McKinsey & Company

The telco reinvention: How AI can fuel value creation

March 2026

Contents

Insights

Issue briefs

Insights

McKinsey & Company

Technology, Media & Telecommunications Practice

The critical bets on the future of telco value creation

After a challenging decade, telecom operators may finally be poised to spur a new era of growth. But it will require embracing fundamental change and making strategic moves in four key domains.

This article is a collaborative effort by Andrea Travasoni, Dev Patel, Theodor Vendrig, and Tiago Palma, with Alima Ziman and Marco Labianca, representing views from McKinsey's Technology, Media & Telecommunications Practice.

March 2026

Telecom operators (telcos) have played a critical role in laying the foundation for the massive, technology-driven economic and financial gains of the 21st century. Yet for most of that period, the industry hasn't managed to share in much of the spoils. While established tech titans and emerging start-ups alike were profiting from successive generations of wired and wireless infrastructure, the telcos that deployed and operated those networks found themselves weighed down by the crushing capital burden as well as intense new competition. That initial period from the late 2000s to 2015 put great pressure on the industry's revenues and margins, leaving it with an ROIC that fell below the cost of capital.

The industry has spent much of the last decade trying to spur a recovery, but its track record has been mixed. Despite making modest progress by embracing selective consolidation, more disciplined and agile operating models, and adjacent businesses or services, it's often seemed the best the industry could do over the past decade has been to tread water. The market value of telcos has done just that, as the persistent challenges of the prior five to ten years continued to take their toll. Regulatory constraints, a legacy business model, shifting customer expectations, and technological disruption have helped saddle telcos with both TSR and enterprise value multiples that have stayed flat at around 5 to 7 times for the past ten years, even as the tech sector's have continued to grow to 15 or higher.

However, as the industry turns the page on the first quarter of the 21st century there are some signs that telcos may finally be poised to embark on a new era of healthy growth. While the industry experienced notably low TSR growth over the last two decades—at 29 percent versus 235 percent for all sectors globally—it has stabilized, growing 28 percent since the start of 2024, firmly on par with the overall global market (Exhibit 1).[1]

At the same time, in most regions, the longstanding imbalance between revenue growth and investment growth is beginning to narrow. From 2019 to 2021, for instance, investment capital (IC) at North American telcos grew at a CAGR of 1.8 percent while revenues grew just 0.4 percent; by contrast, from 2021 to 2024, IC declined at a CAGR of 0.4 percent compared with revenues' annual decline of 1.2 percent. The industry is showing a more hesitant, wait-and-see perspective on capital outlays, fueled in part by the fact that many costly 5G and fiber rollouts are close to completion. In a survey of 152 telco executives last year, only a third said they expected investment capital growth to accelerate over the coming three to five years, a 50 percent decline from the prior year's outlook.[2]

Still, driving healthy, long-term growth for telcos will require more than taking a sharper approach to infrastructure investments or reaping productivity improvements. It will take the industry embracing fundamental change and making strategic moves in four key domains:

— **AI-native transformation.** Comprehensive, AI-driven reinvention of the end-to-end processes and operating model across all functions requires the capability to partner strategically with hyperscalers, system integrators, and solution and change management specialists. At the same time telcos will need to harness the technology and data to help automate workflows and break down legacy silos and fragmented systems.

— **Growth beyond the core.** Achieving the goal of healthy, long-term growth requires using AI to redefine consumer and enterprise value propositions through hyper-personalized, automated engagement, both with core offerings as well as an expanded, diverse portfolio of low-cost adjacent products and services.

[1] Industry TSR growth as reflected in the MSCI Total Shareholder Yield index.
[2] McKinsey Telco CxO Survey was conducted online in February 2025, with a sample of 152 respondents representing Africa, Asia, Australia, Europe, Latin America, the Middle East, North America, and Oceania.

Exhibit 1

After more than a decade of lagging most sectors, telcos' market performance has more recently begun to approach global averages.

