



AIGW Policy Observatory
Global AI Governance & Workforce Transformation

Global AI Governance and Workforce Transformation Policy Observatory

FROM AI ADOPTION TO WORKFORCE ARCHITECTURE — 2026 ENTERPRISE EVIDENCE REPORT

Workforce Architecture for the AI Enterprise

What Nine Enterprise Transitions Reveal About
Redesigning Work, Roles, Capability, Authority
and Careers

October 2026



CONTENTS

Contents

Executive Summary	03	What the Nine Cases Reveal	37
AI Adoption Is No Longer Enough	07	Five Enterprise Dilemmas	41
Nine Enterprises, Nine Transition Episodes	09	The AIGW Workforce Architecture Framework	43
Task Architecture — Deciding What AI Should Do	13	The Enterprise Diagnostic	46
Workflow Architecture — Redesigning How Work Moves	15	An Illustrative 12-Month Enterprise Roadmap	48
Role Architecture — Rebundling the Job	17	For CEOs and Boards	51
Capability Architecture — What Becomes Scarce When Execution Becomes Cheap?	20	For CHROs	52
Learning Architecture — How Expertise Is Reproduced	23	For CIOs and Chief AI Officers	53
Authority Architecture — Who May Decide, Act and Override?	26	For Learning and Development Leaders	54
Career Architecture — Where Does the Worker Go Next?	29	For Policymakers and Public Institutions	55
Workforce Flow Architecture — How People Move Through the Transition	31	What We Still Do Not Know	56
Measurement Architecture — What Counts as Evidence of Transformation?	34	Conclusion — From Technology Deployment to Institutional Design	58
		Notes and References	59
		Source and Method Note	65

SECTION 01

Executive Summary

Enterprise AI is entering a different phase.

For several years, the dominant management question was adoption: Which tools should employees receive? How many people use them? How many copilots, assistants or agents have been deployed? How much time might they save? Those questions still matter. But they become incomplete once AI starts doing more than helping an individual produce an answer.

In the cases examined in this report, AI is already answering customers, drafting professional work, inspecting industrial products, supporting credit decisions, creating clinical documentation, coordinating operational workflows, reallocating network resources and acting as a digital participant inside financial-services teams. In some organizations it changes what a role contains. In others it changes how people learn, how they move through the organization, or who retains the right to make a consequential decision.

At that point, AI is no longer only a technology-adoption problem.

It becomes a workforce-architecture problem.

This report defines AI workforce architecture as the institutional system through which an organization allocates work between people and AI, structures workflows and roles, builds and renews human capability, assigns decision authority, creates career pathways, manages workforce flows and measures the outcomes of transition.

The framework is derived from nine enterprise transition cases researched by the AIGW Policy Observatory: Salesforce, DBS, PwC US, LONGi, JD.com, SimplifyNext, Kaiser Permanente / The Permanente Medical Group, Deutsche Telekom and BNY. These cases span enterprise software, banking, audit, advanced manufacturing, e-commerce and customer

service, AI services, healthcare, telecommunications and financial infrastructure across the United States, Singapore, China, Germany and global operations.

They are not a representative sample of the economy. They were selected because they make particular transition mechanisms unusually visible. The purpose is therefore not to estimate how common each outcome is. It is to identify what organizations have to redesign when AI becomes part of how work actually gets done.

Across the nine cases, the evidence points to nine interdependent architecture domains:

1. Task Architecture — what people do, what AI does, and what is shared.
2. Workflow Architecture — how work moves, escalates, coordinates and completes.
3. Role Architecture — how responsibilities are rebundled after tasks and workflows change.
4. Capability Architecture — what people must be able to do in the redesigned system.
5. Learning Architecture — how judgment and expertise are formed when AI changes the work through which people used to learn.
6. Authority Architecture — who may review, approve, override, decide and remain accountable.
7. Career Architecture — where people enter, move and progress as roles change.
8. Workforce Flow Architecture — who is hired, redeployed, reskilled, not replaced, or exited.
9. Measurement Architecture — how institutions distinguish adoption from capability, productivity from workforce change, and technological activity from worker outcomes.

SECTION 01

These domains are not a sequence. An organization can have advanced workflow automation but weak career pathways. It can have explicit human authority without a major workforce transition. It can build extensive learning infrastructure before it can demonstrate a clean productivity or headcount outcome. The architecture matters precisely because changes in one domain can create liabilities in another.

Salesforce makes that interaction visible. Agentforce absorbed a large share of recurring customer-support interactions, while human support shifted toward more complex, consultative and proactive work. Salesforce also redeployed hundreds of support engineers and reported a movement in support headcount from roughly 9,000 to roughly 5,000. The evidence supports lower support staffing requirements, redeployment, attrition and non-backfill as material mechanisms; it does not support the simplified claim that AI “laid off 4,000 people.” [1] [2]

DBS shows a different path. The bank connected AI adoption to formal Operating Model Transformations that combine process redesign, role redesign, reskilling and redeployment. More than 11,000 employees in roles expected to be significantly changed by AI were identified for deeper upskilling or reskilling, and 96% had commenced training at the reporting point. The bank completed nine operating-model transformations in 2025. [5][6]

PwC exposes a less visible cost of automation: the possibility that AI compresses the work through which juniors historically learned. PwC is using AI across audit research, data preparation and specialized audit tasks while preparing early-career auditors to review, supervise and exercise judgment earlier. At the same time, the firm's US entry-level hiring plans have come

under pressure for multiple reasons, including technology, low attrition and delivery-model change. The challenge is not merely to preserve junior jobs; it is to preserve the capability-formation function of apprenticeship. [9][10][11]

LONGi and JD show why domain expertise can remain valuable after the original task is automated. At LONGi's Jiaxing Lighthouse Factory, AI performs substantial quality inspection, humans double-check flagged defects, and a documented former quality inspector moved into AI-model training using frontline expertise. At JD, customer-service experience has been repurposed into AI training, agent configuration and related AI-operations work. [13][14][16][17]

SimplifyNext illustrates the other side of the apprenticeship problem. Through Singapore's Company-Led Training programme, employment, structured learning, mentoring and real AI project work are joined into one pathway. IMDA reported close to 50 trainees hired, trainee participation in nearly half of more than 35 AI-powered client projects, and more than 90% of trainees progressing into some form of project-leading responsibility, although the precise authority associated with “leading” is not publicly defined. [20][21]

Kaiser Permanente shows that automation can be scaled while professional authority remains explicit. Ambient AI creates draft clinical documentation, while physicians review and edit the draft and remain responsible for medical decision-making. In The Permanente Medical Group, 7,260 physicians used ambient AI across more than 2.5 million encounters; estimated documentation-time savings were nearly 16,000 hours. Those savings do not establish better clinical outcomes or physician headcount reduction. [23][24]

SECTION 01

Deutsche Telekom provides the contrasting authority case. Its RAN Guardian Agent can identify network situations, analyze conditions and autonomously execute corrective actions, including reallocating network resources and changing configurations. The production evidence is strong on agentic action and cycle-time compression. It is much less explicit on action-level human approval, veto and override rights in that specific use case. [26][27]

BNY shows what happens when AI building itself becomes a workforce capability. By the end of 2025, BNY reported that all employees had access to Eliza, nearly half were building AI agents, 160 enterprise AI solutions were in production and 134 “digital employees” were live. The company embedded role-based learning and early-career AI pathways into the operating model. These metrics demonstrate behavioral scale; they do not mean 134 human jobs were eliminated or that training automatically equals proficiency. [29][30]

Together, the cases lead to a more demanding view of AI transformation.

AI adoption changes tools. AI institutionalization changes the architecture of work.

That shift creates several cross-case findings.

First, AI transformation propagates unevenly. Technology and workflow change can move faster than role redesign, learning systems, career pathways or workforce accounting. An organization can therefore look advanced on adoption metrics while carrying unresolved institutional debt.

Second, domain expertise increasingly becomes both operating capital and governance capital. The expert is not always displaced when execution is automated. Expertise may be needed to train the model, identify

exceptions, verify outputs, design workflows and define what the system should be allowed to do. This is what AIGW calls Transition Capital.

Third, learning architecture becomes a constraint on sustainable automation. AI can remove the production value of a junior task faster than an institution can replace its learning value. If organizations automate apprenticeship without rebuilding capability formation, they accumulate Learning Debt.

Fourth, career architecture determines whether productivity becomes mobility or exclusion. Salesforce and DBS show internal movement; SimplifyNext builds an entry pathway; PwC illustrates pressure on entry seats. The technology itself does not determine which of these outcomes occurs.

Fifth, authority becomes a design variable as AI moves from recommendation to action. “Human in the loop” is too imprecise. Human presence does not tell us whether the person has review rights, approval rights, override capability, final decision authority or accountability. Kaiser and Deutsche Telekom make that difference visible from opposite directions.

Sixth, employment effects often appear through flows before they appear as a clean headcount story. Reduced hiring, non-backfill, redeployment, attrition and role creation may all alter labor demand without one discrete “AI layoff” event. Recent external labor-market evidence reinforces the need to distinguish these mechanisms: Stanford researchers find no widespread economy-wide AI displacement in US payroll data through mid-2026, but a widening employment shortfall for workers aged 22–25 in highly AI-exposed occupations, with adjustment appearing primarily through reduced hiring rather than increased separations. [36]

Seventh, distributed AI building creates a new management object: the hybrid human + digital workforce. BNY and Deutsche Telekom show that the issue is no longer only whether employees use AI. Organizations increasingly have to decide how agents are identified, permissioned, supervised, evaluated and integrated with human teams.

The implication for enterprise leadership is that no single function owns the transition.

The CIO or Chief AI Officer cannot solve learning and career pathways alone. The CHRO cannot redesign work without understanding agent capabilities and


system controls. Business leaders cannot treat productivity as a technology metric if the gains depend on verification, exceptions and changes in staffing demand. Boards cannot assess AI risk if human authority is described only in principle. L&D cannot count training hours if capability is never demonstrated in production.

The report therefore proposes an integrated AIGW Workforce Architecture Framework. It is not a maturity model and it is not a validated scoring index. It is a design and diagnostic system for asking whether the institutional structures around AI are keeping pace with the technology.

FIGURE 01

From AI Deployment to Workforce Architecture

Why a technology deployment question becomes an organisational design question



AI adoption changes tools. AI institutionalization changes the architecture of work.

Source: AIGW analysis based on approved enterprise workforce transition cases and related deep-dive evidence.

Note: The figure is an analytical synthesis; enterprise transitions may develop through different pathways.

The central management question is no longer simply:
Where can we deploy AI?
It is:

What kind of organization will our AI deployment create — and have we redesigned the systems through which people work, learn, exercise judgment and build careers to match it?

SECTION 02

AI Adoption Is No Longer Enough

The adoption era is not over. Many organizations are still deciding which models, copilots and agents to use, how to secure them and where to find credible business value. But adoption is no longer a sufficient unit of analysis.

The distinction matters because an enterprise can report impressive AI activity while changing little about the underlying institution.

A large license count may show access. It does not show whether work changed.

A large training count may show participation. It does not show whether capability improved.

A large number of agents may show experimentation. It does not show whether those agents are in production, whether they have decision rights, whether people supervise them effectively or whether roles have changed.

A large time-saving estimate may show technical potential. It does not show whether the released capacity becomes more output, better service, internal mobility, reduced hiring, non-backfill or workforce reduction.

The next phase of enterprise AI therefore requires a distinction between adoption and institutionalization.

Adoption asks whether technology is present and used.

Institutionalization asks whether AI has become part of the formal and informal systems through which the organization operates.

Institutionalization becomes visible when the organization has to make durable decisions about questions such as:

- Which tasks remain human and which can be delegated?
- Where do agents enter and act inside a workflow?
- What remains inside a job after routine execution is compressed?
- What capabilities become more valuable?
- How will inexperienced people learn those capabilities?
- Which AI outputs require review or approval?
- Who can override an autonomous action?
- What happens to people whose roles change?
- What new entry paths or destination roles exist?
- How will the organization measure transition quality rather than just AI usage?

External labor-market evidence points in the same direction. The ILO's 2025 refined global index estimates that one in four workers worldwide is in an occupation with some degree of generative-AI exposure, while emphasizing that most occupations still contain tasks requiring human input and are therefore more likely to be transformed than wholly made redundant. [35] The World Economic Forum's Future of Jobs Report 2025 similarly shows employers expecting to use several workforce strategies at once: 73% planned faster process and task automation, 63% planned to augment workers with technology, 51% planned to move people from declining to growing roles, and 41% anticipated some workforce reduction because of skills obsolescence. [34]

The coexistence of those strategies is important.

There is no single enterprise response called “AI adoption.” The same organization may automate one work unit, augment another, hire new skills, stop hiring

for a declining role, redeploy experienced workers and redesign an entry pathway at the same time.

That is why workforce architecture becomes a more useful management object than tool adoption alone.

