TATA CONSULTANCY SERVICES

TATA

TCS Enterprise AI Study

Intelligence that is real



TCS Thought Leadership Institute

Contents

01 Foreword 03

02 Executive summary 04

03 From productivity to enterprise process transformation 05

04 Building for context, control, and adaptability 12

05 Aligning leadership, skills, and workforce systems for the AI era 23

06 Conclusion 27

07 Next Steps: Priorities for scaling enterprise AI 28

08 Study methodology 29

09 Acknowledgements 31



Amit Kapur
Chief AI and Services Transformation Officer
Tata Consultancy Services

Foreword

A shift in motion: Charting the enterprise AI advantage

We're seeing the evolution of enterprise AI firsthand, both through our client engagements and across several of our large-scale research studies. Once defined by anticipation around its potential, AI has rapidly moved into everyday workflows, decisions, products, and operations. During this time we've seen conversations shift from potential to execution, and from experimentation to value.

Our latest research captures that shift in motion. As organisations have gained more practical experience with AI, respondents see some of its clearest effects in decision-making, changing roles, and responsibilities. At the same time, security, risk, and governance are playing a greater role in shaping AI deployment and investment decisions, while organisations are placing an increased emphasis on maintaining flexibility across the AI stack, including models and platforms.

Truly autonomous action remains the exception rather than the norm. Most organisations are making progress in specific functions, use cases, or business units, yet relatively few have reached a level of enterprise-wide readiness. Business knowledge is often scattered across systems, and many AI initiatives remain tied to individual platforms, environments, or tools.

This is why we believe enterprise AI has reached a decision point. The question now is where and how intelligence can make the greatest difference. That is what we mean by intelligence that is real. Leaders now have to decide where AI should advise, where it should act, and how human judgement will remain part of the system. We also need to protect critical data, intellectual property, and business context while preserving the freedom to choose among a fast-changing ecosystem of models, platforms, and agents.

The opportunity ahead is significant. Organisations that make deliberate choices to lead with connected organisational intelligence will be best positioned to turn today's gains into trusted, resilient, and lasting enterprise advantage.

We hope this report offers a useful perspective as you consider the choices ahead and chart the next stage of your enterprise AI journey.

Executive summary

These study findings are based on a global survey of 1,333 senior business and technology leaders across 25 countries and 12 industries. Survey questions focused on how enterprise AI priorities are evolving across data and AI foundations, scale and deployment, people and work transformation, governance and trust, and value-chain and industry transformation.

Together, the results point to three interconnected areas where leaders must focus their attention: **redesigning processes around AI-enabled decisions, building adaptable and controlled technology foundations, and aligning the workforce with changing responsibilities.** Progress in one dimension cannot compensate indefinitely for gaps in another.

Among the key findings:



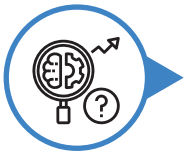
AI is reshaping how decisions get made

Decision-making processes rank as the area experiencing the greatest AI-related change, signaling a shift beyond productivity gains toward AI's growing role in how organisations evaluate options, allocate resources, and act on opportunities.



Security will determine the next phase of enterprise AI

Security, risk, and governance considerations now rank as the leading factor shaping how and where AI systems are developed and deployed, ahead of cost, speed, and ease of deployment.



Practical experience is reshaping expectations for AI's workforce impact

Compared with 2024, respondents are less eager to make predictions on AI's workforce outcomes. The share saying it is too early to know how AI will affect job roles increased by 15 percentage points, suggesting leaders are still assessing how roles, responsibilities, and organisational structures will evolve.



Model flexibility and sovereignty are emerging priorities

Organisations increasingly want the freedom to adapt as AI ecosystems evolve. Nine in ten respondents say model independence is important. More than a third (34%) report that AI sovereignty requirements present a high or very high risk to their organisation, highlighting growing concerns about control, resilience, and long-term flexibility.



AI economics are coming into focus

As deployments mature, organisations are paying closer attention to the costs of intelligence. Thirty-nine percent assess whether AI initiatives generate sufficient value relative to token and compute costs, while 23% actively optimise consumption or continuously improve the value generated per token and unit of compute.

Note: Percentages are rounded to the nearest whole number. As a result, sums of displayed values may differ by one percentage point from summary figures calculated using unrounded data, and displayed percentages may not total to 100%

From productivity to enterprise process transformation

Ever since AI burst into the mainstream, there's been no shortage of talk about its value. While the discussions take different forms, they usually boil down to two questions: is AI creating value and where is that value showing up?

The vast majority of our survey respondents answer in the affirmative and point to productivity gains. Across industries and geographies, the vast majority of organisations identify productivity improvement as the leading source of value realised from AI investments, reflecting the technology's ability to accelerate work, improve efficiency, and augment employee capabilities.

Equally notable is where value has been slower to emerge. Revenue growth ranks lower, putting the onus on the next phase of AI adoption to translate employee productivity gains into top-line enterprise impact.

Employee productivity vs enterprise productivity



84%

of executives surveyed rank productivity improvement as the leading source of value realised from AI investments



But only **10%**

of executives surveyed rank revenue growth as the leading source of value realised from AI investments

> Across industries and geographies, the vast majority of organisations identify productivity improvement as the leading source of value realized from AI investments, reflecting the technology's ability to accelerate work, improve efficiency, and augment employee capabilities

But what happens when AI moves beyond productivity? As AI becomes more deeply embedded in the enterprise, it is starting to influence not just how work is performed, but how decisions are made and how systems operate. According to respondents, that next phase is already beginning to take shape. Respondents ranked decision-making processes as the area experiencing the greatest AI-related change, ahead of many other parts of the business (see Figure 1).



Figure 1. Q: Please rank the top three areas in your organisation that have experienced the greatest change due to AI.

The importance of this shift aligns with earlier research conducted by TCS on *Intelligent Choice Architectures (ICAs).* That work identified decision environments as an emerging source of competitive advantage. Rather than simply automating processes or generating recommendations, AI systems can help surface new options, reveal hidden trade-offs, and continuously refine the context in which decisions are made.

Industry and region insights

Globally, decision-making processes, job roles and responsibilities, and workforce planning are the three areas seeing the greatest AI-driven change. Regional priorities vary: **operating models** enters the top three in LATAM and UKI; and **talent development and career growth** rises into the top three in North America, APAC, and ANZ. Among industries, **insurance diverges most sharply**, ranking operating models first and enterprise workflows third.

AI's path from recommendation to action runs through security

AI may be changing decision processes, but most organisations are not handing over broad authority. Figure 2 shows that two-thirds of respondents either do not permit autonomous AI decisions at all or in only a few specific use cases.

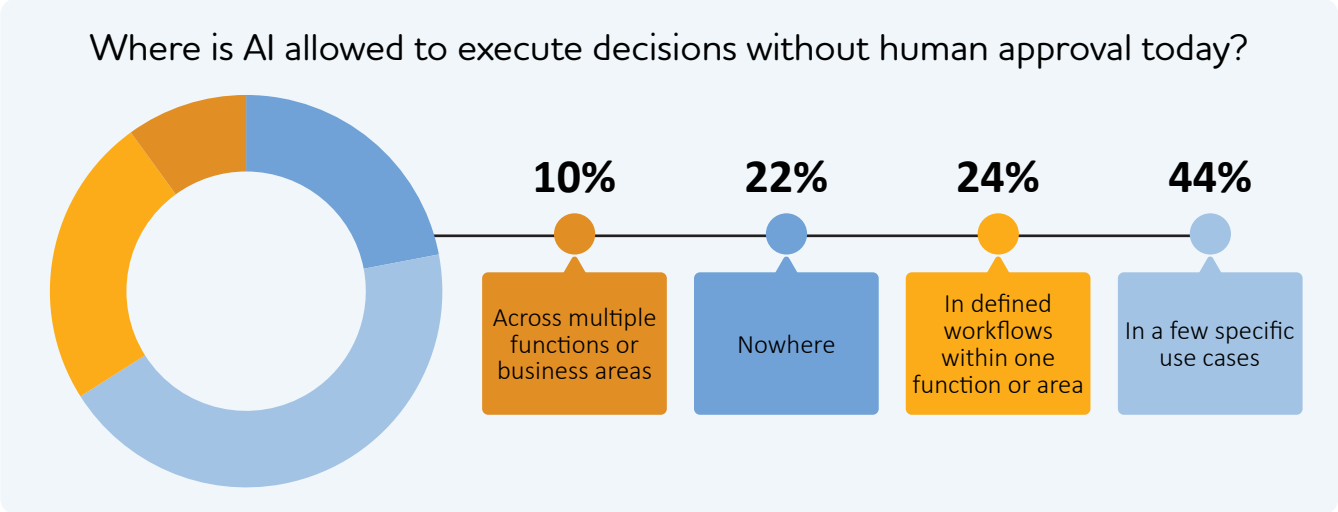


Figure 2. Q: Where is AI allowed to execute decisions without human approval today?

This distribution separates influence from execution. AI may recommend a course of action, flag conditions, or evaluate possible responses while a person retains final approval. Organisations are therefore expanding AI's influence on decisions more quickly than its authority to execute them. This distinction becomes increasingly important as enterprises move toward more agentic forms of AI, where systems can act within defined boundaries rather than simply generate insights or recommendations.

The same bounded progression appears in physical operations. Nearly half of respondents use AI to guide or improve operational decisions, and another third use it to increase visibility. Far fewer use it for task automation or dynamic real-time adjustment (see Figure 3).

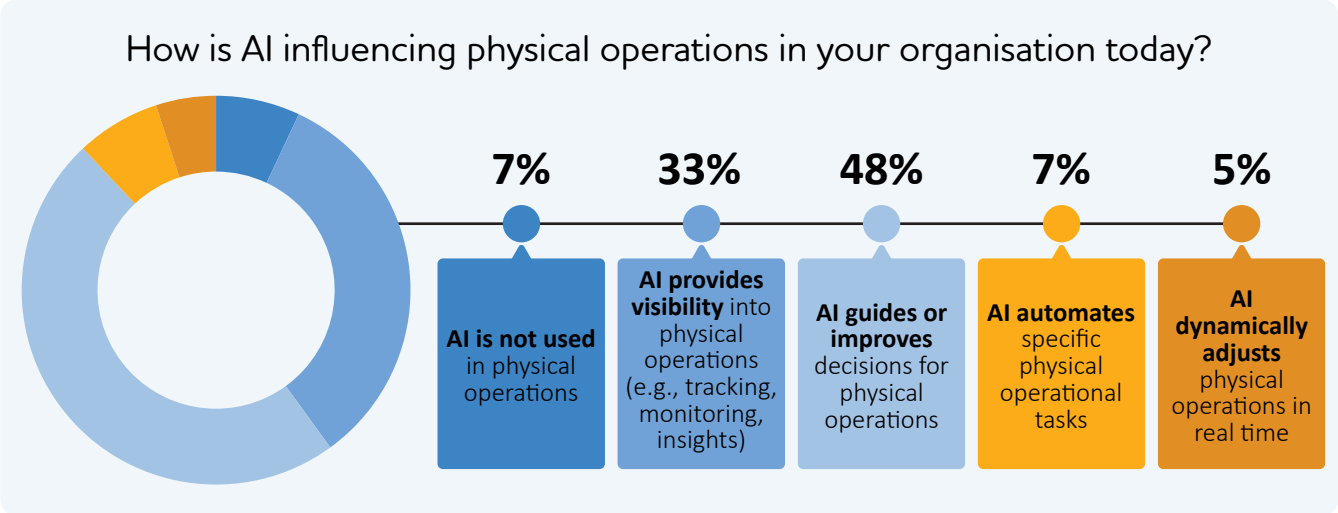


Figure 3. Q: How is AI influencing physical operations in your organisation today?

Industry insights

The use of AI in physical operations varies by sector, yet follows the same general trends across all industries. All the life sciences executives surveyed reported some use of AI, with 96% using it to improve operating technology decisions or to provide visibility into physical operations; however, far fewer (4%) said that AI dynamically adjusts their physical operations in real-time or automates specific tasks. At the opposite end of the use spectrum, banking, financial services and insurance firms reported the least use of AI overall in physical operations (13% said it was not used), yet among the highest uses (20%) for dynamic, real-time adjustments or automation of specific tasks. Other industries with the most advanced use of AI for physical operations are retail (20%), manufacturing (18%), and healthcare (16%). Consumer packaged goods is perhaps the most advanced industry of all for operating technology AI, with 12% of executives saying they use AI to adjust physical operations dynamically in real-time, and another 1% using it to automate specific tasks.

The limits on AI autonomy are closely tied to the limits on what AI can securely access and do across the enterprise. As organisations move from AI that informs decisions to AI that can carry them out, secure connectivity, governed permissions, and oversight become increasingly important. AI can generate recommendations with relatively limited access to enterprise systems. But acting reliably on those recommendations requires secure application connections, governed permissions, monitoring, oversight, and a coordinated pathway across tools and processes.