MSCI total shareholder return,[1] index (Jan 2024 = 100)

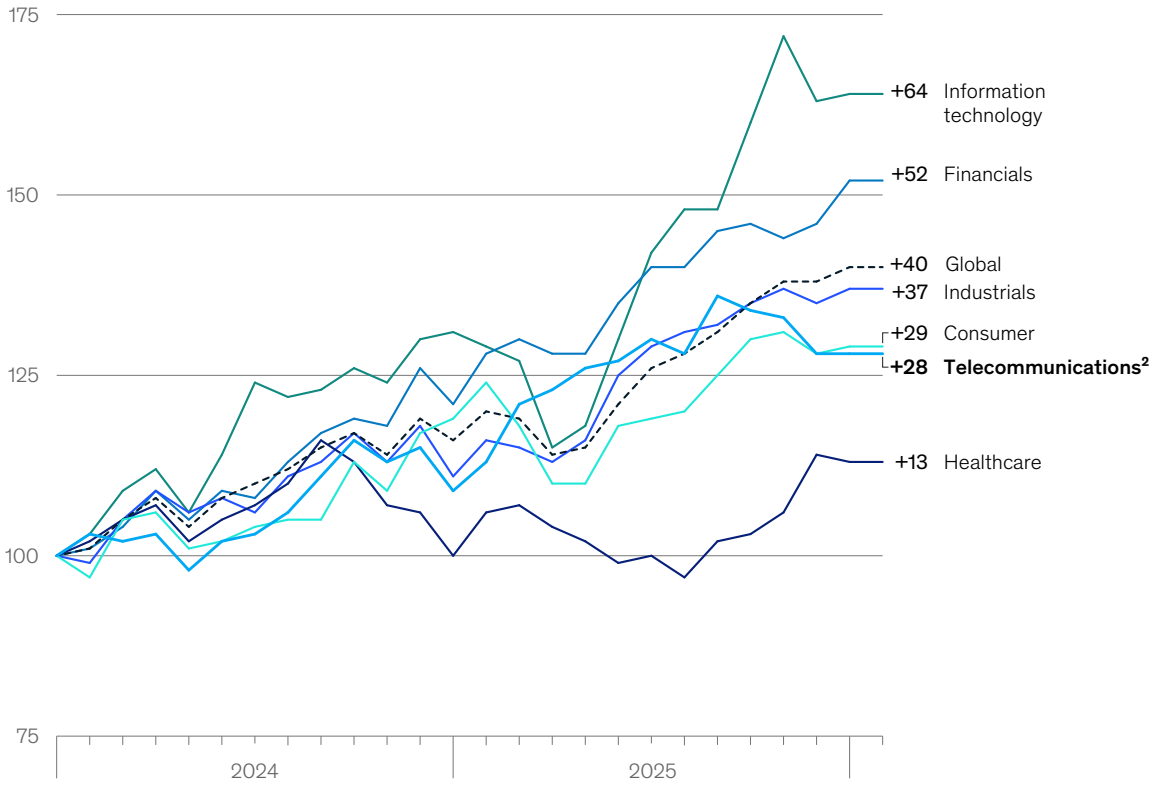

[1]Based on last sale price on Jan 30, 2026 of MSCI All Country World Index (ACWI) and sub-sector indices.
[2]Based on MSCI ACWI Select Telecommunication Services Screened 35/20 Capped Index ($).
Source: Capital IQ

McKinsey & Company

— **Infrastructure innovation.** As sovereign AI and data infrastructure become strategic differentiators in select markets and ownership models evolve to unlock capital and increase asset utilization, rethinking and upgrading legacy systems and technologies is critical.

— **Market shaping.** In an industry becoming increasingly commoditized, carefully navigating market dynamics and regulatory environments is more essential than ever; structural transformation such as the delayering of infrastructure or carveouts of non-core assets and continued consolidation remain critical, but only if pursued as part of a broader holistic strategy and transformation.

