Why accounting must follow the worker

The cross-case evidence suggests a new reporting framework.

Enterprise AI reporting should connect six layers.

Most current public disclosures are strongest in Layers A and B.

The five cases offer meaningful evidence in C, D and E.

They remain much weaker in F.

That is where future research and disclosure need to improve.

Hiring restraint, non-backfill, attrition and involuntary exits can all reduce staffing demand, but a headcount snapshot cannot distinguish them.

Without linked role-transition and worker-outcome data, the same reported efficiency gain may represent redeployment, fewer future hires or incumbent job loss. The framework therefore follows change from deployment through worker outcomes; adoption and training counts alone are not transition outcomes.

WORKFORCE ACCOUNTING

AIGW Workforce Transition Accounting

FIGURE 3.1 proposed disclosure framework

AIGW Workforce Transition Accounting

Layer	What organizations should report
Technology	System deployed; deployment stage; user population; autonomy level. [2] [5] [10] [16] [17]
Work	Tasks automated, compressed or created; workflow handoffs; exception volume. [2] [5] [10] [16] [17]
Role	Roles augmented, streamlined, redesigned or created; role descriptions before and after. [2] [5] [10] [16] [17]
Capability	Skills required; training offered, started and completed; proficiency demonstrated. [2] [5] [10] [16] [17]
Workforce	Headcount by affected role; internal redeployment; new hiring; hiring avoided; non-backfill; attrition; involuntary exits. [2] [5] [10] [16] [17]
Worker outcome	Pay; grade; career progression; retention; time to productivity; worker choice and agency. [2] [5] [10] [16] [17]
Decision rights	Who reviews, approves, overrides and remains accountable in each AI-enabled workflow. [2] [5] [10] [16] [17]

Source: AIGW cross-case synthesis [\[2\]](#) [\[5\]](#) [\[10\]](#) [\[16\]](#) [\[17\]](#)

Note: A proposed AIGW reporting framework synthesized from these cases; not an existing common reporting standard.

EXTERNAL BENCHMARK

The Workforce Strategy in Wider Context

01

What the employer survey can show

The five cases are not statistically representative.

But they align with a broader change visible in employer expectations.

These figures are expectations, not realized outcomes.

They should not be used to validate the five enterprise cases as representative.

Their value is different.

They show that the organizational choices visible in the cases — automate, upskill, redeploy, hire differently, reduce capacity — are also appearing in broader employer strategy.

The five cases provide the institutional detail that a survey cannot.

Employers expect redeployment.

Salesforce, DBS, LONGi and JD show what redeployment can look like in practice.

Employers expect automation.

PwC shows why automation creates a learning-design problem.

Some employers expect staff reduction.

Salesforce and LONGi show why the mechanism needs to be decomposed rather than reduced to a headline number.

02

Institutions choose among the options

A second useful benchmark comes from UNDP's Human Development Report 2025, which argues that AI's developmental effects depend on institutional choices and emphasizes a complementarity economy in which human capability, agency and AI are designed together rather than treated as substitutes by default.[24]

The five cases are consistent with that framing.

Technology capability does not determine the workforce outcome alone.

Organizations choose:

- whether to redeploy; • whether to train; • whether to retain entry pathways; • whether to reduce hiring; • whether humans keep meaningful authority; • whether new roles create credible careers.

The technology changes the option set.

Institutions choose among the options.

