

STEWARDING THE NEXT GENERATION OF ORGANISATIONS

Where people and agents work in sync

Copyright © 2026 Accenture. All rights reserved.

A point of view for Singapore's C-suite

AI

is restructuring the relationship of people to technology, what Project Stonehenge defines as a shift from adopting agents as a technological tool to creating value with agents, with humans in the lead. This puts CEOs at the forefront of re-imagining how work should be designed, presenting them with a new set of operational choices that will shape the future of their organisations for the generations to come.



Accenture Singapore

In collaboration with the Trustworthy Artificial Intelligence (TRACE) Lab at the University of Cambridge's Centre for Human-Inspired Artificial Intelligence (CHIA).

Copyright © 2026 Accenture. All rights reserved.

CONTENTS

	The critical decisions most leaders do not know they're making	3
1	Project Stonehenge: Our answer to what's next for Singapore	5
2	The shift to a hybrid workforce	11
3	The implication: Leaders need to change how they think about the nature of work	17
4	Stonehenge's three-step solution: Re-architect, re-allocate and re-calculate	23
5	Where the rubber meets the road: Banks and hospitals	34
6	How organisations can move forward, and what Accenture is doing	43
7	Stewardship for Singapore	51
	Appendix: Banking workforce economics illustration	53
	About the research	55
	Authors	56
	References	57



Copyright © 2026 Accenture. All rights reserved.



The critical decisions most leaders do not know they're making

Every chief executive officer in Singapore investing in AI is making two sets of critical decisions. The problem is that most are only aware of one. The second is happening by default.

The first concerns the technology: how much to spend, which tools to deploy and how quickly to scale. Leaders discuss and debate these decisions openly at the board level. They explicitly set expectations and budgets, and measure outcomes.

Hidden under the technology decision, however, is a more consequential decision about the workforce. In focusing so intently on the technology, many leaders are implicitly ceding control over the shape of their future workforce: which capabilities their people will retain and build, how they develop and grow, and who will be able to answer for the work when something goes wrong.

That is a critical problem, because the key to capturing the transformative potential of AI is people. Leaders who allow their technology decisions to drive their workforce decisions risk eroding their people's ability to create value, and with it, their organisation's competitiveness.

As reported in our earlier publication, *The CIO Agenda in Singapore: Agentic labour and the CIO's new mandate*, 90% of the 70 Singapore technology leaders Accenture surveyed in partnership with the IT Management Association (ITMA) have moved beyond awareness of AI into implementation. Yet 46% have not begun redesigning roles or responsibilities and only 10% said job redesign was well established.[1]

Copyright © 2026 Accenture. All rights reserved.

90%

have moved into implementation, beyond awareness of AI

46%

have not begun redesigning roles or responsibilities

10%

say job redesign is well established



This gap reveals a paradox. Executives intellectually recognise agentic AI as a fundamentally transformative, categorically different technology. Yet they continue to apply the same frame of reference used for previous technology transformations, leading them to solve the wrong problems and creating consequences that may be increasingly difficult to reverse.

What is needed is an approach that reflects what makes agentic AI different. Agentic AI is a new producer of work, with its own economics, its own failure modes and the inability to be held accountable for anything it does. As a consequence, CEOs face a series of hard questions that most have never had to consider, much less answer.



What does it mean to measure the outcomes of a new type of labour that can be scaled almost infinitely?

Which tasks should stay with people and on what basis?

Who answers when the work goes wrong?

What is the right cost balance between the two types of labour?

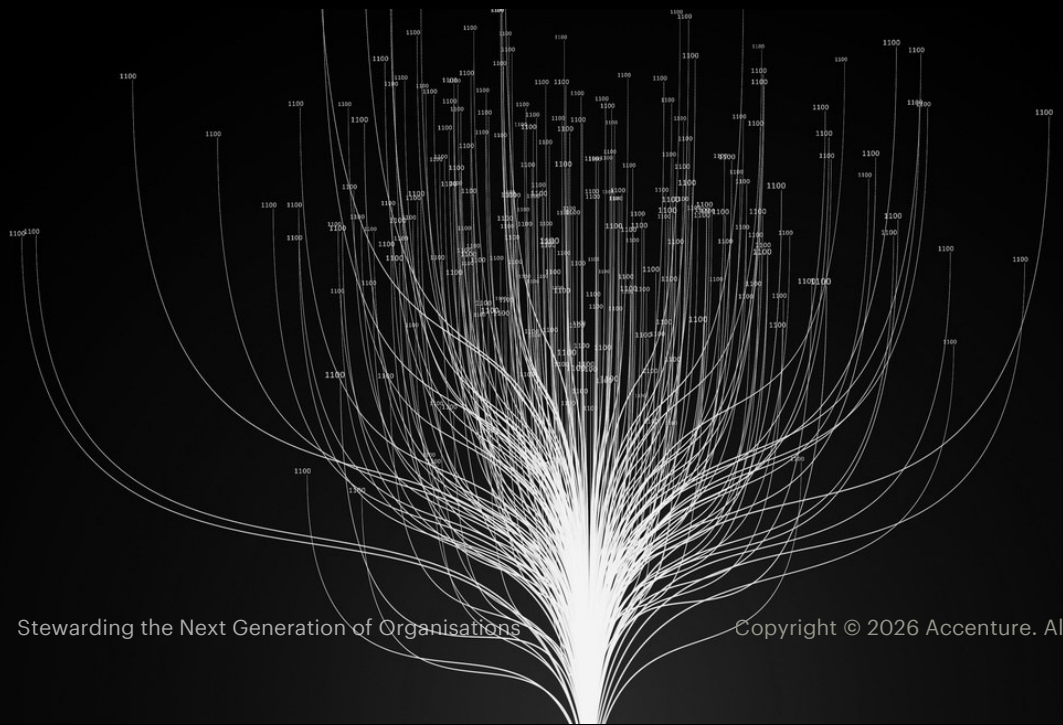
We wrote this paper to help executives across Singapore consider the challenges and opportunities before them with clarity and confidence. It is written for Singapore first, because a country whose growth has always rested on its people has the least room to get this wrong. The lessons, however, apply to all organisations navigating the generational shift to a hybrid workforce.

Copyright © 2026 Accenture. All rights reserved.

01 Project Stonehenge: Our answer to what's next for Singapore

The question our clients in Singapore ask most often is this:

How do we run our organisation well when an increasingly meaningful part of the work is done by AI, not people?



Copyright © 2026 Accenture. All rights reserved.



The name Stonehenge marks a defining threshold between the old ground and the new.

In response, we created Project Stonehenge, an applied approach to redesigning work, where humans and agents work in sync as one hybrid workforce.

It assembles distinct components into a new operating model, one that puts humans in the lead for accountable outcomes and builds the capabilities people need to work alongside technology and pursue new opportunities so that the organisation keeps growing.

To ground the thinking, we conducted three primary research streams and drew on a substantial body of secondary research.

We also collaborated with the Trustworthy Artificial Intelligence Lab (TRACE) at the University of Cambridge's Centre for Human-Inspired Artificial Intelligence (CHIA) on deployment mechanisms discussed in Section 6, including the Human API, the idea that human judgement can be formally called by an agent rather than informally escalated to.[2]

1,008

Singapore workers surveyed across financial services and healthcare[3]

70

Technology leaders surveyed with the IT Management Association (ITMA)

Review panel

Ongoing tests of the emerging model by trusted Singapore and global practitioners

Cambridge

TRACE Lab, Centre for Human-Inspired Artificial Intelligence (CHIA)

Copyright © 2026 Accenture. All rights reserved.

2 of 4

priority industries in Singapore's National AI Strategy are financial services and healthcare[4]

Importantly, we developed the Stonehenge model through the lens of financial services and healthcare, where Singapore's regulatory expectations are especially explicit and the consequences of getting the human dimension wrong are particularly visible.

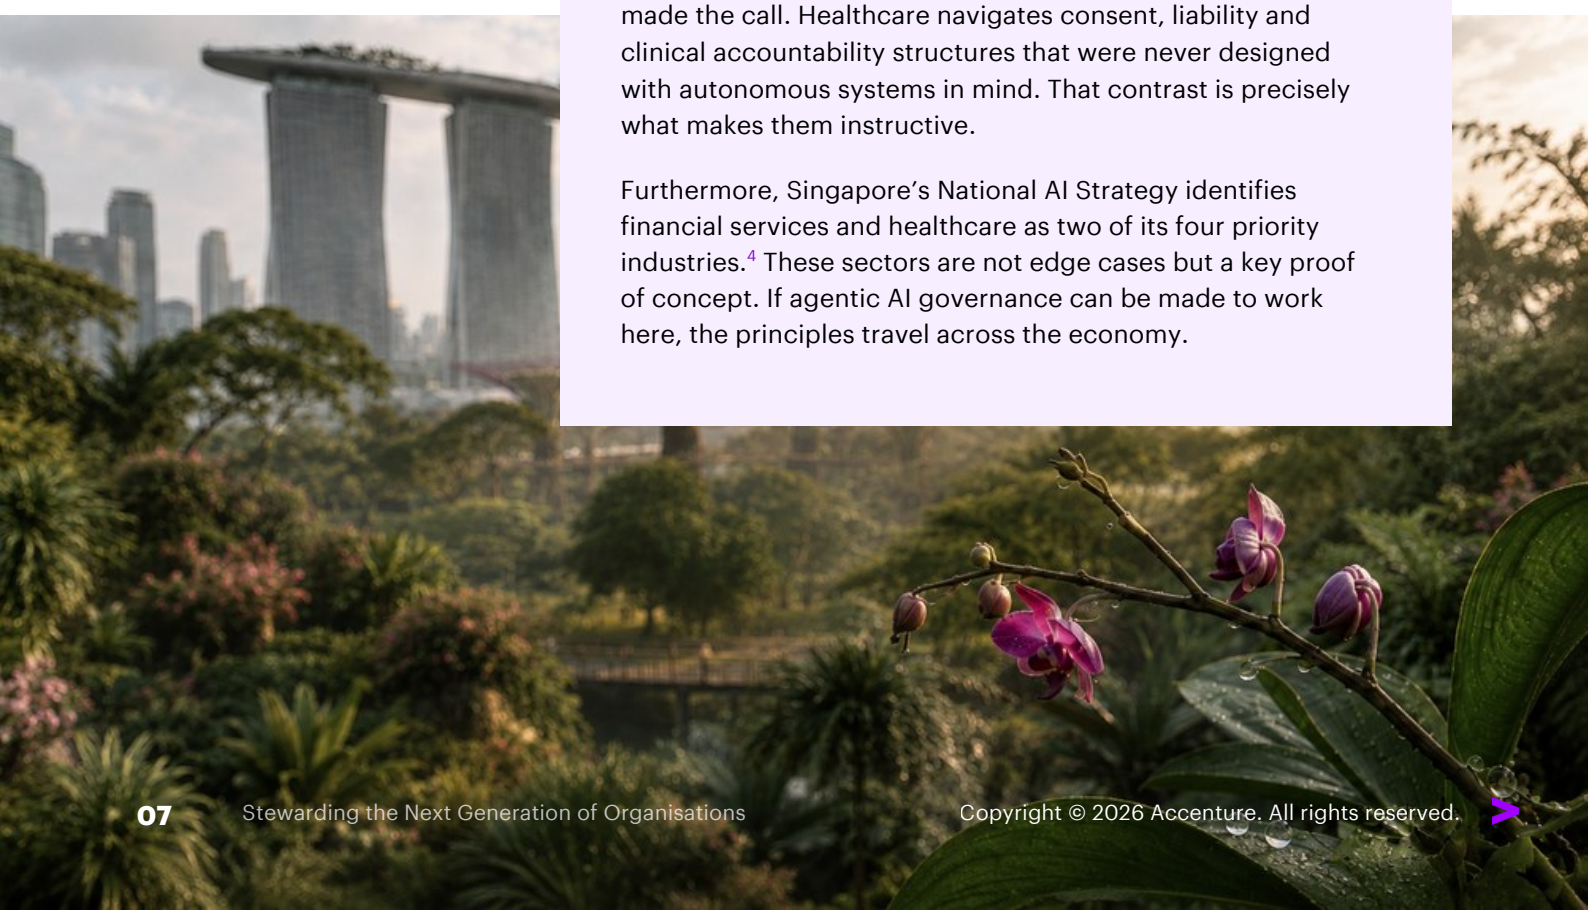
However, the model's relevance extends much further. The same principles can help any organisation use AI to reinvent work, create new opportunities and build the capabilities needed for sustainable growth, in alignment with Singapore's growth agenda.

Why financial services and healthcare?

Few industries face higher stakes for getting AI governance wrong. In financial services, errors compound across interconnected systems. A misaligned model doesn't just affect one decision, but propagates through risk, compliance and customer trust simultaneously. In healthcare, the consequences are even more direct: decisions touch patient safety, clinical outcomes and some of the most sensitive data a person can share.

Both sectors operate under demanding regulatory frameworks, but the nature of those demands differs sharply. Financial services contends with auditability, outcome validation and conduct obligations that assume a human made the call. Healthcare navigates consent, liability and clinical accountability structures that were never designed with autonomous systems in mind. That contrast is precisely what makes them instructive.

Furthermore, Singapore's National AI Strategy identifies financial services and healthcare as two of its four priority industries.[4] These sectors are not edge cases but a key proof of concept. If agentic AI governance can be made to work here, the principles travel across the economy.



Copyright © 2026 Accenture. All rights reserved.



1.1 Why this matters for Singapore: People-led growth

Since independence, Singapore's growth model has depended on three factors: capital, openness to trade and a steady supply of talent moving up the value chain. The first two are intact. The third, however, is breaking down.

67.9%

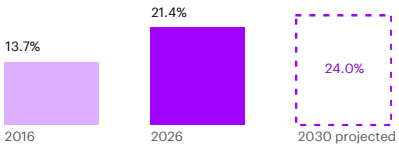
Resident labour force participation in 2025, its fourth consecutive annual decline[5]

0.87

Resident total fertility rate in 2025, a record low, against 1.24 a decade earlier[7,8]

21.4%

of citizens aged 65 and over in June 2026, rising to 24.0% by 2030[6]



64.2%

of employed residents held professional, managerial, executive or technical roles in 2025, precisely the cognitive work most exposed to agentic systems[9]

Fewer people working relative to the population, an ageing workforce and the birth rate hitting a record low are not symptoms of a temporary hiring problem. These figures point beyond a labour shortage; they are signals of a structural shift.

This matters to the state as much as to individual organisations because the roles most exposed to AI are exactly the skilled roles that make up the majority of the workforce.

Copyright © 2026 Accenture. All rights reserved.

15%

of financial services workers describe themselves as very satisfied with their working lives

22%

of healthcare workers say the same[11]



That exposure can foster or exacerbate anxiety, deepening the disengagement that is already entrenched in Singapore's workforce.

As it is, Singapore's engagement numbers are lower than the regional and global average. Just 14% of Singapore employees were engaged at work in 2025, against a Southeast Asia average of 25% and a global mean of 20%.[10]

Employees engaged at work, 2025



Our recent research is more specific where satisfaction is low: only 15% of financial services workers and 22% of healthcare workers describe themselves as very satisfied with their working lives and roughly four in ten sit in a pragmatic middle, neither clearly satisfied nor dissatisfied.

> "It is hard to say [how the organisation can support me] when I know the point of a company opting for AI is to reduce headcount."

Financial services worker, *The Humanness of Organisations in an Age of AI*, survey of 1,008 Singapore workers

Copyright © 2026 Accenture. All rights reserved.

Despite the challenges, Singaporeans show high intent to learn.

43%

of Singapore employees cite lack of time to experiment or learn as one of their top three barriers to AI usage[13]

Training participation, resident labour force

40.7% to 54.8%

An all-time high in 2025, up from the year before[12]

Digital trainees taking AI-related courses

39.0% to 58.1%

Share of digital training participants, 2024 to 2025[12]

However, willingness to re-train does not automatically translate into capacity. 43% of Singapore employees cite lack of time to experiment or learn as one of their top three barriers to AI usage.[13] Interpreting high training participation as evidence of readiness misreads what that participation is costing people.

Overall, this ties to the government's national commitment to an AI transition with no jobless growth: raising productivity without raising unemployment, helping workers move into new or redesigned roles and ensuring growth comes with good jobs and clear career progression.[14] That is the boundary condition every organisation here operates inside, whether or not it has read the strategy.



The challenge, then, is not about preserving existing jobs in their current form. It is about enabling a shrinking pool of highly skilled workers to transition and grow in capability, so that organisations can keep developing the accountable senior people they will need within the next decade.

Copyright © 2026 Accenture. All rights reserved.

02



The shift to a hybrid workforce

To get the most out of AI, C-suite executives need to consider it through a much broader lens than they have thought about for prior technologies.

Copyright © 2026 Accenture. All rights reserved.



Agents are not human employees, but neither are they simply another piece of software.

Unlike machinery or technology designed to complete a single task, AI is versatile and dynamic. It can be designed to deploy in many different capacities, completing workflows end-to-end and learning as it goes. Agentic AI, in particular, needs to be thought of as a producer of work, with implications for accountability, the kinds of people the organisation needs and what work means at all.

Once agents are doing a meaningful share of the work, organisations have to make many of the same decisions they make about any other source of productive capacity:

1. What work to assign them
2. How much capacity to deploy
3. How they work with others
4. How to monitor their performance
5. How to improve them
6. When they should be withdrawn or reassigned

2.1 How humans and agents are distinct

As part of the new hybrid workforce, humans and agents are complements rather than substitutes, but the differences between them run deeper than task allocation. They diverge across every dimension that defines an operating model: scaling capacity, judgement and context, reach and responsiveness, reliability, accountability and authority, organisational capability, continuity and resilience and knowledge source.

Copyright © 2026 Accenture. All rights reserved.

Two types of labour, compared on the dimensions that drive allocation.