FIGURE 02

From Adoption to Institutionalization

Technology presence is not the same as organizational transformation

01	AI ADOPTION Evidence that AI is being used	02	AI INSTITUTIONALIZATION Evidence that the organization itself is changing
01	Licences deployed AI tools and platforms are made available to employees.	01	Workflows redesigned Core processes are restructured around AI-enabled ways of working.
02	Active users People are using AI tools in their day-to-day work.	02	Roles rebundled Jobs are redesigned, combined or split as tasks and responsibilities change.
03	Training completed Employees have completed introductory or advanced AI training.	03	Capabilities changed New capabilities are built at scale and embedded in how work is performed.
04	Pilots launched Initial use cases are tested in parts of the business.	04	Authority reassigned Decision rights are reallocated between people, AI systems and oversight functions.
05	Agents built Custom or configured AI agents are deployed for specific tasks.	05	Career pathways rebuilt Progression routes are updated to reflect new AI-enabled roles and skill requirements.
06	Productivity estimates Initial estimates of time, cost or output improvements are recorded.	06	Workforce flows altered Hiring, mobility, reskilling and attrition patterns change as the operating model evolves.

Adoption metrics tell us whether technology is present. Institutional evidence tells us whether work has changed.

Source: AIGW analysis based on approved enterprise workforce transition cases and related implementation evidence.
Note: Institutionalization is an analytical distinction, not a maturity score.

What changes when AI becomes institutional

Three shifts are especially important.

The first is a shift from individual productivity to workflow productivity. A copilot may make one employee faster. A production agent changes who performs a step, how work is handed off and where exceptions go. Deutsche Telekom's RAN Guardian is a clear example: the system does not merely help an engineer analyze the network; it can carry part of the operating loop itself. [26]

The second is a shift from skills as training content to capability as operating infrastructure. BNY's effort to move employees from AI users toward AI builders

matters because the organization's AI capacity depends on domain employees being able to identify and redesign work, not simply on a central team deploying generic tools. [29]

The third is a shift from human participation to human authority. Kaiser Permanente can automate documentation while keeping clinical decision authority with the physician. Deutsche Telekom can permit certain agentic actions on infrastructure. These are both “human-AI” systems, but the authority structure is fundamentally different. [24][26]

Once these shifts appear, workforce strategy can no longer remain a downstream response to technology deployment. It becomes part of the design of the system itself.

SECTION 03

This report uses nine purposively selected enterprise cases to identify and compare mechanisms of workforce transition.

The cases were selected because each makes a different part of the transition unusually observable. They are not intended to represent the distribution of AI outcomes across all companies, industries or countries.

These cases are purposively selected to identify and compare transition mechanisms. They do not estimate how common each mechanism is across the economy or establish that the observed outcomes will generalize to other organizations.

FIGURE 3.1

Cross-case comparison

The nine transition archetypes

Distinct enterprise settings and evidence boundaries

Organization	Primary transition	Strongest evidence	Main boundary
Salesforce	Support automation + role redesign + lower staffing requirement	Millions of AI-handled conversations; hundreds redeployed; reported support headcount ~9,000 → ~5,000	Exact split among redeployment, attrition, non-backfill and involuntary exits is not public [1] [2] [3]
DBS	AI-linked operating-model transformation	Nine OMTs; >11,000 employees identified for deeper reskilling/upskilling; 96% commenced	Bank-wide value and model counts cannot be causally attributed to specific workforce transitions [5] [6] [7]
PwC US	Entry-level audit role and apprenticeship redesign	AI embedded in audit; >17,000 assurance users of ChatNational; structured AI training; planned pressure on entry intake	Hiring contraction is multi-causal; AI alone does not explain the planned decline [9] [10] [11] [12]
LONGi	Frontline inspection automation + expert transition	AI inspection across nine procedures / 198 items; human double-check; documented inspector → AI trainer	One named worker does not establish a universal pathway; lower labor cost does not prove layoffs [13] [14] [15]

Source: AIGW synthesis; row-level sources are identified in the table.

SECTION 03

FIGURE 3.1

Cross-case comparison

The nine transition archetypes

Distinct enterprise settings and evidence boundaries

Organization	Primary transition	Strongest evidence	Main boundary
JD.com	Basic service automation + AI operations / role transition	4.2bn AI service inquiries; thousands reported job transitions/upgrades; named customer-service → AI trainer pathway	“Thousands transformed” is not thousands of net-new jobs; current denominator and destination mix remain unclear [16] [17] [18] [19]
SimplifyNext	Employment-first AI apprenticeship	Close to 50 trainees hired; real project participation; >90% reported project-leading responsibility	Long-term retention, pay and precise authority associated with “leading” are not public [20] [21] [22]
Kaiser Permanente / TPMG	Documentation automation with explicit professional authority	7,260 physicians; 2.57m encounters; ~15,791 estimated documentation hours saved	Does not establish clinical outcome improvement or physician staffing reduction [23] [24] [25]
Deutsche Telekom	Agentic network operations	Autonomous corrective actions; selected processes compressed from ~1 hour to minutes; >100 remediation actions reported after launch	Use-case-level approval, veto and override architecture is not fully public [26] [27] [28]
BNY	Employee-built agents + digital workforce	100% access; 171,000 AI learning hours; nearly 50% building agents; 134 digital employees live in 2025	Training is not proficiency; digital employees are not equivalent human FTEs; full workflow authority is not public [29] [30] [31] [32] [33]

Source: AIGW synthesis; row-level sources are identified in the table.

FIGURE 3

Nine Enterprise Transitions, Nine Different Workforce Problems

The current AIGW case base shows that enterprise AI transition does not take one form

Enterprise case	What AI changed	What happened to human work	Why the case matters
Salesforce	Customer support execution	Support work shifted toward escalation, complex cases and internal mobility as AI handled more routine interactions	Shows how AI can reshape workforce flows without reducing the story to simple “AI layoffs”
DBS	Banking workflows and service operations	Employees were reskilled and redeployed as work moved toward higher-value tasks and new operating models	Shows workforce transformation as operating-model redesign, not just tool adoption
PwC	Entry-level audit work	Junior professionals moved toward AI-assisted drafting and earlier judgment-intensive review	Shows how AI compresses apprenticeship tasks and changes capability formation
LONGi	Frontline quality inspection	Inspectors shifted from manual checking toward verification, exception handling and AI-supported production roles	Shows how industrial AI transition depends on transition capital, not only automation
JD.com	Customer-service operations	Workers moved from service execution into AI trainer and AI operations roles	Shows a visible role transition from doing the work to training the system
SimplifyNext	Workforce entry into AI delivery	Trainees entered AI work through structured, employment-linked capability building and project exposure	Shows how a new first rung can be designed rather than assumed
Kaiser Permanente	Clinical documentation workflow	AI drafted notes while physicians retained review and edit authority	Shows that “human in the loop” is really a question of authority design
Deutsche Telekom	Network operations and corrective action	The human role shifts from reviewing recommendations toward supervising more autonomous system action	Shows the governance challenge when AI moves from recommendation to action
BNY	Enterprise work orchestration through digital employees	Employees increasingly build, manage and work alongside digital workers	Shows hybrid workforce design as an organization-level operating model

Source: AIGW analysis based on nine approved enterprise workforce transition cases and related deep-dive materials.

Note: The cases illustrate distinct transition problems. They are not intended as a representative sample of all enterprise AI transitions.

Evidence discipline

The evidence base combines several source types.

First-party enterprise sources are the strongest source for internal operating design, training systems, product

deployment and management intent. They are also vulnerable to positive selection and corporate framing.

Annual and sustainability reports provide more formal organizational evidence but often aggregate outcomes at enterprise level, making causal attribution to a specific AI workflow difficult.

SECTION 03

Nine Enterprises, Nine Transition Episodes

Government and public-institution sources are particularly important in the LONGi and SimplifyNext cases because they independently document aspects of worker transition and programme design.

Attributable journalism is used when it adds information not present in company publications, such as PwC's internal hiring plans or executive descriptions of staffing change. Such evidence remains attributable rather than treated as independently audited outcome data.

External institutional research — including the ILO, OECD, World Economic Forum, NIST, Stanford Digital Economy Lab, EU regulatory materials and UNDP — is used for context, triangulation and boundary testing. It does not substitute for case evidence.

The report separates five analytical categories:

- Observed: directly supported by a source.
- Reported: stated by the organization or an attributable actor.
- AIGW Interpretation: a cross-case mechanism inferred from the evidence.
- Implication: a plausible consequence that remains under-evidenced.
- External Benchmark: macro or institutional evidence used to contextualize the enterprise cases.

This distinction matters because enterprise AI evidence is often asymmetric. Organizations publish

adoption and productivity metrics more readily than worker outcomes, pay effects, failed transitions or detailed headcount flows. The absence of those data is itself part of the research finding.

What the portfolio can and cannot support

The nine cases collectively cover the architecture well enough to construct an integrated framework. But the portfolio has clear selection limits.

It overrepresents large or sophisticated organizations, visible AI adopters and institutions willing to disclose operational change. It underrepresents small and medium-sized enterprises, failed deployments, worker resistance, collective bargaining, pay deterioration, long-term career outcomes and organizations that adopted AI without generating measurable value.

The report therefore uses the cases to answer:

What transition mechanisms are becoming visible, and what institutional questions do they create?

It does not answer:

How common is each mechanism across the whole economy?

That boundary should remain visible throughout the report.

SECTION 04

Task Architecture — Deciding What AI Should Do

The most immediate unit of AI transformation is the task.

A task is a discrete activity inside work: answer a recurring customer question, inspect a product defect, summarize a consultation, prepare a research note, reconcile information, draft an audit document, analyze a network condition, validate a payment exception.

At this level, the design problem appears simple:

Can AI perform the activity faster, cheaper or more consistently?

But the nine cases show that this is not the only question that matters.

A task can have at least two kinds of value:

- production value — the direct value of completing the activity;
- developmental value — the value of the activity as a way for people to build context, judgment, pattern recognition or accountability.

Those two values do not always move together.

A repetitive task can be expensive to perform and add little developmental value. Automating it can improve both productivity and worker experience.

But a relatively junior task can also look repetitive while quietly performing an apprenticeship function. Researching source material, checking data, drafting a first analysis or reviewing recurring exceptions can expose a worker to the patterns from which later judgment is built.

That is why task architecture cannot be reduced to a binary choice between “automate” and “keep human.”

A more useful set of design choices is:

Automate · Augment · Delegate · Protect · Redesign

Automate

The organization may choose to remove a low-learning, repeatable task from the human workload entirely.

LONGi provides a clear example. Automated visual-inspection systems now perform a substantial share of quality inspection across nine production procedures and 198 inspection items. Human work is no longer organized around visually inspecting every item from scratch. [13][14]

Salesforce similarly delegates a large share of recurring support conversations to Agentforce. The important point is not that every conversation is automated, but that repetitive support volume is no longer treated as a purely human capacity problem. [1]

Augment

The task remains human-owned, but AI compresses part of the work.

Kaiser Permanente's ambient documentation is a strong example. The system listens, transcribes and drafts. The physician still reviews the note and retains medical decision authority. [23][24]

The task is changed, but it is not transferred wholesale.

Delegate

In agentic systems, the organization can go further and delegate not only analysis but action.

Deutsche Telekom's RAN Guardian Agent can identify network events, evaluate conditions and autonomously initiate certain corrective actions. [26]

SECTION 04

At this point, task architecture and authority architecture begin to overlap. The question is no longer only whether AI can do the activity, but whether it may do so without prior human approval.

Protect

Some tasks should remain meaningfully human because they perform a developmental or governance function even if AI could accelerate them.

PwC's entry-level audit evidence makes this problem visible. If junior professionals stop doing all forms of research, data preparation, testing and first-pass analysis, the organization may save time while weakening the process through which future reviewers learn how evidence behaves in practice. [9] [10]

“Protect” does not necessarily mean preserve the old task unchanged. It means preserve the learning function.

Redesign

Often the right answer is neither to automate the task fully nor to keep it intact.

The organization may redesign the task so that AI performs routine execution while the human performs interpretation, verification, exception handling or communication.

LONGi's quality workflow is a good example: AI detects, while humans double-check flagged defects and some worker expertise moves into AI-model training. JD follows a similar pattern in customer service, where routine service is increasingly handled by AI while human work expands into complex scenarios, domain-specialist service and AI operations. [13][16][17]

The task-allocation mistake

A common enterprise error is to automate based only on technical feasibility.

The better question is:

What is the production value of this task, what is its learning value, and what human capability would disappear if the task disappeared?

This distinction becomes especially important for early-career work.

The ILO's 2025 global exposure index reinforces the need to think at task level. It estimates that one in four workers globally is in an occupation with some generative-AI exposure, but argues that transformation is more likely than wholesale redundancy because most occupations still contain tasks requiring human input. [35]

That macro finding aligns with what the enterprise cases show: the most visible changes often occur inside roles before they appear as complete job elimination.

Decision rule

For every materially affected role, institutions should create a task map with at least five fields:

- task;
- AI treatment: automate / augment / delegate / protect / redesign;
- production value;
- developmental value;
- human control requirement.