EXTERNAL BENCHMARK

External Benchmark: The Cases Match a Wider Workforce Strategy Shift

FIGURE 3.2

macro benchmark only

WEF employer expectations for training and workforce strategy

Measure	Reported expectation	Evidence boundary
Workers expected to need training by 2030	59 of 100	Employer expectation; not a realized outcome. [23]
Upskilled in current roles	29 of 100	A portion of the 59 expected to need training. [23]
Reskilled and redeployed internally	19 of 100	Employer expectation; not a case-level estimate. [23]
May need training but not receive it	11 of 100	A portion of the 59 expected to need training. [23]
Employers prioritizing upskilling	85%	Surveyed employer strategy expectation. [23]
Employers accelerating process/task automation	73%	Surveyed employer strategy expectation. [23]
Employers transitioning staff to growing roles	51%	Surveyed employer strategy expectation. [23]
Employers foreseeing staff reductions	41%	Surveyed employer strategy expectation. [23]

Source: World Economic Forum, Future of Jobs Report 2025 [\[23\]](#)

Note: The survey offers macro context; it does not validate the five enterprise cases as representative or describe realized outcomes.

MANAGEMENT CHOICES

The Management Choices That Shape the Employment Outcome

01

15. The Management Choices That Shape the Employment Outcome

The five cases suggest that the workforce effect of AI is mediated by a relatively small set of recurring management decisions.

02

Choice 1 — Automate the task or redesign the workflow?

A tool can be inserted into an existing process.

Or the process can be rebuilt around the tool.

DBS's Operating Model Transformation is the clearest example of the second approach.[5][6]

The workforce consequence is likely to be larger when organizations redesign the whole workflow rather than automate one isolated task.

03

Choice 2 — Preserve the worker or preserve the expertise?

These are not the same decision.

An organization may still need:

- product knowledge; • process knowledge; • customer knowledge; • quality knowledge;

even if it needs fewer people performing the original task.

Salesforce, LONGi and JD show different ways of moving that expertise into adjacent roles.

04

Choice 3 — Build mobility before reducing capacity?

Internal mobility is easier when:

- destination roles already exist; • managers know which skills transfer; • employees have paid learning time; • transitions begin before the old role disappears.

Salesforce's Career Connect and structured role transitions illustrate this.[2]

DBS's workforce-planning sequence does something similar at larger scale: identify affected roles, define future roles, then design learning pathways.[5]

MANAGEMENT CHOICES

01

Choice 4 — Replace the learning that automation removes?

PwC makes this the most explicit.

If AI removes routine work from the first rung, the firm must deliberately replace:

- exposure; • practice; • observation; • feedback.

Without that, productivity today can become expertise shortage tomorrow.

02

Choice 5 — Use attrition, non-backfill or layoffs?

These mechanisms all reduce staffing but have different worker consequences.

Public reporting often collapses them.

Leaders should not.

A transition that occurs mostly through:

- redeployment; • voluntary attrition; • non-backfill;
- is institutionally different from one driven by immediate involuntary exits.

03

Choice 6 — Define the human control layer?

If humans remain involved, organizations must decide what that means.

Do humans:

- receive escalations; • review; • approve; • override; • remain accountable?

Governance design is workforce design because it determines what human roles remain necessary.

04

Choice 7 — Create a career or create a temporary bridge?

A new AI role may exist only because the system is immature.

JD's second-order automation evidence makes this visible.

Leaders should ask whether a destination role has:

- progression; • durable capability; • portability; • increasing responsibility.

A job title is not a career architecture.

INSTITUTIONAL IMPLICATIONS

Designing the Workforce Response

The workforce outcome depends on choices made across executive leadership, HR, technology, learning, policy and education.

INSTITUTIONAL IMPLICATIONS

Implications for Institutions

01

16. Implications for CEOs and Boards

Boards often receive AI updates in three categories:

- adoption; • productivity; • risk.

A fourth category is needed:

- workforce transition quality

Useful board questions include:

- Which roles are materially changing? • Which tasks are leaving those roles? • Which new tasks are appearing? • How many people are affected? • What percentage are being redeployed? • What happens to hiring? • Which roles are no longer being backfilled? • What new roles are being created? • What capabilities are becoming scarce? • What happens to entry pathways? • Who retains decision authority? • Are worker outcomes improving or deteriorating?

Without these questions, leadership may discover the workforce consequence after the technology program has already locked in the role architecture.