Today, relatively few respondents say they have reached that stage. The most common maturity stage, reported by 44%, is integration with several enterprise applications and tools; only 9% say they have AI agents that securely orchestrate workflows using governed tools and permissions (see Figure 4).

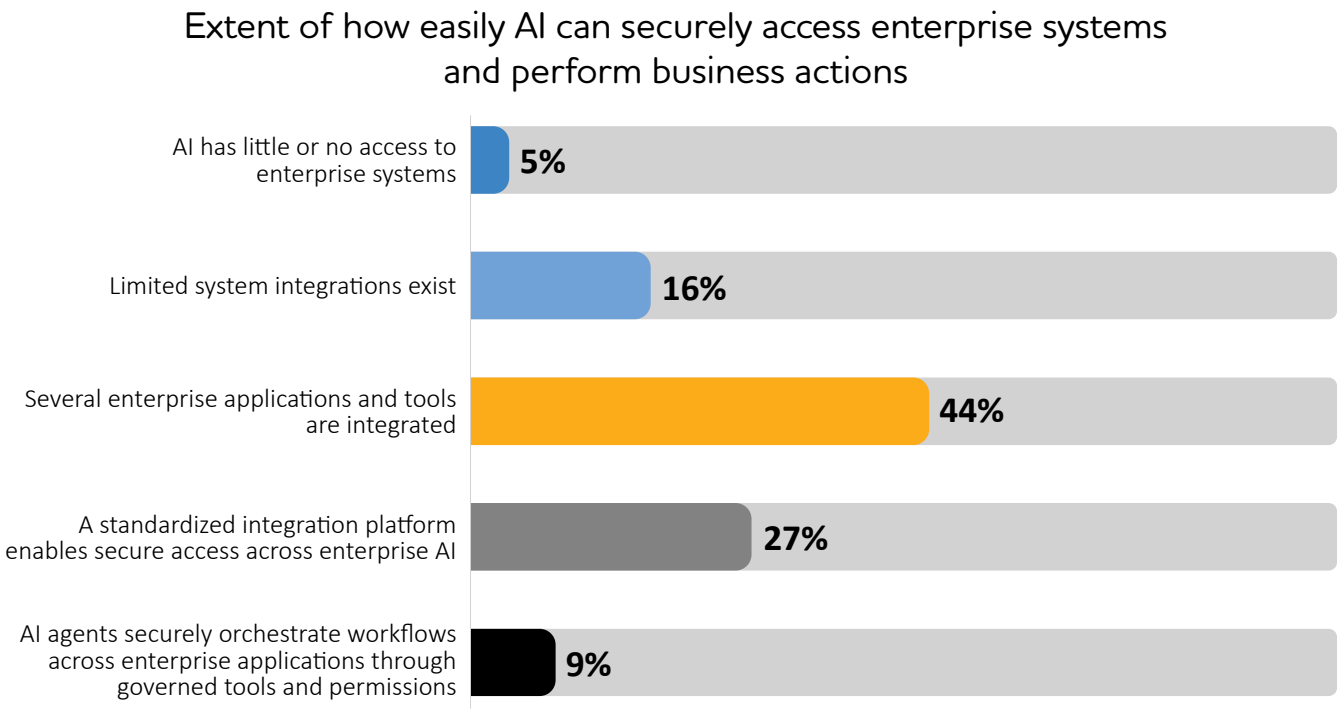


Figure 4. Q: How easily can AI securely access your enterprise systems and perform business actions? (e.g., CRM, ERP, PMS, HR, finance, collaboration tools)

Industry insights

Banking, financial services and insurance has an 11-point advantage in agent orchestration: 20%, compared with 9% globally.

Closing this infrastructure gap will be critical as enterprises move toward more agentic AI. And as AI gets closer to taking action, security and governance are becoming more consequential to deployment decisions. Security, risk, and governance considerations now rank as the leading factor shaping how and where AI systems are developed and deployed, ahead of cost, speed, and ease of deployment.

1

Security, risk, or governance considerations

LEADING DRIVER FOR AI SYSTEM AND DEPLOYMENT DECISIONS

> As AI gets closer to taking action, security and governance are becoming more consequential to deployment decisions.



Among industries, AI's impact has yet to become structural transformation

AI is already influencing how organisations compete, innovate, and engage customers. How far does AI's impact extend beyond the enterprise itself? When asked where AI is having the greatest impact on their industries, respondents rank customer engagement first, followed by product and service innovation and internal enterprise functions (see Figure 5.)

Ranked industry areas where AI is creating the greatest impact on business transformation



Figure 5. Q: Where is AI creating the greatest impact on business transformation in your industry?

Industry and geo results

Customer engagement, product or service innovation, and internal enterprise functions are the three leading areas of AI impact globally. The composition changes across several countries. Supply-chain and ecosystem coordination enters the top three in Japan and France, while both supply-chain coordination and industry business models enter Australia's top three. Industry results are similarly varied. Supply-chain coordination enters the top three in manufacturing, retail, consumer packaged goods, life sciences, and healthcare. Industry business models enter the top three in travel, hospitality and logistics, communications and media, consumer packaged goods, and utilities.

These are meaningful areas of change, but they remain concentrated within individual organisations rather than across entire industries. More structural forms of transformation such as changes to industry business models and revenue streams rank lower, while shifts in competition, market structure, and industry boundaries rank even lower. AI is improving how organisations participate in their industries before it fundamentally reshapes industry structures, boundaries, and revenue models.

Deeper industry transformation depends on coordination beyond a single enterprise. Respondents identify difficulty coordinating data, workflows, and processes across organisations as the leading barrier, followed by limited trust in AI outputs across organisational boundaries and difficulty translating insights into action (see Figure 6). Until organisations can connect these foundations across participants, AI's impact is likely to remain more visible within enterprises than across their broader ecosystems.

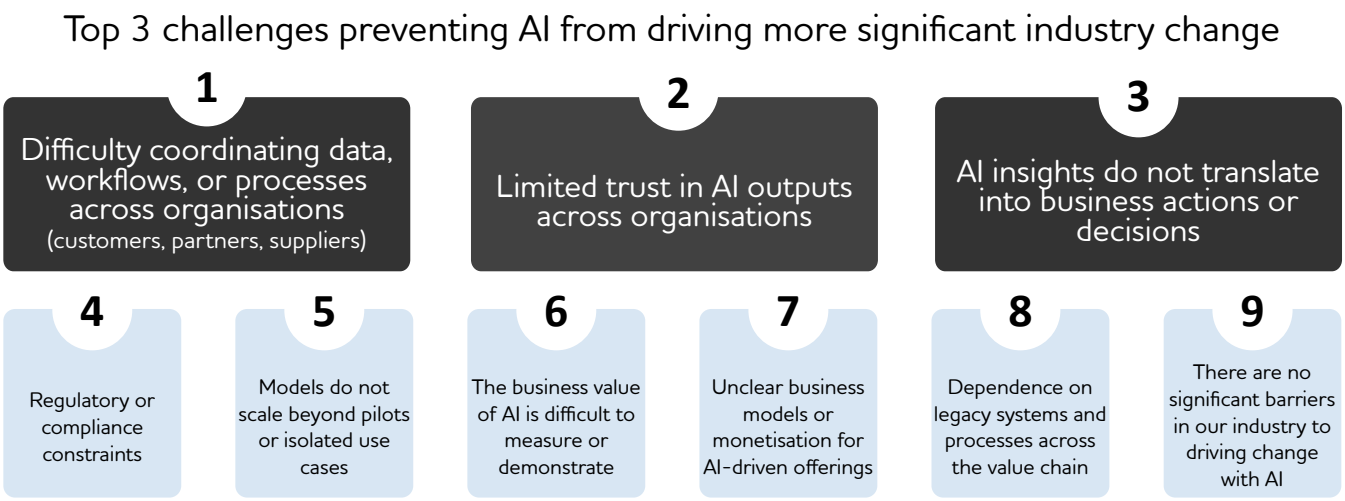


Figure 6. Q: What is preventing AI from driving more significant change in your industry?

Most AI activity still stops at the enterprise boundary

The maturity of external AI adoption helps explain why broader transformation remains limited. Nearly half of respondents (45%) report using AI only within internal operations. Another 31% use AI to support interactions with customers, partners, or suppliers, but without integrating AI into external systems or processes. Only 19% report integrating AI with partner systems and workflows, with fewer than 5% using AI to coordinate activities or enable shared workflows across multiple external stakeholders (see Figure 7).

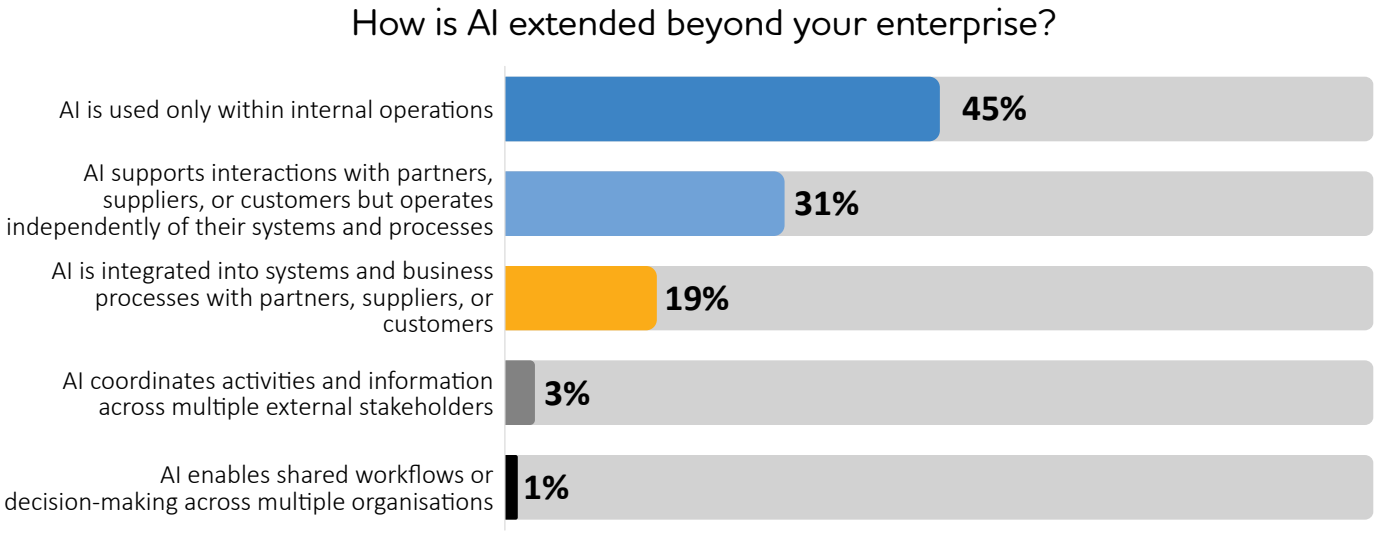


Figure 7. Q: How is AI extended beyond your enterprise?

Industry and geo results

While 45% of global respondents say AI is used only within internal operations, Japan is even more internally focused, at 60%. This pattern also appears in consumer packaged goods (61%), energy and resources (59%), travel, hospitality and logistics (56%), and manufacturing (56%). Insurance points in a different direction: internal-only use falls to 26%, while integration with partners, suppliers or customers rises 13 points above the global aggregate to 32%.

Taken together, the findings in this section show that AI is influencing decisions, workflows, and operations, but broader authority and ecosystem coordination remain limited. The next section examines the context, integration, controls, and architecture that determine whether organisations can move further.

Building for context, control, and adaptability

Organisations are steadily advancing the value they generate from data and AI. The long-term ambition is even greater as enterprises mature toward agentic systems that can continuously interpret changing business conditions, draw on enterprise context, and take appropriate action.

However, AI's growing influence on decisions has not yet translated into broad authority to act. Reliable action depends on business context, connected data, secure access to enterprise systems, governed permissions, and controls that operate consistently across models and platforms. While 29% use AI to uncover previously inaccessible insights and another 29% generate predictions and recommendations, only 16% provide contextual intelligence in the flow of work, and just 8% continuously interpret enterprise-wide data to anticipate needs and adapt actions in real time (see Figure 8).