The map should be reviewed not only by technology leaders, but also by the managers responsible for capability formation and workforce planning.

SECTION 05

Workflow Architecture — Redesigning How Work Moves

Task automation becomes institutionally consequential when it changes the workflow around the task.

Workflow architecture is the sequence through which work moves from trigger to completion. It includes handoffs, escalation, sequencing, exception paths, system coordination and the points at which humans or AI systems are allowed to act.

This distinction matters because two organizations can automate the same task and create very different operating models.

One organization may use AI to generate a recommendation that a human must review before anything changes.

Another may allow AI to execute automatically unless a threshold is breached.

A third may use several agents that coordinate with one another before a human becomes involved.

The productivity and governance implications are different in each case.

DBS: from use cases to operating-model transformation

DBS is the clearest example in the portfolio of treating workflow redesign as an enterprise discipline rather than an accumulation of AI pilots.

Its Operating Model Transformation approach combines process re-engineering, human-AI collaboration, role redesign, reskilling and organizational simplification. Nine OMTs were completed in 2025. [5][6]

The implication is important:

A workflow should not be considered transformed merely because one step has become faster.

The organization has to ask whether the whole process now has a different structure.

In KYC, for example, DBS reports significant efficiency improvement in name screening. In credit preparation, later agentic work uses multiple agents to gather information and draft a credit memo while a human credit manager retains the final call. [7][8]

The workflow is not “human replaced by AI.” It is a redistribution of analytical work, drafting work and decision authority across a new operating sequence.

JD: basic service, complex service, AI operations

JD provides a different workflow pattern.

Basic customer interactions can be handled by AI. Complex cases remain with people. Behind the front line, humans maintain knowledge, diagnose response failures, configure agents and tune the system. [16][17]

The workflow therefore has at least three layers:

1. automated basic service;
2. human handling of complex or sensitive scenarios;
3. backstage AI operations that improve the system itself.

That third layer is easy to miss.

When AI becomes a service channel, the organization creates new work around the maintenance of that channel.

This is why task substitution can coexist with role creation.

SECTION 05

Deutsche Telekom: when the workflow closes the loop

Deutsche Telekom pushes the logic further.

RAN Guardian can identify an event, assess network conditions and implement corrective actions. In later reporting, Deutsche Telekom said the system autonomously triggered more than 100 remediation actions in its first month around Christmas-market events. [26][27]

The workflow therefore changes from:

human observes → human analyzes → human decides → human executes

into something closer to:

agent observes → agent analyzes → agent acts within defined system capabilities → human monitors / intervenes where required

The public evidence does not fully specify the action-level human-control matrix, but the workflow redesign itself is clear.

BNY: AI as a team participant

BNY's "digital employees" make workflow architecture visible in yet another way.

The bank describes these systems as multi-agentic AI participants with distinct identities, credentials and supervisors. In payment workflows, digital employees perform routine operational work alongside human teams. [29][30]

The organization therefore has to design not only human-to-human handoffs, but human-to-agent and agent-to-human handoffs.

That creates a new operating question:

Who owns the workflow when neither the human nor the AI performs every step?

A workflow is more than a process map

A mature AI workflow description should specify:

- trigger;
- AI action;
- human action;
- escalation condition;
- exception path;
- approval point;
- rollback or correction path;
- logging;
- ownership.

Without these elements, an organization may know that "AI is in the workflow" without understanding how accountability is actually distributed.

AIGW interpretation

The cases suggest a progression in workflow intensity:

AI beside the workflow — an optional aid to a human task.

AI inside the workflow — a routine step becomes AI-enabled.

AI coordinating the workflow — AI organizes or routes multiple steps.

AI acting within the workflow — AI can execute consequential actions under specified conditions.

This is not a maturity ladder. Different contexts should stop at different points.

The relevant question is whether the workflow's technical autonomy and its governance architecture are aligned.

SECTION 06

Role Architecture — Rebundling the Job

Jobs rarely disappear in one clean motion.

More often, AI changes the bundle of tasks, responsibilities, interactions and decision rights inside the job.

That is why the concept of role architecture is more useful than the binary question of whether a job exists.

AIGW defines a role through six components:

1. Task portfolio — what the person actually does.
2. Workflow position — where the person sits in the process and with whom or what they interact.
3. Decision rights — what the person may decide, approve, verify or override.
4. Capability bundle — what the role requires the person to know and do.
5. Learning pathway — how a person becomes competent in the role.
6. Staffing envelope — how much human capacity the organization requires in that role.

FIGURE 04

The Six Components of Role Architecture

A role changes when more than a job title changes

Task Portfolio

What activities remain in the role?

Workflow Position

Where does the person sit in the AI-enabled process?

Decision Rights

What may the person decide, approve or override?

Capability Bundle

What must the person know and be able to do?

Learning Pathway

How is capability built inside the role?

Staffing Envelope

How many people and what capacity does the role require?

**A role changes when one or more of these components change—
not only when the job title changes.**

Source: AIGW analysis based on approved enterprise workforce transition cases.

Note: Role architecture can change even when titles remain unchanged.

Salesforce: role content and staffing change together

Salesforce is valuable because several components move at once.

Agentforce takes a larger share of recurring support interactions. Human support work shifts toward complex cases, adoption, implementation health and proactive customer engagement. Some employees move into Success, Renewals and Forward Deployed

Engineering. At the same time, the total staffing requirement for the support function declines. [1][2]

The role is therefore not simply “the same support engineer with AI.”

The task portfolio changes.

The workflow position changes.

The destination roles for some workers change.

And the staffing envelope changes.

SECTION 06

PwC: the role becomes more senior before the worker becomes senior

PwC exposes a different form of rebundling.

As AI performs more research, data preparation and specialized execution, junior auditors are expected to review and exercise judgment earlier. [9][11]

The role acquires more senior-like responsibilities before the employee has accumulated the experience that historically came before them.

This creates a design challenge:

How do you increase role complexity without pretending the worker has already acquired the judgment that normally comes from experience?

The answer cannot be only more training content. It requires changes in supervision, feedback and the progression of responsibility.

LONGi and JD: role redesign through retained expertise

LONGi and JD show how roles can be rebundled around the expertise that remains scarce after automation.

At LONGi, inspection expertise becomes useful for verifying AI results and training the system. At JD, customer-service experience becomes useful for identifying knowledge gaps, configuring agents and maintaining service quality. [13][16]

In both cases, the worker's value shifts from performing the entire original task toward helping the system perform it better.

BNY: the role includes AI building

BNY broadens the role further.

A domain employee may be expected not only to use AI, but to build agents, supervise digital employees or identify workflows that should be redesigned. [29]

The role therefore acquires an explicit system-design dimension.

That suggests a wider organizational shift:

The workforce is becoming a partial designer of the operating system in which it works.

Role redesign must include the staffing envelope

One reason role-redesign programmes become misleading is that they describe enriched work while ignoring how many people the organization expects to perform it.

Salesforce shows why this matters.

A role can become more complex and consultative while the total number of people required in that role falls. [2]

PwC shows a similar tension from the entry side: the work expected from junior professionals can become more demanding while the number of entry seats becomes smaller. [11]

Therefore a credible role architecture must answer two questions together:

What will this role contain?

and

How many people will the organization require in it?

Ignoring the second question produces a role-design story without a workforce strategy.

SECTION 07

Capability Architecture — What Becomes Scarce When Execution Becomes Cheap?

As AI reduces the cost of some forms of execution, the value of human capability does not disappear uniformly.

Instead, some capabilities become more important because they are needed to direct, verify, contextualize or govern AI-enabled work.

Across the nine cases, five recurring capability families stand out:

- AI Fluency — understanding how to use AI tools effectively and responsibly.
- Judgment & Verification — evaluating whether output is correct, appropriate and sufficiently supported.
- Domain Expertise — understanding the business, client, production or professional context.
- Human Interaction — communication, trust, relationship management and negotiation.
- AI Orchestration — configuring, supervising, improving or coordinating AI systems and agents.

A sixth capability increasingly connects them:

- Business Translation — turning a real organizational problem into an AI-enabled workflow while respecting

operational and governance constraints.

The OECD's 2026 work on skills in the AI age similarly emphasizes that AI changes skill requirements rather than simply reducing demand for human capability, including greater importance for data use, interpretation and adaptation to changing work. [38]

Transition Capital

AIGW uses the term Transition Capital to describe existing human knowledge that retains value after the original task is automated.


Transition Capital can include:

- product knowledge;
- customer context;
- tacit process knowledge;
- professional judgment;
- pattern recognition;
- institutional memory;
- knowledge of exceptions and failure modes.

FIGURE 05

Transition Capital

Where existing expertise can move when tasks are automated or redesigned



Interpretation

Expertise can remain valuable even after the task that originally used it is automated.

Illustrative case anchors

LONGi | JD.com | Salesforce | BNY

Source: AIGW analysis based on approved enterprise workforce transition cases.

Note: Transition capital refers to capabilities that can be redeployed into new forms of work after automation or workflow redesign.

LONGi: tacit inspection expertise becomes model knowledge

The LONGi case is particularly clear.

A former frontline quality inspector moved into AI-model training, using experience from visual

inspection to teach the system how to recognize defects. [13]

The expertise was not obsolete when the inspection task was automated.

It became an input to the automation system itself.

That is a profound shift in the economic role of experience.

SECTION 07

JD: customer-service knowledge becomes AI-operations capability

JD's customer-service transition follows the same mechanism in a different environment.

Experienced workers can identify whether an AI response problem comes from configuration, knowledge content or more fundamental system logic. They update content, configure agents and help improve the service system. [16][17]

The capability is not generic "AI skill." It is domain expertise translated into system improvement.

Salesforce: product knowledge travels across roles

Salesforce shows Transition Capital through redeployment.

Support engineers possess product and customer knowledge that can be used in Success, Renewals and Forward Deployed Engineering. [1][2]

The transition is easier because the organization can identify the part of the old role that remains valuable after routine support volume falls.

BNY: domain expertise becomes build capability

BNY's employee-building model generalizes the mechanism.

When employees outside engineering build agents, the value comes precisely from their knowledge of payments, operations, risk, clients or other domains. [29]

The AI platform becomes a channel through which domain knowledge can be formalized into workflow logic.

Capability architecture is not a training catalogue

A common enterprise mistake is to respond to role change with a generic AI curriculum.

But capability architecture should begin with the redesigned role.

For each role, institutions should specify:

- which old capabilities remain valuable;
- which capabilities lose importance;
- which new capabilities become essential;
- which capabilities require real-world practice rather than classroom exposure;
- how proficiency will be demonstrated.

This last point is critical.

Access ≠ use ≠ training ≠ demonstrated capability ≠ role transition ≠ worker outcome.

BNY's scale of AI training is important evidence of institutional investment. It should not be treated as proof that every employee has equivalent proficiency. [30]

DBS's figure that 96% of identified employees had commenced deeper training is similarly significant, but "commenced" is not "completed" and neither is the same as a successful role transition. [6]

Capability architecture therefore needs evidence beyond participation.

SECTION 08

Learning Architecture — How Expertise Is Reproduced

AI can improve the productivity of experienced workers while simultaneously weakening the mechanisms through which inexperienced workers become experienced.

This is one of the least visible risks in enterprise AI adoption.

The problem is not that all routine work is educational. Much routine work is genuinely low-value. The problem is that organizations often do not know which repetitive activities are also carrying learning value until the activities disappear.

AIGW calls the resulting gap Learning Debt:

AI can remove the production value of a task faster than the institution replaces the task's learning value.

The debt accumulates when productivity improves today but the pipeline for future judgment weakens.

PwC: the apprenticeship paradox

PwC is the strongest case in the portfolio for this problem.

Audit has historically relied on an apprenticeship model. Junior professionals perform structured work, receive review and gradually build the judgment required for more senior responsibility.

AI changes that sequence.

Research becomes faster. Data preparation becomes more automated. AI can assist specialized audit execution. The junior professional is expected to move more quickly toward review and judgment. [9][10]

The immediate opportunity is obvious: less time on low-level execution and earlier exposure to more valuable work.

The risk is equally important: if the organization removes the experience that previously created judgment, it has to build a new way to create judgment.

PwC's response includes technical and leadership training, small-group reviewer development, coaching, revised professional frameworks and greater emphasis on human and AI skills. [10]

This is not simply “upskilling.” It is a redesign of the learning system around a changed job.

SimplifyNext: an AI-native first rung

SimplifyNext starts from the opposite direction.

Instead of preserving an old apprenticeship after AI changes it, the company uses an employment-first model in which participants are hired, trained, mentored and exposed to real AI projects from the beginning. [20][21]

The relevant design principle is guided autonomy.

The entrant is not protected from real work. Nor are they expected to operate without support.

They receive meaningful production exposure inside a supervised environment.

AIGW calls this Demand-Backed Apprenticeship:

employer demand → employment → structured learning → mentored production work → increasing responsibility.