Figure 1 Two types of labour, compared on the dimensions that drive allocation

	Humans How humans best contribute	Agents How agents best contribute
Production economics		
Scaling capacity	Capacity is added in fixed headcount blocks, constrained by labour supply, time and utilisation	Capacity can scale more elastically without equivalent increases in headcount, subject to technology and governance constraints
Achieving outcomes		
Judgement and context	Contextual judgement, relationships and consequential trade-offs where the right answer cannot be fully specified in advance	Rapid synthesis and execution where objectives and acceptance criteria can be sufficiently specified
Reach and responsiveness	Adapting to novel circumstances, with breadth constrained by individual time, availability and acquired expertise	Broader service coverage across demand peaks, languages, knowledge and customer needs
Reliability	Strong at novel situations that are low frequency but high impact	Strong at learning from past patterns to apply to large volumes of similar situations
Sustaining the organisation		
Accountability and authority	Exercising authority and being held accountable and answerable for consequential decisions and outcomes	Executing delegated work and providing evidence and traceability, but not itself bearing accountability
Organisational capability	Building judgement, tacit knowledge and the future pipeline of experienced talent through doing the work	Capturing and applying knowledge at scale and augmenting human performance
Continuity and resilience	Distributed human capacity and judgement that can provide fallback when AI-enabled production is disrupted	Persistent and scalable capacity, but shared models, infrastructure or providers can create correlated failures across multiple workflows
Knowledge source	Operates on both explicit and tacit knowledge and transmits knowledge that cannot be fully specified in advance	Constrained to codified knowledge (digitised, retrievable or formally specifiable)

Sources: Accenture, *The CIO agenda in Singapore: Agentic labour and the CIO's new mandate*, May 12, 2026. Infocomm Media Development Authority, *Model AI Governance Framework for Agentic AI*, June 5, 2026. Michael Polanyi, *The Tacit Dimension* (New York: Doubleday, 1966).

Copyright © 2026 Accenture. All rights reserved.



Consider how human and agentic capacity scale very differently.

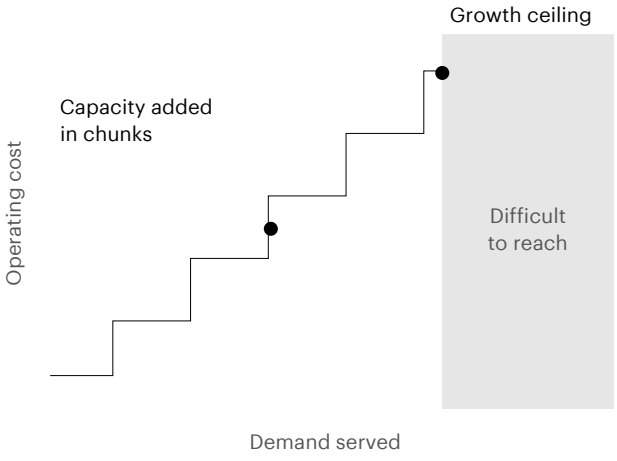
Human capacity is added one person at a time, and each new hire requires months of recruiting, training and support before reaching full productivity. Agentic capacity behaves differently. Once the underlying technology and integration are in place, it can scale more elastically with demand and at a much lower marginal cost for suitable work. This makes additional volumes, peaks in demand and some forms of work economically attractive to shift to agents where adding another human worker may not be.

That creates a powerful economic pull towards using agents to get more work done. But lower marginal production costs does not make what agents produce equivalent to what humans produce. Some uniquely human and critically important properties do not transfer when production moves to an agent.

Figure 2 The scaling economics of traditionally human versus agentic labour

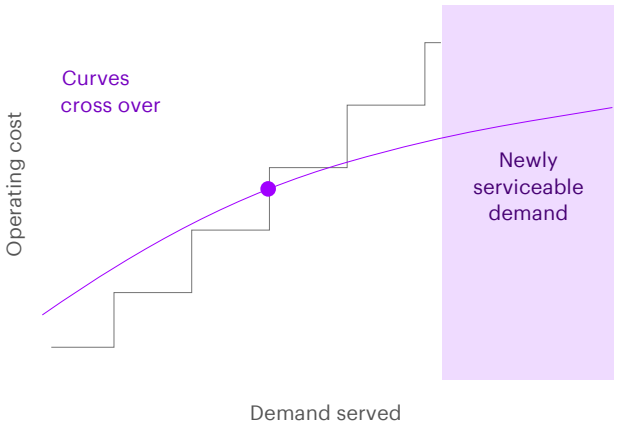
Traditional workforce

Capacity grows in steps, and becomes difficult to expand beyond the workforce limit.



Agentic workforce

A lower, smoother cost curve, with the ability to serve new, higher levels of demand.



Source: Accenture, *The CIO agenda in Singapore: Agentic labour and the CIO's new mandate*, May 12, 2026.

Copyright © 2026 Accenture. All rights reserved.



2.2 Accountability: Only humans can ultimately be accountable

Of all the ways that agents differ from humans, accountability is a key dimension. Responsibility, accountability and liability are used interchangeably in many organisations but must be distinguished in an operating model where people work with AI and agentic AI, because only one of the three can move to an agent.

Figure 3 Three terms the operating model keeps distinct

Term	What it means	Can it move to a machine?
Responsibility	The operational duty to carry out or oversee a task	Yes It can be shared with, or delegated to, an agent
Accountability	The obligation to explain, justify and answer for a decision and its outcome	No A named human remains accountable, whatever executes
Liability	The legal, professional and commercial exposure when things go wrong	No It rests with the human, the professional and the organisation, singly or together

Every accountability structure organisations rely on today assumes that a person did the work. Employment law assumes an employee. Professional liability assumes a registered professional whose name is on the file. Corporate governance assumes officers who can be identified, questioned and, where necessary, removed. Agents satisfy none of these. They have no legal personhood, no professional registration and no capacity to be held liable.

Copyright © 2026 Accenture. All rights reserved.

Moreover, there is a deeper asymmetry beneath the legal one.

An agent can be stopped, corrected, retrained or withdrawn. It cannot be deterred. It cannot feel the weight of having been wrong, and a system that cannot feel a consequence cannot carry the obligation that consequence enforces. Most human control systems rely on deterrence, which is simply absent on one side of a hybrid workforce.

In Singapore, this is not a philosophical position. The Monetary Authority of Singapore's (MAS) guidelines on AI Risk Management, issued in October 2026,[15] keep accountability with the institution and its people: the board and senior management oversee AI risk, deployed AI needs a designated accountable person, and a firm that buys AI from a vendor remains primarily accountable for it. These build on the fairness, ethics, accountability and transparency (FEAT) principles, which since 2018 have expected firms to place automated decisions under an appropriate internal approving authority.[16] Health sector guidance treats AI as an aid that should "enhance, not replace, professional judgement" and makes human oversight mandatory for all clinical use.[17] Both are versions of the same requirement: a named person behind a consequential decision.

Accountability must therefore be a choice made before deployment, not a name assigned after it. Meaningful accountability depends on whether:

1. The right authority is involved
2. That person has the information the decision requires
3. They are capable at the moment of the decision
4. They can realistically review, intervene or supply accurate input

A person who cannot do those things has been named as accountable without being placed in a position to be.

That requirement, in turn, introduces a new asymmetry: agent capacity can expand rapidly, but human judgement and accountability do not scale in the same way.

An organisation can therefore increase the amount of work being produced much faster than human capacity to assess that work, make decisions on it and answer for the outcome. This is why agents cannot simply be added to the existing operating model.

Copyright © 2026 Accenture. All rights reserved.

03

The implication: Leaders need to change how they think about the nature of work

Integrating agents into the workforce has profound implications for how leaders should think about the nature of work itself.

What is needed is a new way to measure value.



Copyright © 2026 Accenture. All rights reserved.



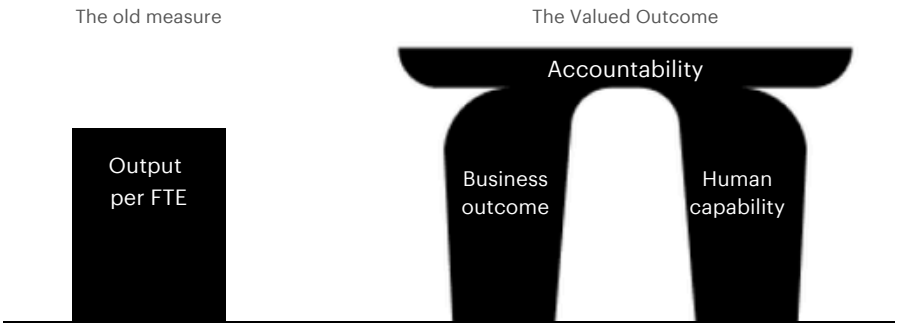
For a century, the unit of account for work was the hour of human effort and its proxy was the full-time equivalent (FTE). Output per FTE-hour worked well because work and human effort were the same thing. That is no longer true after the introduction of agents, and the measure now actively misleads. What is needed is a new way to measure value.

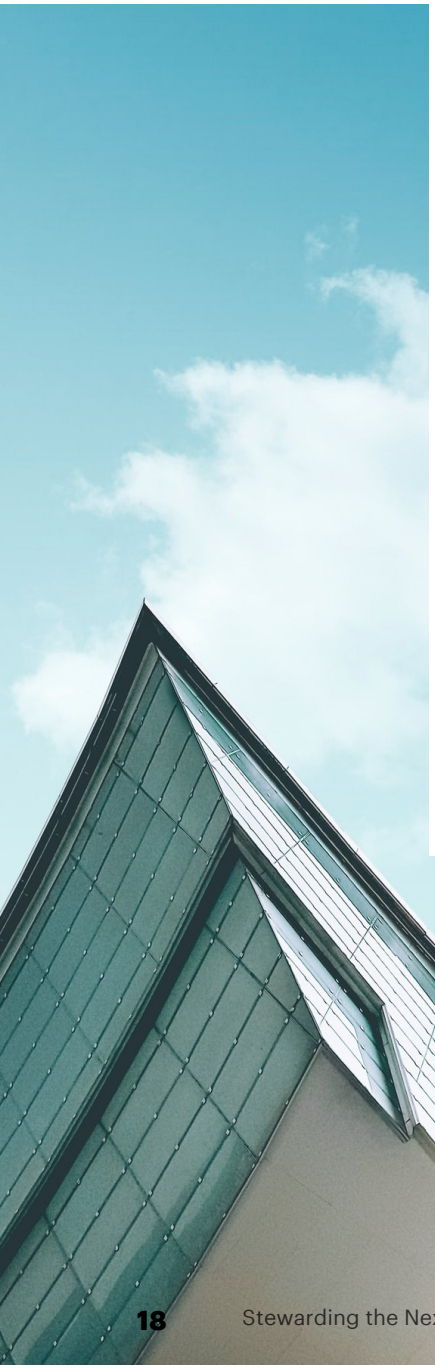
3.1 From throughput to valued outcomes

The existing approach of measuring productivity fails to capture the true value of a hybrid workforce made of both humans and agents. It quantifies the stock of input a company has and then asks how much output can be produced by that stock. Agentic production shatters this equation, seeming to scale output at near zero marginal cost but without the essential accountability that makes that output worthwhile.

This unbundling of productivity from other outcomes such as accountability, organisational capability and business continuity makes throughput an almost arbitrary measure, rendering it misleading as a management signal on its own. The replacement is a measure of work that describes flow rather than stock. We call it the Valued Outcome.

Figure 4 The change in the measure of work





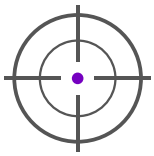
Copyright © 2026 Accenture. All rights reserved.

The old equation	The new equation
Output per FTE-hour	Valued Outcome = (Achieved business outcome ∧ human capability grew), gated by accountable provenance
Counts input, held by one kind of producer. Assumes work and human effort are the same thing. Rewards volume. Silent on who answers for the result and on what the arrangement does to the people inside it.	Measures value flowing through a human and agent pairing to a result someone owns. Holds its meaning in both directions of the pairing and across all five tiers of work discussed in the next section.

Valued Outcome = (Achieved business outcome ∧ human capability grew), gated by accountable provenance

The symbol ∧ is the Boolean AND: the outcome is valued only when both conditions are met together.

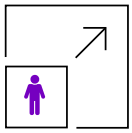
The Valued Outcome has two flow dimensions and one gate, and the distinction between them matters.



Business outcome achieved

Flow dimension

The first flow dimension is whether the work met the intended business outcome to the standard expected, deliberately leaving out the quantity of work produced. For example, if a contact centre handles a hundred times the volume at the same standards such as the same service level agreements, Achieved is satisfied and nothing further is claimed about the volume itself. The Achieved measure fails when output is generated but the real outcome is not met.



Human capability grew

Flow dimension

The second flow dimension is whether the people involved end up more capable than they began. This is the reinforcing loop that expands the system's binding constraint, which is human judgement. This becomes most relevant when leaders have to make decisions on where the freed capacity of the humans is redirected towards. Over a period of time, a human worker's job description should demonstrate elevation towards higher value density tasks. For example, from coordinating, writing, maintaining type of tasks to analysing, managing, evaluating type of tasks. When tracked longitudinally it becomes the capability trajectory: whether a team's judgement is rising or falling.



Accountable provenance

The gate

Finally, an outcome is considered valued only if there is a traceable line of human accountability across the human–agent pairing. This is not a third dimension but a gate, the precondition for achieving the valued outcome. Regardless of whether it is the human or the agent performing the work, there always needs to be meaningful human accountability.

Copyright © 2026 Accenture. All rights reserved.

There are two important considerations. First, the Valued Outcome must be defined at the right level of proximity to the work (Level 1, Level 2 process). For example, a contact centre should measure whether the customer's problem was resolved rather than how many calls were handled, or whether the need for the call could have been removed altogether. Second, a comprehensive set of defined Valued Outcomes is not a prerequisite. An organisation can start with the measures its sector already accepts, refine them around the capabilities being transformed and let the outcome definition and agentic transformation mature together.

Measuring a Valued Outcome means measuring its two flow dimensions, which move on different timescales. The measurement of human capability grew is inspired by Kirkpatrick's model developed to evaluate the effectiveness of training programmes.[18]

Measuring whether human capability grew

Capability level	What it asks	How it maps to the Valued Outcome
1 **Engagement**	Does the person engage with AI-enabled work, and is there confidence and buy-in?	A leading signal but not a complete proof of capability growth. Measured by pulse surveys, adoption and sentiment analysis.
2 **Preservation of current capability**	Has the person preserved their proficiency in the capabilities required to do the work?	Demonstrated competence, not self-rated. Measured by objective instruments to assess business outcomes (such as service-level agreements met), practice tests and recertifications.
3 **Demonstrated enhanced capability**	Did the person learn new capabilities that allow them to perform higher value density work?	Demonstrated competence for performing higher value work. Measured by supervisor feedback, review of the work performed and new or advanced certifications.



Organisations may combine several measures as appropriate. Preference should be given to objective and observed measures. Measurements can be aggregated at the level of the team or cluster of tasks. Measures should also be read longitudinally as capabilities develop or degrade over time.

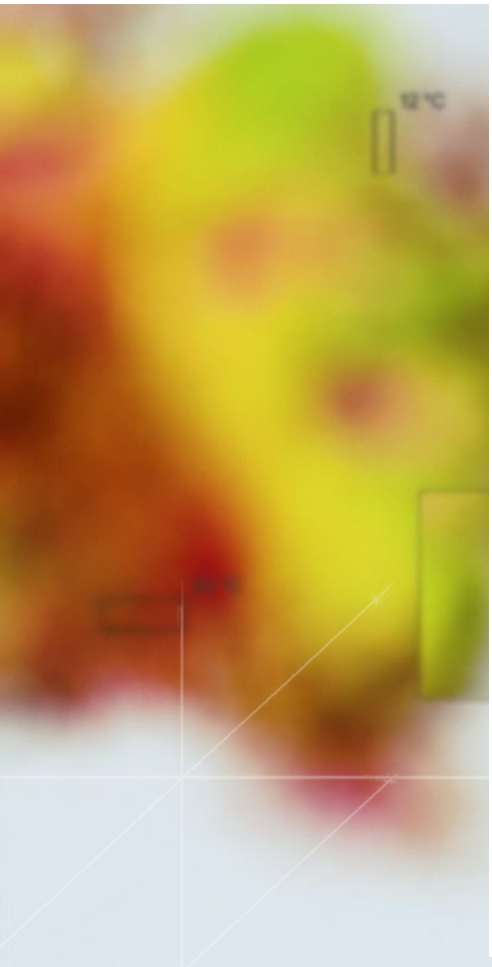
Copyright © 2026 Accenture. All rights reserved.



3.2 Why this matters: Throughput trades longevity for the current period

A leadership team that optimises only for throughput is not just making an error of measurement. It is trading off long-term organisational capability for a short-term throughput increase. A simple case might look something like this: automate the junior analyst's work and book the saving for the quarter, only to discover a decade later that you have stopped producing senior managers, because the junior work was the training ground.

Beyond the upstream shrinking of the talent pipeline, prioritising throughput also comes at the expense of the organisation's existing skills. There is increasing evidence that skills atrophy is a predictable consequence of optimising human–AI interactions purely for productivity. In a study reported in *The Lancet Gastroenterology and Hepatology*, endoscopists routinely exposed to AI-assisted colonoscopy saw their unaided adenoma detection rate fall from 28.4% to 22.4%.[19] That is measured deskilling of experienced practitioners after months of routine support. It is not a story about novices failing to learn. It is about experts unlearning. If it can happen to endoscopists, it can happen to credit officers, analysts and advisers. MAS' latest guidelines explicitly name skill degradation from overreliance on generative AI as one of the human-factor risks that AI can introduce.[20]



Higher confidence in a generative AI tool is associated with less critical thinking by the person using it.[21] The automation-complacency literature has warned for four decades that automating the routine portion of a task leaves the operator with a decayed skill set at precisely the moment the rare, hard case arrives. [22,23]

Even when companies recognise that over-reliance on AI can cause skill atrophy, they might still choose to buy talent instead of building it. This seems like a rational choice on paper. However, it is impractical to scale because tenure, context and institutional understanding cannot be bought but only retained or cultivated, and an organisation that only buys eventually stops producing the senior people that every other role depends on.