02

17. Implications for CHROs

The CHRO's unit of analysis should move from:

- job family

toward:

- task + role architecture + worker flow

Build a role-impact map — For each AI deployment, identify:

- tasks removed; • tasks compressed; • tasks added; • decisions shifted; • capability changes; • staffing implication.

Build a transition-capital map — Ask:

- What expertise does the affected worker possess? • Where else is it valuable? • What new capability would make it transferable?

Build destination roles before training programs — A worker should be able to see:

- current role → capability gap → destination role → progression

rather than:

- current role → generic AI course.

Measure successful movement — Track:

- started training; • completed training; • moved role; • remained after 12 months; • pay change; • grade change; • time to productivity.

The distinction between training and transition should become non-negotiable.

INSTITUTIONAL IMPLICATIONS

01

18. Implications for CIOs and AI Transformation Leaders

Technology teams frequently optimize for:

- model performance; • latency; • automation rate; • cost; • throughput.

Those metrics are necessary.

They are insufficient.

Every scaled automation should also have a human operating model.

That model should specify:

- exception routing; • escalation; • review; • override; • failure recovery; • knowledge feedback; • staffing implications.

A high-performing model with a weak human operating model can create:

- hidden review burden; • degraded learning; • unclear accountability; • poor customer experience; • future skill shortages.

The system boundary should include the people who operate around the AI.

02

19. Implications for L&D and Workforce Development

The five cases suggest a move away from generic AI literacy as the end goal.

AI literacy remains foundational.

But transition requires role-specific capability formation.

Customer support — AI fluency + product expertise + customer judgment + escalation.

Banking — AI/data capability + domain knowledge + relationship skills + change management.

Audit — AI use + accounting fundamentals + skepticism + review.

Manufacturing — AI tools + process knowledge + defect judgment + system verification.

AI service operations — Domain expertise + configuration + error diagnosis + evaluation.

The learning design question — For every task removed, ask:

- What did the worker learn by doing this task?

Then decide how that learning will be replaced.

This is especially important for:

- early-career work; • apprenticeship occupations; • technical roles built through repeated exposure.

INSTITUTIONAL IMPLICATIONS

01

20. Implications for Policymakers

Employment policy needs to observe the transition before it appears in unemployment statistics.

Track hiring as well as layoffs — If companies stop hiring into a role, entry opportunities can shrink without incumbent displacement.

DBS and PwC make this visible.

Track internal movement — Redeployment can protect employment while still producing major occupational change.

A worker may remain employed but move into:

- a different function; • a different skill level; • a different career trajectory.

Track transition quality — Useful indicators include:

- pay retention; • occupational progression; • training access; • completion; • worker choice; • portability of new skills.

Support skill portability — China's formal recognition of Artificial Intelligence Trainer as an occupation illustrates one institutional response to new AI work. [22]

Occupational standards can help make emerging roles:

- legible; • trainable; • portable.

But standards need to evolve because the work itself will continue changing.

Protect entry pathways — PwC raises a wider labor-market problem.

If AI removes the routine work through which young workers gain experience, policy needs to pay attention to:

- apprenticeships; • work-integrated learning; • employer incentives; • supervised practice; • new entry-role design.

The labor market cannot rely on senior workers forever if the pipeline producing senior workers weakens.

02

21. Implications for Education

AI changes the meaning of “entry ready.”

Employers may increasingly expect new entrants to combine:

- domain fundamentals; • AI fluency; • data literacy; • source verification; • judgment; • communication.

The danger is to respond by replacing deep knowledge with shallow tool familiarity.

The cases suggest the opposite.

The more AI performs routine execution, the more valuable it becomes for humans to understand:

- why a result is correct; • what evidence supports it; • where it can fail; • what to do when it fails.

Education therefore needs to design for:

- AI-enabled judgment

not:

- AI tool usage alone.