BNY: learning becomes part of the operating model

BNY adds another variation.

AI learning is not a separate academy running beside the business. It is connected to role-based pathways, agent building, early-career development and production deployment. [29][30]

The important shift is from:
learn AI, then perhaps use it

into:

learn by identifying, building and operating real AI-enabled workflows.

That makes learning architecture part of the production architecture.

FIGURE 06

Old Apprenticeship vs AI-Native Capability Formation

Capability must still be built even when AI changes the work juniors once learned from

Traditional pathway		AI-native pathway	
How juniors built capability in the pre-AI model		How juniors can build capability in an AI-enabled model	
01	Routine execution Juniors perform well-defined tasks as directed.	01	AI-assisted execution Juniors use AI to perform and iterate on tasks, with human direction.
02	Repetition Tasks are repeated to build familiarity and fluency.	02	Verification Outputs are checked for accuracy, completeness and appropriate use.
03	Senior review Work is checked, corrected and refined by more experienced colleagues.	03	Explanation Juniors explain how the output was produced, including assumptions and limitations.
04	Context accumulation Juniors build situational understanding across varied cases and edge conditions.	04	Guided practice Juniors apply AI with increasing complexity under supervision, including edge cases.
05	Judgment Over time, experience supports independent judgment and decision-making.	05	Bounded ownership Juniors take ownership of defined problems or workstreams with appropriate guardrails.
		06	Judgment Over time, experience with AI-supported work enables independent judgment and decision-making.

The developmental function must survive even when the old task does not.

Illustrative case anchors: **PwC** | **SimplifyNext** | **BNY**

Source: AIGW analysis based on approved enterprise workforce transition cases.
Note: The pathways are analytical models, not universal sequences.

SECTION 08

What good learning architecture requires

The WEF's 2026 entry-level report identifies job access, job design, talent pipelines and education-system alignment as central pressure points in AI-era early-career work. [37] The ILO's Quality Apprenticeships Recommendation similarly emphasizes structured and compensated learning, qualified supervision, recognized occupational competence and an appropriate balance between apprentices and workers. [39]

Those external benchmarks reinforce several design requirements:

Real work

Capability forms under real constraints, not only simulations.

Qualified supervision

If AI allows juniors to take on higher-context tasks earlier, the senior review layer becomes more important, not less.

Explicit progression

A first rung is only successful if it creates the second rung.

Verification and explanation

Workers should not be evaluated only on whether they can produce an output with AI. They should be able to explain, challenge and modify it.

Protected reflection

Faster work can reduce the time in which people notice patterns, discuss mistakes and understand why a decision was made. Learning architecture has to deliberately restore those moments.

The capability test

A useful learning ladder is:

Access → Use → Verify → Explain → Judge → Modify
Independently

The first two levels can make a worker look productive quickly.

The last four are closer to professional capability.

That distinction becomes increasingly important as AI makes polished output easier to produce.

Authority Architecture — Who May Decide, Act and Override?

As AI moves from recommendation toward action, authority becomes a first-order design problem.

The phrase “human in the loop” is no longer precise enough.

A human can be present in a workflow without holding meaningful authority. A person may see an AI output but lack the right to change it. A human may review a decision after the system has already acted. A manager may be accountable for an autonomous agent without approving every action it takes.

For governance, these are not equivalent arrangements.

AIGW therefore distinguishes several layers of human authority:

Permission → Scope → Monitoring → Verification → Review → Approval → Override / Rollback → Final Authority → Accountability

Not every workflow requires every layer. The point is to specify which ones exist.

FIGURE 07

The Human Authority Stack

Human involvement in AI workflows can occur at different layers of control, review and accountability

AI-enabled workflow	01	Permission May the AI system act in this domain at all?	Case evidence anchors
	02	Scope Setting What is the system allowed to do, and within what boundaries?	
	03	Monitoring Who watches system behaviour as it operates?	
	04	Verification / Review Who checks the output, recommendation or draft?	
	05	Approval Who must approve before an action or output is accepted?	
	06	Override / Rollback Who can stop, reverse or overrule the system?	
	07	Final Authority Who ultimately decides in consequential cases?	
	08	Accountability Who remains answerable for the outcome?	
			Kaiser Permanente Physician review/edit authority
			LONGi Human defect verification
			DBS Manager final decision
			Deutsche Telekom Autonomous action; override rights not yet evidenced

Figure 07. The Human Authority Stack

Human involvement in AI workflows is not a single condition. It can occur at multiple layers of authority, from permission and scope setting to review, approval, override, final decision-making and accountability.

Source: AIGW analysis based on approved enterprise workforce transition cases, especially Kaiser Permanente, LONGi, DBS and Deutsche Telekom.

Note: The stack is an analytical framework, not a universal process sequence. Case anchors illustrate publicly evidenced authority layers; missing evidence for a layer does not imply that the layer is absent in practice.

SECTION 09

Kaiser Permanente: bounded AI authority

Kaiser Permanente provides one of the clearest examples of explicit professional authority inside an AI-enabled workflow.

Ambient AI listens to the clinical encounter and produces a draft note. The physician reviews and edits the draft before it becomes part of the medical record, and the tool does not make clinical-care decisions or recommendations. [23][24]

The authority boundary is therefore clear in at least three ways:

- the AI may draft;
- the physician must review the clinical documentation;
- the physician retains clinical decision authority.

Patient permission also matters. Kaiser Permanente's implementation descriptions state that clinicians seek patient permission before recording the encounter. [24]

This creates multiple human authority holders in one workflow: the patient has a consent-related right, the physician has professional review and clinical decision authority, and the organization retains governance responsibility over the deployed system.

DBS: AI can prepare; the human can decide

DBS's later agentic credit-memo workflow provides another clear boundary.

AI agents gather information and help prepare a credit memo, but the credit manager retains the final call. [8]

This distinction is central.

The system can do meaningful analytical and drafting work without owning the decision that creates the

financial commitment.

That is not the same architecture as a fully autonomous credit decision.

LONGi: verification is visible; final authority is not

LONGi's quality-inspection case provides a narrower form of human control.

AI detects potential defects and human workers double-check flagged items. [13]

Verification is explicit.

Broader final production authority is not sufficiently detailed in the public evidence.

The correct conclusion is therefore not that “humans remain fully in control.”

It is that human verification is directly evidenced.

Deutsche Telekom: the Authority Gap

Deutsche Telekom makes the opposite side of the problem visible.

RAN Guardian can autonomously reallocate network resources or change configurations. [26]

The company also describes interaction between AI and network experts. But the use-case-level public material does not fully specify whether a human must approve each action, what thresholds trigger intervention, who can veto an action, or how rollback works.

AIGW calls this an Authority Gap:

autonomous action is publicly evidenced, but the corresponding human approval, override and accountability architecture is not described at the same level of specificity.

SECTION 09

Later Deutsche Telekom governance material describes a risk-based agent model with identities, permissions, logged actions and the possibility of requiring human approval for particularly critical actions. [28]

That is relevant governance evidence.

But enterprise policy should not be silently treated as proof of the exact configuration of an earlier operational use case.

The unit of governance is the action

One of the most important implications is that “autonomy” should not be assigned to an entire system as if it were a single property.

A system may:

- read data autonomously;
- classify events autonomously;
- make low-risk changes autonomously;
- require approval for higher-risk changes;
- prohibit certain actions entirely.

The relevant governance question is therefore:

Which action is autonomous, under what conditions, with what consequence, and with what human recourse?

NIST's AI Risk Management Framework supports this more differentiated approach. It explicitly calls for

organizations to define and differentiate human roles and responsibilities in AI decision-making and oversight, recognizing that human-AI configurations range from fully manual to highly autonomous. [40]

The EU AI Act's Article 14 similarly states that human-oversight measures for high-risk systems should be commensurate with the risks, level of autonomy and context of use. [41]

These benchmarks should not be read as validating every case in this report. They reinforce the broader principle that oversight should be designed at a level more precise than “a human is involved.”

Verification economics

Authority also has a productivity cost.

If AI generates an output in seconds but a qualified person must spend substantial time verifying it, the gross time saving overstates the reliable benefit.

A more realistic relationship is:

gross automation benefit – verification burden – exception handling – control infrastructure – error remediation = reliable net benefit

Kaiser Permanente's ambient AI case makes this visible. The system saves documentation time, but the physician still has to inspect and correct the draft. [23]

In high-stakes domains, this is not inefficiency. It is part of the operating cost of reliable automation.

Authority architecture therefore belongs inside productivity measurement, not outside it.

SECTION 10

Career Architecture — Where Does the Worker Go Next?

A role transition is incomplete if the organization can describe the new work but cannot describe the worker's next move.

Career architecture is the longitudinal layer of workforce design.

It asks:

Where can people enter, where can they move, what capabilities unlock the next step, and what happens to workers whose old role contains less human work?

A useful transition structure is:

Current Role → Transition Capital → Capability Gap → Destination Role → Next Rung

This structure matters because reskilling is not a destination.

A worker can complete a course and still have nowhere credible to go.

Salesforce: mobility as a transition mechanism

Salesforce provides unusually direct evidence of named destination roles.

Support employees moved into Success, Renewals, Forward Deployed Engineering and other growing functions. [1][2]

The transition is possible because product and customer knowledge remain valuable outside the compressed support queue.

The important infrastructure is therefore not only training.

It is the ability to identify transferable expertise, match people to adjacent demand and create roles that can

absorb that expertise.

DBS: role change supported by career infrastructure

DBS similarly combines reskilling with internal mobility infrastructure.

Its public reporting describes workers moving from customer service into relationship and wealth-planning pathways, while tools such as iGrow and career-advisory support help identify opportunities and capability gaps. [5][6]

This is an important distinction:

workforce transition requires both capability formation and a labor market inside the organization.

Without internal demand, reskilling can become a credentialing exercise disconnected from real work.

SimplifyNext: building the first rung

SimplifyNext addresses the career problem from the entry side.

Its employment-first model brings fresh graduates and mid-career entrants into paid work, structured learning and real AI delivery. [20][21]

The case is useful because it reduces a common gap between learning and employment.

Rather than:

training → certificate → job search

it is closer to:

job → training → mentored work → increasing responsibility.

SECTION 10

This does not prove long-term career success. Public evidence remains limited on retention, pay progression and promotion.

But it shows how a first rung can be institutionally constructed around real demand.

BNY: AI enters the formal early-career pathway

BNY's Analyst Program illustrates another model.

Early-career analysts are introduced to structured AI pathways such as AI Explorers, AI Builders and AI Pioneers alongside rotations, mentoring and progression toward associate roles. [31]

This matters because AI becomes part of professional formation rather than an optional skill layer added after the employee has “learned the real job.”

But the model also raises a risk:

Can a young worker learn to build AI before fully understanding the domain being automated?

That question links career architecture back to learning architecture.

LONGi and JD: transition examples, not universal pathways

LONGi and JD provide worker-level examples of movement into AI training and AI operations. [13][16]

They are valuable because they show concrete destination roles.

They should not be generalized into claims that all affected inspectors or customer-service workers can follow the same path.

The evidence is stronger on the existence of the pathway than on its scale or accessibility.

Moving Safe Harbor

AIGW uses the term Moving Safe Harbor for a destination role that appears protected because it is “an

AI job” but can itself become a target for automation.

JD illustrates the risk. AI trainers and knowledge-base roles help maintain service automation, but tools can also automate parts of knowledge extraction, FAQ mining and configuration. [18]

BNY's employee-builder model creates a similar question over time: as agent-building platforms become easier, the valuable capability may move away from basic agent construction toward workflow design, governance, exception handling and domain judgment.

Therefore career architecture should not ask only:

What AI job can this worker move into today?

It should ask:

Which capability is likely to remain valuable as the technology itself improves?

Career quality matters

A career transition should be measured not only by whether a person moved.

High-quality transition evidence would include:

- destination role;
- pay;
- grade;
- responsibility;
- retention;
- promotion;
- portability of capability;
- worker choice;
- time to proficiency.

Most public enterprise AI reporting does not yet provide this level of detail.

That gap should be treated as a measurement problem, not a reason to assume that all mobility is positive.

Workforce Flow Architecture — How People Move Through the Transition

Headcount is a stock.

Workforce transition happens through flows.

That distinction is essential for understanding AI employment effects.

A company's total workforce can remain stable while major changes occur underneath:

- fewer junior hires;
- more senior hiring;
- internal redeployment;

- attrition without replacement;

- new AI-role creation;

- elimination of a narrow role;

- workforce reduction in one function and growth in another.


AIGW therefore distinguishes at least eight workforce-flow mechanisms:

Hire · Reskill · Redeploy · Create Roles · Reduce Hiring · Non-backfill · Attrition · Reduce

FIGURE 08

Workforce Stocks vs Workforce Flows

A stable headcount can conceal major workforce transition



Headcount alone can miss how AI reshapes workforce movement, role composition and internal mobility.

Source: AIGW analysis based on approved enterprise workforce transition cases.