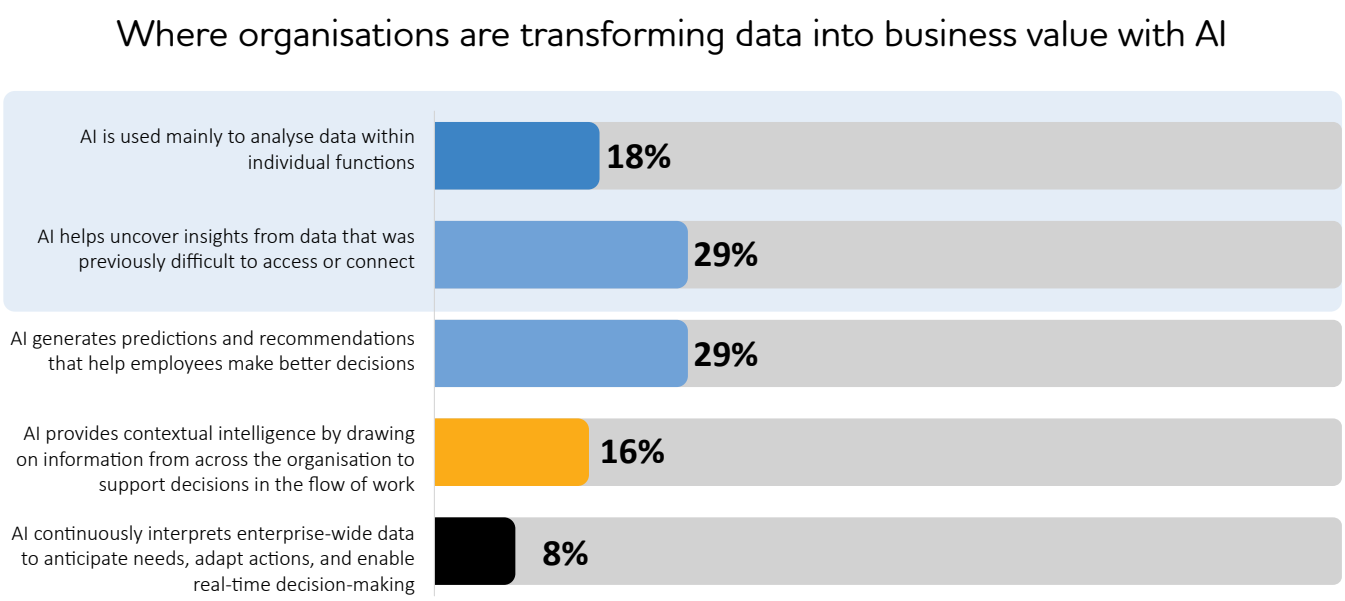


Figure 8. Q: Which statement best describes the value your organisation is creating from data through AI?

Industry and region insights

Industry results reveal different paths forward. Prediction and recommendation generation soars for a number of industries: energy and resources (48%); utilities (44%); consumer packaged goods (38%) and manufacturing (35%). Insurance stands out for continuous enterprise-wide interpretation at 36%, compared with 8% globally.



Many organisations are encountering the practical limits of fragmented business context. Forty-four percent still rely on employee knowledge or fragmented documentation, while insufficiently integrated, accessible, or trusted data is the leading challenge when scaling beyond initial use cases (see Figures 9 and 10). Similar to the coordination challenge at the industry level, here it is visible inside the enterprise. Accessible business context is the connective layer that allows AI to produce consistent intelligence across functions and, eventually, across ecosystems. Agentic AI turns this already important layer into a critical one. And as AI systems gain the ability to take actions within defined boundaries, organisations must determine how much authority to delegate and where human approval remains necessary.

How organisations are structuring and standardising business context for AI

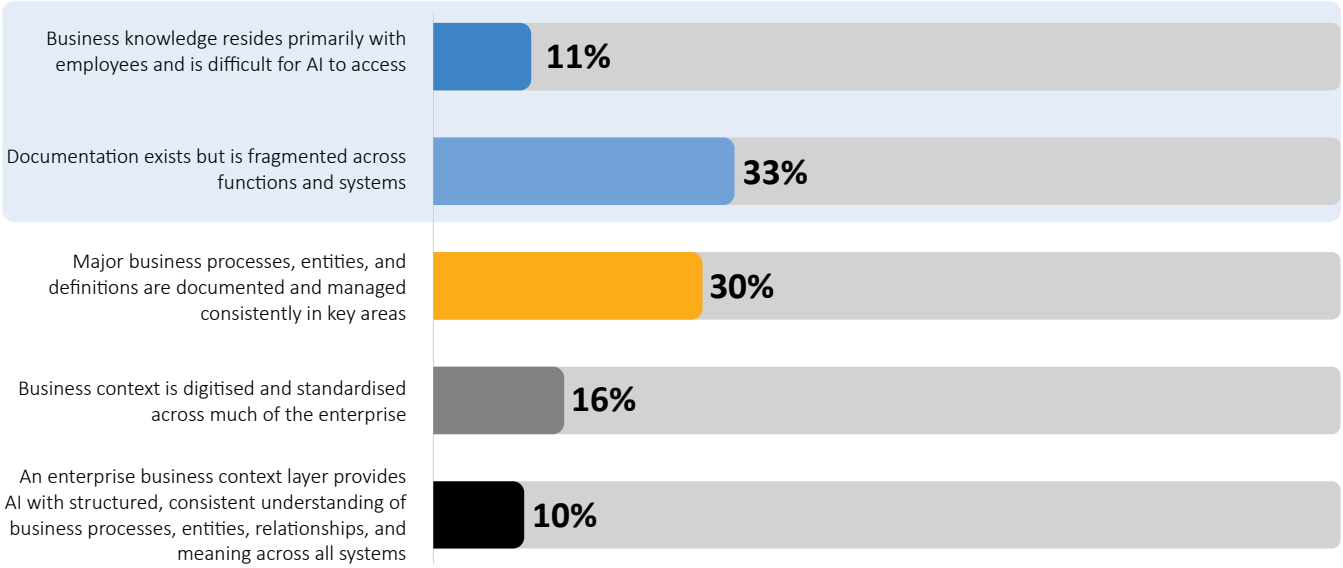


Figure 9. Q: Which statement best describes how well your organisation structures and standardises business context so AI can consistently operate across business functions? (e.g., finance, procurement, HR, marketing, revenue management, customer operations)

The consequences become visible when organisations attempt to scale beyond initial use cases. Data that is insufficiently integrated, accessible, or trusted ranks as the leading challenge to broader AI deployment (see Figure 10). Business teams struggling to adopt, trust, or operationalise AI ranks second, followed by faster-than-expected cost increases.

> “
> Many organisations are encountering the practical limits of fragmented business context.

Top challenges scaling AI beyond initial use cases

Data is not sufficiently integrated, accessible, or trusted to support broader deployment	01
Business teams struggle to adopt, trust, or operationalise AI	02
Costs increase faster than expected	03
Existing systems cannot integrate or scale AI workloads	04
AI models and tools become difficult to manage or maintain	05

Figure 10. Q: What are the top challenges that occur when you try to scale AI beyond initial use cases?



Production readiness is selective rather than enterprise-wide

Organisations have made meaningful progress in building the technology foundations required for enterprise AI. The majority of respondents say they have moved beyond experimentation and can support selected production use. However, readiness remains uneven across the technology stack. Across platforms, models, applications, APIs and integration, and governance and controls, only 5% to 7% describe their environments as fully scaled (see Figure 11). Targeted production readiness is now common; enterprise-wide maturity is not.

How well can your environment support scaling AI?

Capability	Not supported	Ready for limited pilots and experimentation	Ready for select production use	Ready for widespread production use	Fully scaled across the enterprise
Platforms Using AI capabilities across business functions	**2%**	**25%**	**42%**	**25%**	**7%**
Models Deploying, managing, and updating AI models and tools	**1%**	**25%**	**46%**	**21%**	**7%**
Applications Embedding AI into day-to-day business processes	**2%**	**20%**	**48%**	**24%**	**6%**
APIs/integration Connecting systems, data, and workflows across the organisation	**2%**	**31%**	**42%**	**19%**	**5%**
Governance and controls Ensuring oversight, compliance, and responsible use of AI systems	**4%**	**24%**	**40%**	**28%**	**5%**

Figure 11. Q: How well can your technology environment support scaling AI in the following areas?



Governance is present, but fragmented controls weaken consistency and trust

As AI becomes more deeply integrated into business processes and decision-making, trust becomes a prerequisite for scale. In an agentic environment, organisations need to govern not only what AI generates, but what it is permitted to do, under what conditions, and who remains accountable for the outcome. Organisations may invest in context layers, secure integrations, and production-ready environments, but those capabilities create value only if people are willing to rely on the outputs they produce.

Figure 12 shows that respondents identify unpredictable, inconsistent, or inaccurate outputs as the greatest factor undermining trust in AI. Lack of accountability follows, while limited transparency or explainability ranks third. Unpredictable, inconsistent, or inaccurate outputs are the leading source of trust breakdown, followed by unclear accountability and limited transparency or explainability.

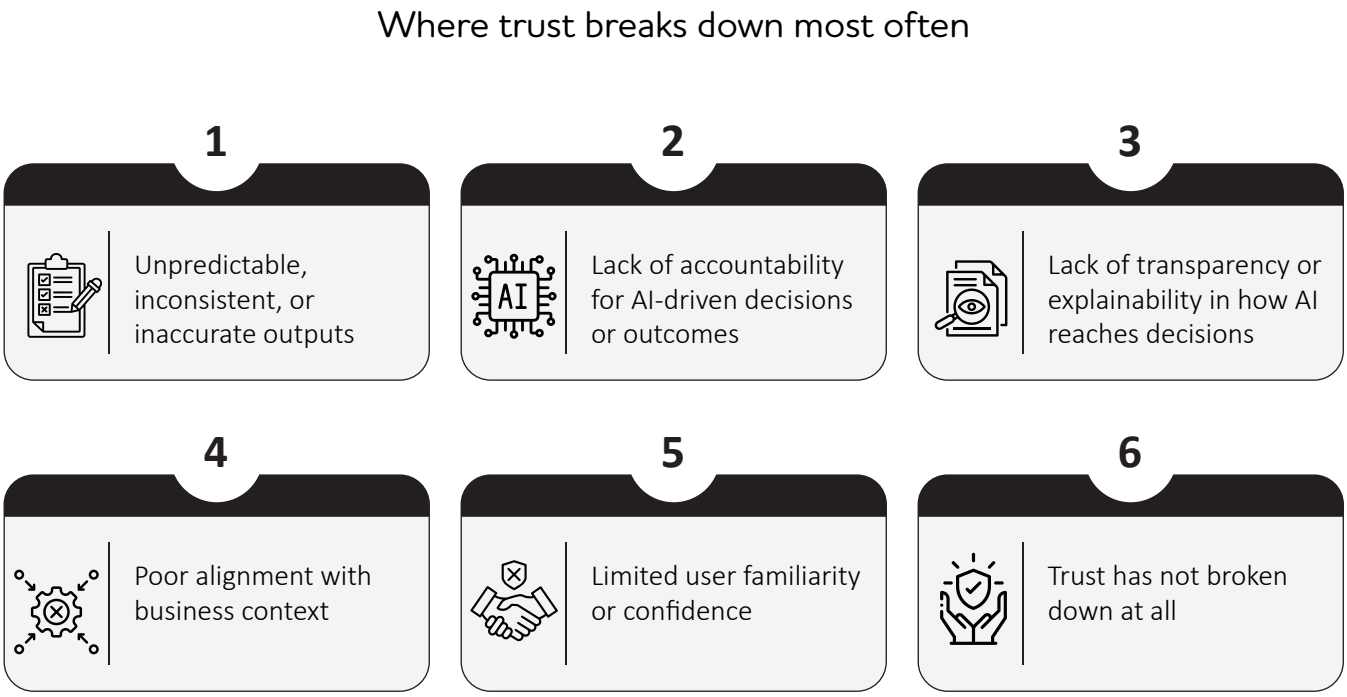


Figure 12. Q: Where does trust in AI break down most often?

> In an agentic environment, organisations need to govern not only what AI generates, but what it is permitted to do, under what conditions, and who remains accountable for the outcome.

Organisations are not starting from an absence of controls, but the largest share rate them only moderately effective (see Figure 13). Confidence is highest in exception handling and human intervention, while real-time monitoring receives the lowest highly effective rating. That gap matters: intervention mechanisms depend on the ability to detect when intervention is needed.

How effectively is your AI environment controlled?

Capability	Not effective	Limited effectiveness	Moderately effective	Highly and consistently effective	Don't know
Monitoring AI performance in real time	**2%**	**31%**	**58%**	**9%**	**0%**
Identifying and managing risk	**2%**	**24%**	**47%**	**27%**	**0%**
Enforcing policies and compliance requirements	**2%**	**24%**	**50%**	**24%**	**1%**
Managing exceptions, escalation, and human intervention	**3%**	**23%**	**42%**	**32%**	**1%**
Ensuring transparency and explainability for AI outputs	**3%**	**18%**	**49%**	**29%**	**1%**
Defining and maintaining accountability for AI-driven decisions and outcomes	**5%**	**17%**	**47%**	**29%**	**1%**

Figure 13. Q: How effectively are the following controlled in your AI environment?

Multi-platform governance remains a work in progress

Governance becomes more difficult as AI expands across multiple models, platforms, applications, and vendors. Consistency depends not only on the quality of individual controls, but also on the ability to coordinate them across the broader AI environment.