Accountability in a regulated economy is not a fixed assignment made once. It has to be renewed continuously as people retire, move on or step up, and it can only be renewed if a pipeline has been developed through years of progressively harder work.

Copyright © 2026 Accenture. All rights reserved.

This is why **growth** needs to be part of the equation.

1 in 5

organisations with heavy AI adoption report that nothing fundamentally new has been created with the time AI has saved them[24]

Growth is not top-line revenue alone. In a private enterprise it means competing and defending market share, elevating the customer and employee proposition, entering new arenas, building ecosystems, resetting the cost base in order to reinvest and deepening trust. In a mission-led institution it means reach, quality and safety, access and equity, capacity and resilience for an ageing population, and public legitimacy for how AI is used on citizens and patients. Singapore's healthcare system often has to deliver both commercial sustainability and public mission.

AI releases time, cognitive bandwidth, operating capacity and capital as a capacity dividend. A leader can choose to harvest the capacity dividend as cost, refill it with more of the same work, use it to relieve unsustainable workload or reinvest it in capability, service and new demand. Choosing to convert it into cost savings improves a period's efficiency. Doing more with the same human labour is work intensification. Easing unsustainable workloads protects the capability base; only reinvesting the capacity dividend converts productivity into growth.

Yet one in five organisations with heavy AI adoption report that nothing fundamentally new has been created with the time AI has saved them.[24]

The trap is not productivity itself, but productivity without a named destination and without a fair account of who captures the gain.



Copyright © 2026 Accenture. All rights reserved.

04

Stonehenge's three-step solution



01

RE-ARCHITECT

work by tier, not title

02

RE-ALLOCATE

work across the nine types of human worker

03

RE-CALCULATE

the cost of both human and agentic labour

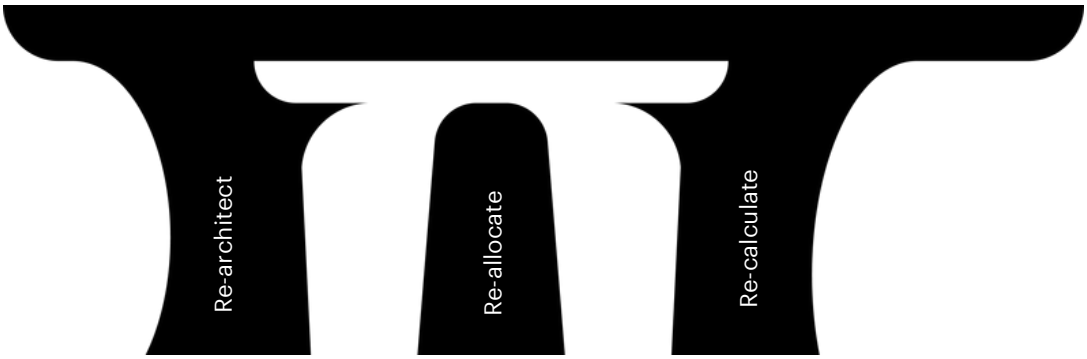
Copyright © 2026 Accenture. All rights reserved.

If leaders accept the premise that agents are not merely a tool but a new type of labour, and are willing to redefine work around the Valued Outcome, the question then becomes how to redesign work around this new reality.

A radically new operating model is required, one that goes beyond the familiar territory of where to deploy AI or which tasks to automate and into terra incognita, an unmapped landscape with no precedent.

Built on research and our direct experience with clients, Stonehenge offers a way forward. The Stonehenge operating model is a three-step decision flow: re-architect work by tier, not title; re-allocate work across the nine types of human worker; and re-calculate the cost of both human and agentic labour. Together, they turn a new understanding of labour into a practical system for designing, deploying and governing the hybrid workforce.

Figure 5 The three steps of Project Stonehenge



01

RE-ARCHITECT

work by tier, not title

02

RE-ALLOCATE

work across the nine types of human worker

03

RE-CALCULATE

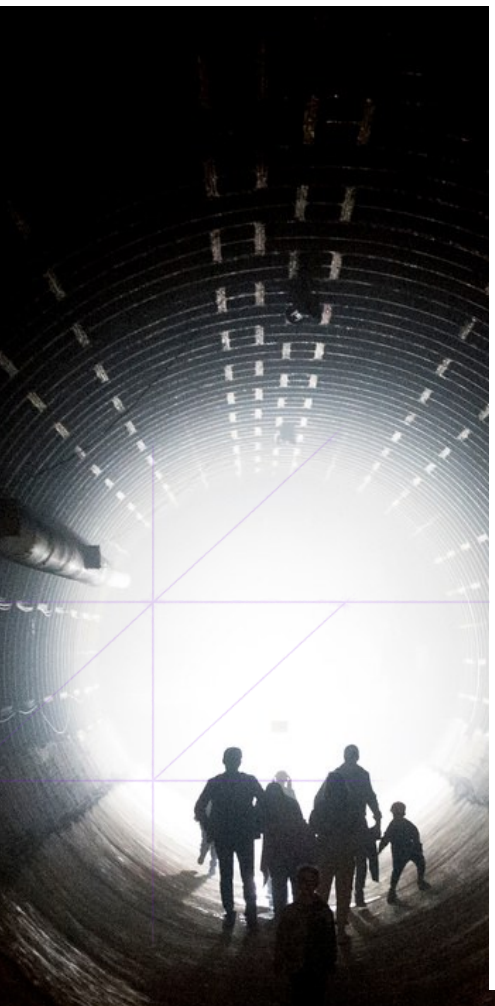
the cost of both human and agentic labour

Copyright © 2026 Accenture. All rights reserved.

4.1 Re-architect work by tier, not title

Work has traditionally been organised around titles: a relationship manager advises clients while a nurse coordinates care, for example. In both cases, the title acts as a shorthand for a bundle of very different kinds of work that belong to an individual human worker. However, as agents produce an increasing share of work, the title-based definition of work collapses.

Instead, Stonehenge adopts a more precise definition: Work is the intentional application of capability, human or machine, to produce a valued outcome within a system of constraints. This definition changes the design question; it is no longer a binary choice between whether a human or an agent should do the work. Rather, leaders now have a spectrum of choices, shaped by the outcome they seek. Our definition of capability deliberately makes no distinction between cognitive and physical work. Rather, the key shift is from measuring output per FTE to measuring the Valued Outcome.



To support leaders in making the right choices, we have defined a set of design principles:

1. **Human-centred.** Humans continue to derive meaning from the work they perform. The model should avoid skill atrophy and enable humans to grow professionally.
2. **Designed for business continuity.** A human fallback path is built in should the need arise for humans to take back control, and meaningful human accountability is maintained.
3. **Optimised as a whole system, not its parts.** Leaders should invest in elevating the key binding constraint in future, which will be human judgement capacity. Keeping the minimum human share is critical.

Making these questions answerable requires a way to classify work that is finer than the job title. Stonehenge organises cognitive work into five tiers, ordered by the nature of the decision each contains rather than by seniority. A single role routinely spans several tiers, and the purpose of the work taxonomy is to unbundle the role so each part can go to the producer best suited to it.

Copyright © 2026 Accenture. All rights reserved.

Figure 6 The five-tier work taxonomy and the allocation pattern at each tier. T3 is the frontier, where the two types of labour are most tightly paired and where deskilling does the most damage.

Tier	What it is	Who leads
T1	Procedural. Structured, rule-based, repeatable work	Agent leads and executes at scale subject to acceptable risk. Human sets the rules and audits exceptions. Humans continue to perform a minimum share of work to avoid skill atrophy.
T2	Analytical. Structured cognitive production	Agent produces. Human reviews before the output is used. Humans continue to perform a minimum share of work to avoid skill atrophy.
T3	Adaptive judgement. Contextual, ambiguous, partly irreversible	Human leads and decides. Agent enriches with synthesis, options and evidence, but never decides.
T4	Relational and ethical. Trust, care, ethical deliberation, named accountability	Human only. Agents confined to logistics and coordination.
T5	Creative and strategic. Novel synthesis, institutional reinvention	Human authors and remains in command. Agent partners in synthesis only.

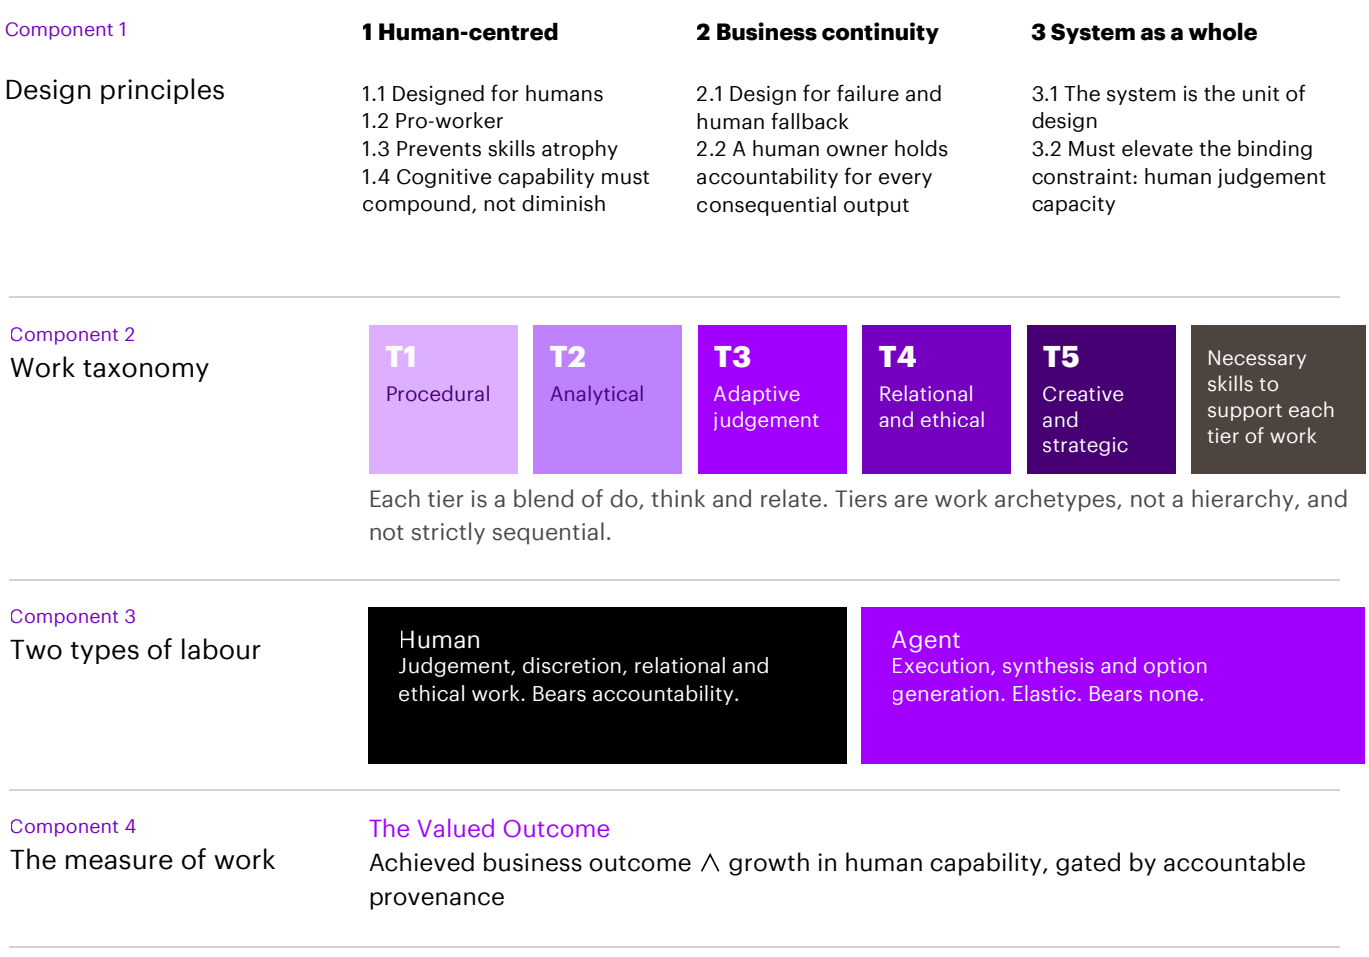
Two properties of the tiers carry most of the design weight. Value density rises as the tier rises, which is where scarce human judgement creates disproportionate value. This is where humans should concentrate as agents absorb the lower tiers. And the time it takes to verify a decision lengthens as the tier rises. A T1 error is caught by a quality check almost immediately. The quality of a T5 decision may only be visible over quarters or years. Any measure of work has to accommodate outcomes whose success cannot be confirmed at the moment of production.

The mapping to tiers is designed to be bidirectional, meaning that agents can call on human expertise and vice versa. For example, a human at T3 may call an agent to retrieve and synthesise information, and an agent running a T1 process may call a human when it hits an exception. Tiers are intended as a snapshot of a dynamic process rather than a static allocation. However, what must not move is who answers for the result, which should always be a named human with meaningful accountability. In addition, adequate safeguards need to be built into T1 and T2 tasks primarily handled by agents to avoid skill atrophy for humans. These include seeded errors, practice quotas, blind reviews and minimum dedicated review capacity for each human.

Copyright © 2026 Accenture. All rights reserved.

In the full Stonehenge operating model, the five-tier work taxonomy is the second of four components. It sits beneath the design principles, which establish the rules by which work is designed, and above the two types of labour, which define which can perform it. The final component is the Valued Outcome discussed in Section 3, the measure the system ultimately counts and optimises. Together, the four components move from the principles that govern work, to how work is classified, who performs it and, ultimately, the value it is intended to create.

Figure 7 The four components of the work architecture: design principles, work taxonomy, the two types of labour and the measure of work.



Copyright © 2026 Accenture. All rights reserved.

4.2 Re-allocate work across the nine types of human worker

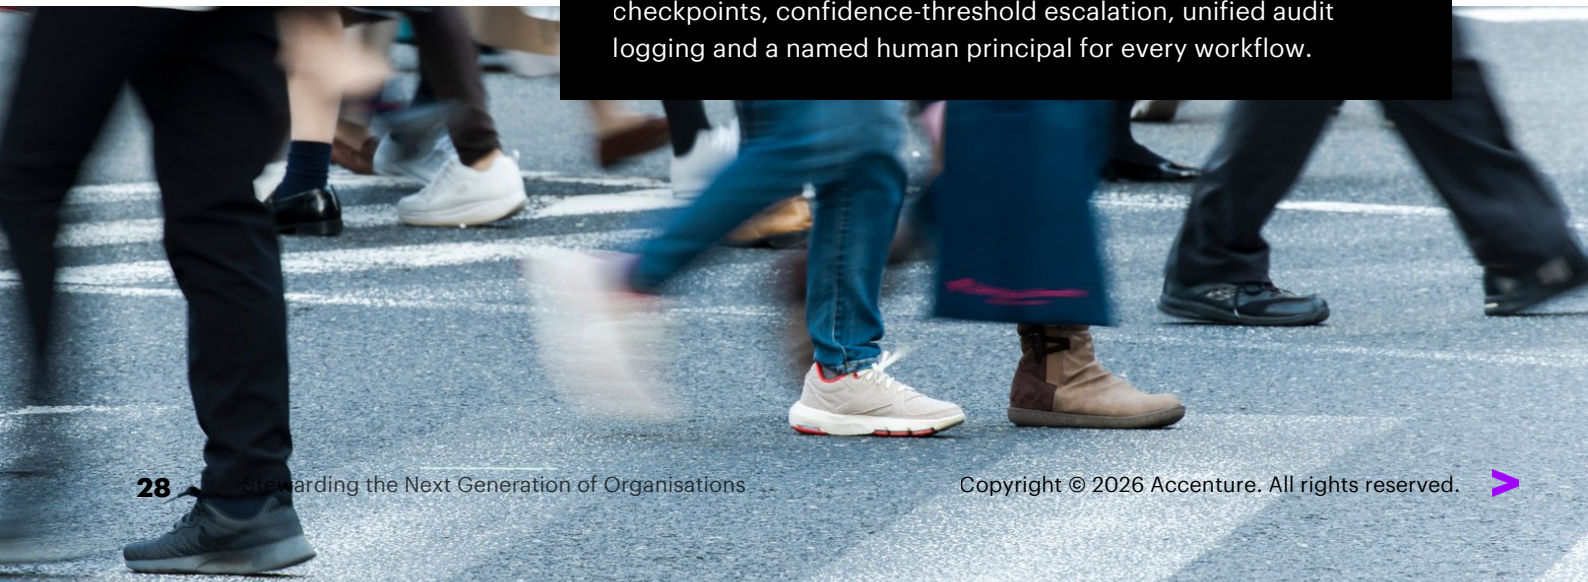
The next consequence of a hybrid workforce comprising both humans and agents is that it creates seams: the points of handoff where work passes from one type of labour to the other. When one kind of producer did all the work, coordination happened inside a single human workflow among people who shared context. A second producer introduces a seam that did not previously exist, and it is at that seam that accountability, context and trust are either preserved or dropped.

It is thus the design of the seam itself that creates, or destroys, value. However, organisations currently focus on the capabilities on either side of the seam instead, investing in the capability of the agent or the training of the person and leaving the handoff between them to look after itself.

Designing the seam requires four moves: allocate at the task level rather than the role level; classify oversight by reversibility and consequence rather than abstract risk labels; design for escalation rather than universal review; and redefine human roles around orchestration. Each move is sequenced deliberately. Task allocation determines what humans and agents do. If the human share falls below the minimum determined by leaders, the system increases the cost of assigning work to agents, thus shifting more tasks back to people. Oversight classification determines how humans stay in the loop. Escalation design determines where human review concentrates. Orchestration determines what the manager becomes. Together they constitute a redesign of the operating model around two types of labour.