Note: Workforce transition can occur even when the total workforce stock appears stable.

SECTION 11

Salesforce: why “jobs cut” is an incomplete category

Salesforce is the strongest case in the portfolio for the importance of flows.

The company reported a support-headcount movement from roughly 9,000 to roughly 5,000 and directly linked lower staffing need to the Agentforce transition. It also documented hundreds of redeployments and later described attrition and non-backfill as mechanisms. [1][2]

What remains unknown is the exact decomposition.

That uncertainty matters.

A worker who is redeployed, a vacancy that is not refilled and an involuntary exit all reduce staffing in different ways and have different implications for transition quality.

Therefore the safe conclusion is:

support staffing requirements fell materially, with redeployment, attrition and non-backfill visibly contributing, while the exact contribution of each mechanism remains unresolved.

DBS: mobility before exit

DBS makes a different strategic choice visible.

The bank says it identifies roles likely to be materially changed by AI, maps new or enhanced roles, and reskills or redeploy existing employees. [5][6]

CEO Tan Su Shan also said the bank would not hire new people into some jobs it expected AI to replace, while focusing on retraining people already in affected roles. [7]

That is a workforce-flow signal even without a layoff count.

It indicates reduced future hiring demand in some roles combined with an internal-mobility strategy for incumbents.

PwC: entry seats can change before the role disappears

PwC illustrates why entry-level labor demand needs separate measurement.

The junior auditor role still exists and is being redesigned around AI-enabled work. At the same time, US entry-level hiring plans were expected to decline, with multiple factors cited, including technology, low attrition and greater use of delivery centers. [11]

It would be incorrect to describe the entire decline as AI displacement.

It would also be incorrect to say the workforce effect is irrelevant because no whole occupation disappeared.

Fewer entry seats are themselves a workforce-flow change.

SimplifyNext: new hiring can be part of AI transition

SimplifyNext is a useful counterexample to narratives that treat AI workforce effects only as reduction.

The company used the CLT model to hire close to 50 trainees into AI and automation work and planned additional intake. [21]

The transition mechanism is Hire and Entry-Path Redesign.

This reminds us that AI can simultaneously reduce demand for one type of work and create demand for another.

SECTION 11

JD: transition is not the same as net job creation

JD reported that thousands of people had completed job transformation or upgrading over roughly two years. [17]

That is meaningful transition evidence.

It is not evidence that thousands of net-new jobs were created.

Without a denominator, exit data and destination distribution, the net employment effect remains uncertain.

External labor-market evidence: flows matter outside the case portfolio

Stanford Digital Economy Lab's revised 2026 payroll study finds no evidence of widespread economy-wide AI displacement, but identifies a widening employment shortfall for workers aged 22–25 in highly AI-exposed occupations. Crucially, the adjustment appears primarily through reduced hiring rather than increased separations. [36]

A separate Stanford study using 1.25 billion job postings and 154 million employment records across 41 countries finds that foreign affiliates of AI-adopting firms reduce the junior share of their workforce relative to comparable controls. The decline is driven mainly by growth in senior employment rather than a

simple fall in junior employment, with suggestive evidence of modest overall employment growth. [42]

These are not direct causal estimates for the nine companies in this report.

They reinforce the analytical need to measure composition and flows rather than relying on a single stock measure.

Workforce-flow accounting

A credible enterprise AI transition should be able to report, by affected role or work unit:

- beginning headcount;
- hires;
- internal transfers in;
- internal transfers out;
- non-backfilled vacancies;
- voluntary attrition;
- involuntary exits;
- new-role entries;
- ending headcount.

It should also identify which flows are directly linked to the AI transition and which reflect wider business conditions.

Without that accounting, claims about “jobs created” or “jobs lost” remain too coarse for serious workforce strategy.

Measurement Architecture — What Counts as Evidence of Transformation?

Most enterprises can measure AI activity.

Far fewer can measure AI workforce transition.

The distinction is visible across the case portfolio.

BNY can report access, training hours, agent builders and digital employees. DBS can report models, use cases, employees commencing training and operating-model transformations. Kaiser Permanente can report users, encounters and estimated documentation-time

savings. Salesforce can report AI-handled conversations, resolution rates, redeployment and staffing change.

These are valuable metrics.

But they live at different levels of the system.

Comparing them without a measurement architecture creates false equivalence.

AIGW proposes an eight-part evidence system.

FIGURE 09

AIGW Workforce Architecture Measurement Stack

What institutions need to measure when AI changes work, roles and workforce outcomes

Measurement domain		Core question	Illustrative evidence or example indicators
01	Technology	What was deployed?	access • adoption • agents • autonomy
02	Work	What tasks and workflows actually changed?	cycle time • task compression • automation • handoffs
03	Role	How did human responsibilities change?	tasks removed • tasks added • decision rights • exceptions
04	Capability & Learning	What can people now do, and how was it learned?	demonstrated proficiency • supervised practice • verification capability • learning progression
05	Authority	Who can decide, approve, override and remain accountable?	approval rights • override rights • escalation • accountability
06	Workforce Flow	Who entered, moved, stayed or exited?	hires • redeployment • attrition • non-backfill • reduction
07	Worker Outcome	What happened to the people affected?	pay • grade • progression • retention • agency
08	Business Outcome	What happened to organizational performance?	productivity • quality • service • risk • cost

Usage is an input metric. Workforce transformation requires outcome evidence.

Access ≠ Use ≠ Training ≠ Capability ≠ Role Transition ≠ Worker Outcome

Source: AIGW analysis based on approved enterprise workforce transition cases and related implementation evidence.

Note: The indicators are illustrative. The stack is an analytical measurement framework, not a validated scoring model.

SECTION 12

1. Technology

Measure:

- deployment stage;
- users;
- models / agents;
- production status;
- degree of autonomy;
- system access.

Technology metrics answer:

Is the system there, and what can it do?

They do not answer:

What happened to work or people?

2. Work

Measure:

- tasks automated;
- tasks augmented;
- cycle time;
- handoffs;
- exception rate;
- escalation rate;
- verification burden;
- work remaining with humans.

Work metrics answer:

How did the operating process change?

3. Roles

Measure:

- before / after task portfolio;
- responsibilities added;

- responsibilities removed;
- workflow position;
- decision rights;
- staffing requirement.

Role metrics answer:

What is the job becoming?

4. Capability and learning

Measure:

- training offered;
- training started;
- training completed;
- assessed proficiency;
- production application;
- mentor access;
- time to independent performance.

This is where enterprise reporting most often needs improvement.

Training participation is a learning input.

It is not a capability outcome.

5. Authority

Measure:

- actions AI may take;
- actions requiring approval;
- human review points;
- override / rollback capability;
- escalation rules;
- final authority;
- accountable owner.

SECTION 12

Authority metrics are especially important in agentic systems because productivity can rise while the governance boundary becomes less visible.

6. Workforce flows

Measure:

- hires;
- redeployments;
- role creation;
- reduced hiring;
- non-backfill;
- attrition;
- involuntary exits.

This answers:

What happened to the population of workers?

7. Worker outcomes

Measure:

- pay;
- grade;
- promotion;
- retention;
- job quality;
- workload;
- worker agency;
- access to transition opportunities.

This is the weakest public evidence category across the nine cases.

The absence matters because an enterprise can deliver business value while producing poor worker outcomes.

8. Business outcomes

Measure:

- productivity;
- quality;
- service;
- risk;
- revenue;
- cost;
- customer experience;
- resilience.

Business value is necessary for sustainable adoption.

But business outcomes should not be silently attributed to a specific workforce intervention unless the causal link is clear.

A measurement architecture, not a score

This framework should not be treated as a validated maturity index.

Its purpose is to force analytical separation.

A company with 100% AI access may still have weak role redesign.

A company with high productivity may still have weak career mobility.

A company with extensive reskilling may still have no evidence that workers moved into better roles.

A company with autonomous agents may still have incomplete public evidence of human authority.

The discipline is simple:

Measure the transition at the layer where the claim is being made.

SECTION 13

What the Nine Cases Reveal

The nine cases do not converge on one workforce outcome.

That is precisely why they are useful.

Some cases show lower staffing requirements. Some show large-scale reskilling and redeployment. Some show new hiring. Some show role creation. Some show productivity gains without a documented employment effect. Some show explicit human authority. Others expose gaps in how that authority is publicly described.

The value of the portfolio is not that it tells one clean story.

It reveals where different parts of workforce architecture move together — and where they do not.

Seven cross-case findings stand out.

Finding 1 — AI transformation propagates unevenly across architecture layers

Technology can move faster than institutions.

Deutsche Telekom can place autonomous agents into network operations before the public evidence describes the corresponding workforce transition in detail. Kaiser Permanente can scale ambient documentation and create a strong authority boundary without producing a major career redesign. BNY can create widespread agent-building capability before it can show long-term worker outcomes. PwC can redesign the junior role before the profession has resolved how entry-level hiring and apprenticeship will settle.

This means enterprise AI maturity cannot be inferred from one visible layer.

High adoption does not imply high capability.

High productivity does not imply strong career pathways.

Agent autonomy does not imply clear human authority.

Extensive training does not imply successful redeployment.

The architecture therefore needs to be assessed as a system.

Management implication

Organizations should expect architecture lag: one domain will often advance faster than another.

The practical goal is not perfect synchronization. It is to identify where the lag creates risk.

Examples include:

- automation without a learning system;
- role redesign without destination roles;
- agent autonomy without explicit authority;
- reskilling without demand;
- productivity without workforce-flow accounting.

Finding 2 — Domain expertise becomes both operating capital and governance capital

The cases repeatedly show that the expertise embedded in an affected role often remains valuable after the original task is automated.

SECTION 13

LONGi's former inspector uses frontline knowledge to train AI. JD's customer-service expertise becomes useful in diagnosing response failures and configuring agents. Salesforce moves product and customer knowledge into adjacent customer-facing roles. BNY relies on domain employees outside engineering to build AI agents. [13][16][1][29]

This extends the idea of Transition Capital.

Domain expertise is not only useful for performing residual human work.

It becomes useful for:

- deciding what should be automated;
- defining exceptions;
- training models;
- verifying outputs;
- designing escalation;
- governing autonomous systems.

That makes domain experts part of AI governance, not just AI adoption.

Management implication

Before reducing a role, organizations should ask:

Which knowledge inside this role will still be needed to train, verify or govern the system that replaces part of it?

If that knowledge exits before it is captured or redeployed, automation can destroy institutional capability while improving short-term efficiency.

Finding 3 — Learning architecture becomes a constraint on sustainable automation

The strongest form of AI productivity can remove the very work that helped build future experts.

PwC makes this visible in professional services. AI can compress entry-level execution, but future audit quality still depends on experienced human judgment. [9][10]

SimplifyNext demonstrates that the learning function can be rebuilt around employment, mentorship and real AI-enabled work. BNY shows that AI learning can be integrated directly into role development and production. [20][29]

Together, the cases suggest:

Automation is sustainable only if capability formation keeps pace with task compression.

This is the institutional meaning of Learning Debt.

Management implication

For each automated task with developmental value, leaders should identify a replacement learning mechanism.

If the answer is “the employee will learn by using AI,” the design is incomplete.

Finding 4 — Career architecture determines whether productivity becomes mobility or exclusion

The same underlying technology can create very different worker outcomes depending on whether career architecture exists around it.

Salesforce provides internal destination roles.

DBS links reskilling to internal mobility.

SECTION 13

SimplifyNext creates an entry route.

BNY integrates AI into the analyst pathway.

PwC illustrates a harder outcome: the first rung can become more demanding while the number of entry opportunities tightens. [2][6][21][31][11]

The difference is not explained by AI capability alone.

It is explained partly by institutional choices about hiring, mobility, apprenticeship and workforce planning.

Management implication

Every material AI use case should eventually connect to a people-flow question:

If this work changes, what happens to the people who performed it, and where can the next cohort enter?

Finding 5 — Authority becomes a design variable as AI moves from recommendation to action

The portfolio shows several distinct human-AI authority structures.

Kaiser Permanente preserves physician review and clinical authority.

DBS preserves a human final decision in the credit-memo example.

LONGi preserves human verification of flagged defects.

Deutsche Telekom allows autonomous corrective action but publishes less use-case-specific information about approval and override. [24][8][13][26]

This variation matters because two systems with similar technical capability can have very different governance architectures.

Management implication

“Human in the loop” should be replaced by explicit action-level authority descriptions.

Finding 6 — Employment effects appear through flows before clean headcount causality becomes visible

AI's labor effect is often a composition problem before it is a layoff event.

Salesforce shows lower staffing requirements and multiple flow mechanisms.

PwC shows pressure on entry hiring.

DBS signals reduced future hiring in some roles alongside incumbent retraining.

SimplifyNext shows new hiring into AI-enabled work.

JD shows role transition without a clean net-employment number. [2][11][7][21][17]

This pattern is consistent with emerging labor-market research showing reduced hiring and changes in workforce composition without broad economy-wide displacement. [36][42]

Management implication

Organizations and policymakers need flow accounting, not only layoffs and ending headcount.