EVIDENCE BOUNDARIES

What the Five Cases Establish

01

22. What the Five Cases Establish

Taken together, the evidence supports several conclusions.

The cases establish that: — • AI is changing task composition in scaled production environments. • Workflow redesign can occur before a clear net employment effect is measurable. • Human work can shift toward exceptions, judgment, relationships and system oversight. • Real worker redeployment into new or redesigned roles is occurring. • Existing domain expertise can become an input to AI operations. • Training and internal mobility can be part of AI deployment rather than after-the-fact remediation. • Lower labour requirements can emerge alongside redeployment. • Hiring demand can change even when incumbents are retained. • Early-career pathways can be redesigned as routine work compresses. • AI-support roles themselves can become exposed to further automation. • Human control structures differ materially across workflows.

The cases do not establish that: — • role redesign always precedes job loss; • every worker affected by AI can be reskilled successfully; • redeployment fully offsets reduced staffing demand; • new AI roles create net employment growth; • training guarantees a transition; • “higher-value work” is necessarily higher paid or more secure; • AI is the sole cause of every headcount or hiring change observed; • the five cases are representative of their sectors or countries; • the same transition path is available to smaller firms with fewer resources.

The synthesis therefore supports a transition framework, not a universal forecast.

02

23. The Research Gaps That Matter Most

The five cases repeatedly expose the same missing evidence.

1. Worker-flow accounting — We need:

• source role; • destination role; • number moved; • number retained; • number exited.

2. Transition quality — We need:

• pay; • grade; • job quality; • retention; • progression; • worker agency.

3. Learning effectiveness — We need:

• not only training started; • but proficiency; • performance; • time to autonomy.

4. Entry-path impact — We need:

• graduate hiring; • apprenticeship slots; • reduced replacement hiring; • internal entry pathways.

5. Decision rights — We need:

• escalation; • review; • approval; • veto; • override; • accountability.

6. Longitudinal durability — We need to know whether workers who move into new AI roles:

• remain there; • progress; • face second-order automation; • move again.

7. Independent evaluation — Most detailed enterprise evidence remains first-party.

Independent research should test:

• workforce outcomes; • productivity attribution; • service quality; • distributional effects.

EVIDENCE BOUNDARIES

01

24. A Research Agenda: From AI Adoption to Workforce Architecture

The next generation of AI-and-work research should move beyond exposure scores alone.

Exposure tells us which tasks AI could affect.

It does not tell us what the organization does next.

A stronger research program would follow five questions longitudinally.

Question 1 — What work moved? — Track specific task portfolios.

Question 2 — What role replaced the old bundle? — Track role descriptions and decision rights.

Question 3 — What capability became necessary? — Track actual workplace competence.

Question 4 — What happened to people? — Track worker flows and outcomes.

Question 5 — What happened over time? — Track whether the destination role persists.

This would allow researchers to distinguish:

- temporary transition jobs; • durable new occupations; • hidden displacement; • successful redeployment; • stalled learning; • workforce upgrading.

DECISION FRAMEWORK

The AIGW Decision Framework

01

1. What task is changing?

Be specific.

02

2. What happens to the workflow after the task changes?

Do not assume the rest of the process stays the same.

03

3. What happens to the role?

Which tasks leave, remain and appear?

04

4. What expertise must be preserved?

Identify transition capital.

05

5. What new capability is required?

Define a destination, not a course.

06

6. What happens to the worker?

Track movement and employment mechanism.

07

7. Who keeps authority?

Specify the human decision layer.

A deployment is not workforce-ready until all seven questions have an answer.

CONCLUSION

Conclusion

The five enterprise cases in this brief do not resolve the debate over whether AI will create or destroy more jobs.

They reveal why that debate is often framed at the wrong level.

AI changes work before its full employment effect becomes visible.

It removes tasks.

It changes queues.

It rewrites handoffs.

It changes what counts as expertise.