The implication is clear

Design the seam deliberately, with variable-engagement checkpoints, confidence-threshold escalation, unified audit logging and a named human principal for every workflow.



Copyright © 2026 Accenture. All rights reserved.

When considering how work should be allocated across the seam, a second taxonomy is needed. While re-architecting uses five tiers to describe what kind of work each task is, it does not say which human should do it. That requires a second classification, of workers rather than the work. In the Stonehenge Work Architecture, there are nine types of human workers, defined by where each person stands in relation to the seam, rather than by function or seniority.

Figure 8 The nine human worker typologies

1 · Direct & design the seam

- **1.1 Work architect** — Designs the split · T5
- **1.2 Conductor** — Directs the work · T3–T5
- **1.3 Agent manager** — Owns the agent fleet · T3–T5

2 · Across the seam with agents

- **2.1 Co-architect & Judgement specialist** — The hard-call expert · T3–T4
- **2.2 Augmented practitioner** — Keeps a craft core · T2→T3

3 · Beyond the seam Primarily human

- **3.1 Relational or ethical anchor** — Relationships or ethics · T4

4 · The pipeline build & sustain

- **4.1 Apprentice** — Builds org memory · T1–T2
- **4.2 Developing practitioner** — Some T2–T3 · prevents atrophy

5 · Retained reserve

The ninth type · cross-cutting spine

- Retains institutional knowledge
- Trains, corrects & grades the AI
- Mentors the pipeline
- Failure fallback & takeback · T1–T3

Kept operationally warm: practised, not standby.



Copyright © 2026 Accenture. All rights reserved.

First, there are those required to direct and design the seam. These are the **Work Architect**, who decides what the work should be based on the Valued Outcome; the **Conductor**, who allocates work between humans and agents and owns the exceptions; and the **Agent Manager**, who owns, runs and maintains the agentic workforce: provisioning, identity, monitoring, retraining, cost and decommissioning.

Second, humans are also required to work across the seam. The **Co-Architect and Judgement Specialist** is the expert an agent calls at pre-defined moments that matter: uncertainty, irreversibility or an ethical line. The **Augmented Practitioner**, on the other hand, takes ownership of the workflow but makes calls on agents throughout the process.

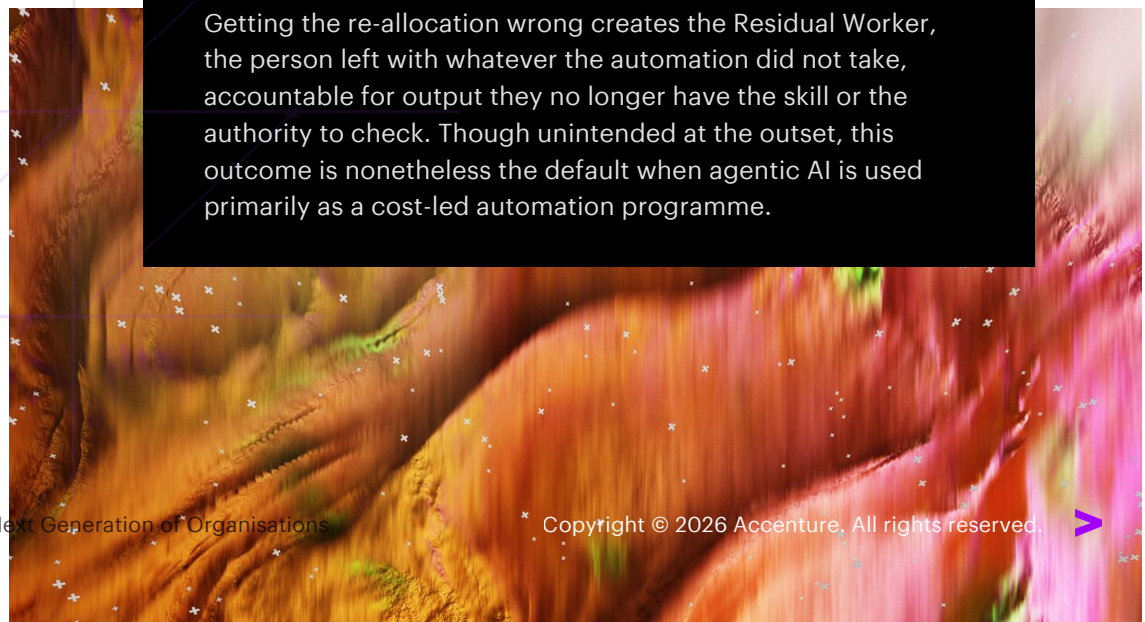
Then there is the **Relational or Ethical Anchor**, who operates beyond the seam altogether. These people hold care, trust and negotiation where human presence is itself the value, and moral judgement where a regulated decision must be made by a person.

While the first six types of workers are defined by their current capabilities, the next cluster is where those capabilities are developed within the organisation. **Apprentices** are assigned foundational work for the express purpose of acquiring basic skills, while **Developing Practitioners** focus on mastery and the critical thinking required to question and not just check AI outcomes.

Finally, the ninth type of worker is the **Retained Reserve**, cutting across all the other types but playing a unique four-fold role: retaining tacit knowledge, training and assessing AI, mentoring the pipeline and providing failure fallback. Singapore has an unusually strong pool of candidates for the Retained Reserve: experienced and retiring workers. The statutory retirement age rose to 64 in July 2026 and the re-employment age to 69, with the government committed to raising them to 65 and 70 respectively by 2030,[25] with the Part-Time Re-employment Grant, the Senior Employment Credit and SkillsFuture credits supporting that path. In an economy that cannot hire its way out of a shortage, a worker who both preserves institutional knowledge and cascades it is valuable.

The Residual Worker

Getting the re-allocation wrong creates the Residual Worker, the person left with whatever the automation did not take, accountable for output they no longer have the skill or the authority to check. Though unintended at the outset, this outcome is nonetheless the default when agentic AI is used primarily as a cost-led automation programme.



Copyright © 2026 Accenture. All rights reserved.

4.3 Re-calculate the cost of both human and agentic labour

Once leaders know the two variables within their control, what work is done and by which kind of labour, they next need to understand what each combination of those variables actually costs their organisation. To make a sound decision about their target hybrid workforce structure, the Chief Financial Officer needs to be able to rigorously defend the chosen combination of human and agentic labour, with a business case that can stand up to board scrutiny.

35%

of organisations can calculate the cost per business outcome, even for their largest AI use cases[26]

Companies are already building the business case for agentic labour, but two blind spots remain. First, measuring inputs such as tokens and salaries is not the same as attributing costs to the business work performed and outcomes achieved. Today, only 35% of organisations can calculate the cost per business outcome even for their largest AI use cases.[26] Second, even where production savings are visible, organisations can overlook what human labour previously supplied alongside production: capability development, accountability and continuity. These costs may only become apparent later, when critical capabilities or processes need to be restored.

Emerging approaches are helping organisations address the first challenge. Accenture's AI tokenomics for enterprises provides greater visibility into AI consumption and costs across models and business workloads, supporting more informed attribution and management of AI expenditure.[27] Similarly, the Tokenomics Foundation, an initiative under the Linux Foundation, has proposed a Total Cost of AI (TCA) approach that extends beyond token and model costs to encompass the broader resources required to create and operate AI capability, linking these costs to realised business value.[28] Both represent important steps towards understanding the full economics of AI production.

While the underlying principle of assessing AI economics against the whole cost required to produce measurable business value is a necessary starting point, it remains incomplete. It does not capture the full economic consequences of replacing human production, distinguish immediate value from more contingent future value, or reflect how changing labour type changes the risk profile of the workforce.



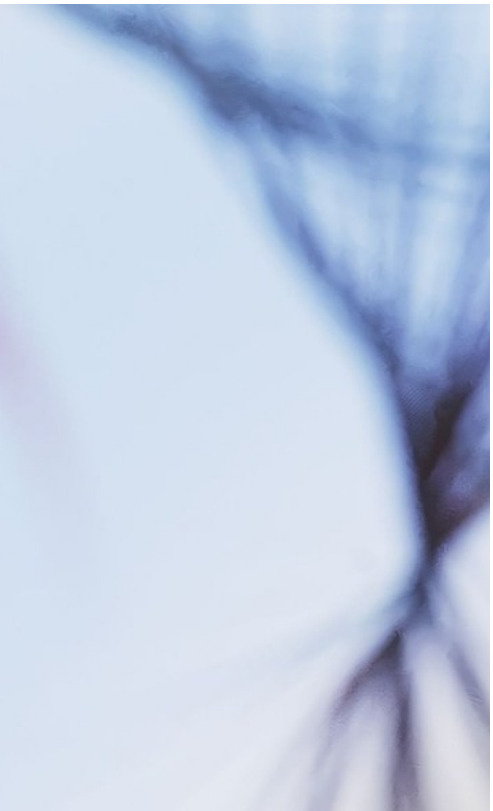
Copyright © 2026 Accenture. All rights reserved.

First, the total cost of AI needs to reflect the fact that human labour has always delivered more than immediate output. Through doing the work, people developed organisational capability, carried accountability for consequential outcomes and provided continuity when normal production was disrupted. These contributions were largely bundled into the production workforce and therefore rarely appeared as separate economic lines.

Agentic labour separates these requirements from production. The organisation's need for human capability, accountability and continuity, however, does not disappear when tasks move across the seam from humans to agents. As a result, capability, accountability and continuity must be deliberately designed—and priced—into the target operating model as what we call retained requirements.

Retained requirements refer to the capability, accountability and continuity the organisation must preserve after production shifts to AI. They may be provided through normal human roles, a Retained Reserve, controls or other mechanisms. The aim is not to duplicate costs already captured in TCA, but to make explicit requirements that were previously bundled into human production.

Second, AI unit economics need to be able to distinguish between immediate and incremental value. Most approaches already account for immediate value in the form of measures like productivity gains or cost savings. Stonehenge goes further to capture incremental value, which reflects the upside when organisations successfully redeploy freed up capacity or capture new opportunities.



Including incremental value can change the business case. For example, retaining the human labour needed to ensure accountability and long term talent development may seem like a cost driver but it actually delivers better overall economics if the incremental value created is sufficiently large. The discipline is to keep the two sources of value visible rather than combining them prematurely into one headline number.

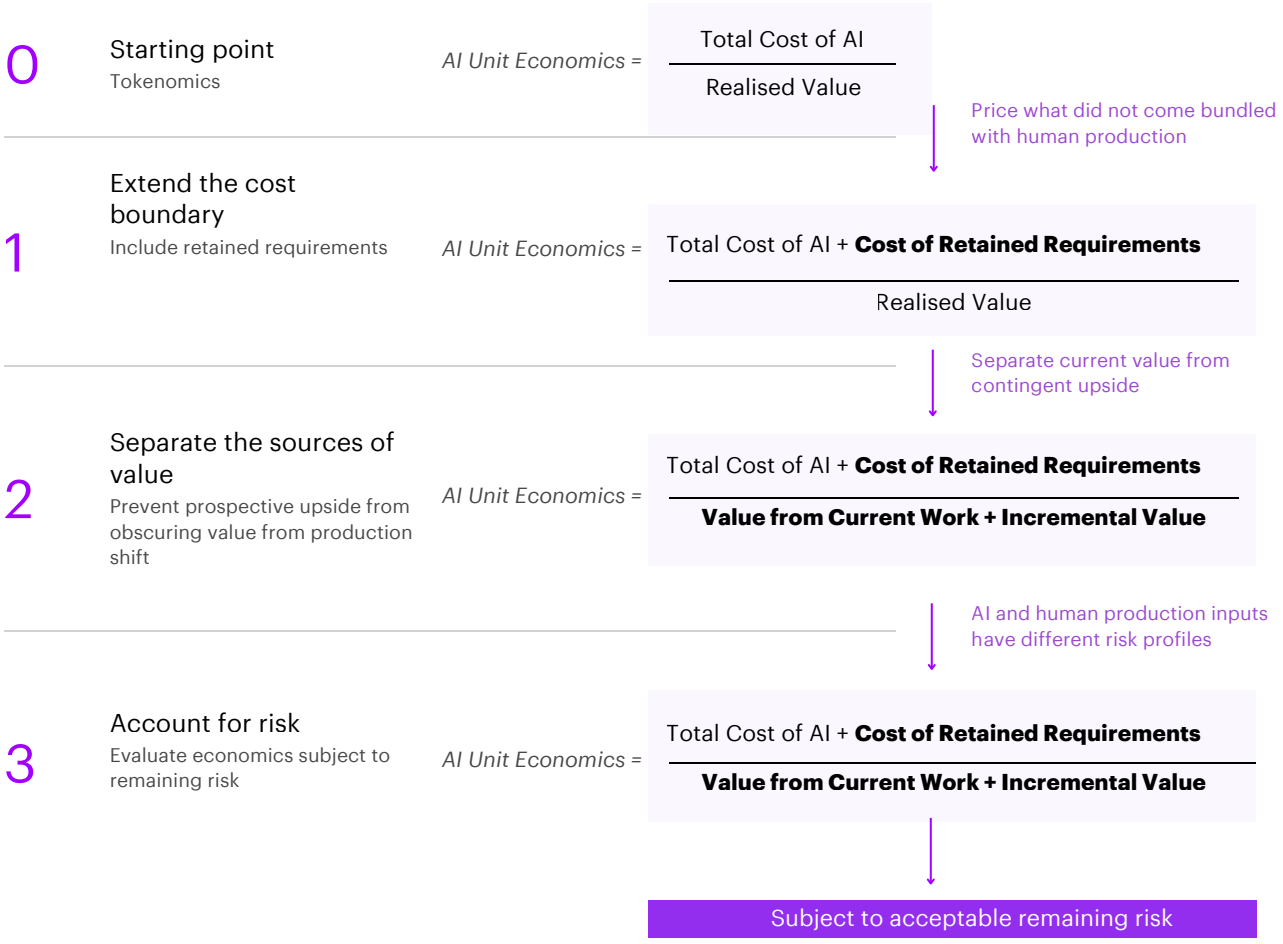
Finally, a comprehensive assessment of the rewards of a hybrid workforce must also take risks into account. Meeting the retained requirements can reduce risk, but changing the balance of humans and agents also introduces different risks. Controls can reduce this exposure, but risk never reaches zero. Retaining too much human capacity can undermine the economics of switching, while retaining too little can leave material capability, accountability or continuity exposure. Where remaining risks can be credibly quantified, they should be reflected in the economics; otherwise, they become a management gate: is the remaining risk acceptable given the economics of the target workforce?

Stonehenge extends AI unit economics in three ways: it makes the retained requirements of the production system explicit, separates immediate value from incremental value, and considers the resulting economics subject to acceptable remaining risk.

Copyright © 2026 Accenture. All rights reserved.

Interdependent steps, built on the foundation of stewardship.

Figure 9 From AI unit economics to the full economics of the hybrid workforce



Re-architect, re-allocate and re-calculate are interdependent steps of the Stonehenge approach. All three are built on the foundation of stewardship and designed to help leaders make the critical decisions that will shape the future of their organisations and the fate of the people they lead.

In this pivotal moment, leaders are also inadvertently making decisions that, on aggregate, will determine whether Singapore successfully grows its economy while bringing its people along. The stakes have never been higher.

Copyright © 2026 Accenture. All rights reserved.

05

Where the rubber meets the road: Banks and hospitals

Two industries demanding high levels of human accountability, but with very different types of work, workers and underlying economics.



Copyright © 2026 Accenture. All rights reserved.

73%

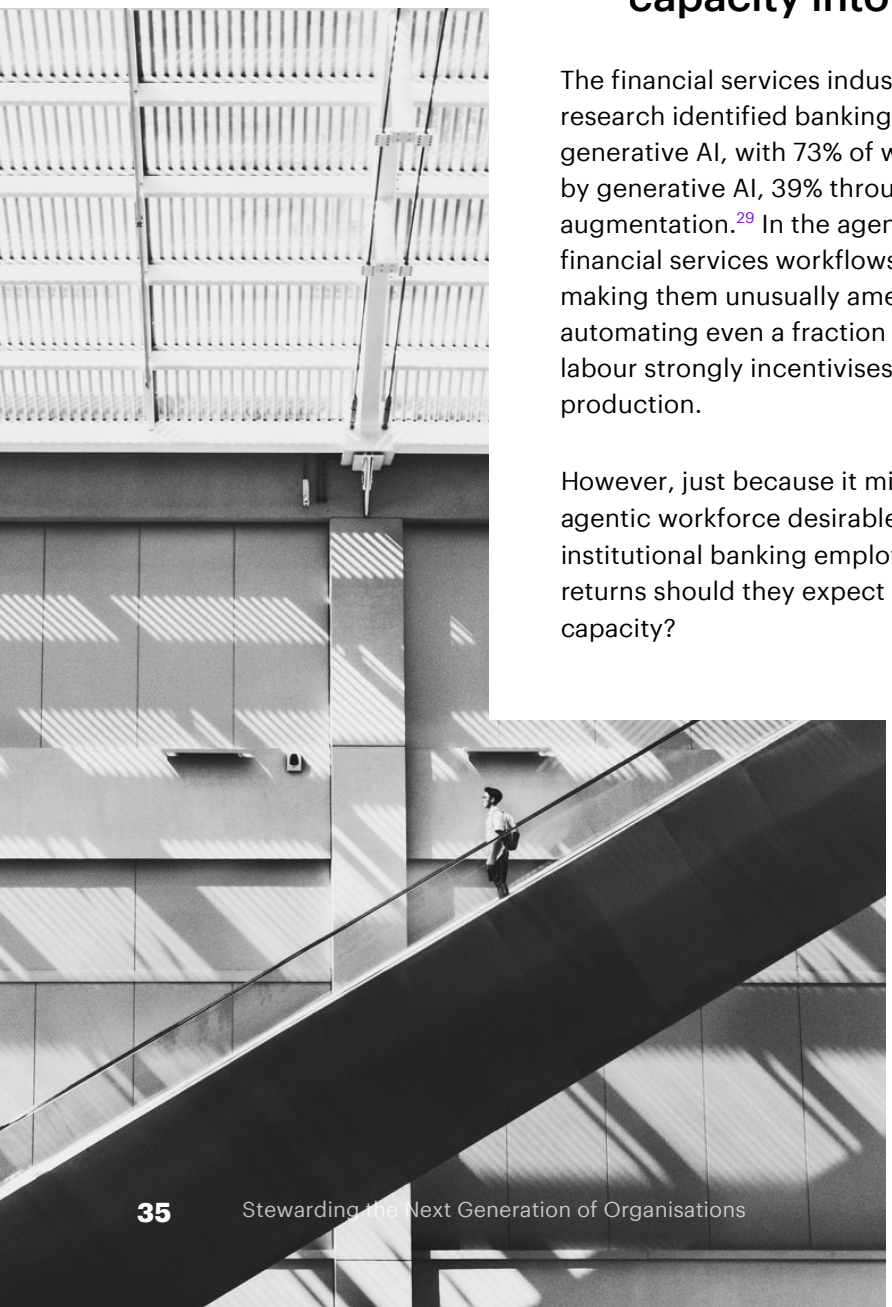
of working time in US banking potentially affected by generative AI: 39% through automation, 34% through augmentation[29]

Meeting this generational shift in the definition of labour requires organisations to critically examine their own unique circumstances and the values reflected in their decision making.