SECTION 13

Finding 7 — Distributed AI building creates a hybrid human + digital workforce

BNY makes this development especially visible, but Deutsche Telekom and JD point in the same direction.

As employees build agents and agents become operational participants, managers increasingly oversee mixed systems of people and software actors.

The management object changes from:

team of people using software

into:

team of people + agents + automated workflows with different identities, permissions and escalation rules.

This creates new questions about span of control, performance management, access rights, system

ownership and accountability.

Management implication

Enterprises need to decide how digital workers are governed as organizational actors without confusing them with human employees.

A digital employee is not an FTE.

But it can still have:

- an identity;
- a manager;
- permissions;
- a workflow role;
- performance expectations;
- monitoring requirements.

That is an operating-model shift, not merely a software deployment.

SECTION 14

Five Enterprise Dilemmas

A good architecture does not eliminate trade-offs.

It makes them visible.

The nine cases reveal five dilemmas that cannot be solved through technology choice alone.

Dilemma 1 — Productivity vs Capability Formation

AI can make current workers more productive while weakening the mechanism through which future workers acquire expertise.

PwC is the clearest warning case.

Routine junior work may be economically compressible but developmentally useful. If the work disappears, the organization has to rebuild exposure, feedback and practice deliberately. [9][10]

SimplifyNext offers a possible alternative: start the career with AI already present, but combine it with real work, mentorship and increasing ownership. [20][21]

The dilemma is not whether productivity or learning should win.

It is whether the organization can redesign learning quickly enough that productivity gains do not create a future expertise shortage.

Decision test

For every automated junior task:

What capability did this task help build, and where will that capability now come from?

Dilemma 2 — Automation vs Internal Mobility

Automation can lower labor demand in one role while the organization still wants to retain the people who

hold useful domain knowledge.

Salesforce and DBS show two versions of this problem.

Salesforce reduced support staffing requirements while redeploying hundreds of employees. [2]

DBS uses reskilling and redeployment as part of formal operating-model change. [6]

But internal mobility has limits.

The number of destination roles may not equal the number of affected workers. Workers may not want the destination. The new role may require different capability. Geography, compensation and manager incentives can block movement.

Decision test

Are we building mobility before displacement becomes urgent — and is there enough destination demand to make the pathway credible?

Dilemma 3 — Democratized AI Building vs Governance

BNY's model suggests that AI transformation can accelerate when domain employees become builders rather than passive users. [29]

This is powerful because the people closest to the work often understand where automation opportunities and operational risks actually lie.

But distributed building creates a control problem.

More builders mean more:

- agents;
- use cases;
- permissions;
- data connections;

SECTION 14

- production pathways;
- failure modes.

Deutsche Telekom's agent-governance work illustrates why autonomy needs identities, access rights, logging and action boundaries. [28]

Decision test

Can we democratize building without democratizing unrestricted deployment authority?

The answer usually requires a governed path from experimentation to production.

Dilemma 4 — Higher-Value Human Work vs Fewer Human Seats

Corporate AI narratives often say automation removes routine work so people can focus on higher-value work.

The cases show that this can be true and still hide a distributional problem.

Salesforce's remaining support work becomes more complex while total support staffing need falls. [2]

PwC expects junior work to become more judgment-intensive while entry-level hiring plans tighten. [11]

Therefore:

Higher-value human work does not imply more human jobs, better pay or broader access.

It may instead mean a smaller number of more demanding roles.

Decision test

As work becomes more sophisticated, who gains access to it – and who loses the opportunity to enter?

Dilemma 5 — Autonomy vs Accountability

AI systems can create value precisely because they do not require a person to approve every step.

That is the promise of agentic operations.

But if nobody can explain who remains accountable for an autonomous action, the operating model is incomplete.

Kaiser Permanente demonstrates one end of the design space: AI drafts, the physician reviews, and medical decision authority remains human. [24]

Deutsche Telekom demonstrates a more autonomous model in which agents can execute network changes. [26]

The correct answer is not that every autonomous action needs prior human approval.

It is that control should match consequence.

Decision test

Which actions can the AI take independently, which require approval, who can intervene, and who remains accountable after the action is taken?

These five dilemmas matter because they prevent a false conclusion from the architecture framework.

The goal is not to optimize every layer independently.

It is to make the trade-offs visible enough that leaders can choose deliberately.

The AIGW Workforce Architecture Framework

The nine cases suggest that enterprise AI transformation should be designed across nine interdependent domains.

The framework is not a maturity staircase. It is an institutional system.

FIGURE 10

AIGW Workforce Architecture

Nine interdependent design domains for managing AI-driven work transition

AI adoption changes tools. AI institutionalization changes the architecture of work.



1. Task Architecture

Decide what people, AI or both should do.

Key questions:

- Which tasks should be automated?
- Which should be augmented?
- Which can be delegated to agents?
- Which should be protected because they build capability or preserve control?
- Which need redesign rather than removal?

SECTION 15

2. Workflow Architecture

Design how work moves after task allocation changes.

Key questions:

- Where does AI enter?
- What triggers escalation?
- Where are the handoffs?
- Which steps can close automatically?
- What happens when the system is uncertain or wrong?

3. Role Architecture

Rebundle responsibilities around the new workflow.

Key questions:

- What disappears from the role?
- What becomes more important?
- What new responsibilities are added?
- What decision rights change?
- How many people are still required?

4. Capability Architecture

Define what people must be able to do.

Key questions:

- Which domain expertise remains valuable?
- Where is judgment required?
- Who needs AI fluency?
- Who needs AI orchestration capability?
- Which capabilities must be demonstrated in production?

5. Learning Architecture

Design how those capabilities will be built and renewed.

Key questions:

- Which learning experiences disappear with automation?
- What replaces them?
- Who mentors?
- How does responsibility increase?
- How will capability be assessed?

6. Authority Architecture

Specify who may decide, act, approve and override.

Key questions:

- What may the AI do without approval?

SECTION 15

- What must a human verify?
- Where is approval required?
- Who can override or roll back?
- Who retains final authority and accountability?

7. Career Architecture

Create credible movement through the redesigned system.

Key questions:

- Where can people enter?
- What destination roles exist?
- What transition capital can be reused?
- What is the next rung?
- Are the new capabilities portable?

8. Workforce Flow Architecture

Manage population change.

Key questions:

- Who is being hired?
- Who is reskilled?
- Who is redeployed?
- Which roles stop hiring?
- Where does non-backfill occur?
- Where does workforce reduction occur?

9. Measurement Architecture

Create an evidence system across the other eight domains.

Key questions:

- Are we measuring usage or changed work?
- Training or demonstrated capability?
- Gross time savings or reliable net benefit?
- Headcount stocks or workforce flows?
- Business value or worker outcomes?
- Human presence or human authority?

The framework's central principle

Each domain constrains the others.

A workflow cannot be safely autonomous if authority is undefined.

A role cannot be sustainably redesigned if learning is absent.

A reskilling programme is not credible if career destinations do not exist.

A productivity gain is not a complete workforce outcome if flows are not measured.

The architecture therefore provides a way to diagnose institutional incompleteness.

The question is not whether every layer is “mature.”

It is whether the layers that matter for a specific AI transition fit together well enough for the transition to be sustainable.

SECTION 16

The Enterprise Diagnostic

The framework can be translated into a practical diagnostic without pretending that every organization can be scored on one universal maturity scale.

The purpose of the diagnostic is to expose unanswered questions.

Task Architecture

- Which tasks are materially changed by AI?
- Which are automated, augmented, delegated, protected or redesigned?
- Which tasks have high developmental value?
- Which tasks carry high decision or safety consequence?

Workflow Architecture

- What is the trigger for the AI-enabled process?
- Where does AI hand work to a person?
- Where does a person hand work to an AI system?
- What happens when the system is uncertain?
- What is the exception route?

Role Architecture

- What is removed from the role?
- What is added?

- Does the role move closer to judgment, relationships, verification or orchestration?
- Does the staffing envelope change?

Capability Architecture

- Which capabilities become more valuable?
- Which existing domain capabilities can be reused?
- Which capabilities require formal training?
- Which require supervised practice?
- How will proficiency be demonstrated?

Learning Architecture

- What did people learn from the old work?
- Where will that learning now happen?
- Who provides feedback and review?
- How does the worker move from assisted performance to independent judgment?

Authority Architecture

- Which AI actions require no human intervention?
- Which require monitoring?
- Which require verification or approval?
- Who can override or roll back?
- Who is accountable after the action?

SECTION 16

Career Architecture

- Where does an affected worker go next?
- Which destination roles have real demand?
- What is the entry pathway for new workers?
- What is the second rung?

Workforce Flow Architecture

- What happens to hiring?
- What happens to vacancies?
- How many workers move internally?
- How many leave voluntarily?
- Are there involuntary exits?
- Are new roles being created?

Measurement Architecture

- Which metrics show adoption?
- Which show work change?

- Which show capability?
- Which show worker outcomes?
- Which show workforce flows?
- Which show business value?

A strong diagnostic conversation should end with a small number of architecture gaps, not a generic “AI readiness score.”

Examples might include:

- strong automation, weak learning architecture;
- strong reskilling, weak destination demand;
- strong agent deployment, unclear authority;
- strong productivity evidence, weak workforce-flow accounting;
- strong early-career recruitment, weak progression evidence.

These gaps are more decision-useful than a single maturity number because they point directly to what has to be redesigned.

An Illustrative 12-Month Enterprise Roadmap

No enterprise will redesign workforce architecture in exactly twelve months.

The sequence below is therefore illustrative, not universal.

Its purpose is to show how an organization can move from technology deployment toward institutional design without waiting for every answer to be perfect.

FIGURE 11

Illustrative 12-Month Workforce Architecture Roadmap

A practical sequence for beginning AI-driven workforce redesign



The sequence is illustrative. Organizations may enter at different points depending on AI maturity, risk and workforce context.

Source: AIGW synthesis based on approved enterprise workforce transition cases and the Workforce Architecture framework.
Note: The roadmap is not a universal implementation sequence.

0–90 days — Map the transition

Build the task and workflow map

Identify the work units where AI has already changed execution or where material change is expected.

For each use case, document:

- tasks;
- workflow;
- role impact;
- current human authority;
- affected population.

SECTION 17

Establish the evidence baseline

Capture:

- current staffing;
- hiring rate;
- cycle time;
- quality;
- training state;
- current decision rights.

Without a baseline, later workforce claims will be difficult to interpret.

Identify architecture gaps

Flag use cases where:

- work is changing but roles are not;
- roles are changing but learning is not;
- agents can act but authority is unclear;
- reskilling exists without destination roles;
- productivity metrics exist without worker metrics.

3–6 months — Pilot redesigned work and transition paths**Redesign selected roles**

Choose a small number of high-value work units and define the before/after role architecture.

Build learning into production

Create:

- supervised project work;
- verification checkpoints;
- review routines;
- mentor capacity;
- proficiency criteria.

Design authority at action level

For agentic workflows, classify actions by consequence and specify:

- autonomous;
- monitored;
- approval required;
- prohibited.

Test mobility

Identify destination roles for affected workers and test whether real internal demand exists.

6–9 months — Integrate the systems

SECTION 17

Connect CHRO, CIO and business ownership

Workforce transition should become part of the AI deployment governance process rather than a downstream HR reaction.

Integrate learning and career pathways

Link role-specific capability to:

- internal jobs;
- progression;
- manager development;
- early-career pathways.

Formalize hybrid-team management

Where digital workers or agents operate in production, define:

- identity;
- permissions;
- manager / owner;
- escalation;
- performance;
- monitoring;
- accountability.

9–12 months — Measure and institutionalize**Report workforce flows**

Move beyond ending headcount.

Track:

- hires;
- redeployments;
- new roles;
- non-backfill;
- attrition;
- exits.

Report worker outcomes

Where possible, add:

- promotion;
- pay;
- retention;
- workload;
- worker experience.

Refresh the architecture

AI capability changes quickly.

A role architecture designed for today's copilot may be obsolete when agents can execute tomorrow.

The architecture should therefore be reviewed as part of normal workforce and technology governance.

The goal is not a one-time transformation programme.

It is an institutional capability to redesign work repeatedly as AI changes.

SECTION 18

For CEOs and Boards

Enterprise AI should no longer be governed only as a technology portfolio.

Once AI changes work allocation, authority and workforce demand, boards and executive teams need to treat it as an operating-model and institutional-design issue.

The board-level question is not:

How many AI use cases are live?

It is:

What organizational dependencies, workforce effects and decision-right changes are being created by those use cases?

Three responsibilities follow.

First, leaders should require a workforce consequence view alongside major AI investments. That does not mean demanding a layoff forecast for every pilot. It means asking which roles, skills, career pathways and

staffing assumptions are likely to change if the use case scales.

Second, boards should ask whether the organization is creating hidden institutional liabilities. A productivity gain that removes apprenticeship, an autonomous agent with unclear override rights, or a reskilling programme without destination roles may look successful on a technology dashboard while creating future risk.