It can reduce staffing requirements, create new roles, narrow entry pathways, move employees, and change the skills through which a career is built.

Those effects do not occur in isolation.

They form a transition architecture.

Across Salesforce, DBS, PwC, LONGi and JD, the recurring pattern is not:

- AI arrives → job disappears.

It is closer to:

- AI changes the task → the workflow is redesigned → the human role is re-bundled → capabilities shift → the organization reallocates people and staffing.

Sometimes the final outcome includes lower headcount.

Sometimes it includes redeployment.

Sometimes it includes fewer future hires.

Sometimes it includes new roles.

Often it includes several of these at once.

That is why the most important AI workforce question for an organization is not:

- How many jobs will AI replace?

It is:

- How are we redesigning roles now — and what workforce are those redesigned roles creating?

By the time the answer appears in the headcount, much of the future of work has already been decided.

METHOD AND SOURCE QUALITY

Method and Source Quality

01

Case-selection logic

This Evidence Brief synthesizes five enterprise cases previously researched as AIGW Workforce Transition Deep Dives.

They were selected because each has evidence beyond a speculative statement about AI exposure.

At minimum, the cases provide evidence of:

- deployed AI in real work; • identifiable task or workflow change; • observable role or workforce implications; • attributable source material.

Several also provide:

- named worker transitions; • quantitative operating outcomes; • training or mobility mechanisms; • staffing or hiring evidence.

The cases were not selected randomly.

They should not be used to estimate the probability that an organization will redeploy, reduce, reskill or create roles.

02

Unit of analysis

The unit of analysis is not:

- the company; • the entire sector; • the technology vendor.

It is the AI-enabled work transition episode.

For example:

- Salesforce customer support; • DBS AI-affected roles under Operating Model Transformations; • PwC US entry-level audit; • LONGi quality inspection; • JD customer service and AI operations.

A company can contain many other AI transitions with different workforce effects.

03

Evidence hierarchy

The synthesis prioritizes:

- regulatory filings and annual / sustainability reports; • official company operational reporting; • government and public-institution sources; • attributable executive or worker interviews; • high-quality journalism; • institutional benchmark research.

First-party sources are necessary because companies often hold the most detailed workforce data.

They are not treated as independent validation.

METHOD AND SOURCE QUALITY

01

Fact, interpretation and implication

The brief distinguishes three levels.

Fact — Directly supported by a source.

Example: Salesforce says hundreds of support engineers were redeployed.

Interpretation — AIGW synthesis of directly supported facts.

Example: Domain expertise can function as transition capital.

Implication — A decision-relevant lesson.

Example: Employers should map transferable expertise before automating a role.

The latter two should not be presented as if the source itself made the claim.

02

Causal discipline

The synthesis does not infer:

- layoffs from automation; • headcount reduction from productivity; • redeployment from training; • successful career transition from course participation; • AI causality from a coincident workforce change.

Where multiple causes are plausible, the brief preserves them.

Examples include:

- Salesforce's mixed redeployment / attrition / non-backfill mechanisms; • PwC's AI + low attrition + Acceleration Center hiring context; • LONGi's automation + severe photovoltaic-industry pressure.

03

Comparability

Metrics are kept at their original analytical level.

The brief does not directly compare:

- conversations with cases; • training starts with completions; • factory-wide cost reductions with one inspection workflow; • company-wide AI value with workforce-transition ROI.

This is particularly important because enterprise AI reporting often changes metric definitions as deployments scale.

04

Publication discipline

Safe headline:

- Across five enterprise cases, AI's first visible workforce effects often appear as task, workflow and role redesign before they appear as a clean net employment outcome.

Avoid:

- AI never replaces jobs; it only redesigns them.

Also avoid:

- These five companies prove that reskilling will solve AI displacement.

Neither conclusion is supported by the evidence.