An operating model like Stonehenge is only as useful as the clarity it provides and the decisions it surfaces. How organisations choose to apply Stonehenge inside their organisations is how the operating model comes to life.

That application can look wildly different depending on the goals of the organisation and on who it serves. Nowhere is the difference clearer than in financial services and healthcare, two industries demanding high levels of human accountability but with very different types of work, workers and underlying economics.

5.1 Financial services: Converting released capacity into economic value



The financial services industry is heavily exposed to automation. Our prior research identified banking as the sector most likely to be disrupted by generative AI, with 73% of working time in US banking potentially affected by generative AI, 39% through automation and 34% through augmentation.[29] In the agentic era, this exposure is likely to intensify; financial services workflows are largely digital and information-based, making them unusually amenable to agentic labour. And the economics of automating even a fraction of the work performed by highly-paid human labour strongly incentivises the substitution of agentic for human production.

However, just because it might be feasible does not make a heavily agentic workforce desirable. Consider a Singapore bank with 1,200 institutional banking employees planning an agentic AI pilot. What kind of returns should they expect and what should they do with their freed-up capacity?

Copyright © 2026 Accenture. All rights reserved.

Recent AI initiatives at local banks report or target preparation-time reductions of around 30% for corporate credit assessments. Similarly, industry case studies suggest that 30–50% of banker capacity can be absorbed by credit assessment and adjacent internal and administrative workflows handled by agents. Assuming a conservative lower bound of 30% workflow exposure, handing work off to agents results in the release of approximately 108 FTE capacity.

1,200 FTE × 30% workflow exposure × 30% AI effort reduction = 108 FTE-equivalent released

At a market-based fully loaded cost of approximately S$150,000 per FTE, the gross annual economic value of released production capacity is therefore S$16.2 million. This represents capacity that can potentially be removed or redeployed.

108 FTE × S$150,000 loaded cost per FTE = S$16.2m

A conventional business case might treat the full S$16.2 million as value available to fund AI. While it might be tempting to bank the S$16.2 million and shed 108 roles, this approach to accounting fails to capture the true cost to the bank in terms of lost capability, accountability and continuity. When both retained requirements and incremental value are considered, the potential annual technology-cost headroom could reach approximately S$14.3 million, subject to the additional value being realised.

Using the share of managers and administrators in Singapore's employed resident workforce, approximately 17%, as a broad proxy for the judgement-bearing workforce of the bank, that translates to 204 FTE. Applying an approximately 11% annual replenishment rate based on Singapore PMET resignation rates[30] and a three-year formation period to develop judgement, the retained requirement becomes 67 FTE.

204 judgement-bearing FTE × 11% annual replenishment rate × 3 years ≈ 67 FTE

For this illustration, we assume that 67 FTE of human production capacity must be retained to develop the next generation of judgement-bearing workers. In practice, the requirement will depend on how much developmental experience remains after automation and how effectively AI can support learning.

Copyright © 2026 Accenture. All rights reserved.

At the same S$150,000 loaded cost per FTE, the retained formation requirement costs approximately S$10.1 million annually. This illustration prices capability formation explicitly. Any additional costs required to maintain accountability or continuity would further reduce the economic headroom.

67 FTE × S$150,000 = S$10.1 million of retained requirements annually

The economics of the current work therefore support approximately:

S$16.2m human production displaced – S$10.1m retained requirements = S$6.1m

In other words, after providing for retained requirements, the maximum annual attributable AI technology cost is approximately S$6.1 million for the current-work economics to remain positive, before considering incremental value.

After providing for the assumed retained requirement of 67 FTE, approximately 41 FTE-equivalents of released capacity remain available for redeployment. As the destination roles are not predetermined, the illustration adopts S$300,000 in annual operating contribution per productive FTE, informed by the 2025 bank-wide operating economics of Singapore's three local banking groups. It assumes that redeployed employees realise an average of 67% of mature operating contribution during a three-year ramp-up. Actual incremental value will depend on the roles created, demand for their services and the additional resources needed to make them productive.

This produces:

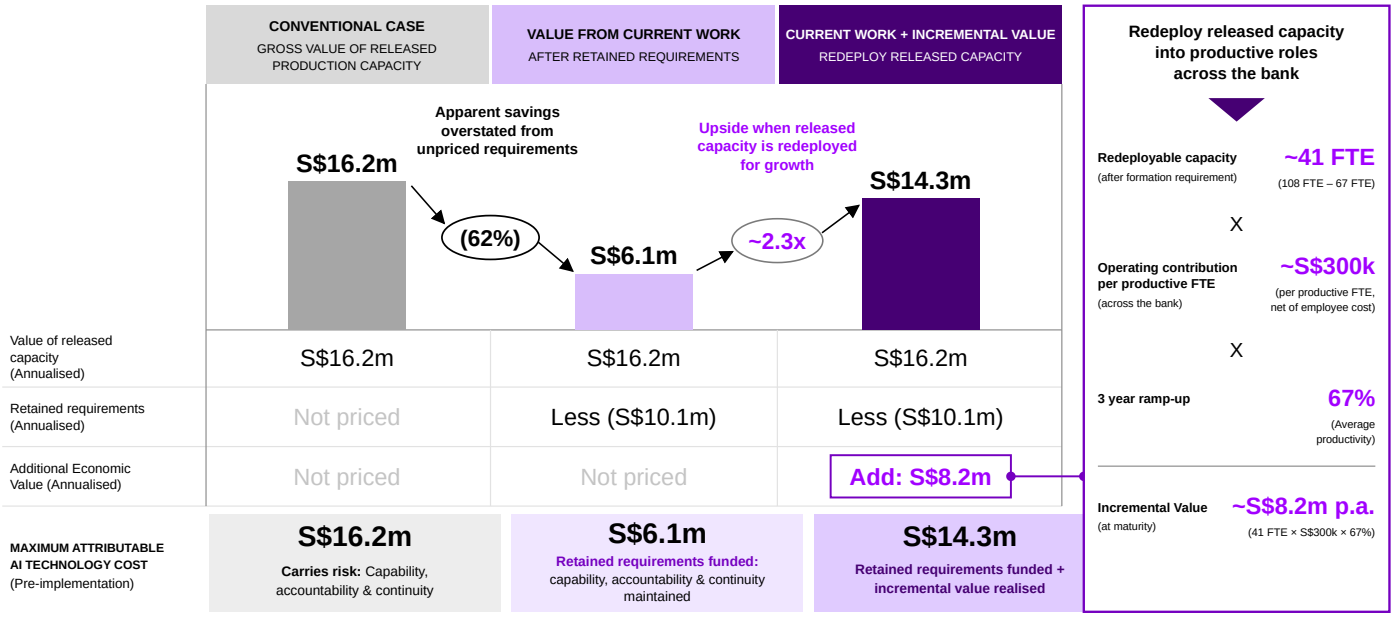
41 FTE × S$300,000 × 67% ≈ S$8.2 million of incremental value annually

Under the current-work case, the S$6.1 million represents potential savings from the 41 FTE no longer needed for production. If these employees are redeployed instead, their salaries remain on the bank's payroll. As the S$8.2 million operating contribution is calculated after staff costs, the two figures together provide approximately S$14.3 million in average annual technology-cost headroom over the three-year ramp-up, assuming the additional value is realised.

The full assumptions and calculations are set out in the Appendix.

Copyright © 2026 Accenture. All rights reserved.

Figure 10 How retained requirements and additional economic value reshape the business case for agentic AI



Finally, these economics must be considered alongside risk, which is reduced by funding the retained requirements but never completely eliminated. A smaller human production layer may increase dependence on AI infrastructure, reduce capacity for unusual cases or independent challenge, or affect responsiveness and client relationships.

Not every risk needs to be converted into a financial number. Where exposure can be credibly quantified, such as operational losses, remediation costs or regulatory penalties, its expected economic effect can be estimated using an appropriate probability and impact assessment, and reflected in the economics where relevant. Other exposures, including brand, trust, relationship quality or strategic flexibility, may be better assessed qualitatively. In either case, quantification does not replace the management judgement over whether the remaining risk is acceptable. That judgement is also becoming a regulatory expectation. MAS' AI Risk Management Guidelines, issued in October 2026 and taking effect in stages from October 2027, set expectations for board oversight of AI risk appetite and for banks to assess and manage risks at the individual use-case level. [31]

Rather than a poorly defined return on investment metric, the decision for the bank crystallises into whether immediate and incremental value, after accounting for attributable AI technology costs and retained requirements, remains attractive subject to an acceptable level of remaining risk.

Copyright © 2026 Accenture. All rights reserved.

5.2 Healthcare: Better care in the hospital and beyond

18% vs 12%

of an in-charge nurse's shift spent on documentation versus direct patient care at Singapore General Hospital. On the surgical ward, 21% versus 20%[33]



While the economic case for a hybrid workforce in financial services is clear, the calculus changes for healthcare, as the goal is not profit maximisation alone but the sustainable delivery of high-quality care. The situation is particularly urgent in Singapore, which officially became a super-aged society earlier this year, with over 21% of the citizen population aged 65 and older.[32] The demand for care is continually rising and the shortage of skilled human labour is a key limiting factor.

In a hospital setting, agentic labour is not about displacing already scarce human labour but enabling the existing workforce to meet the demand. A time-motion study at Singapore General Hospital (SGH) found in-charge nurses spent 18% of a shift on documentation against 12% on direct patient care. On the surgical ward it was 21% versus 20%.[33] Those hours are not idle but they are bundled into clinical roles and defended as clinical work. Agents have the capability to absorb certain tasks, but should only be trusted to do so when the work is deliberately re-architected and re-allocated in a way that preserves critical decision making.

Figure 11 The existing care pathway

Today, the circle of care is sparse when the patient goes home

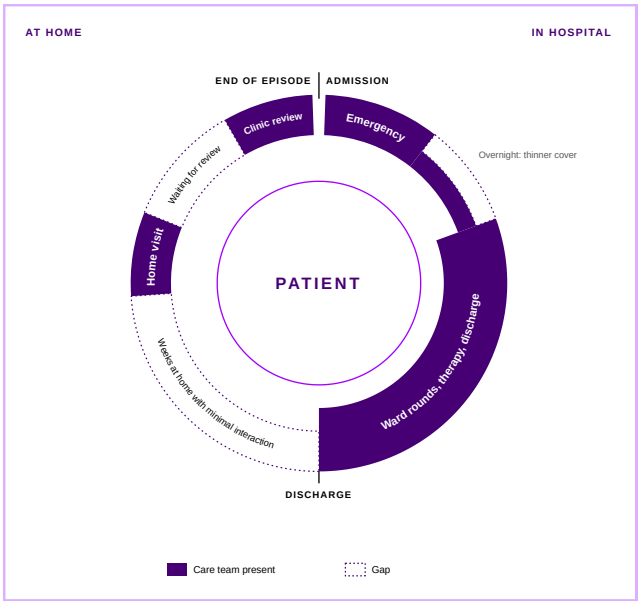
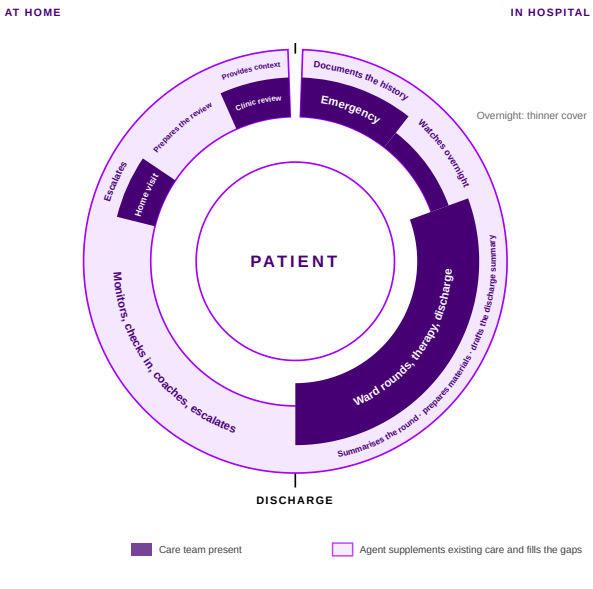


Figure 12 The agent-enhanced care pathway

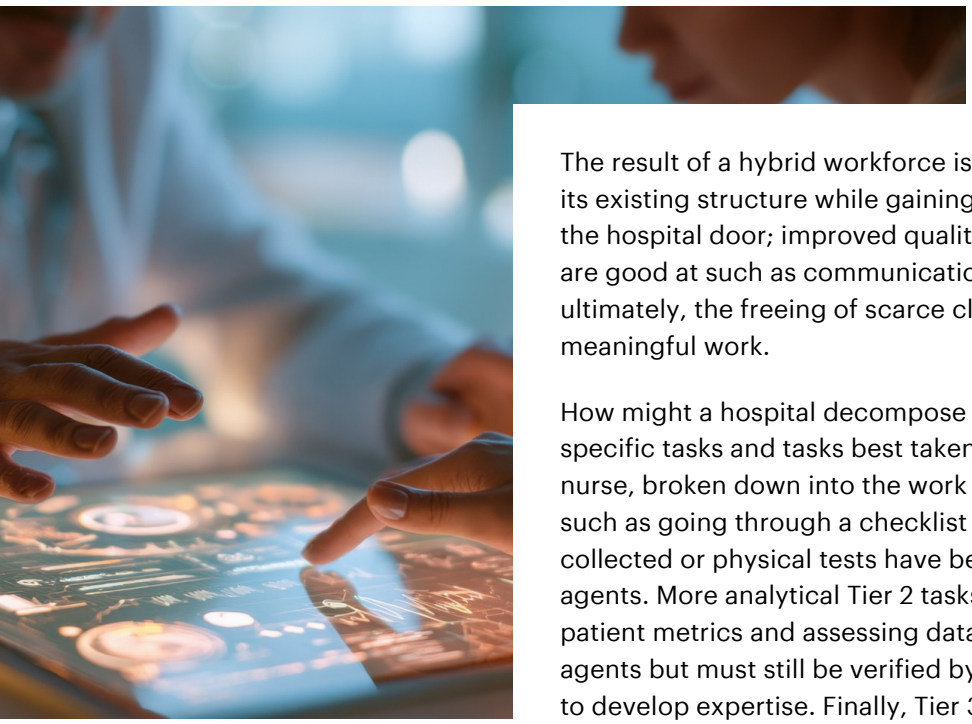
With agents, the circle of care stays continuous



The current care pathway faces limits of time and place; staffing thins on overnight shifts, for example, and clinicians have less contact with, and visibility of, patients once they are discharged to their homes and communities. It is impractical to fill these gaps using human labour alone, but agentic labour is well-placed to supplement existing workflows, available round the clock and across any healthcare setting

.