If telcos successfully execute a number of moves across these four domains, they could increase the industry's lagging ROIC rates and lay the foundation for renewed top- and bottom-line performance, which could in turn lead to more favorable market valuations. Over the last five years, fewer than one in five large telcos across the globe has achieved above-industry-average growth in both revenue and profits, and those were the only players rewarded with enterprise value (EV) growth, delivering around 11 percent annual EV growth on average.[3] Players that fail to make these significant shifts could face growing competitive pressure in an AI-driven environment. This article explores what it takes for telcos to avoid that fate.

Four strategic moves key for a new era of telco growth

The range of strategic moves telcos can make to reignite growth is broad, spanning four primary domains, but the pace of adoption varies widely throughout the industry, and not all moves carry the same importance (Exhibit 2).

AI-native transformation

Like many other industries, telcos haven't yet seen much top- or bottom-line impact from their experiments with gen AI across various parts of their organization. As with some of their previous technological initiatives, telcos have suffered from a siloed, piecemeal approach to incorporating gen AI or more recently agentic AI. They have tended to favor narrow use-case implementations or cost-reduction pilots over a comprehensive redesign of end-to-end domains, which have so far yielded little financial upside.

Industry leaders increasingly recognize that capturing AI's full potential requires a more holistic transformation, supported by two critical enablers.[4] First, data: access to a broad and integrated data foundation is essential to redesign processes end-to-end, and prevents data from becoming the source of a new generation of invisible silos. Second, the operating model: scaling AI requires cross-functional teams, rapid test-and-learn cycles, and continuous improvement embedded into day-to-day execution.

Though they remain relatively early in their AI journeys, with very few initiatives fully implemented, many telcos have started to make progress. Sixty-one percent of telco executives in a recent McKinsey gen AI–focused survey said they are focused on scaling gen AI use cases across functions and less on isolated pilots. And, 47 percent reported experiencing some impact from those efforts, up 20 percentage points from a couple of years ago.[5] That impact is limited so far, with AI driving less than two percent of total revenues for more than half the operators we surveyed that have tested AI use cases. But the industry remains optimistic about its potential. Fully 64 percent of leaders said they expected the technology to contribute more than five percent of revenues in the near future, and 40 percent anticipated AI-driven cost reductions to exceed 10 percent once the technology is fully scaled across their organizations.[6]

As they attempt to achieve those goals, telcos will need to adapt to the fact that data is the core competitive asset in an AI environment. That gives even greater urgency to the industry's pressing need to standardize its fragmented systems, break down silos, and make structured data easily accessible at all levels of the company. At the same time, standing out from the crowd could become even more challenging. The rise of AI-powered agents risks further commoditizing the telco market, with all players offering similar efficiency and service quality, a situation that could force telcos to be even more innovative and creative in communicating a distinctive value proposition.

[3] Based on a sample of 166 publicly traded, large telecom operators, including tower companies, across Africa, Asia, Europe, the Middle East, Latin America, and North America.
[4] These two enablers are part of *Rewired: The McKinsey Guide to Outcompeting in the Age of Digital and AI*, McKinsey's broader, comprehensive digital and AI transformation framework. The second edition of *Rewired* will be available April 14, 2026. For more information, see, Eric Lamarre, Kate Smaje, and Rodney Zemmel, "Rewired to Outcompete," *McKinsey Quarterly*, June 20, 2023.
[5] McKinsey Gen AI Telco CxO Survey, December 2025 (n = 49).
[6] McKinsey Telco CxO Survey, January 2026 (n = 125).

Exhibit 2

The strategic moves telcos can make to reignite growth vary by level of importance and implementation stage.