The results indicate that such coordination remains a challenge. More than one-third of respondents (34%) rely primarily on governance capabilities provided by individual vendors, with limited coordination across platforms. Just 8% operate a federated model with enterprise-wide visibility, automated controls, and coordinated oversight across vendor and internally developed ecosystems (see Figure 14).

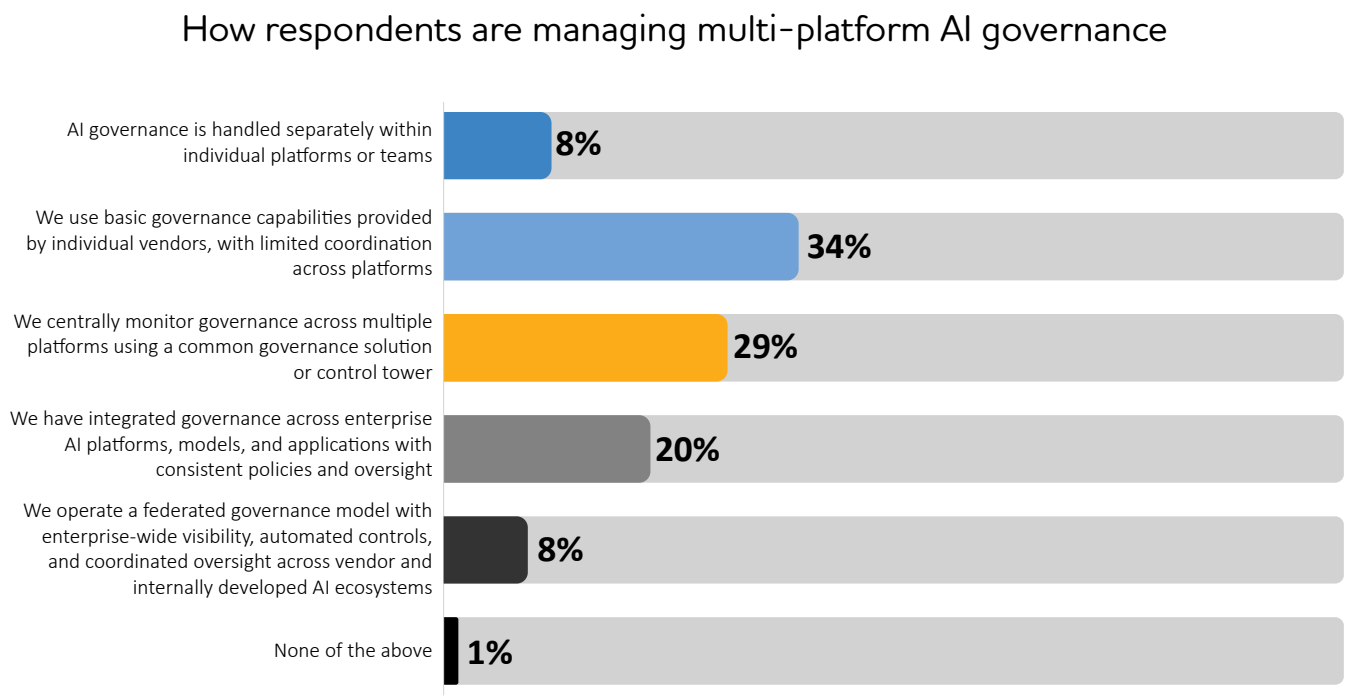


Figure 14. Q: Which best describes how your organisation manages AI governance across multiple AI platforms, models, and applications?

Architecture choices now extend from deployment to resilience and return

Multi-platform governance is only one dimension of managing AI at scale. As organisations deploy AI across multiple models, providers, and environments, attention is also shifting to the architectural decisions that determine how much control, flexibility, resilience, and value those environments deliver over time. As seen earlier, security, risk, and governance considerations now rank as the leading factor shaping how and where AI systems are developed and deployed. The need to retain control over data and intellectual property follows closely behind (see Figure 15).

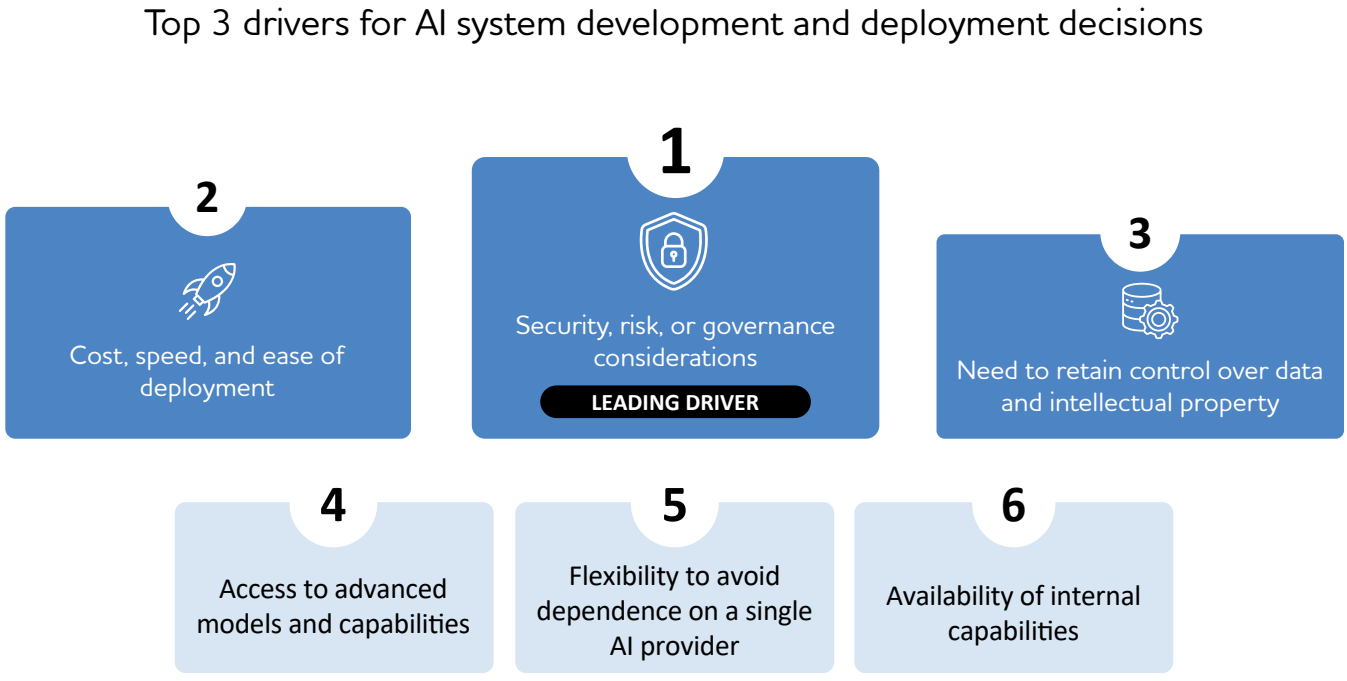


Figure 15. Q: What is driving decisions about where and how AI systems are developed and deployed?

Model independence is one reflection of that shift. Figure 16 shows that 90% consider avoiding dependence on a single model or platform important to some degree, including 27% who view it as a strategic architecture and investment requirement. Concerns about control extend beyond vendors and platforms.

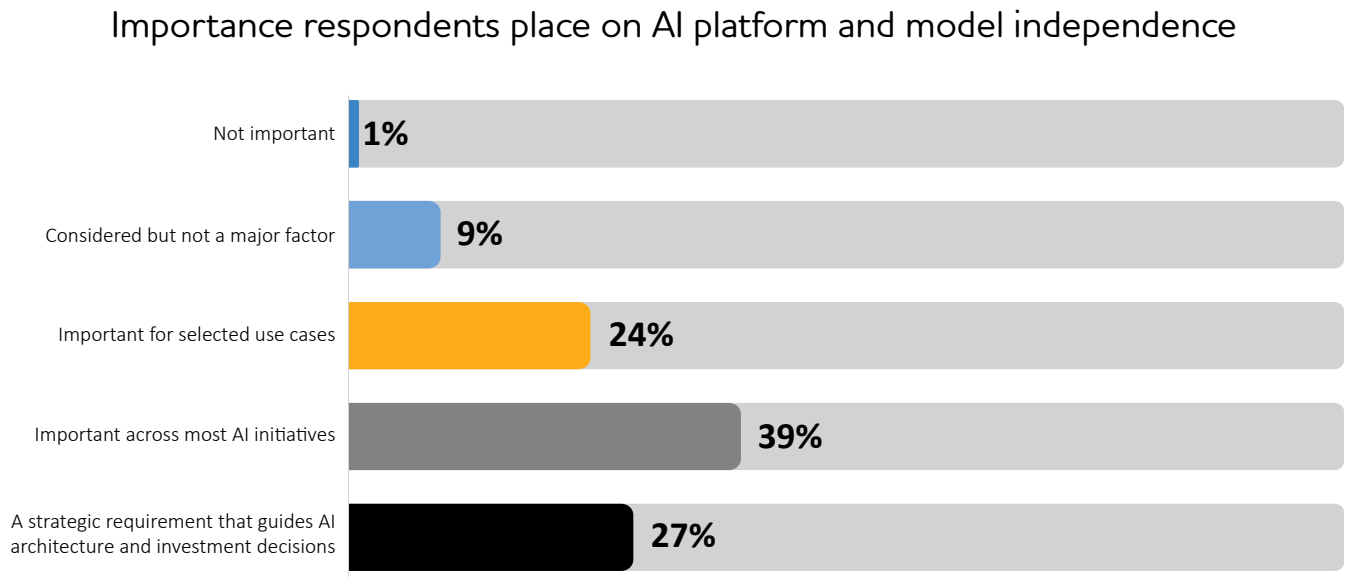


Figure 16. Q: How important is maintaining independence from any single AI platform or model when making AI investment decisions?

90% of respondents say model independence is an important factor in AI investment decisions

Industry and region insights

France places greater emphasis on model independence across most AI initiatives, at 56% vs 39% for the global respondents. Aerospace gives independence greater strategic weight, with 39% identifying it as an architectural and investment requirement vs 27% globally.



Sovereign AI requirements are also emerging as an enterprise risk consideration. More than a third (34%) describe requirements for control over AI data location, model choice, and regulatory or data-residency compliance as a high or very high organisational risk (see Figure 17).

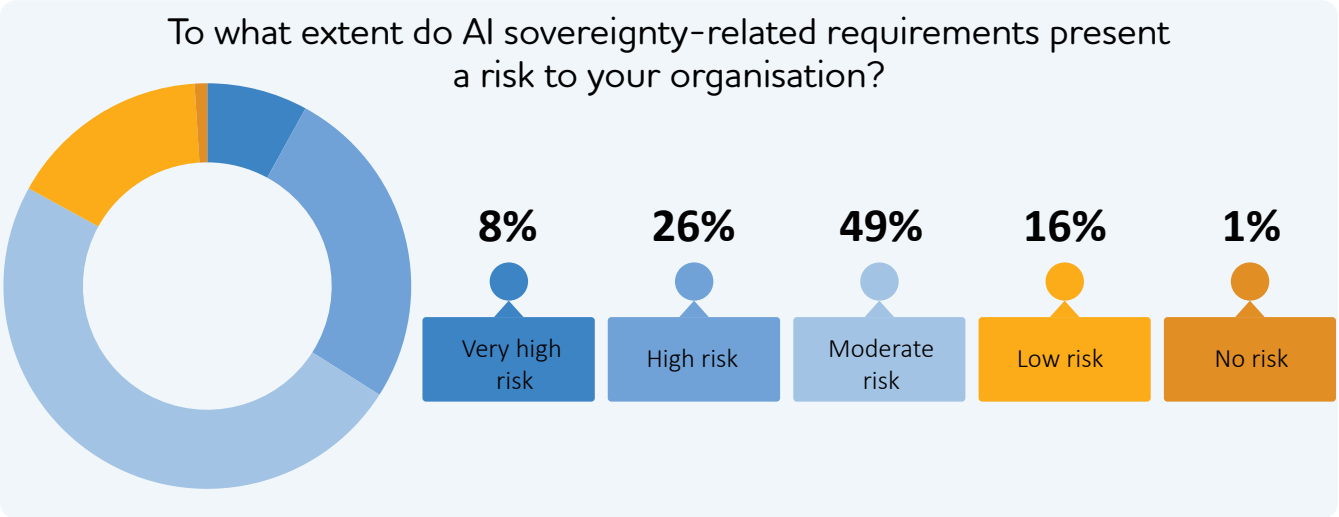
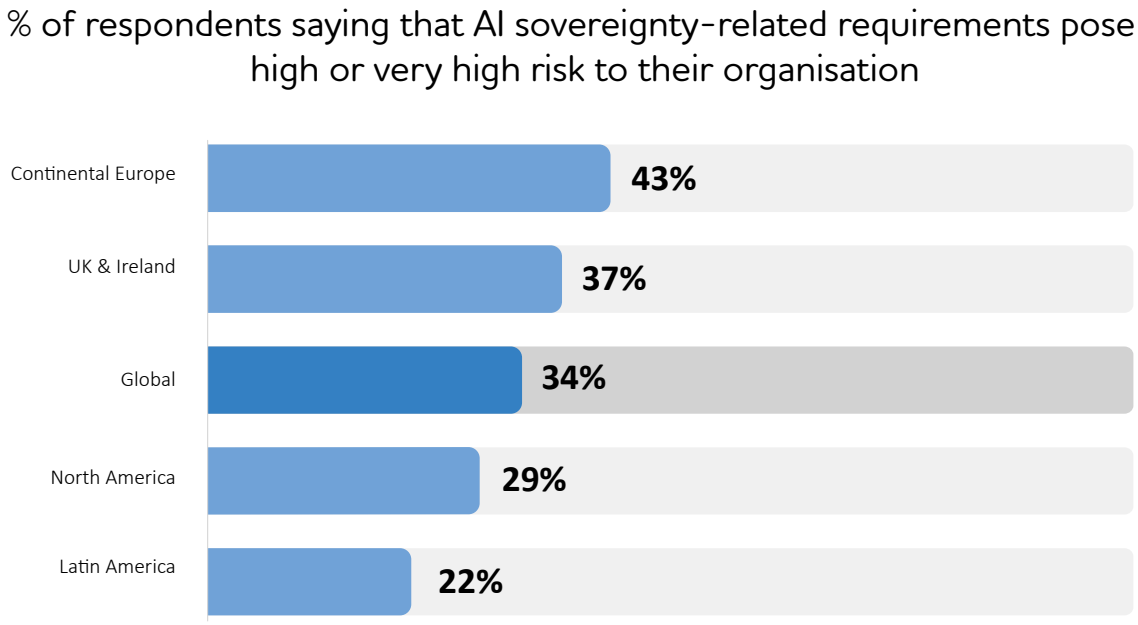