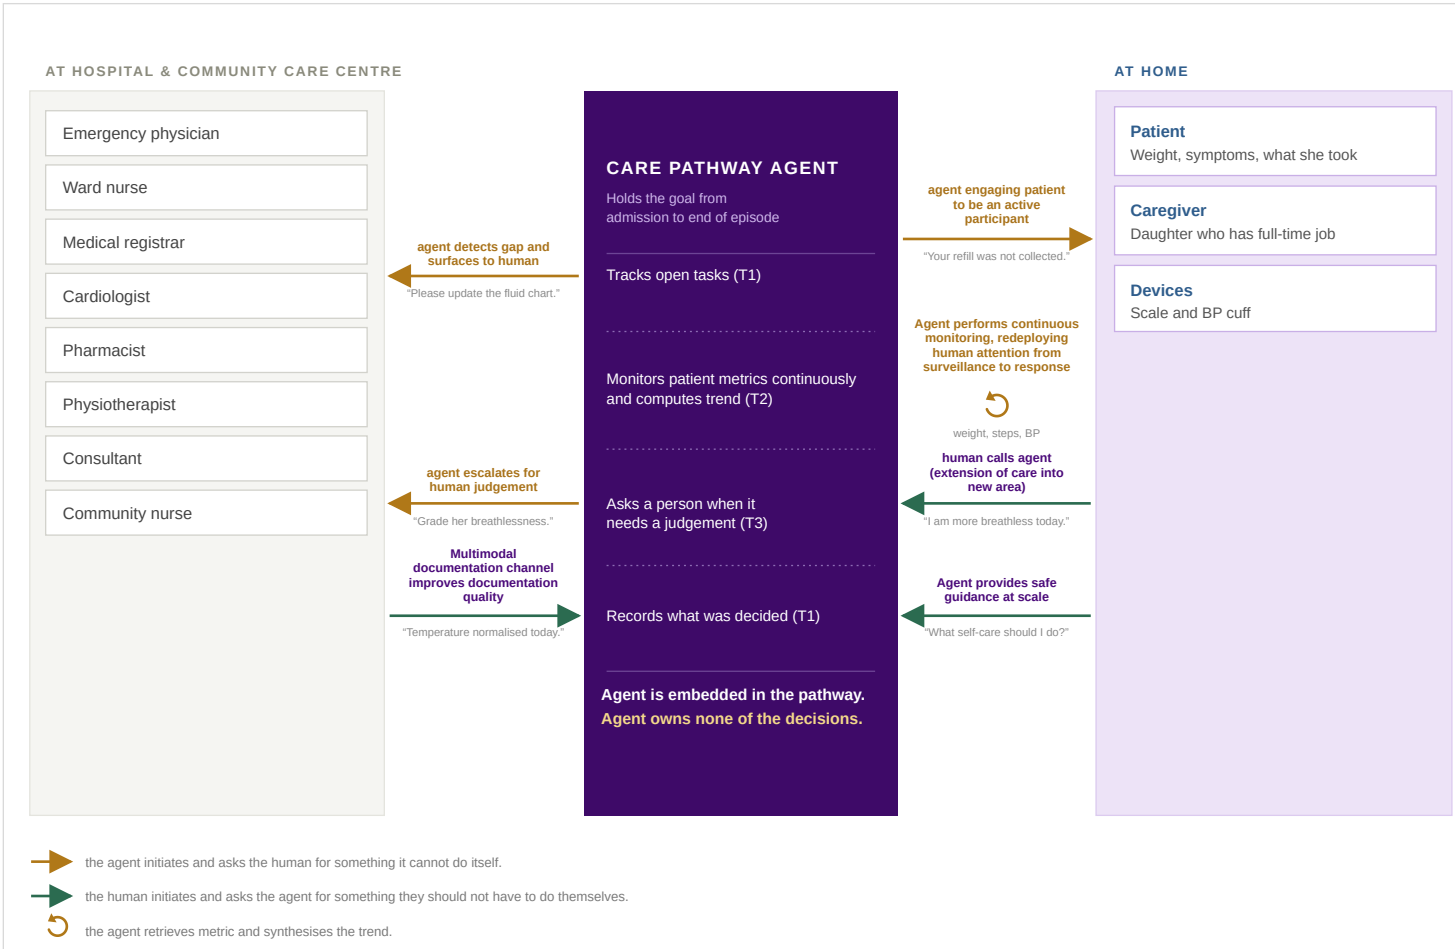
Copyright © 2026 Accenture. All rights reserved.



The result of a hybrid workforce is an enhanced care pathway that retains its existing structure while gaining three things: the ability to extend past the hospital door; improved quality of care, particularly for tasks agents are good at such as communication, synthesis and monitoring; and, ultimately, the freeing of scarce clinical hours for other high-value, more meaningful work.

How might a hospital decompose existing clinical tasks into human-specific tasks and tasks best taken on by agents? Consider the work of a nurse, broken down into the work taxonomy tiers. Tier 1 procedural tasks such as going through a checklist of whether a medicine refill has been collected or physical tests have been conducted can be absorbed by agents. More analytical Tier 2 tasks, such as continuously monitoring patient metrics and assessing data to recognise patterns, can be moved to agents but must still be verified by human labour and routinely practised to develop expertise. Finally, Tier 3 tasks that call for adaptive judgement must remain with human labour; clinical reasoning, weighing an ambiguous patient presentation and interpreting the investigation results.

Figure 13 How agents can be embedded in a care pathway



Copyright © 2026 Accenture. All rights reserved.



Once a job has been decomposed into tasks of different tiers, those tasks can be re-allocated across the different types of workers. In a hybrid workforce, a nurse can play the role of a Co-Architect and Judgement Specialist, using agents to monitor patients across time and care settings as well as interacting directly with the patient, intervening only when there is a significant issue such as an unexplained spike in blood pressure. Similarly, agents can enable doctors to become Augmented Practitioners, absorbing the task of synthesising information via ambient documentation and proposing a list of relevant diagnoses.

In healthcare, the larger economic opportunity is not necessarily reducing labour cost, but expanding the capacity of scarce clinical labour.

Applying agents to nursing documentation in the SGH example above is unlikely to produce a simple reduction in headcount. Nursing is not a single role; a ward team combines patient care assistants, enrolled nurses, registered nurses and advanced practice nurses. Each grade carries a different mix of cognitive and physical work, so documentation relief reaches these grades unevenly. The more likely effects are a change in the skill mix of each role and a change in the proportion of nurses of each grade.

For advanced practice nurses, routine documentation shrinks and the role draws more on clinical judgement, oversight of agent output and coaching of junior staff. For patient care assistants and enrolled nurses, the role remains centred on hands-on care while agents handle their records and prompts. These roles therefore draw less on administrative skill and more on direct care and on recognising when to escalate.

Hospitals may then choose to hold or increase the proportion of bedside nurse grades, since agents do not replace physical care, and to rebalance the senior tier toward oversight and complex cases.



Copyright © 2026 Accenture. All rights reserved.

The value also extends beyond the hospital. Agentic monitoring and coordination can support more care at home and in the community, allowing lower-acuity patients to be managed outside hospital and reducing avoidable deterioration or readmission. But care is more than treatment. Patients and their families also value reassurance, personal attention and human connection. For home-based care to become a meaningful alternative, patients and caregivers must be comfortable with how that care is delivered.

Care is more than treatment. Patients and families also value reassurance, personal attention and human connection.

The economic benefit is therefore twofold: greater capacity to serve an ageing population and better allocation of scarce clinical resources. Hospital capacity can be concentrated on patients who genuinely require higher-acuity care, while others receive an appropriate level of attention closer to home. The value of agentic labour is therefore not simply the nurses removed or hours saved, but the ability to expand the effective capacity of the health system and direct scarce human clinical capability to where it creates the greatest value. The economic value comes when this additional capacity translates into more patients served, better care or lower resource needs, after accounting for the technology and human oversight required.



This is particularly important within Singapore's operating envelope. A super-aged society will continue to increase healthcare demand, while slower population and workforce growth limits the ability to meet that demand simply by adding more healthcare workers. The economic imperative is therefore to increase the capacity of the overall care system without requiring clinical manpower to grow one-for-one with demand or compromising the **quality of care.**

Copyright © 2026 Accenture. All rights reserved.

06

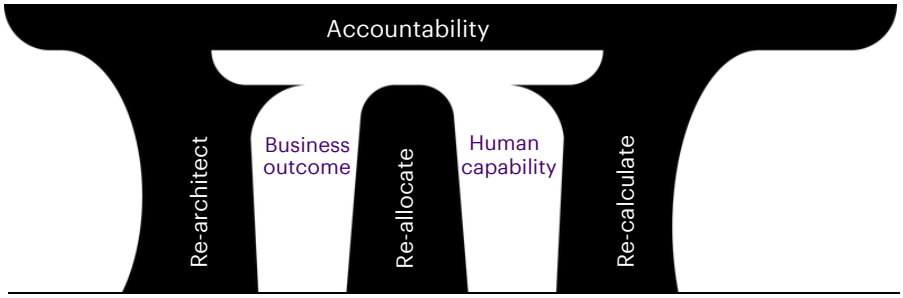
How organisations can move forward, and what Accenture is doing

A structured route from leadership intent to an operating hybrid workforce, built continuously on the human conditions for change: clarity, connection, confidence, capability and commitment.



Copyright © 2026 Accenture. All rights reserved.

Figure 14 The Project Stonehenge method



Stonehenge offers leaders a structured method for addressing an opportunity, albeit one that comes with enormous attendant challenges, unlike any other they have faced before: the once-in-a-generation introduction of a new form of labour. It is a new frame of reference that continually raises new questions and forces critical decisions. The right execution will differ by organisation, shaped as much by its values, culture and ambitions for the future as by the business case and competitive landscape.

What we have observed across years of AI implementation in real-life client environments is that the points of failure are consistent even when the circumstances are not. Change needs to begin with the operating model, where frameworks meet friction, pilots meet production and ambition meets a legacy operating environment. We have also seen what works and have built a body of work that includes points of view, accelerators and deployable solutions that have been tested against reality, not just theory.

For organisations keen to apply these approaches, we offer practical guides and solutions to get started.

6.1 From leadership decisions to a path for action

Project Stonehenge provides a practical playbook for moving from leadership decisions to action. It helps leaders work through the choices and trade-offs involved in designing a hybrid workforce, while giving practitioners a structured methodology for translating those choices into what needs to be designed, tested, built and operated.

The approach moves through five stages: Diagnose, Design, Prove, Scale and Run. It begins by establishing the organisation's ambition and whether it has the maturity to pursue it; develops the target hybrid workforce design; tests that design in practice; builds the capabilities required to scale it; and then operates and continuously recalibrates the model. The intent is not to prescribe a single end state, but to provide a structured route from leadership intent to an operating hybrid workforce.

Copyright © 2026 Accenture. All rights reserved.

Figure 15 The Stonehenge approach to designing and operating a hybrid workforce

01

Diagnose

Where should we focus first?

Identify where agentic AI can materially advance the organisation's priorities, and select the opportunities to take forward.

Scope · Enterprise-wide or within a self-contained business unit

KEY DECISION

Which opportunities should we pursue, and at what scope?

AI maturity gate

Can our current AI maturity support the transformation we intend in our selected ambition and scope?

YES PROCEED

02

Design

What should the future human–AI workforce look like?

Design how humans and AI work together, including the work split, retained human requirements, authority, controls and economics.

Scope · Same scope selected in Diagnose

KEY DECISION

What should AI do, what must remain human, and under what authority and controls?

03

Prove

Does the design work in practice?

Test the target design in a bounded part of the organisation, validate the assumptions and refine it before wider deployment.

Scope · Selected product, workflow or business unit

KEY DECISION

Has the design been proven well enough to scale?

04

Scale

How do we put it into operation at scale?

Industrialise the proven design: put production capabilities and controls in place, transition people and workflows, and expand deployment.

Scope · Enterprise-wide or business unit-wide

KEY DECISION

Is the operating model ready to enter normal operations?

05

Run

How do we keep it effective over time?

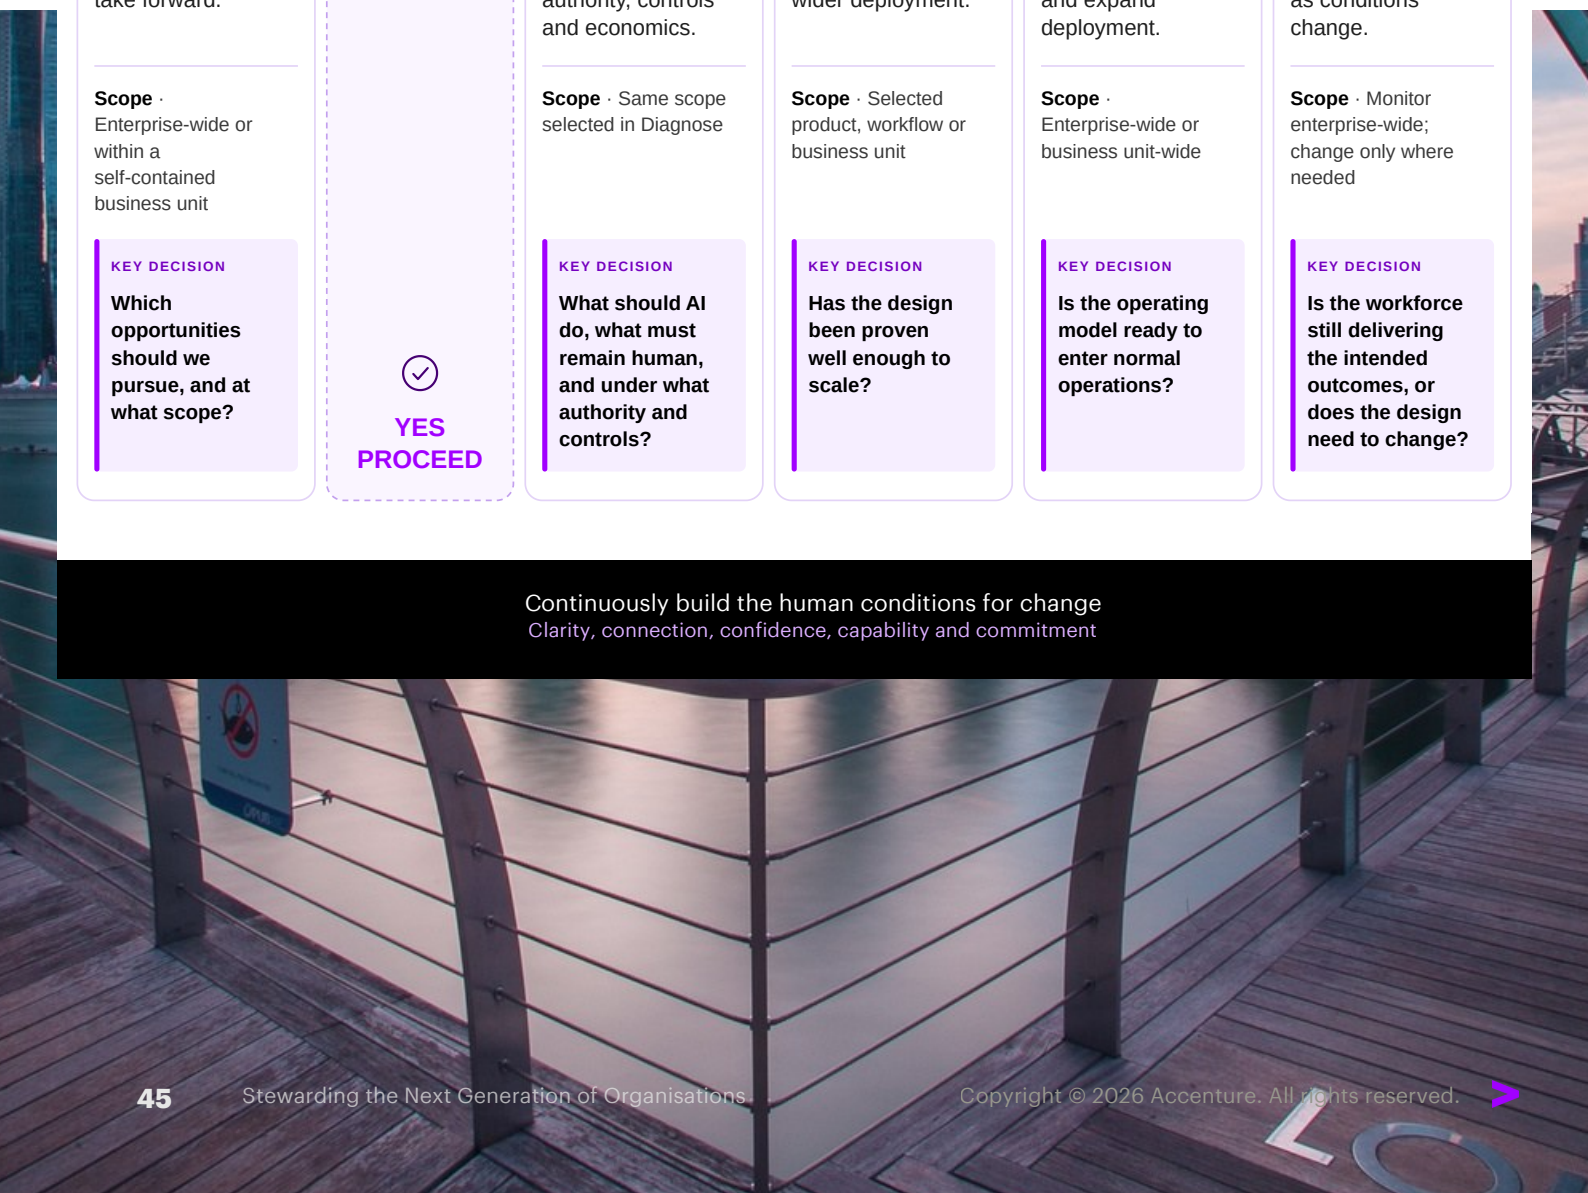
Operate the hybrid workforce, monitor outcomes, capability and controls, and recalibrate the human–AI design as conditions change.

Scope · Monitor enterprise-wide; change only where needed

KEY DECISION

Is the workforce still delivering the intended outcomes, or does the design need to change?

Continuously build the human conditions for change
Clarity, connection, confidence, capability and commitment



Copyright © 2026 Accenture. All rights reserved.

Adoption is a behaviour. Readiness is a state.

6.2 Conditions for change readiness

Whatever operating model is designed ultimately needs to be implemented for human beings. Our prior research shows that five conditions determine whether people choose to lean in and participate in the change or are simply subjected to it: clarity, connection, confidence, capability and commitment.[34] They are conditions that organisations need to fulfil, rather than programmes they need to run. Adoption is a behaviour. Readiness is a state, and an organisation can score well on the first while failing on the second.

Figure 16 The five human conditions needed for change readiness. The outer ring conditions are set by leadership on an organisational level; the inner ring conditions determine how employees experience change individually



Source: Accenture, *The Humanness of Organisations in an Age of AI*, October 28, 2026.

The research reveals a concerning gap between senior executives' confidence in AI and their actual fluency with the technology. Senior workers in our sample were most likely to be overconfident about their AI fluency. If leaders are miscalibrated on what the technology can and cannot do, this may result in underinvestment in applied development, or overinvestment in awareness programmes that workers consistently describe as generic and disconnected from their roles.