SOURCES

Notes and References

The five underlying AIGW Deep Dives contain a broader evidence base than the selected references. This Evidence Brief cites the sources most directly necessary for its cross-case claims.

1. Salesforce. *How Salesforce Is Reimagining Its Workforce for the Agentic Enterprise*. 15 January 2026. Company / first-party workforce reporting.
Supports: 2.6 million customer conversations; 63% resolution; customer satisfaction roughly comparable with humans; hundreds redeployed; shift toward strategic work.
Boundary: First-party company reporting; not independent validation.
<https://www.salesforce.com/news/stories/agentic-enterprise-workforce-evolution/>
2. Salesforce. *How Salesforce Is Reshaping Its Workforce in the Age of AI*. 29 April 2026. Company / first-party workforce reporting.
Supports: redeployment process; Career Connect; cohort moves; targeted matching; learning during work time; named transitions; non-backfill; shift toward proactive customer work.
Boundary: First-party company reporting; reported transitions do not quantify all workforce flows.
<https://www.salesforce.com/news/stories/salesforce-reshaping-workforce-in-age-of-ai/>
3. The Logan Bartlett Show. *Marc Benioff interview*. 29 August 2025. Attributable executive interview.
Supports: reported support-headcount movement from roughly 9,000 to roughly 5,000 and CEO statement that fewer support staff were required.
Boundary: does not provide a decomposition of involuntary exits, redeployment, attrition and non-backfill.
<https://podscan.fm/podcasts/the-logan-bartlett-show/episodes/ep-149-marc-benioff-ceo-salesforce-predicts-half-of-conversations-will-be-with-ai-agents-next-year>
4. Salesforce. *Workforce Innovation Playbook: The 4Rs to Building the Human + Agent Workforce*. n.d.. Company / workforce playbook.
Supports: support-engineer-to-Forward-Deployed-Engineer reskilling example; retained capabilities; new technical and consultative skill requirements.
Boundary: Company playbook examples; not a representative estimate of transition outcomes.
<https://www.salesforce.com/en-ap/wp-content/uploads/sites/14/documents/pdf/agentforce/workforce-innovation-playbook-ap.pdf>

SOURCES

DBS References

5. DBS Group Holdings. *Sustainability Report 2025*. n.d.. Company sustainability report.

Supports: >11,000 employees identified for deeper upskilling/reskilling; 96% commenced; role-impact assessment; augmented, streamlined and new roles; customer-service → relationship/wealth-planning pathway.

Boundary: Company-reported training commencement does not establish completion or role transition.

https://www.dbs.com/annualreports/2025/i/pdf/DBS_SR2025.pdf

6. DBS Group Holdings. *Annual Report 2025*. n.d.. Company annual report.

Supports: >2,000 models; >430 AI use cases; nine Operating Model Transformations; bank-wide AI/analytics value; OMT design and workflow context.

Boundary: Company-wide adoption and value metrics do not quantify case-specific workforce outcomes.

<https://www.dbs.com/iwov-resources/images/investors/annual-report/dbs-annual-report-2025.pdf>

7. Bloomberg. *Singapore's DBS Girds for AI as Tellers Morph Into Bankers*. 7 November 2025. Attributable journalism.

Supports: CEO Tan Su Shan's statement that DBS would not hire new people into some jobs expected to be replaced by AI and would retrain people already in affected roles.

Boundary: An announced future hiring restraint; not a quantified incumbent reduction.

<https://www.bloomberg.com/news/articles/2025-11-07/singapore-s-dbs-girds-for-ai-as-tellers-morph-into-bankers>

8. DBS. *CEO Tan Su Shan on how agentic AI is transforming the bank*. n.d.. Company / executive interview.

Supports: agentic credit-memo workflow; AI performs research/drafting; human credit manager reviews and makes the final decision.

Boundary: An executive-described workflow example; it does not establish authority design across every credit decision.