Figure 17. Q: To what extent does AI sovereignty (your organisation's control over AI data location, model choice, and regulatory/data-residency compliance) present a risk to your organisation?

There are considerable regional differences when it comes to perceived risk related to AI sovereignty requirements.



But could organisations act on that desire for independence if a critical provider or model became unavailable? While 47% are prepared to switch providers for a few critical use cases and 34% for most business-critical use cases, only 7% report being prepared to switch critical AI workloads across the organisation (see Figure 18). Among organizations with a continuity plan, 35% have not tested it.

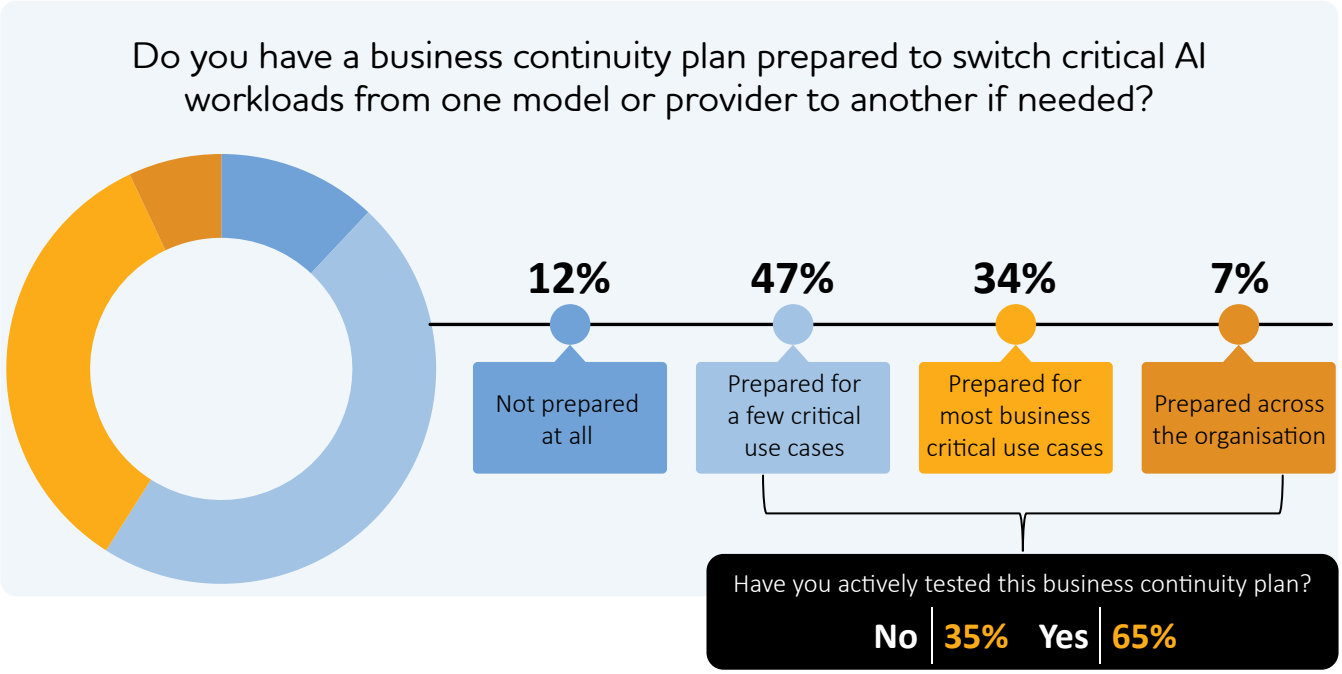


Figure 18. Q: Do you have a business continuity plan prepared to switch critical AI workloads from one model or provider to another if needed? Q. Have you actively tested this business continuity plan?

Open-source adoption is similarly selective, favoring workload fit over broad standardisation. Figure 19 shows that 46% report targeted deployment for selected use cases, compared with 13% reporting broad deployment across multiple functions. Just 7% describe open source as a strategic priority involving significant customisation or deployment.

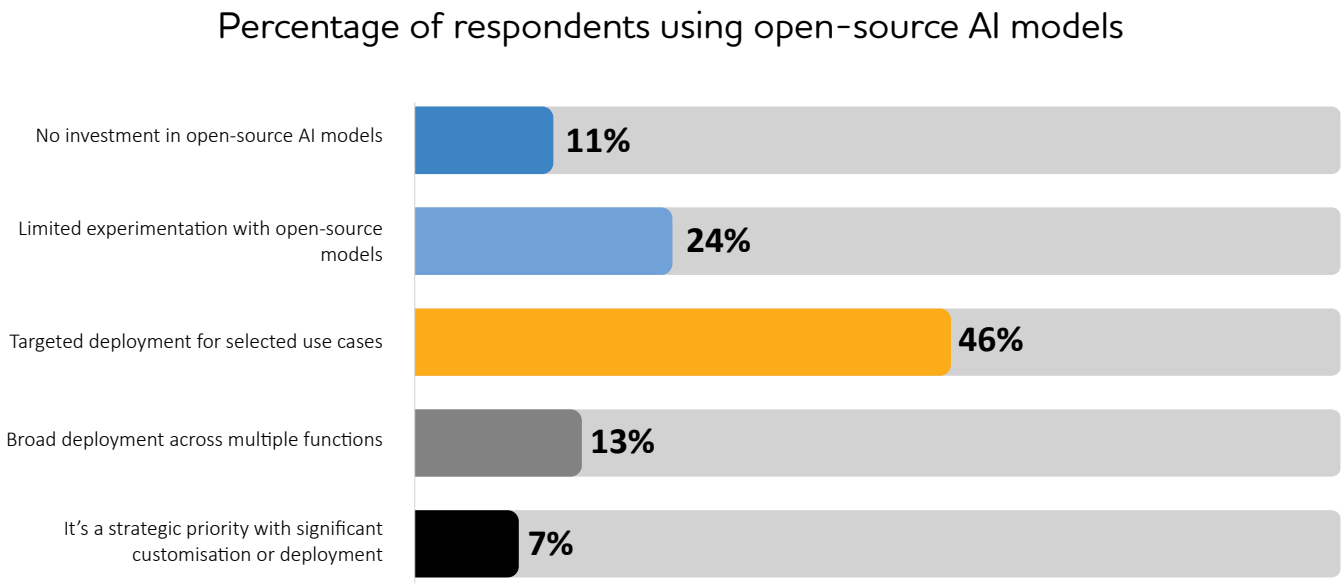


Figure 19. Q: To what extent is your organisation invested in developing, customising, or deploying open-source AI models (including small language models) instead of relying exclusively on frontier AI models?

Industry and region insights

Banking, financial service, and insurance reaches 17% for strategic priority, compared with 7% globally. Technology 19% for broad deployment and CMI 12% for strategic priority.

Economic discipline is developing alongside architectural flexibility. Figure 20 shows that 39% assess whether AI initiatives generate sufficient value relative to token and compute costs, while 31% track those costs without consistently evaluating returns. Only 9% continuously optimise value generated per token or unit of compute through ongoing refinement of prompts, architectures, and models.

The result is a broader definition of AI architecture. Beyond governance, success increasingly depends on maintaining the ability to choose, switch, control, recover, and continuously improve the economics of AI over time.

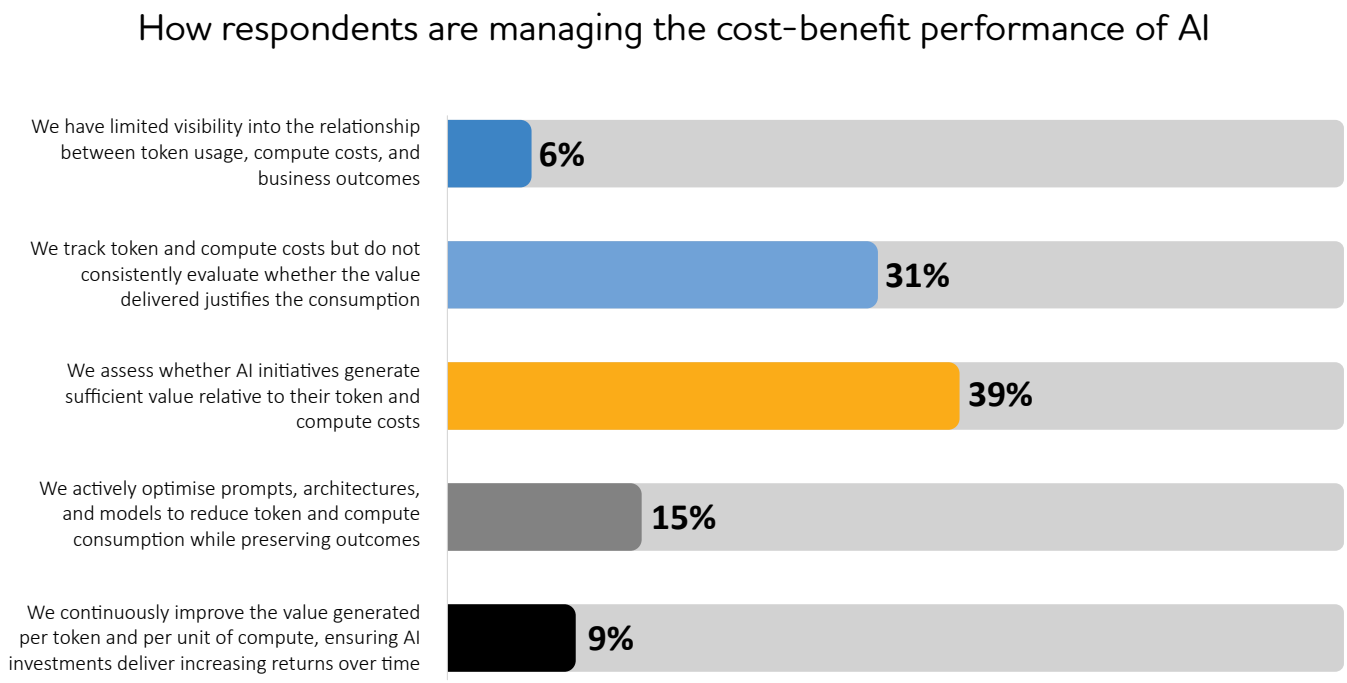
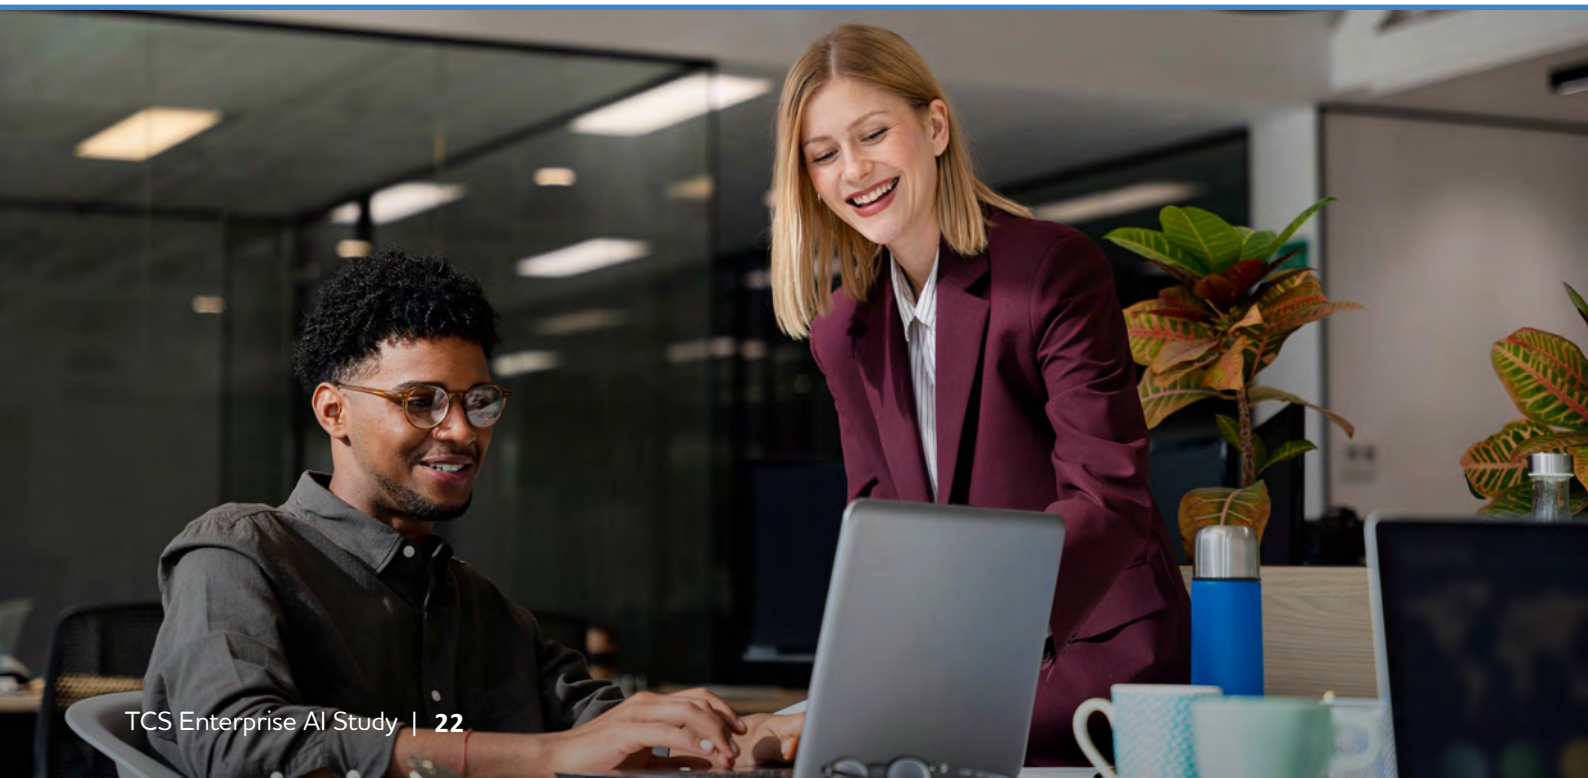