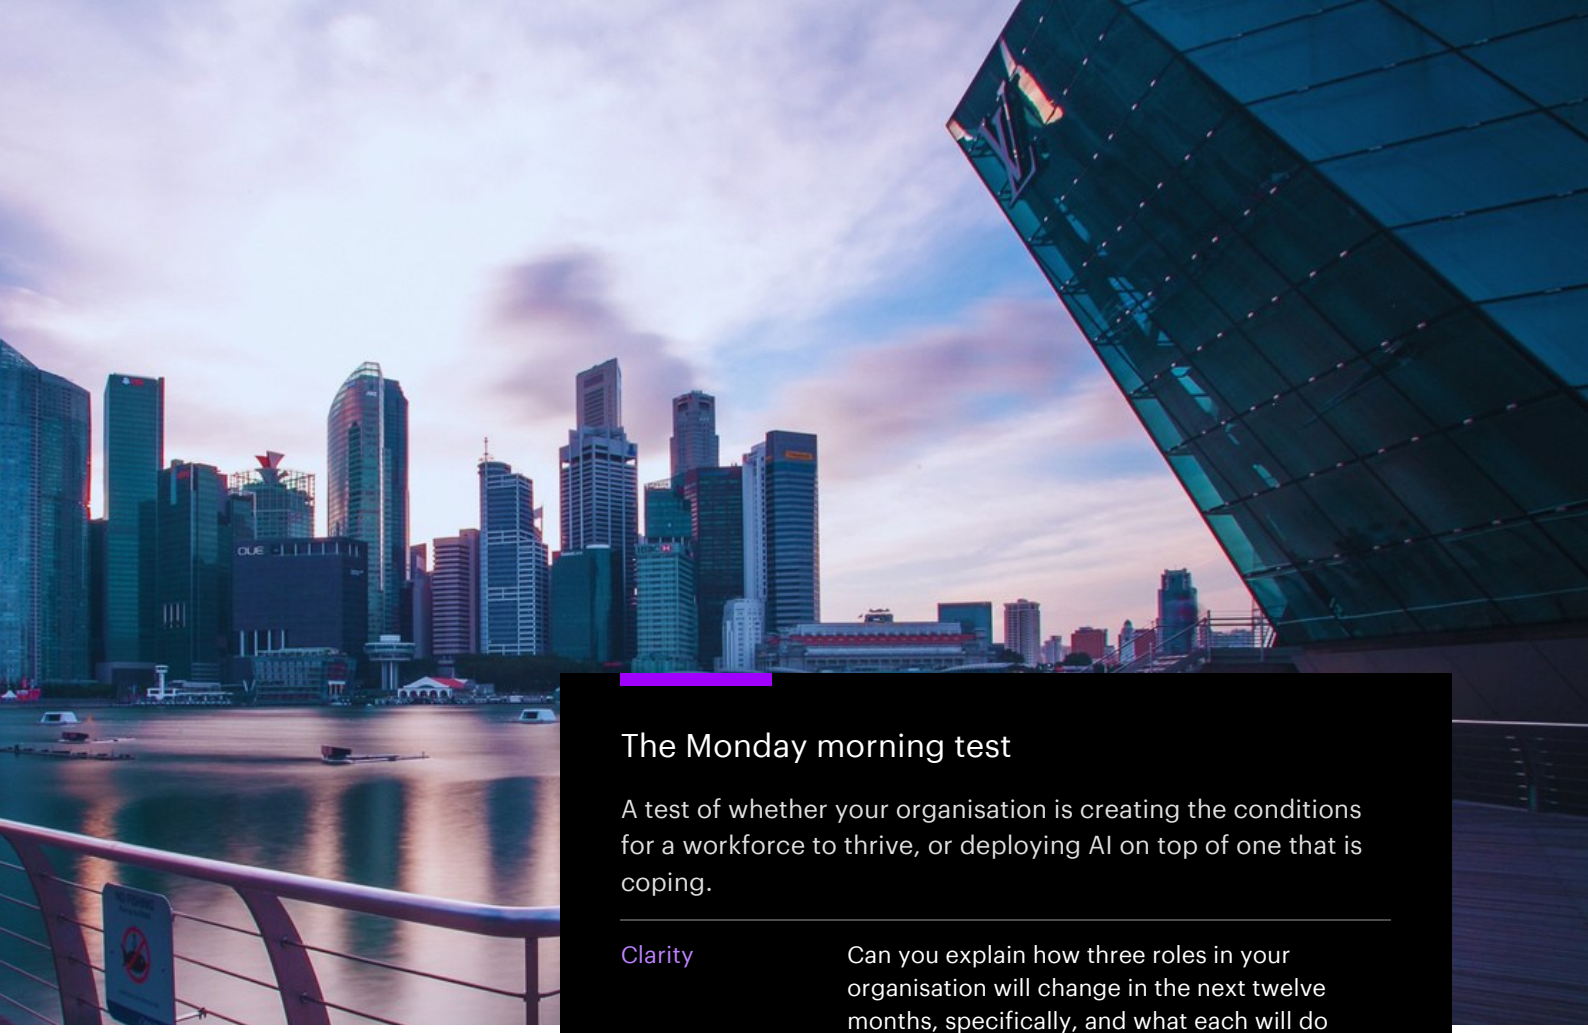
Copyright © 2026 Accenture. All rights reserved.

Figure 17 The five conditions of change readiness

Condition	The question the employee is asking	What is needed for the condition to be met	What our research found
Clarity	Where is our organisation heading with AI?	Employees understand where the organisation is heading with AI, what will change, and what will stay human-led.	From Accenture's Pulse of Change survey, 53% of Singapore employees, one of the highest proportions in any APAC market surveyed, cite clear guidance on how their role is evolving as one of their three most helpful forms of support.
Connection	What does this mean for me?	Workers can see how the AI transformation supports what they personally value in their work and sense of purpose, and what AI means for their role, tasks and career.	Only 18% of financial services and 24% of healthcare workers describe AI as extremely relevant to their current role, despite 62 to 78% of financial services workers using it at least three days a week. Frequent use and perceived relevance have separated entirely.
Confidence	Can I engage with AI safely and credibly?	Workers have the tools, governance and support to engage with AI safely, while knowing their judgement and expertise still matter.	Workers in both sectors are broadly optimistic and willing to engage with AI. However, while training is provided, not all of it translates into meaningful capability. Three in ten who received employer training found it unhelpful. Workers express frustration about lack of applied, role-specific training, and unavailability of useful, sanctioned tools.
Capability	What skills do I need in my actual role?	Before this condition can be addressed, work and roles will need to be redesigned. After work has been redesigned, workers develop the skills required for AI-enabled work through real tasks and workflows, not just awareness training.	29% of workers are overconfident about their AI capability, rising to 46% of senior financial services leaders and 55% of senior healthcare leaders.
Commitment	Can I trust that human value still has a place?	Workers can see evidence that the organisation is using AI to strengthen human work, not simply extract more from it.	This condition sits beneath all the others. In healthcare, AIHGle 2.0 explicitly states that AI should augment and empower healthcare professionals. Without similar reassurance in the financial sector, workers are more likely to view AI through a productivity lens: 62% of financial sector workers say AI makes them question how valuable parts of their work are.

Sources: Accenture, *The Humanness of Organisations in an Age of AI*, October 28, 2026. Accenture Employee Pulse of Change survey, 2026, N = 100 Singapore respondents. Ministry of Health and Health Sciences Authority, *Artificial Intelligence in Healthcare Guidelines (AIHGle) 2.0*, March 2026.

Copyright © 2026 Accenture. All rights reserved.



The Monday morning test

A test of whether your organisation is creating the conditions for a workforce to thrive, or deploying AI on top of one that is coping.

Clarity	Can you explain how three roles in your organisation will change in the next twelve months, specifically, and what each will do differently once AI is embedded in those workflows?
Connection	Can your employees explain, in their own words, why the transformation matters to them personally and what it means for their own future at work?
Confidence	Does a worker in your organisation know which tools they may use, what the governance around them is, where they can experiment safely, and who to ask?
Capability	Is your development investment tied to the work people actually do, and do you know how calibrated your own leadership is about what AI can and cannot do?
Commitment	Is there an explicit, written statement of what remains distinctively human in your organisation, and can an employee point to it?

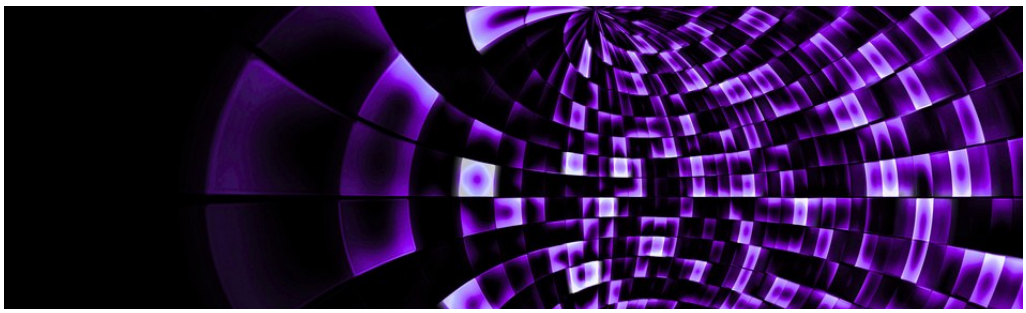
Copyright © 2026 Accenture. All rights reserved.

6.3 Talent Navigator

Talent Navigator is Accenture's workforce planning and skills platform. For Stonehenge we have extended it with a responsible AI constraint layer and an accountability layer, so that an allocation decision can be modelled against the design principles, the tier boundaries and the accountability requirements set out in this paper rather than against cost alone.

We have used it to model Singapore organisations and to test what these choices actually mean: what changes in a workforce when a leadership team decides one way rather than another. The purpose is not to produce an answer. It is to make the trade-off visible before it is taken, which is the thing that is currently missing from most AI business cases.

A framework without deployment mechanisms remains a position paper, not an operating model.



6.4 Deployment mechanisms

Three deployment mechanisms, developed in collaboration with the TRACE Lab at Cambridge's Centre for Human-Inspired Artificial Intelligence (CHIA), turn the design into something an organisation can adapt to operate, measure and defend.[35]

Copyright © 2026 Accenture. All rights reserved.

Figure 18 The three deployment mechanisms

Mechanism	What it does	Why it matters
The Human API	Establishes a standard for an AI agent to ask a person for information, judgement or approval that it cannot get from software or databases. Each request specifies what is needed, who is being asked, its expected cost and reliability, and whether the person accepts or declines. All requests and responses are captured and auditable, showing how much an agent depends on people and who made each decision.	It replaces ad hoc escalation with a record. Every consequential input has an owner, a cost and a timestamp. Workers can decline or modify a request, and that action stays on the record. Oversight is enforced by the interface rather than assumed.
Skill inference	A method for estimating what humans and agents are good at based on their task performance. Each task is linked to a set of required skills. Task results provide evidence about those skills, while relationships between skills allow one task to provide information about several skills. These skill profiles are used to assign future tasks.	It distinguishes low proficiency from insufficient evidence, which are different states requiring different action. Every dimension is a named skill, so an allocation decision can be traced to the estimate that drove it and challenged.
Dynamic controls	A method for maintaining a minimum share of work for people. Without this control, more work may gradually move to agents, reducing opportunities for people to practise and develop their skills. The system sets a minimum human-work share, called a floor. If the human share falls below the floor, the system increases the cost of assigning work to agents, thus shifting more tasks to people. This cost shows how much output the organisation gives up to maintain the chosen human-work level.	A quota reports only whether a target was met. A priced human share of the work reports what the commitment cost in output, which allows a board to set the level deliberately, defend it, and revise it on evidence.

Source: Kostadin Devedzhiev et al., *Designing for Hybrid Human-AI Workforces: a human-centred architecture for the division of labour between humans and agents*, Centre for Human-Inspired Artificial Intelligence, University of Cambridge, November 9, 2026.

These deployment mechanisms work together in a closed loop. The Human API enables calls to people and generates outcome data. That data updates capability profiles through the skill inference mechanism, which in turn informs routing. Finally, dynamic controls give organisations a way to determine the share of work that stays human and report what holding it costs.

A note on evidence informing these mechanisms

The three deployment mechanisms were evaluated in simulation, not in production. The population data in the experiments was drawn from published occupational skill ratings. All inputs were synthetic rather than empirically derived: the individual capabilities, the agents and their advantages, the arriving tasks, the graded scores and whether a task was cleared. No real individual's learning or skill atrophy is modelled, and the observation channel has not yet been integrated into a live process. For more detail, refer to the University of Cambridge CHIA TRACE Lab paper, *Designing for Hybrid Human-AI Workforces*, by Kostadin Devedzhiev et al.

Copyright © 2026 Accenture. All rights reserved.

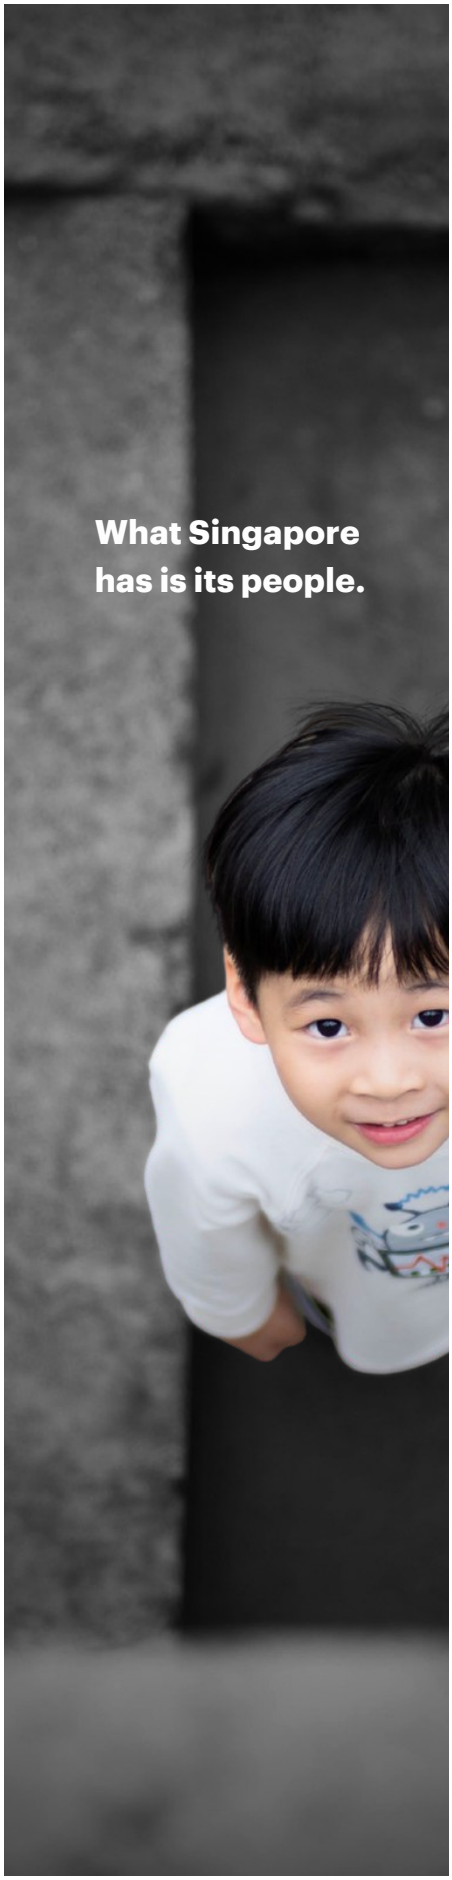
07

Stewardship for Singapore

Human labour has always created value beyond immediate production through accountability, capability, continuity, judgement and relationships.



Copyright © 2026 Accenture. All rights reserved.



Human labour has always created value beyond immediate production through accountability, capability, continuity, judgement and relationships. The first wave of AI transformation has begun to expose what is lost when organisations optimise for production efficiency without deliberately preserving these wider sources of value.

Every organisation is already making a set of trade-offs it has never had to weigh before: what the organisation is prepared to invest to sustain the growth of human capability versus what it is prepared to risk by not investing. These are not abstract hypotheticals; they affect each individual in the workforce today.

Other economies may have alternatives, but Singapore is constrained by land and natural resources. What the country has is its people. It has always grown because of wise decisions that enable Singapore's prosperity and what its people could become next. A country whose strength is in its human capability cannot afford the unintended consequence of organisations that quietly stop building it. This is why the national commitment to no jobless growth is of great importance, and why an organisation that hollows out its own capability is not simply making a commercial choice but a decision that impacts the future of the nation as a whole.

Just like Singapore, Accenture regards stewardship of future generations as core to our belief. Like Singapore, our business is built primarily on people. We developed the Stonehenge approach as much for ourselves, to guide how we may make these decisions and trade offs for our own organisation, as we would want to help our clients think through theirs. We do not claim to have solved the problem completely, but we believe this blueprint will help us engineer a more deliberate and intentional hybrid workforce design that will enable the next generation of organisations to thrive.

Future prosperity, for individual organisations and Singapore as a whole, depends on decisions made today. Between productivity and growth. Between temporary gain and lasting value. Between material wealth and immeasurable worth. Those choices are being made across Singapore right now, in budget papers and deployment decisions, whether or not anyone names them.

Stewardship is deciding, then acting with purpose.

Copyright © 2026 Accenture. All rights reserved.

Appendix: Banking workforce economics illustration

The tables below set out the assumptions and calculations behind the 1,200-person institutional banking example in Section 5.1 and Figure 10.

A Released production capacity and current-work value

Component	Value or build-up	Sources and assumptions
Illustrative workforce	1,200 FTE	DBS's announced rollout to 1,500 bankers, scaled using the average FY2025 institutional and wholesale banking income of Singapore's three local banks as a proxy for workforce size.
Workflow exposure	30%	Assumed share of time spent on credit assessment and related internal workflows.
AI effort reduction	30%	OCBC reports up to 30% preparation-time reduction;[36] DBS targets at least 30% in credit assessment.[37]
Capacity released	108 FTE	1,200 × 30% × 30% = 108 FTE-equivalents.
Fully loaded labour cost	S$150,000/FTE	Randstad Singapore's 2025 salary guide reports a median S$9,000 monthly base salary for corporate banking relationship managers.[38] Assumes two months' bonus, statutory employer CPF and additional employment-related costs.
Gross annual capacity value	S$16.2m	108 × S$150,000 = S$16.2m in released production capacity.