Third, leadership should explicitly decide what released capacity is for. If AI creates time, the enterprise can use it for growth, quality, service, risk reduction, redeployment, reduced hiring or cost reduction. Leaving that choice implicit makes the workforce outcome accidental.

The board does not need to manage individual workflows.

It does need confidence that the organization has a credible architecture for how AI changes work and people.

SECTION 19

For CHROs

The CHRO's role in AI transformation should extend beyond training and change management.

The strongest cases in this report show that the central HR task is workforce transition architecture.

That includes:

- identifying affected roles before disruption becomes urgent;
- mapping transition capital;
- designing destination roles;
- connecting learning to real demand;
- tracking workforce flows;
- rebuilding early-career pathways;
- preparing managers to supervise hybrid human + AI teams.

DBS offers the clearest example of workforce planning being connected to operating-model change.

Salesforce shows why internal mobility infrastructure matters when staffing demand falls. SimplifyNext

shows why learning should connect to employment.

BNY shows why AI capability can become part of role and career design rather than a generic training programme. [6][2][20][29]

CHROs also need a better evidence standard.

Participation in AI learning should not be the endpoint.

The relevant sequence is:

learning offered → learning started → learning completed → capability demonstrated → role changed → worker outcome observed

The CHRO should be able to explain where the organization has evidence along that chain — and where it does not.

SECTION 20

For CIOs and Chief AI Officers

Technology leaders increasingly design work even when that is not how the project is described.

A system that changes who performs a task changes role content.

A system that closes a workflow changes handoffs and escalation.

A system that acts autonomously changes authority.

A platform that lets employees build agents changes capability requirements and governance surface.

Therefore the CIO / Chief AI Officer should not treat workforce design as a downstream HR workstream.

Every production AI use case should have a corresponding human operating model.

That model should specify:

- who uses the system;

- who supervises it;

- where errors go;

- what people must know;

- what decisions remain human;

- what role content changes;

- what metrics will show whether the new workflow works.

Deutsche Telekom and BNY are particularly useful here. The first shows why autonomous action requires action-level control. The second shows why distributed building requires a governed path from experimentation to deployment. [26][28][29]

Technology architecture and workforce architecture are becoming interdependent.

SECTION 21

For Learning and Development Leaders

AI creates a risk that L&D becomes very busy while capability evidence remains weak.

Organizations can rapidly generate:

- courses;
- prompt libraries;
- certifications;
- office hours;
- AI communities;
- training-hour totals.

Those assets are useful, but they are not the same as professional formation.

The L&D mandate should shift from content delivery toward capability formation.

That requires four design changes.

First, learning should be tied to redesigned roles rather than generic AI literacy.

Second, capability should be demonstrated through work. SimplifyNext's live project exposure and BNY's builder pathways illustrate the value of connecting learning to production. [21][29]

Third, L&D should identify where automation removes learning opportunities. PwC's apprenticeship problem shows that some capability can no longer be assumed to emerge naturally from the old work sequence. [9]

Fourth, managers and mentors need capacity to teach. If more junior workers take on higher-context responsibilities earlier, review and coaching become part of production infrastructure.

The right L&D question is therefore not:

How many employees completed AI training?

It is:

What are employees now capable of doing — and how do we know?

SECTION 22

For Policymakers and Public Institutions

The enterprise cases also expose limits in public labor-market measurement.

AI workforce effects can occur through reduced hiring, internal mobility, non-backfill and changes in role composition before they appear as mass layoffs.

Stanford's 2026 evidence on reduced hiring among younger workers in highly exposed occupations and changing junior workforce share across 41 countries illustrates why public statistics need more than layoff counts. [36][42]

Policymakers should improve measurement in at least five areas:

1. Entry-level hiring — not only unemployment.
2. Internal mobility — whether workers move into new roles.
3. Training-to-work conversion — whether reskilling leads to employment and progression.

4. Worker outcomes — pay, grade, job quality, retention and agency.

5. AI decision rights — especially in high-stakes or public-facing systems.

Public policy can also reduce coordination failure.

SimplifyNext's use of Singapore's Company-Led Training programme demonstrates the value of connecting employer demand, employment, training, mentorship and career progression in one institutional mechanism. [20]

The ILO's Quality Apprenticeships Recommendation reinforces the importance of structured, remunerated learning, qualified supervision and recognized occupational competence. [39]

The policy goal should therefore be broader than “more AI skills.”

It should be to help labor-market institutions create credible pathways from changing work to durable capability and employment.

SECTION 23

What We Still Do Not Know

The nine cases provide enough evidence to construct a workforce-architecture framework.

They do not provide enough evidence to answer several of the most important worker-level questions.

These gaps should remain visible.

Pay

Very few of the cases disclose whether a worker who moves into an AI-enabled role earns more, less or the same.

A move can look successful while leaving compensation unchanged or worse.

Promotion and long-term progression

We often know that a transition occurred.

We rarely know what happened 12, 24 or 36 months later.

Did the worker progress?

Did the AI-enabled role create a durable career?

Did the capability remain valuable outside the firm?

Worker voice and agency

Most evidence in this report comes from institutions, executives, annual reports and public case studies.

There is much less systematic evidence on how workers experience:

- increased monitoring;
- changing workloads;
- higher expectations;

- reduced autonomy;
- career uncertainty;
- AI-assisted management.

Collective bargaining and worker representation

The portfolio contains almost no detailed evidence on how unions or worker representatives shape AI deployment, role redesign or transition support.

That is a major gap for understanding transition quality.

Failed reskilling

Enterprise reports often describe successful participation and movement.

We know much less about:

- people who could not complete training;
- workers who were not selected;
- roles for which no credible destination existed;
- transitions that failed.

Full headcount reconciliation

Salesforce provides unusually direct staffing evidence, but even there the exact decomposition among redeployment, attrition, non-backfill and involuntary exits is incomplete.

Most cases provide less.

Without workforce-flow reconciliation, it remains difficult to compare transition mechanisms across firms.

SECTION 23

Small and medium-sized enterprises

The cases are concentrated in organizations with substantial capability, capital and management infrastructure.

SMEs may face a different architecture problem:

- fewer mentors;
- less internal mobility;
- fewer destination roles;
- greater dependence on external tools;
- weaker governance capacity.

Unequal access to new pathways

A new AI-enabled role may exist but remain accessible only to workers with strong education, technical confidence, managerial sponsorship or geographic mobility.

The portfolio is thin on who gains access and who is screened out.

The long-term economics of verification

As AI systems improve, verification burden may fall.

Or it may shift into more complex, high-stakes exception work.

The economics of human oversight over time remain poorly measured.

When the safe harbor moves again

AI trainers, agent builders and workflow supervisors are currently growth roles in several cases.

Some of those tasks will themselves become easier to automate.

Longitudinal evidence is needed to identify which capabilities remain durable as AI systems advance.

These gaps do not invalidate the framework.

They define the next research frontier.

SECTION 24

Conclusion — From Technology Deployment to Institutional Design

The central question of enterprise AI is changing.

For the first phase of adoption, the key problem was access: which models, copilots and agents should organizations deploy, and where could they create value?

The next problem is institutional.

Once AI performs tasks, coordinates workflows, drafts professional work, manages customer interactions, trains on human expertise or takes operational action, organizations have to redesign the systems around the technology.

The nine cases in this report show that the resulting workforce transition is not one thing.

It can involve lower staffing requirements and redeployment, as in Salesforce.

It can involve formal operating-model transformation, as in DBS.

It can weaken apprenticeship and force a redesign of early-career learning, as in PwC.

It can turn frontline expertise into model-training and verification capability, as in LONGi and JD.

It can create a new employment-first pathway, as in SimplifyNext.

It can preserve explicit professional authority inside an automated workflow, as in Kaiser Permanente.

It can cross from recommendation into autonomous operational action, as in Deutsche Telekom.

It can turn employees into builders and supervisors of a digital workforce, as in BNY.

The common thread is not one employment outcome.

It is the need for architecture.

Organizations have to decide:

- what work AI should do;
- how workflows should change;
- how roles should be rebundled;
- what human capabilities remain scarce;
- how those capabilities will be learned;
- who retains authority;
- where workers can move;
- how workforce flows will be managed;
- how transition quality will be measured.

That is the difference between deploying AI and institutionalizing it.

AI adoption asks whether people are using the technology.

Workforce architecture asks what kind of organization the technology is creating.

The institutions that navigate the transition best will not simply deploy the most capable models.

They will build the organizational capacity to repeatedly redesign work, capability, authority and careers as the technology changes.

That is the deeper workforce challenge of the AI era.

REFERENCES

Notes and References

1. SALESFORCE

Salesforce. *Agentic enterprise workforce evolution*.

Supports: Supports the changing customer-support workflow, role redesign, internal mobility and the broader workforce-evolution model around Agentforce. Used for the transition from recurring support work toward more complex and proactive human work.

<https://www.salesforce.com/news/stories/agentic-enterprise-workforce-evolution/>

2. Salesforce. *How Salesforce Is Reshaping Its Workforce in the Age of AI*.

Supports: Supports actual redeployment, reskilling, named destination roles and Salesforce's later description of workforce mechanisms including attrition and non-backfill. Used together with attributable CEO statements on lower support staffing requirements.

<https://www.salesforce.com/news/stories/salesforce-reshaping-workforce-in-age-of-ai/>

3. The Logan Bartlett Show. *Marc Benioff interview*.

Supports: Supports the attributable executive statement that support headcount moved from roughly 9,000 to roughly 5,000 as the company needed fewer support staff after AI deployment. The interview does not provide a complete decomposition of the reduction mechanism.

<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*.

Supports: Supports the organization's 4Rs framing and workforce-transition mechanisms. Used as institutional context rather than independent outcome validation.

<https://www.salesforce.com/en-ap/wp-content/uploads/sites/14/documents/pdf/agentforce/workforce-innovation-playbook-ap.pdf>

5. DBS

DBS Group Holdings. *Annual Report 2025*.

Supports: Supports enterprise AI scale, operating-model transformation, business context and formal company reporting on AI-enabled change.

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

6. DBS Group Holdings. *Sustainability Report 2025*.

Supports: Core workforce-transition source. Supports more than 11,000 employees identified for deeper upskilling/reskilling, 96% having commenced training, role redesign, redeployment and the Operating Model Transformation process.

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

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

Supports: Attributable interview evidence on changing labor demand, including CEO statements that the bank would avoid hiring into some roles expected to be replaced by AI while retraining incumbents. Used as changing-hiring-demand evidence, not proof of a quantified workforce reduction.

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

REFERENCES

8. **DBS**
DBS. *CEO Tan Su Shan on how agentic AI is transforming the bank.*
Supports: Supports the later credit-memo example in which AI agents assist research and drafting while the human credit manager retains the final call. Used specifically as follow-on authority evidence rather than generalized across every DBS workflow.
<https://www.dbs.com/media/features/ceo-tan-shu-shan-on-how-agentic-ai-is-transforming-the-bank.page>
9. **PWC US**
PwC US. *No more pyramids: Rethinking your workforce for the agentic AI era.*
Supports: Supports PwC's workforce-design argument, AI use in audit, changing junior work and the need to preserve a strong early-career capability pipeline rather than simply removing the base of the traditional professional-services pyramid.
<https://www.pwc.com/us/en/tech-effect/ai-analytics/agentic-ai-workforce-redesign.html>
10. PwC US. *2025 Audit Quality Report.*
Supports: Supports AI training, professional capability development, audit technology and the continuing importance of judgment, review, quality and traditional professional foundations.
<https://www.pwc.com/us/en/services/audit-assurance/assets/pwc-2025-audit-quality-report.pdf>
11. Business Insider. *PwC hiring fewer junior associates / internal hiring plans.*
Supports: Supports reported internal planning for lower US Tax and Assurance associate hiring through FY28 and PwC's confirmation that campus hiring goals were decreasing. The report identifies multiple contributors including technology, low attrition and Acceleration Centers; the decline should not be attributed to AI alone.
<https://www.businessinsider.com/pwc-hiring-fewer-junior-associates-ai-offshoring-big-four-2025-8>
12. PwC US. *ChatNational first-year evidence.*
Supports: Supports more than 17,000 assurance professionals using ChatNational and more than 350,000 questions in the first year. Used as operational-scale evidence, not an employment outcome.
<https://www.pwc.com/us/en/about-us/newsroom/press-releases/chatbot-chatnational-anniversary.html>
13. **LONGI**
Xinhua. *Smart tech, better lives: AI empowers China's common prosperity drive.*
Supports: Core worker-level source. Supports AI inspection across nine procedures and 198 inspection items, human double-checking of flagged defects, the named transition of former inspector Zhang Chen into AI training, and a reported reduction in inspection labor cost. It does not establish that all inspectors followed the same path or that lower labor cost equals layoffs.
<https://english.news.cn/20260718/82d978e634d74a018b437dfd668a6ce0/c.html>
14. LONGi. *Technology and Innovation / Jiaxing Lighthouse Factory.*
Supports: Supports the broader scaled AI inspection and digital-manufacturing context at the Jiaxing Lighthouse Factory. Used for technology and operating context rather than worker-level transition counts.
<https://www.longi.com/en/sustainability/esg-society-technology/>

REFERENCES

15. LONGI

LONGI. *Sustainability Report 2023*.