Figure 20. Q: Which statement best describes your organisation's ability to improve the cost-benefit performance of AI over time?



Aligning leadership, skills, and workforce systems for the AI era

The findings throughout this study point to a common conclusion. Enterprise AI has definitively moved past technical capability. Organisations are redesigning workflows, introducing new forms of decision support, establishing governance models, and building increasingly complex AI portfolios.

As AI becomes embedded in decisions and operations, organisational alignment becomes the next limiting factor. Scaling requires leadership structures that coordinate initiatives across the business, operating models that connect accountability with execution, and workforce systems that adapt as roles, responsibilities, and skills evolve.

AI scale requires a defined enterprise operating model

How organisations structure AI initiatives provides an indication of their maturity and readiness to scale. Centres of Excellence remain the most common model, cited by 38% of respondents, while 29% report some form of shared governance (see Figure 21). Only a small minority treat AI as a strategic enterprise capability supported by standardised architecture, governance, reusable assets, and measurable value.

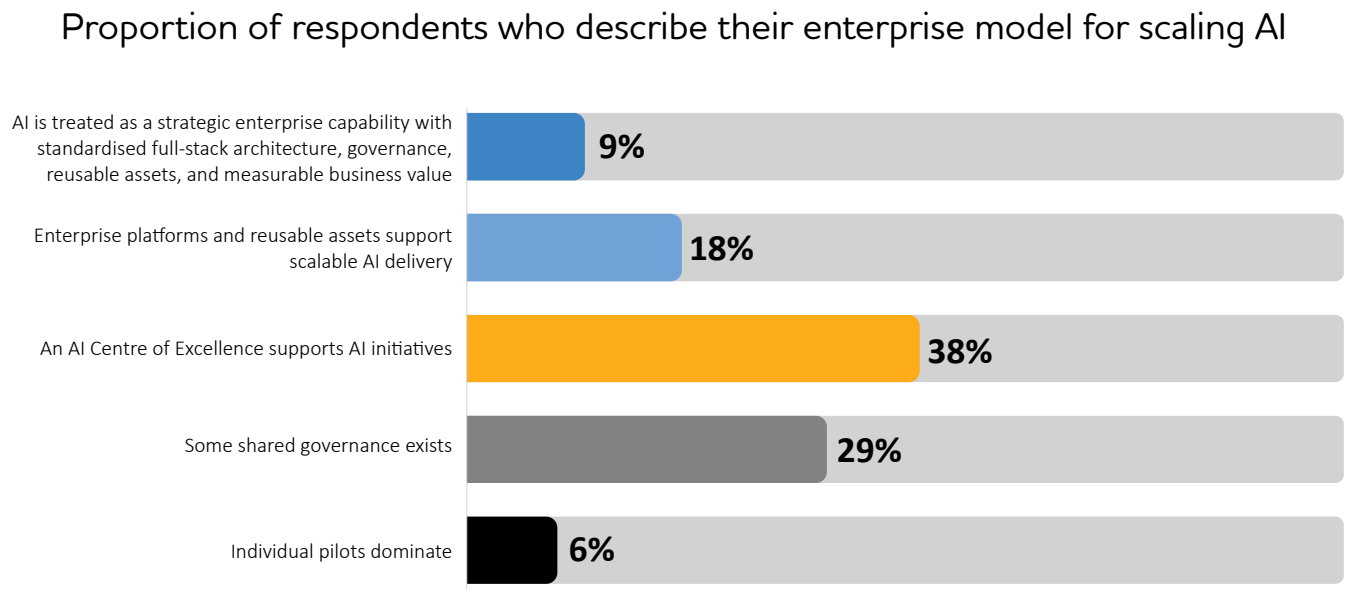


Figure 21. Q: Which statement best describes your enterprise operating model for scaling AI?

These findings suggest that organisations are building coordination in stages. A Centre of Excellence concentrates expertise and accelerates experimentation. Shared governance distributes responsibility while maintaining oversight. Enterprise platforms create foundations for repeatable delivery across functions. Reaching enterprise scale requires these capabilities to operate as a connected system rather than as isolated initiatives.

Where CAIOs exist, they lead scaling

As enterprises formalise AI leadership, the Chief AI Officer (CAIO) is emerging as a focal point for scaling, consolidating responsibility that otherwise remains distributed across executive roles. The Chief AI Officer (CAIO) role is taking hold globally. However, adoption varies considerably by region (see Figure 22).

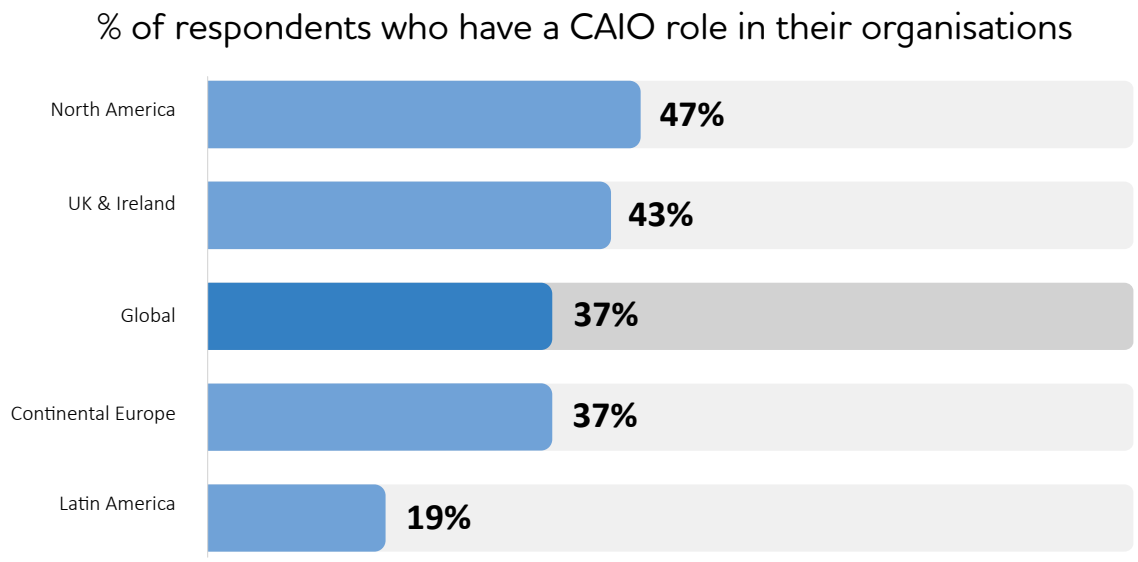


Figure 22. Q. Do you have a Chief AI Officer?

Where a CAIO exists, executive responsibility for scaling AI is highly concentrated: 96% identify the CAIO as having primary responsibility. Without a CAIO, ownership is more distributed across the C-suite (see Figure 23).

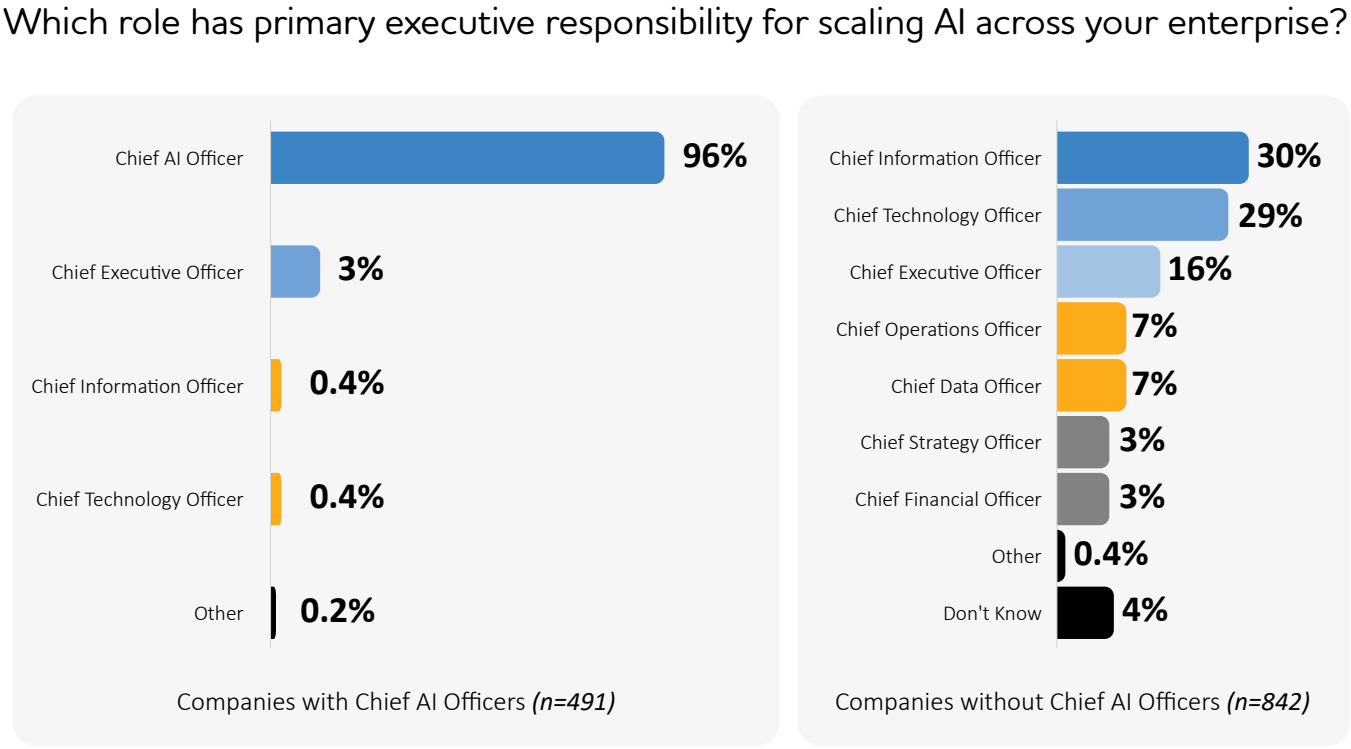


Figure 23. Q. Which role has primary executive responsibility for scaling AI across your enterprise?