<https://www.dbs.com/media/features/ceo-tan-shu-shan-on-how-agentic-ai-is-transforming-the-bank.page>

SOURCES

PwC US References

9. PwC US. *No more pyramids: Rethinking your workforce for the agentic AI era*. n.d.. Company workforce strategy / first-party analysis.
Supports: entry-level audit redesign; apprenticeship risk; hourglass workforce model; need to preserve an AI-literate early-career base.
Boundary: Company workforce strategy; not a measured employment outcome.
<https://www.pwc.com/us/en/tech-effect/ai-analytics/agentic-ai-workforce-redesign.html>
10. PwC US. *2025 Audit Quality Report*. n.d.. Company audit-quality / transparency report.
Supports: >50,000 AI-use training hours for Audit and DAT; apprenticeship model; first-year training; new-reviewer development; professional skepticism; review; CPA progression; entry-hiring context.
Boundary: Company-reported training and audit-quality evidence; training hours do not establish proficiency or future hiring effects.
<https://www.pwc.com/us/en/services/audit-assurance/assets/pwc-2025-audit-quality-report.pdf>
11. Business Insider. *Getting a job at PwC out of college will be a lot tougher*. 21 August 2025. Journalism based on internal PwC presentation + company statement.
Supports: reported plan for lower US Tax and Assurance associate hiring through FY28, including audit-specific targets; PwC confirmation that campus hiring goals were decreasing.
Boundary: PwC did not publicly confirm the specific internal figures; the change is multi-causal.
<https://www.businessinsider.com/pwc-hiring-fewer-junior-associates-ai-offshoring-big-four-2025-8>
12. PwC US. *PwC Learning Collective Reimagines Development for the AI Age*. 5 February 2026. Company workforce / learning release.
Supports: 30 human + AI skills; Associate Discovery; simulations, deployment, mentorship and experiential learning.
Boundary: Company-described learning programs; resulting proficiency and worker outcomes are not quantified here.
<https://www.pwc.com/us/en/about-us/newsroom/press-releases/learning-collective-ai-workforce-development.html>
13. Public Company Accounting Oversight Board. *Staff observations on generative AI in audits and financial reporting*. 22 July 2024. Independent regulator / staff outreach.
Supports: human review and supervision remain essential; firms reported privacy, reliability and governance concerns as GenAI enters audit work.
Boundary: Regulator staff observations from outreach; not a prevalence estimate.
<https://pcaobus.org/news-events/news-releases/news-release-detail/pcaob-staff-shares-observations-from-outreach-on-use-of-generative-artificial-intelligence-in-audits-and-financial-reporting>

SOURCES

LONGi References

14. LONGi. *Sustainability Report 2023*. n.d.. Company sustainability report.
Supports: AI quality system; 10% reduction in direct labour at reinspection stations; +4.9% first-pass yield; -30% defect rate.
Boundary: Company-reported process metrics; they do not provide a worker-by-worker transition reconciliation.
https://static.longi.com/2023_longi-sustainability-report.pdf
15. LONGi. *Jiaxing Lighthouse Factory / Technology and Innovation*. n.d.. Company / first-party operational reporting.
Supports: >90% automation; -33% single-line manpower allocation; 30+ AI models/algorithms; broader Lighthouse operating outcomes.
Boundary: factory-wide metrics should not be attributed solely to the quality-inspection transition.
<https://www.longi.com/en/sustainability/esg-society-technology/>
16. Xinhua. *Smart tech, better lives: AI empowers China's common prosperity drive*. 18 July 2026. Attributable national news reporting.
Supports: Zhang Chen's quality-inspector-to-AI-trainer transition; human double-check of AI-flagged defects; 198 inspection items across nine procedures; worker-reported inspection labour-cost reduction.
Boundary: A named worker case; it does not estimate workforce-wide transition rates.
<https://english.news.cn/20260718/82d978e634d74a018b437dfd668a6ce0/c.html>