B Retained formation requirement and current-work cost headroom

Component	Value or build-up	Sources and assumptions
Judgement-bearing share of workforce	17%	Singapore managers and administrators share of employed resident workforce, used as a proxy.[39]
Judgement-bearing population	204 FTE	1,200 × 17%.
Annual replenishment	~11%	MOM PMET resignation rate of 0.9% per month in Q4 2025; 0.9% × 12 = 10.8%, rounded to 11%.[40]
Formation period	3 years	Market reference of 3–5 years for AVP credit analyst roles: DBS Midcap Institutional Banking specifies at least five years of experience;[41] MUFG Corporate Banking specifies at least three years.[42] Three years is used as the accelerated development period.
Human production retained	67 FTE	204 × 11% × 3 = 67.32 FTE (shown as 67); matching capacity retained for development.
Retained formation cost	S$10.1m	67.32 × S$150,000 = S$10.098m (~S$10.1m).
Current-work headroom	S$6.1m	S$16.2m − S$10.098m = S$6.102m (~S$6.1m).

Copyright © 2026 Accenture. All rights reserved.

C Redeployment and total technology-cost headroom

Component	Value or build-up	Sources and assumptions
Capacity for redeployment	41 FTE	108 − 67.32 = 40.68 FTE-equivalents (~41).
Mature annual contribution	S$300,000/FTE	Net operating contribution per productive FTE, after staff and other operating expenses; see Table D.
Three-year ramp	67%	Assumed average share of mature operating contribution during the three-year ramp-up.
Additional annual value	S$8.2m	40.68 × S$300,000 × 67% = S$8.177m (~S$8.2m) annual average during ramp-up.
Total technology-cost headroom	S$14.3m	S$6.102m + S$8.177m = S$14.279m (~S$14.3m).
Change over current work	~2.3x	S$14.279m / S$6.102m = ~2.34x (~2.3x).

D 2025 bank-wide operating contribution benchmark

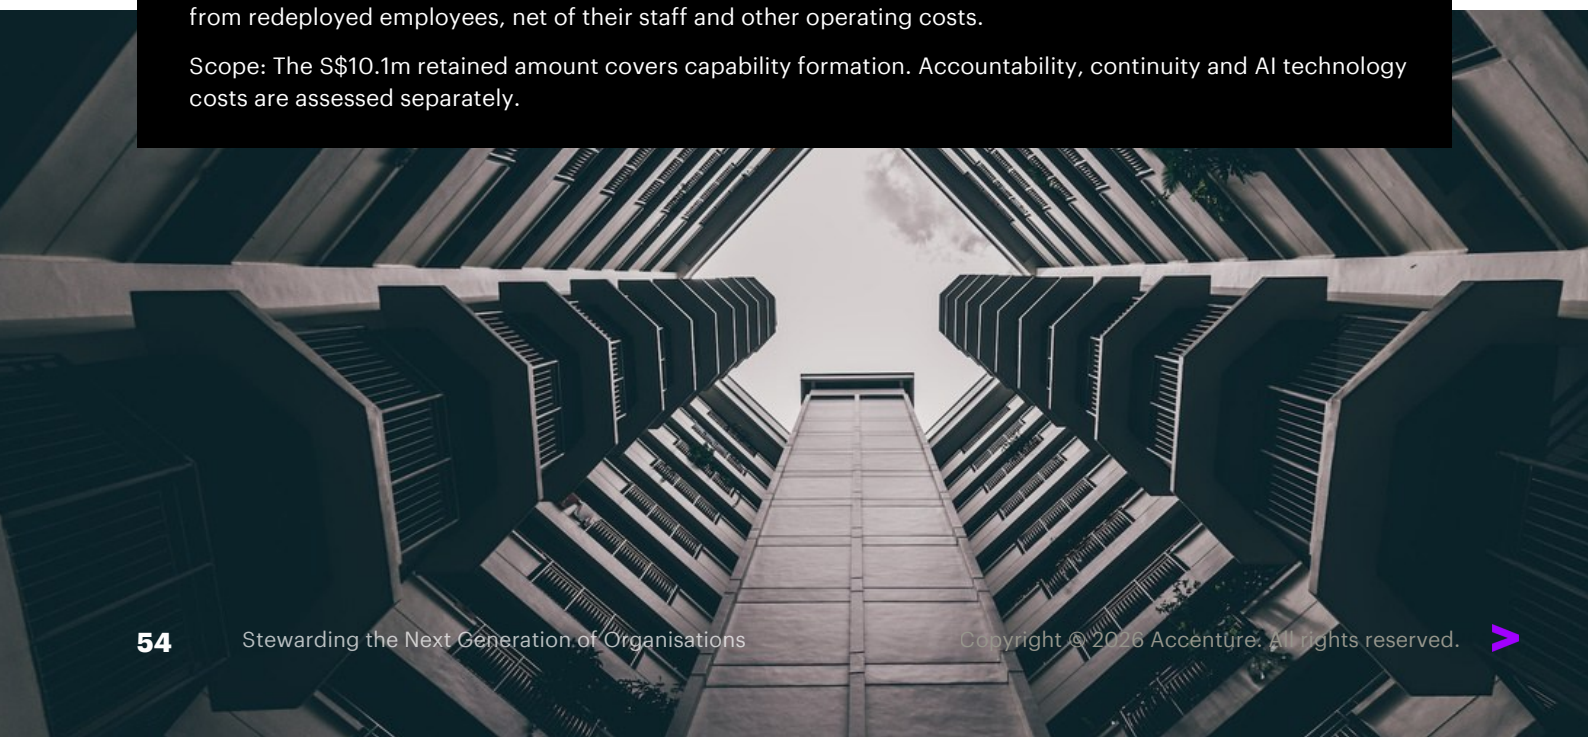
Component	Value or build-up	Sources and assumptions
DBS Group	S$13.70bn / 39,983	2025 profit before allowances / group employee headcount.[43,44]
OCBC Group	S$8.70bn / 33,323	2025 operating profit before allowances and amortisation / group workforce, including Great Eastern.[45,46]
UOB Group	S$7.70bn / 31,095	2025 operating profit / total group headcount.[47,48]
Combined indicative mean	~S$300,000/FTE	S$30.1bn / 104,401 employees = ~S$288k, rounded to S$300k.

The S$300,000 contribution assumption is benchmarked against 2025 operating profit before credit allowances per employee across DBS, OCBC and UOB. These figures are after staff and other operating expenses.

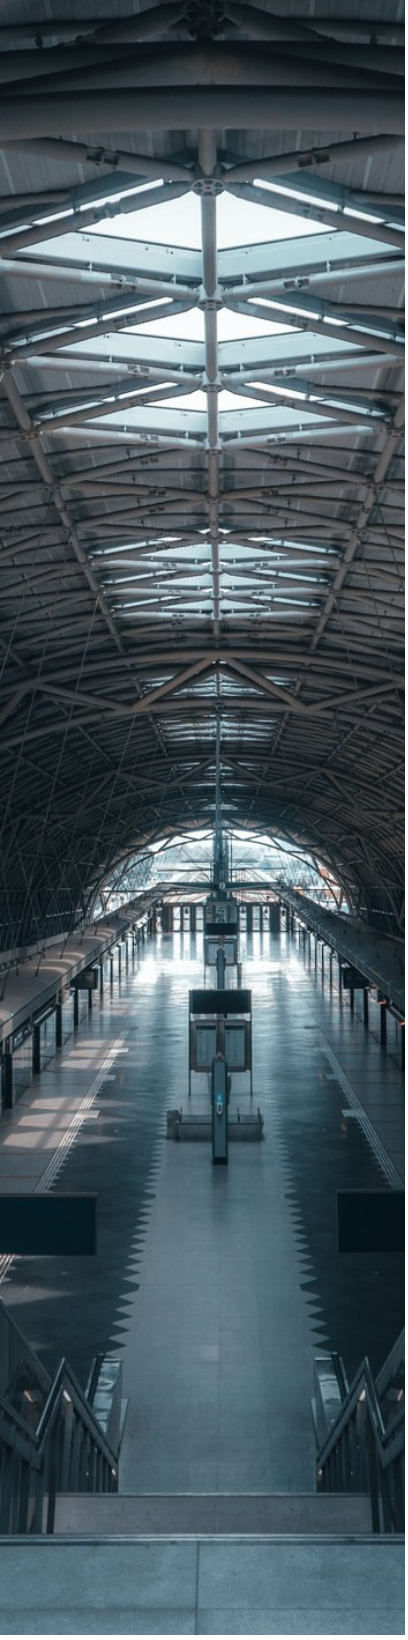
Accounting basis and scope

Accounting basis: S$6.1m is the cost of labour released from existing production. S$8.2m is the contribution from redeployed employees, net of their staff and other operating costs.

Scope: The S$10.1m retained amount covers capability formation. Accountability, continuity and AI technology costs are assessed separately.



Copyright © 2026 Accenture. All rights reserved.



Public data on Singapore's labour market, demography and regulatory frameworks is drawn from the Ministry of Manpower, the Monetary Authority of Singapore, the Ministry of Health and the Infocomm Media Development Authority, as referenced. Generative AI was used in the research production process. All outputs were reviewed and validated by research experts against traditional research methods, in accordance with Accenture's Responsible AI standards.

About the research

This paper draws on Project Stonehenge, an applied Accenture Singapore programme defining a hybrid workforce operating model. Four evidence streams sit behind it.

Worker survey. A survey of 1,008 Singapore workers across financial services and healthcare, covering AI use, perceived relevance, capability, satisfaction and readiness, with free-text responses used as the source of verbatim quotation in this paper.

Technology leader survey. A survey of 70 Singapore technology leaders conducted with the IT Management Association, covering implementation maturity, role redesign and investment intent.

C-suite interviews and expert input. A programme of interviews with Singapore Chief Executives, Chief Human Resources Officers and Chief Information Officers with sector subject-matter input in financial services and healthcare, tested against a trusted review panel of Singapore and global practitioners.

Collaboration with the University of Cambridge. Work with the Trustworthy Artificial Intelligence (TRACE) Lab at the Centre for Human-Inspired Artificial Intelligence (CHIA) on trustworthy AI deployment mechanisms. The three mechanisms in Section 6.4 originate in that collaboration and have been evaluated in simulation only. Their limitations are stated in Section 6.4 and set out in full in the source paper.

The primary evidence streams were informed by an extensive review of the current literature on the future of work, automation and capability formation, organised across five workstreams.

1. Internal corpus synthesis. A review of existing Accenture work on workforce transformation, AI deployment and operating model design.
2. C-suite needs analysis. A mapping of the questions CEOs, CIOs and CHROs are actively asking, drawn from the interview programme, to anchor the literature review in decisions practitioners are facing.
3. External literature review. A review of current academic and practitioner work on the future of work, automation and human capability formation. All sources are referenced and traceable.
4. Singapore context. An analysis of national policy, sector regulation and labour market conditions, ensuring findings are grounded in the specific operating environment under study.
5. Meaning of work. A review of the literature on work as a source of purpose, identity and wellbeing, providing a conceptual grounding for the human dimensions of workforce transformation.

Copyright © 2026 Accenture. All rights reserved.

AUTHORS



Mark Tham
Country Managing Director,
Accenture Singapore



Grace Yip
Managing Director,
Accenture Singapore



Tarun Pardal
Managing Director,
Accenture Singapore



Andrew Eddy
Principal Director,
Accenture Singapore



Chai Chong Keng
Senior Manager,
Accenture Singapore



Yeap Wei Yang
Senior Manager,
Accenture Singapore

RESEARCH & EDITORIAL

Dr Rebecca Tan, Joel Lee, Elly Chiu, Joanna Dorai, Dr Angie Silva

ACKNOWLEDGEMENTS

Thank you to the executives and experts who shaped this hybrid workforce operating model:

Sivaram Krishna Balasubramanian
Managing Director, CEO Advisory Lead, Accenture Singapore

Jimmy Ng
Former Group Chief Information Officer and Group Head, Technology and Operations, DBS

Dr Eric Sandosham
Founder and Partner, Red & White Consulting Partners; Adjunct Associate Professor, Nanyang Technological University

Lee Yan Hong
Managing Director and Head of Group Human Resources, DBS

Dr Kaveh Safavi
Health Senior Advisor, Accenture

Dr Gerald A. (Gerry) Kraines
President and CEO Kraines Consulting

Dr Umang Bhatt
Lab Director, Trustworthy AI Lab, Centre for Human-Inspired Artificial Intelligence (CHIA), University of Cambridge

Kostadin Devedzhiev
Trustworthy AI Lab, Centre for Human-Inspired Artificial Intelligence (CHIA), University of Cambridge

Copyright © 2026 Accenture. All rights reserved.

References

1 Accenture survey of ITMA professional network, January 2026. N = 70 technology leaders.

2 Kostadin Devedzhiev et al., Designing for Hybrid Human-AI Workforces: a human-centred architecture for the division of labour between humans and agents, Centre for Human-Inspired Artificial Intelligence, University of Cambridge, November 9, 2026.

3 Accenture, The Humanness of Organisations in an Age of AI, October 28, 2026.

4 Smart Nation Singapore, Update to NAIS. AI for the Public Good for Singapore and the World, May 20, 2026.

5 Ministry of Manpower, Singapore, Labour Force in Singapore 2025, January 29, 2026.

6 National Population and Talent Division, Strategy Group, Prime Minister's Office, Population in Brief 2026, September 25, 2026.

7 Ibid.

8 Deputy Prime Minister Gan Kim Yong, Speech at the Committee of Supply Debate 2026, Singapore, February 26, 2026.

9 Ministry of Manpower, Singapore, Labour Force in Singapore 2025, January 29, 2026.

10 Singapore Institute of Directors and Gallup, Singapore Workplace Report 2026, June 22, 2026.

11 Accenture, The Humanness of Organisations in an Age of AI, October 28, 2026.

12 Ministry of Manpower, Singapore, Labour Force in Singapore 2025, January 29, 2026.

13 Accenture Employee Pulse of Change survey 2026. N = 100 Singapore respondents.

14 Public Communications Division of the Ministry of Digital Development and Information, How Singaporean Workers Are Supported Through the AI Transition, May 8, 2026.

15 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management for Financial Institutions, October 7, 2026.

16 Monetary Authority of Singapore, Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector, November 2018.

17 Ministry of Health and Health Sciences Authority, Artificial Intelligence in Healthcare Guidelines (AIHGle) 2.0, March 2026.

18 Donald L. Kirkpatrick, Techniques for Evaluating Training Programs, Journal of the American Society of Training Directors, November 1959.

19 Krzysztof Budzyń et al., Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study, The Lancet Gastroenterology & Hepatology, August 2025.

20 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management for Financial Institutions, October 7, 2026.

21 Hao-Ping Lee et al., The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers, CHI Conference on Human Factors in Computing Systems, April 2025.

22 Lisanne Bainbridge, Ironies of automation, Automatica, November 1983.

23 Barry Strauch, Ironies of Automation: Still Unresolved After All These Years, IEEE Transactions on Human-Machine Systems, October 5, 2018.

24 Accenture, Future of Work, October 21, 2026.

25 Minister for Manpower Dr Tan See Leng, Speech at Ministry of Manpower Committee of Supply 2026, Singapore, March 3, 2026.

26 Accenture, Future of Work, October 21, 2026.

27 Accenture Tokenomics Launched to Help Enterprises Manage AI Token Spend, Accenture Newsroom, July 29, 2026.

28 Why Tokens Aren't the AI Bill: Introducing Total Cost of AI (TCA), Tokenomics Foundation, September 7, 2026.

29 Accenture, Banking in the age of generative AI, February 28, 2024.

30 Ministry of Manpower, Singapore, Labour Market Report Fourth Quarter 2025, March 20, 2026.

31 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management for Financial Institutions, October 7, 2026.

32 National Population and Talent Division, Strategy Group, Prime Minister's Office, Population in Brief 2026, September 25, 2026.

33 Mei Ling Lim and Shin Yuh Ang, A time–motion observation study to measure and analyse clinical nursing workload in an acute care hospital in Singapore, Proceedings of Singapore Healthcare, Vol. 28, March 2019.

34 Accenture, The Humanness of Organisations in an Age of AI, October 28, 2026.

35 Kostadin Devedzhiev et al., Designing for Hybrid Human-AI Workforces: a human-centred architecture for the division of labour between humans and agents, Centre for Human-Inspired Artificial Intelligence, University of Cambridge, November 9, 2026.

36 Oversea-Chinese Banking Corporation Limited, Annual Report 2025, March 24, 2026.

37 DBS scales agentic AI to transform way of working for corporate bankers, freeing up time for more strategic client engagements, DBS Bank, August 19, 2026.

38 Randstad Singapore, Singapore's 2025 job market outlook and salary guide, December 18, 2024.

39 Ministry of Manpower, Singapore, Labour Force in Singapore 2025 (Table D), January 29, 2026.

40 Ministry of Manpower, Singapore, Labour Market Report Fourth Quarter 2025, March 20, 2026.

41 DBS Bank, AVP/VP – Specialist, Credit Analyst, Midcap, Institutional Banking Group, LinkedIn Jobs, 2026.

42 Mitsubishi UFJ Financial Group, Assistant Vice President, Credit Analyst, MUFG Careers, 2026.

43 DBS Group Holdings Ltd., Annual Report 2025, March 9, 2026.

44 DBS Group Holdings Ltd., Sustainability Report 2025, March 9, 2026.

45 Oversea-Chinese Banking Corporation Limited, Annual Report 2025, March 24, 2026.

46 Oversea-Chinese Banking Corporation Limited, Sustainability Report 2025, March 24, 2026.

47 United Overseas Bank Limited, Annual Report 2025, March 19, 2026.

48 United Overseas Bank Limited, Sustainability Report 2025, March 19, 2026.

Copyright © 2026 Accenture. All rights reserved.

accenture



Copyright © 2026 Accenture. All rights reserved.