Practical experience is reshaping expectations for AI's workforce impact

As organisations gain practical experience with AI, confidence in predicting workforce outcomes appears to be declining. The most notable change since we posed this question in our ***2024 TCS AI for Business Study*** is the growth in uncertainty.

As Figure 24 shows, the share reporting that it is too early to know how AI will affect job roles increased 15 points from 4% in 2024 to 19% in the current study. Expectations that AI will eliminate more jobs than it creates also fell from 47% to 37%. At the same time, fewer respondents now expect AI to create more jobs than it eliminates.

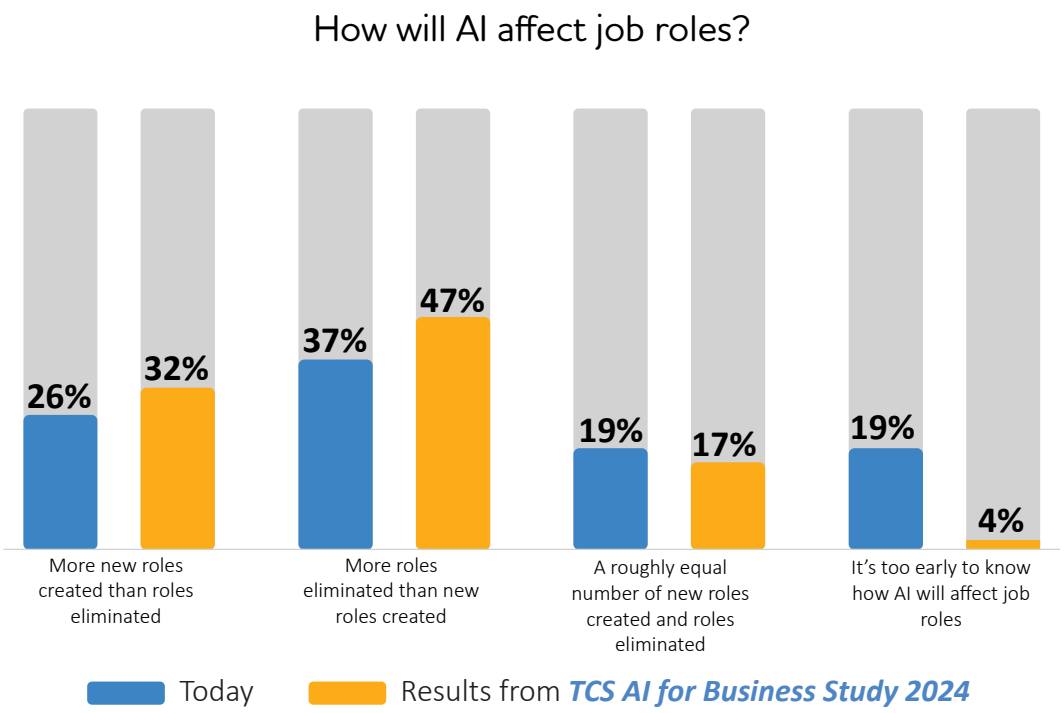


Figure 24. Q: How do you foresee the balance between roles created and roles eliminated as a result of AI's use at your company?

Still, this comparison should not be interpreted as a simple shift toward either optimism or pessimism. Instead, experience appears to be making the impact of AI on work more difficult to predict confidently. Organisations can see that AI is altering tasks, decisions, and responsibilities, but the longer-term workforce implications remain tied to where new forms of work emerge and how processes are redesigned. And as agents take on bounded activities within those processes, organisations will increasingly have to define which responsibilities sit with people, which can be delegated to AI, and where human judgement or intervention remains essential.



Workforce systems lag changes in roles and responsibilities

That challenge becomes more apparent when examining how organisations define the skills needed in an AI-enabled workforce. Progress is visible, but there is limited enterprise integration.

A third have identified important AI-related skills but have not organised them into a formal framework (see Figure 25). Only 6% continuously evolve a comprehensive workforce taxonomy that connects skills, roles, learning, career paths, and workforce planning.

Defining the skills needed in an AI-enabled workforce

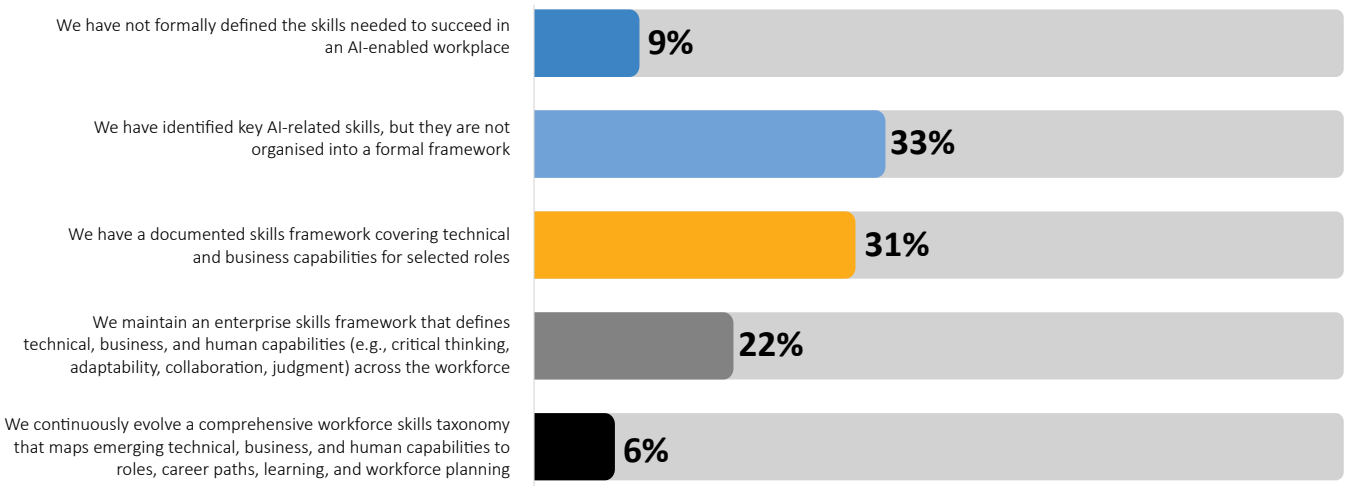


Figure 25. Q: What statement best describes how your organisation defines the skills needed to succeed in an AI-enabled workforce?

Just over a quarter (28%) of respondents maintain an enterprise framework or evolving taxonomy

As AI reshapes workflows and decision-making, there will be greater impetus for skills to evolve beyond stand-alone learning initiatives. Enterprise-wide skills taxonomies become part of a broader workforce system for continuously adapting workforce systems as technology and work evolve.

> As AI reshapes workflows and decision-making, there will be greater impetus for skills to evolve beyond stand-alone learning initiatives.

Conclusion

Organisations have moved beyond proving that AI can improve productivity. Across the study, respondents describe AI reshaping decisions, workflows, operating models, and enterprise technology strategies. The challenge is how organisations will scale that value in ways that are trusted, resilient, and sustainable.

The findings point to a common reality: progress depends on the alignment of process transformation, technology foundations, and workforce adaptation. AI is influencing more decisions and workflows, but broader authority remains constrained by fragmented business context, uneven governance, and technology environments that are still maturing. At the same time, leadership structures, skills frameworks, and workforce systems are continuing to evolve alongside the technology itself.

Perhaps the study's most important finding is that enterprise AI is becoming less about individual technologies and more about the systems that connect them. Competitive advantage will not come from access to models, platforms, agents, or data alone. It will come from an organisation's ability to connect these capabilities to how decisions are made, how work is organised, and how trust is maintained at scale.

Enterprise AI has reached a decision point. Realizing the full potential of enterprise AI depends on an organisation's willingness to redesign work, strengthen the foundations of trust and control, and prepare for new forms of intelligence embedded across the business. The choices made now will determine whether today's gains remain localised improvements or become lasting enterprise advantage.

Next Steps: Priorities for scaling enterprise AI

Four priorities stand out for leaders seeking to move from localised success to trusted enterprise transformation.

Build AI-ready data and security-first foundations

The study repeatedly shows that trust, data readiness, governance, and operationalisation remain among the largest barriers to scale. Organisations should prioritize investments that improve data accessibility, quality, context, and interoperability, while establishing built-in controls and instrumentation that monitor outcomes, manage risk, and maintain accountability. As AI takes on greater responsibility for recommendations and actions, governance must evolve from project oversight to continuous operational control capable of supporting AI at enterprise scale.

Architect for flexibility, security, and choice

AI strategy is increasingly becoming architecture strategy. Leaders should evaluate how architecture decisions affect long-term flexibility, continuity, and risk exposure. This includes reducing unnecessary dependence on individual models or providers, strengthening continuity planning, assessing sovereign AI requirements, and establishing clear approaches to platform and model management. Enterprise AI architectures should be designed for performance today, but also for adaptability as technologies, regulations, and business requirements continue to evolve.

Establish enterprise accountability and business value measurement for AI transformation

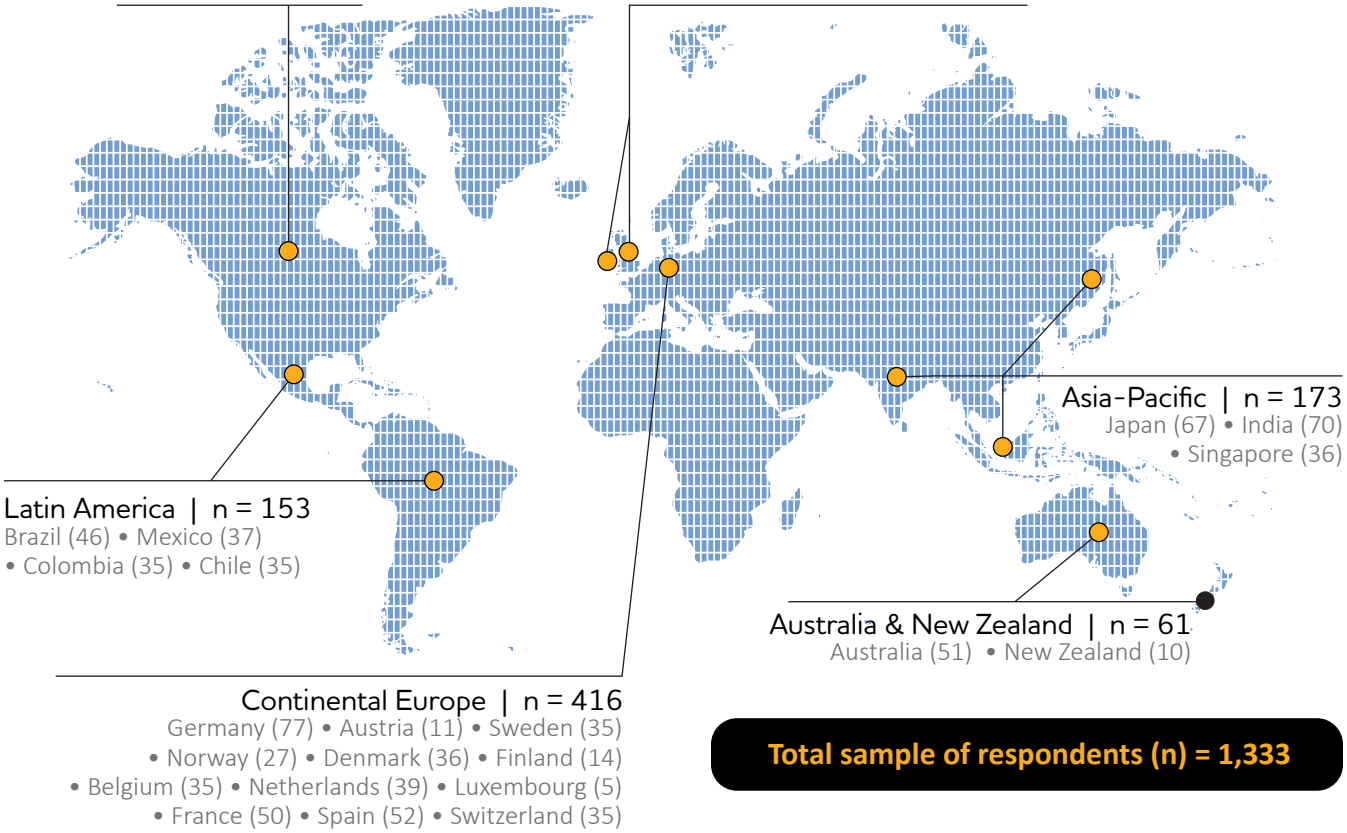
Scaling AI goes beyond technical ownership. Organizations should establish clear executive accountability and operating models that connect technology decisions with process transformation, governance, workforce planning, and business value realisation. Centres of Excellence often provide an effective starting point, but enterprise scale requires coordination across business, technology, operations, risk, and HR functions. Leadership structures should be designed to align decisions across these domains rather than manage them independently.