SOURCES

JD.com References

17. Science and Technology Daily. *AI为他们带来新身份：京东协同推进技术应用与就业转型探索实践*. 19 May 2026.
Attributable national science/technology journalism.
Supports: Yuchen's customer-service-to-AI-trainer transition; error diagnosis; content update; agent configuration; maintenance/tuning; adjacent technical pathways.
Boundary: An attributed individual case; it does not estimate workforce-wide transition rates.
https://www.stdaily.com/web/gdxw/2026-05/19/content_518575.html
18. Xinhua / Science and Technology Daily. *未来五年，如何接住AI“新饭碗”*. 7 September 2026. National media / company-attributed workforce reporting.
Supports: multiple new AI-related customer-service job categories; JD statement that thousands of people completed job transformation or upgrading over roughly two years; AI-basic-service / human-complex-service model.
Boundary: Company-attributed transition counts; current headcount and a comparable transition denominator are not disclosed.
<https://www.news.cn/tech/20260907/eb1d91d961c84dd8b417ba41d649b129/c.html>
19. JD.com. *Fourth Quarter and Full Year 2025 Results*. 2026. Company investor reporting.
Supports: >4.2 billion AI customer-service inquiries during 2025 11.11; >50,000 AI agents integrated into internal systems by Q4 2025.
Boundary: inquiry volume is an adoption/scale metric, not a job-substitution count.
<https://ir.jd.com/news-releases/news-release-details/jdcom-announces-fourth-quarter-and-full-year-2025-results-and>
20. JD.com. *2024 Environmental, Social and Governance Report*. n.d.. Company ESG report.
Supports: >20,000-person professional customer-service workforce at end-2024; 24/7 multi-channel service; training and employee-development context.
Boundary: The end-2024 headcount is a point-in-time stock, not the denominator for later transitions.
https://ir.jd.com/system/files-encrypted/nasdaq_kms/assets/2025/06/19/15-13-13/2024_jd.com_inc_environmental_social_and_governance_report.pdf
21. JD JoyAI. *Intelligent Customer Service*. n.d.. Company / product documentation.
Supports: automated conversation/document analysis, FAQ mining and knowledge enhancement designed in part to reduce AI-trainer knowledge-base configuration effort.
Boundary: Product documentation describes intended capabilities, not independent performance evaluation.
<https://jdai.jd.com/portal/ai-products/customer-service>
22. Ministry of Human Resources and Social Security / Ministry of Industry and Information Technology. *Artificial Intelligence Trainer National Occupational Skill Standard*. 2021. Official national occupational standard.
Supports: formal definition of Artificial Intelligence Trainer; five skill levels; work including data processing, business analysis, algorithm testing, error analysis, human-machine interaction and intelligent-system optimization.
Boundary: An occupational standard defines work and skill levels; it is not evidence of realized employment outcomes.
<https://www.mohrss.gov.cn/wap/zc/zqyj/202106/W020210617509883457681.pdf>

SOURCES

External Benchmarks References

23. World Economic Forum. *Future of Jobs Report 2025*. January 2025. Institutional employer survey / labor-market benchmark.
Supports: 59 of 100 workers expected to need training by 2030; 29 upskilled in current roles; 19 reskilled/redeployed; 11 unlikely to receive needed training; employer workforce-strategy expectations including upskilling, automation, internal transition and staff reduction.
Boundary: Employer expectations are a macro benchmark, not realized outcomes or validation of the five selected cases.
<https://www.weforum.org/publications/the-future-of-jobs-report-2025/>
24. United Nations Development Programme. *Human Development Report 2025: A Matter of Choice*. n.d.. Institutional conceptual benchmark.
Supports: framing of AI outcomes as institutional choices; human agency; complementarity between people and AI; capability-oriented development.
Boundary: Conceptual development framing, not enterprise workforce outcome evidence.
<https://hdr.undp.org/content/human-development-report-2025>