Supports: Supports inspection-performance and workforce-process context. Factory-wide and inspection-system outcomes are kept distinct from the individual role-transition evidence.

https://static.longi.com/2023_longi-sustainability-report.pdf

16. JD.COM

Science and Technology Daily. *AI brings new identities to JD workers*.

Supports: Core named-worker source for customer-service experience moving into AI training and AI-operations work. Supports the mechanism through which domain expertise is used to diagnose AI response problems, configure systems and improve outputs.

https://www.stdaily.com/web/gdxw/2026-05/19/content_518575.html

17. Xinhua / Science and Technology Daily. *How to catch the new AI jobs over the next five years*.

Supports: Supports JD's statement that thousands of people completed job transformation or upgrading over roughly two years and describes the service architecture in which AI handles basic service while people handle complex scenarios and related support roles. "Thousands transformed" is not equivalent to thousands of net-new jobs.

<https://www.news.cn/tech/20260907/eb1d91d961c84dd8b417ba41d649b129/c.html>

18. JD JoyAI. *Intelligent Customer Service*.

Supports: Supports evidence of second-order automation inside AI operations, including automation of knowledge-base configuration and related support tasks. Used to illustrate the Moving Safe Harbor concept, not to claim elimination of AI-trainer roles.

<https://jdai.jd.com/portal/ai-products/customer-service>

19. JD.com. *Fourth Quarter and Full Year 2025 Results*.

Supports: Supports the scale of AI customer-service usage and enterprise adoption context. Used as process/adoption evidence, not a workforce-outcome measure.

<https://ir.jd.com/news-releases/news-release-details/jdcom-announces-fourth-quarter-and-full-year-2025-results-and>

20. SIMPLIFYNEXT

Infocomm Media Development Authority. *Company-Led Training Programme*.

Supports: Authoritative source for the CLT institutional model: employment contract before programme start, salary, structured on-the-job learning, practical projects, experienced mentors, career progression and programme duration. Used to define the employment-first pathway.

<https://www.imda.gov.sg/how-we-can-help/techskills-accelerator-tesa/company-led-training-programme-clt>

21. IMDA. *How Upskilling Talent Powers AI Transformation*.

Supports: Core SimplifyNext outcome source. Supports close to 50 trainees hired, plans for additional intake, more than 35 AI-powered client projects, trainees involved in nearly half, and the reported figure that more than 90% had progressed into project-leading responsibility. The precise authority associated with "leading" is not publicly defined.

<https://www.imda.gov.sg/resources/blog/blog-articles/2026/01/how-upskilling-talent-powers-ai-transformation>

REFERENCES

22. SIMPLIFYNEXT

SimplifyNext. *AI Engineer (Graduate)*.

Supports: Current role-design evidence showing early-career work across AI development, data preparation, solution architecture, validation, deployment and client collaboration. Used as demand-side role evidence, not proof that every CLT trainee occupies the role.

<https://job-boards.greenhouse.io/simplifynext/jobs/4996764007>

23. KAISER PERMANENTE / THE PERMANENTE MEDICAL GROUP

NEJM Catalyst. *Ambient AI scribes: learnings after one year and more than 2.5 million uses*.

Supports: Core scale and workload-outcome source. Supports 7,260 physicians, 2,576,627 encounters, estimated aggregate documentation-time savings of 15,791 hours and differences by usage intensity. Time savings are estimates and do not establish clinical-outcome improvement or staffing reduction.

<https://catalyst.nejm.org/doi/full/10.1056/CAT.25.0040>

24. Permanente Medicine. *Quality assurance informs large-scale use of ambient AI clinical documentation*.

Supports: Supports the draft-note model, physician review/editing, patient permission, quality-assurance process, documented error categories and the distinction between AI documentation assistance and medical decision-making.

<https://permanente.org/quality-assurance-informs-large-scale-use-of-ambient-ai-clinical-documentation/>

25. Kaiser Permanente Division of Research. *AI-assisted notetaking gains steady support*.

Supports: Additional first-party research summary supporting scaled use, documentation burden reduction and physician experience. Used as corroboration rather than a separate employment outcome.

<https://divisionofresearch.kaiserpermanente.org/ai-assisted-notetaking-gains-steady-support-from-kaiser-permanente-physicians/>

26. DEUTSCHE TELEKOM

Deutsche Telekom. *AI agents for mobile network*.

Supports: Core production source for RAN Guardian. Supports multi-agent monitoring, analysis, recommendation and autonomous corrective action, including resource reallocation and configuration changes, with selected processes compressed from roughly one hour to a few minutes. Does not establish a workforce reduction or complete human approval/override matrix.

<https://www.telekom.com/en/newsroom/latest-updates/media-information/2025/11/deutsche-telekom-ai-agents-for-mobile-network>

27. Deutsche Telekom. *Deutsche Telekom and Google Cloud Collaborate for Superior Network Experience with Agentic AI*.

Supports: Supports follow-on operating evidence, including more than 100 autonomous remediation actions and the MINDR direction toward broader cross-domain diagnostics and remediation. MINDR's planned / expanding architecture is kept distinct from the already-live RAN Guardian evidence.

<https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/2/mindr-ai-agents-in-telekom-network>

28. Deutsche Telekom. *How much should an AI agent be allowed to decide on its own?*

Supports: Supports later enterprise agent-governance principles around identity, permissions, logging, risk classification, human approval for particularly critical actions and organizational responsibility. Used as follow-on governance context, not retroactive proof of every RAN Guardian control.

<https://www.telekom.com/en/newsroom/latest-updates/news/2026/09/how-much-should-an-ai-agent-be-allowed-to-decide-on-its-own>

REFERENCES

29. BNY

BNY. *Unlocking Value with BNY's Enterprise AI Platform.*

Supports: Core operating-model source. Supports Eliza, employees building agents across technical and nontechnical functions, firmwide onboarding, role-based learning, analyst bootcamps, digital employees embedded in payment teams, and governance elements including credentials, supervisors, guardrails, telemetry and approvals.

<https://www.bny.com/corporate/global/en/insights/unlocking-potential-enterprise-ai-platform-bny.html>

30. BNY. *Annual Report 2025.*

Supports: Formal company source supporting 100% employee access to Eliza, 171,000 AI learning hours, nearly 50% of employees building AI agents, 160 enterprise AI solutions in production and 134 digital employees live at year-end. These are adoption and operating-model metrics, not human-job displacement counts.

<https://www.bny.com/corporate/global/en/investor-relations/annual-report-2025.html>

31. BNY. *Analyst Program.*

Supports: Current early-career evidence showing the 24-month analyst pathway, rotations, mentoring and AI learning paths including AI Explorers, AI Builders and AI Pioneers. Used to show AI becoming part of formal professional formation.

<https://www.bny.com/corporate/global/en/about-us/careers/students/analyst-program.html>

32. BNY. *Responsible, Ethical Use of Data and AI.*

Supports: Supports enterprise governance principles including approval processes, risk management, human oversight, model review, explainability and lifecycle monitoring. Enterprise policy does not establish the exact decision rights of every production workflow.

<https://www.bny.com/corporate/global/en/about-us/technology-innovation/artificial-intelligence/data-ai-ethical-commitment.html>

33. Microsoft. *BNY customer story on Eliza.*

Supports: Vendor-hosted follow-on evidence supporting later 2026 process metrics such as faster client-settlement inquiries, AI-assisted onboarding and payment-repair activity. These metrics are treated as customer-story evidence rather than audited workforce outcomes.

<https://www.microsoft.com/en/customers/story/27322-bny-microsoft-365-copilot>

34. EXTERNAL BENCHMARKS

World Economic Forum. *Future of Jobs Report 2025, Workforce Strategies.*

Supports: Supports employer-expectation benchmarks used in the report: 73% planning faster process/task automation, 63% planning workforce augmentation with technology, 51% planning internal movement from declining to growing roles and 41% foreseeing staff reductions due to skills obsolescence. These are expectations, not observed enterprise outcomes.

<https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/4-workforce-strategies/>

35. International Labour Organization. *Generative AI and Jobs: A 2025 Update.*

Supports: Supports the global task-exposure benchmark that one in four workers is in an occupation with some GenAI exposure and the conclusion that transformation is more likely than wholesale redundancy for most exposed occupations because human input remains necessary.

<https://www.ilo.org/publications/generative-ai-and-jobs-2025-update>

REFERENCES

36. EXTERNAL BENCHMARKS

Stanford Digital Economy Lab. *Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence*, revised August 2026.

Supports: Supports the descriptive finding that employment among workers aged 22–25 in highly AI-exposed occupations stood about 19% below the path implied by less-exposed peers, with adjustment appearing primarily through reduced hiring rather than increased separations. The authors explicitly caution that these are descriptive patterns rather than causal estimates.

<https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

37. World Economic Forum. *Artificial Intelligence and the Future of Entry-Level Work*, 2026.

Supports: Supports the entry-level framework covering job access, job design, talent pipelines and education-system alignment, and provides external context for the pressure on early-career work.

<https://www.weforum.org/publications/artificial-intelligence-and-the-future-of-entry-level-work-a-framework-for-safeguarding-and-reinventing-early-career-pathways/>

38. OECD. *Skills in the AI Age*, 2026.

Supports: Supports external context on changing skill requirements, labor-market adaptation and the need for new capabilities as AI diffuses across work.

https://www.oecd.org/en/publications/skills-in-the-ai-age_972bd15e-en.html

39. International Labour Organization. *Quality Apprenticeships Recommendation*, 2023 (No. 208).

Supports: Supports structured and compensated learning, qualified supervision, occupational competence and the institutional design of quality apprenticeship systems.

<https://www.ilo.org/resource/other/r208-quality-apprenticeships-recommendation-2023>

40. National Institute of Standards and Technology. *AI Risk Management Framework / Human-AI Interaction*.

Supports: Supports the principle that human roles and responsibilities in AI decision-making and oversight should be clearly defined and differentiated, and that human-AI configurations can vary from highly autonomous to fully manual.

<https://airc.nist.gov/airmf-resources/airmf/appendices/app-c-ai-risk-management-and-human-ai-interaction/>

41. European Union. *Regulation (EU) 2024/1689, Article 14*.

Supports: Used narrowly as a governance benchmark. Article 14 states that human-oversight measures for high-risk AI systems should be commensurate with risk, level of autonomy and context of use. This report does not treat the Article as validating the governance of the enterprise cases.

<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>

42. Stanford Digital Economy Lab. *How Does AI Change Labor Demand? Evidence from 41 Countries*, 2026.

Supports: Supports the cross-country finding that foreign affiliates of AI-adopting companies reduced the junior share of their workforce relative to comparable controls, with the change primarily driven by senior-employment growth rather than a simple decline in junior employment. Used as composition evidence, not as a direct estimate for the nine case organizations.

<https://digitaleconomy.stanford.edu/publication/how-does-ai-change-labor-demand/>

43. United Nations Development Programme. *Human Development Report 2025: A Matter of Choice*.

Supports: Used conceptually for human agency, complementarity and the idea that institutional choices shape how AI affects human development. It is not used as enterprise-outcome evidence.

<https://hdr.undp.org/content/human-development-report-2025>

SOURCE & METHOD

Source and Method Note

This report synthesizes nine AIGW enterprise Deep Dives rather than re-researching each institution from zero inside the flagship manuscript. Each Deep Dive was constructed from attributable public evidence and separated fact, interpretation and implication. The flagship preserves those evidence boundaries and uses selected underlying sources for traceability.

The case set is purposive. It is designed to compare transition mechanisms across sectors and organizational contexts, not to estimate prevalence. The report therefore makes mechanism claims such as “this transition can occur” or “these cases reveal a recurring pattern,” rather than statistical claims about the share of firms experiencing each mechanism.

When quantitative outcomes are company-reported, they are presented as such. When evidence comes from vendor-hosted customer stories, journalism or executive interviews, the source type is kept visible. Enterprise-wide metrics are not silently converted into outcomes of a narrower workforce transition.

The evidence hierarchy used throughout the AIGW research process prioritizes:

1. formal company reports and first-party operational evidence;
2. government, regulatory and public-institution sources;
3. independent institutional research;
4. attributable journalism where it adds material information;
5. vendor customer stories with explicit labelling and appropriate caution.

The report does not infer: - layoffs from automation alone; - reskilling from training access alone; - role transition from tool use alone; - net job creation from movement into new roles; - human authority from the mere presence of a person in the workflow; - clinical, financial or quality outcomes from time savings alone.

The objective is not rhetorical certainty.

It is a more precise account of how AI is changing the architecture through which enterprises organize human work.