Connect workforce strategy to new ways of work

AI is already reshaping decisions, responsibilities, and ways of working, even as long-term workforce impacts remain uncertain. Organisations should move beyond isolated skills initiatives and develop workforce systems that connect emerging capabilities with learning, career progression, performance management, workforce planning, and mobility. The objective is not to predict exactly how AI will affect every role, but to create mechanisms that allow the organization to continuously adapt as work evolves.

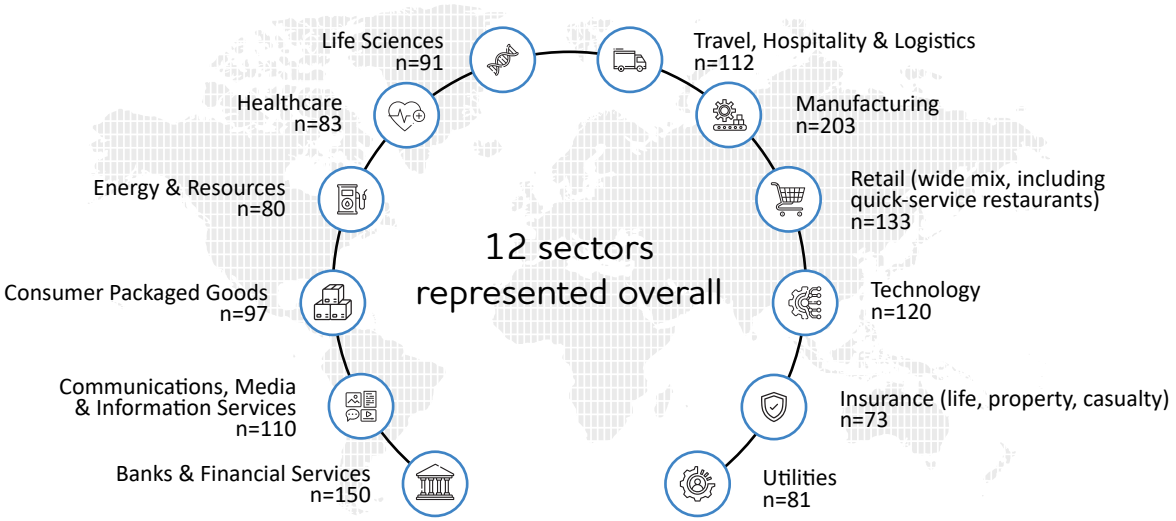
Study methodology

1,300+ Respondents span six regions across 25 countries



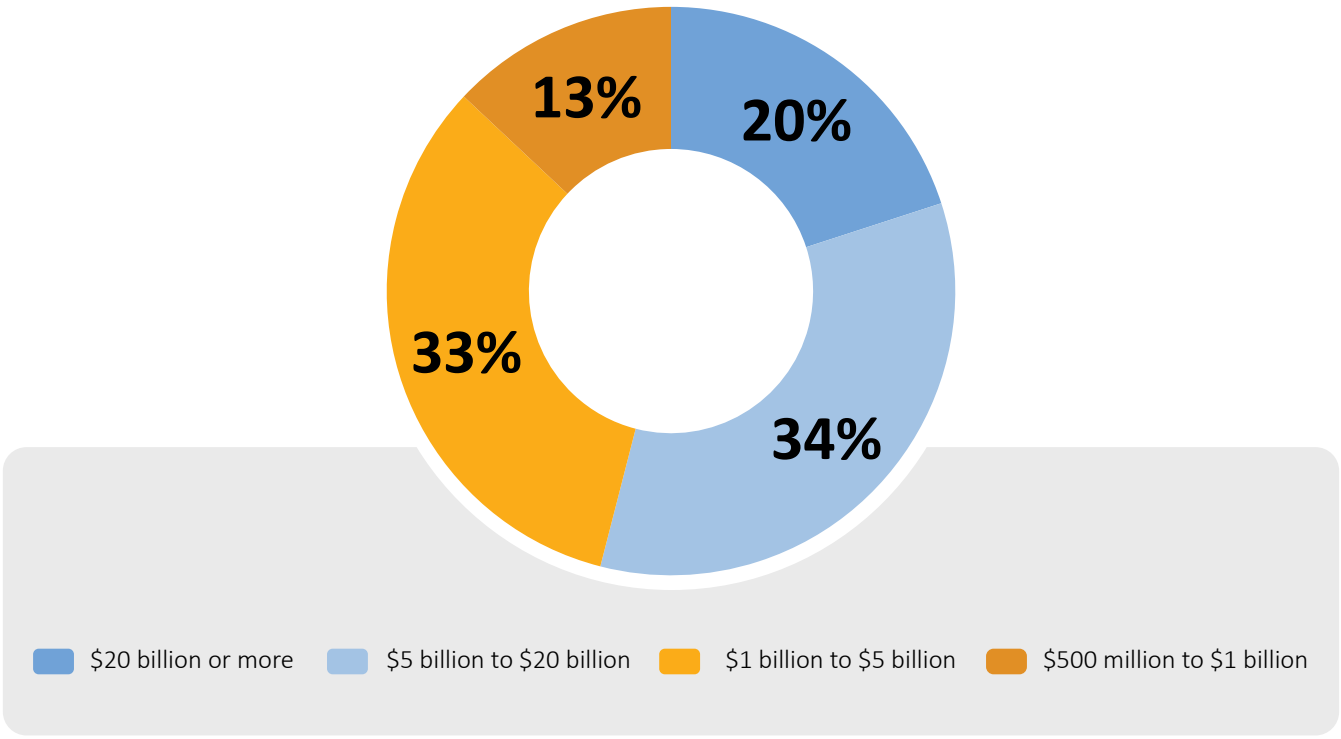
Total sample of respondents (n) = 1,333

Industries surveyed

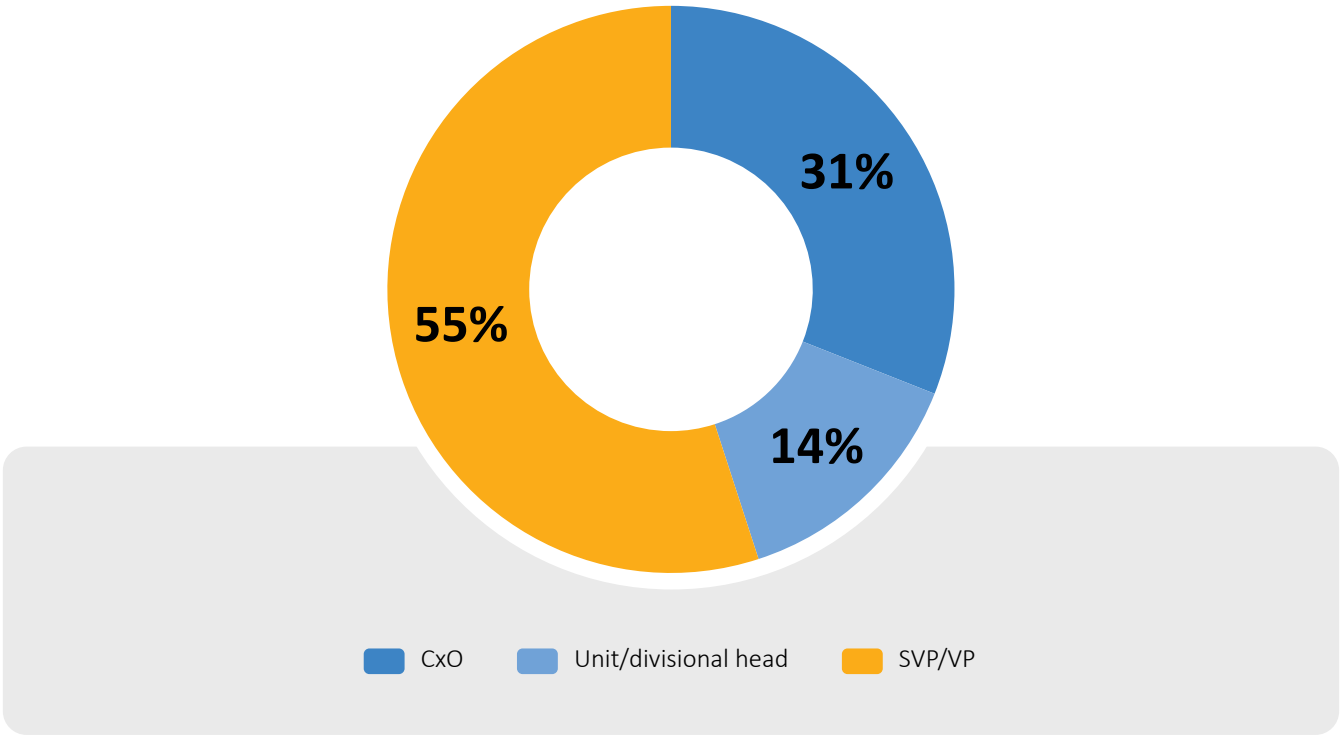


Study methodology continued

Based on annual revenue (in USD)



Sampling focused on senior leadership



Acknowledgements

Study authors



Amit Kapur
Chief AI and Services Transformation Officer, TCS



Ashok Krish
Vice-President and Head- AI Practice, TCS



Spandan Mahapatra
Vice President and Head of AI and Data, North America, TCS



Sankaranarayanan Viswanathan
Vice President and Head- Business Innovation, TCS



Naresh Mehta
Global Chief Technology Officer, Manufacturing, TCS



Priyanka Sethi
GTM Lead, AI and Data, TCS

The authors would like to thank the following contributors:
Saravana P, Binu Jacob, Sai Prasad Parameswaran, Gaurav Syal, Siddhartha Gupta, Sivaram Sivanarayanan, Nirmalya Shome, Srinivasan Rajesh, Murali R, Muralidharan Murugesan, Akash Bhadani, B Subramanian, Sanjeev Sachdeva, Geetha Thiagarajan, Ajay Nandgaonkar, Nitin Kumar, Saswati Collam, Sujatha Gopal, Amitava Das, Bhaskar Bhaumik, Babu Unnikrishnan, Prabhakaran Pitchandi, Bala Peddigari, Grivasan Venkataraman

About the TCS Thought Leadership Institute

Since 2009, the TCS Thought Leadership Institute has initiated conversations by and for executives to advance the purpose-driven enterprise. Led by Serge Perignon, the Thought Leadership Institute conducts primary research to deliver forward-looking and practical insights around key business issues to help organizations achieve long-term, sustainable growth. For more information, visit **tcs.com/insights/global-studies**

TCS Thought Leadership Institute contributors

Editor-in-chief: Serge Perignon
Editorial lead: Laura Rudolph
Data analysis lead: Derek Baker
Research & editorial team: Srilatha Bezgoan, Sanjay D'mello, Somil Bharti, Abhishek Thakur, Supratim Majumdar, Monisa Mukherjee, Clay Hamilton, Priya Prasad, Alice Angom, Kratika Jain, Sneha Gopinath, Michelle Beuscher, Galileo Tignini, Marianna Friscia

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services (BSE: 532540, NSE: TCS) is the technology partner of choice for industry-leading organizations worldwide. Since its inception in 1968, TCS has upheld the highest standards of innovation, engineering excellence, and customer service.

It has set an aspiration to become the world's largest AI-led technology services company and is enabling its clients to transform themselves across the full AI stack, from infrastructure to intelligence.

Rooted in the heritage of the Tata Group, TCS is focused on creating long term value for its clients, its investors, its employees, and the community at large. With a highly skilled workforce spread across 56 countries and 194 service delivery centers across the world, the company has been recognized as a top employer in six continents. With the ability to rapidly apply and scale new technologies, the company has built long-term partnerships with its clients. Many of these relationships have endured into decades and navigated every technology cycle, from mainframes in the 1970s to artificial intelligence today.

TCS sponsors 14 of the world's most prestigious marathons and endurance events, including the TCS New York City Marathon, TCS London Marathon, Tata Mumbai Marathon, and TCS Sydney Marathon with a focus on promoting health, sustainability, and community empowerment.

TCS generated consolidated revenues of over US $30 billion in the fiscal year ended March 31, 2026.

For more information, visit **www.tcs.com**

Follow TCS on **LinkedIn| Instagram | YouTube| X**

Content / information presented here is the exclusive property of Tata Consultancy Services Limited (TCS). The content / information contained here is correct at the time of publishing. No material from here may be copied, modified, reproduced, republished, uploaded, transmitted, posted or distributed in any form without prior written permission from TCS. Unauthorized use of the content / information appearing here may violate copyright, trademark and other applicable laws, and could result in criminal or civil penalties. Copyright © 2026 Tata Consultancy Services Limited