Telco value creation strategic bets by domain

Domain	Strategic bet	Overview	Relevance,[1] % of respondents	Implementation is ongoing or fully live,[2] % of respondents (0–100)
AI-native transformation	AI-driven cost efficiency	Triggering significant cost reduction by harnessing AI to redesign processes across all operations and support functions	58	
	Zero-touch operations	Using AI and automation to enable proactive, autonomous, and faster operations with minimal human intervention across the enterprise	34	
	AI-powered customer life cycle engagement	Driving highly tailored engagement across the entire customer journey, from brand building and conversion to retention and upsell	46	
	AI-driven value proposition	Delivering personalized experiences through intelligent, value-added services, such as vendor-agnostic AI assistants	48	
Growth beyond the core	Vertical ecosystem expansion	Expanding beyond core connectivity by bundling standard telecom services with adjacent offerings (eg, security, entertainment, insurance, energy)	33	
	ICT[3] products and services	Becoming trusted digital transformation partners with end-to-end enterprise tech solutions that span cloud, cybersecurity, and IT managed services	30	
	Network APIs	Fostering the development of new 5G-related products and services by working together to roll out standardized, global API platforms	18	
	Private 5G networks	Monetizing 5G infrastructure through vertical-specific enterprise offerings in such areas as manufacturing, logistics, and public infrastructure	26	
Infrastructure innovation	AI infrastructure	Positioning telcos as trusted providers of AI-capable infrastructure (including sovereign AI) to support national priorities and enable enterprise-scale AI deployment	55	
	Network excellence	Incorporating AI and automation to provide intelligent and resilient networks as network equality reemerges as a key differentiator	40	
	Infrastructure and IT modernization	Updating legacy networks and IT systems with cloud-native, real-time architectures to reduce complexity, cut costs, and unlock AI-enabled service innovation at scale	46	
Market shaping	Delayering	Carving out high-potential telco units into independent entities to enable specialized governance, faster scaling, and stronger performance	6	
	Strategic M&A	Focusing on industry consolidation to benefit from enhanced economies of scale and value proposition differentiation	58	

[1]Question: Which of the strategic bets listed are most relevant to you (please select the 5 most relevant)?
[2]Question: How far is your company in the implementation journey for your selected 5 most relevant bets? Please indicate level and status of implementation. Implementation figures represent the share of respondents who identified the given strategic bet as one of their 5 most relevant and said their implementations are either "ongoing" or "fully live."
[3]Information and communications technology.
Source: McKinsey Telco CxO Survey, Jan 2026 (n = 125)

McKinsey & Company

Overcoming such risks and succeeding at an AI-native transformation will, above all, require telcos to embrace foundational change in the following areas:

— **AI-driven cost efficiency.** Arguably one of the most foundational elements of an AI-native transformation is to drive significant cost reduction by harnessing AI to redesign processes across all operations and support functions. It makes sense then that this area attracted the most consensus in our most recent telco leader survey, with close to 60 percent of executives choosing it as one of their top strategic priorities.[7] Telstra, the leading Australian telco, achieved AU $122 million in cost savings in just the first year of its program to embed AI across all its corporate processes and business workflows. As part of what it dubbed Project T25, Telstra shifted from relying on a roster of 18 different tech vendors to establishing strategic joint ventures with just two, improving capability integration and governance in the process. Another telco developed a gen-AI powered copilot for its financial planning and analytics team that could handle much of the previously manual and time-consuming data collection and research activities, allowing employees to focus on higher value, more strategic analyses, and saving several millions in operating expenses.

— **Zero-touch operations.** Using AI and automation to minimize human intervention is another key piece of the puzzle. Reaching this goal of proactive, autonomous operations has the potential to radically reshape back-office and support functions, but it is no small feat. It requires, among many key enablers, cross-domain data flows that seamlessly link IT, network, and operations with real-time, cross-domain data flows. BT's procurement unit (BT Sourced) uses AI sourcing and analytics tools to automate and quicken tendering, supplier onboarding, and spend analysis, improving decision-making in the process.

— **AI-powered customer life cycle engagement.** Providing tailored engagement for both B2C and B2B customers is increasingly table stakes for any successful telco, and leveraging AI is essential to reaching that goal. If telcos can incorporate the technology into every stage of the customer journey, from brand building and conversion to retention and upsell, they can spur sizeable improvements in revenue (5 to 8 percent), upsell (50 percent), and churn reduction (30 percent). One telco call center struggling with periods of undercapacity and overcapacity, for example, was able to shave its customer wait times from 3 minutes to 30 seconds with no budget increase by using AI to better predict call traffic patterns and adjust staffing as a result. Another telco has used gen AI tools to cut by 70 percent the time it takes to develop and activate a new personalized marketing campaign or product offering.

— **AI-driven value proposition.** A key strategic bet that forms the core of an AI-native transformation is tapping the technology to provide hyper-personalized experiences and services. Such innovative offerings, which in some cases have shown the potential to increase B2C ARPU by as much as 5 to 10 percent, could range from vendor-agnostic AI assistants that help users deal with fragmented LLMs, to localized edge hosting of LLMs closer to the end user. Several telcos have already gone down this route: SK Telekom, for instance, rolled out an AI-powered personal assistant as an integrated part of its mobile plans that helps in such areas as shopping and travel.

This type of bold, multifaceted effort depends on several factors that can make the difference between realizing ambitions and encountering significant setbacks. These key enablers include:

— **Change management.** AI implementation ultimately comes down to helping a group of people, often after years of ingrained routines, to fundamentally reshape how they work day to day. That seismic shift requires time, committed and visible leadership, a comprehensive redesign of core processes and incentives, and a significant investment in upskilling. In a recent gen AI–focused survey of telco top executives, more than

[7] McKinsey Telco CxO Survey, January 2026 (n = 125).

three-quarters identified slow adoption because of weak change management as the biggest obstacle to scaling AI impact.[8]

— **AI-ready systems.** These include systems that integrate and structure data from customer journeys in real time. They also include modular gen AI tools for content creation and targeting, and cross-functional, agile teams to enable personalization sprints and rapidly roll out and refine new AI campaigns.

— **Cross-functional product squads.** These include teams from the product, engineering, and go-to-market functions, and edge-ready infrastructure that can host LLMs locally, securely, and cost-effectively.

— **Ability to form, manage, and scale strategic partnerships.** These need to be made with a broader set of players complementing telcos in the AI and tech ecosystem, including hyperscalers, system integrators, specialist solution players, digital infrastructure specialists, and impact partners.

In those markets where telcos happen to be part of larger diversified conglomerates (such as those in many parts of Asia), building these enablers often presents platforming opportunities to commercialize capabilities both internally and externally.

Growth beyond the core

As traditional connectivity markets mature and competitive intensity continues to rise, telcos face mounting pressure to unlock new sources of sustainable growth. Even data, once the industry's primary value creation engine, is increasingly commoditizing. Amid this challenging environment, telcos seeking a new commercial growth path are already testing and scaling new business models that extend well beyond core connectivity.

Telcos happen to be well-positioned for this pursuit. The sector's position at the center of the digital economy affords it real-time access to data generated across networks, devices, locations, and usage patterns, and at a scale and granularity matched by few other sectors. This unique data advantage, combined with trusted customer relationships and ubiquitous infrastructure, enables telcos to shift from selling discrete products to delivering integrated, value-added services.

By responsibly and securely leveraging network and customer insights, operators can create differentiated, customized offerings in both B2C and B2B markets. These range from personalized consumer experiences (including energy, content, insurance, security, omnichannel retail and other offerings) built around a vertical ecosystem to industryspecific enterprise solutions that enhance security, efficiency, and performance. Leading players are already translating this shift into action by expanding into adjacent digital verticals, end-to-end ICT services, network APIs, and private 5G networks. In the process, they are converting experimentation into scalable growth, unlocking new revenue pools, and building a more resilient, long-term value creation model.

— **Vertical ecosystems expansion.** Expanding beyond core connectivity to offer adjacent services that unlock new revenue pools is a critical component for telcos looking to drive a new era of healthy growth. These offerings could revolve around everything from security and entertainment to insurance, energy, and fintech, tapping the industry's reach and infrastructure to dramatically expand the scope of customer relationships. Forward-thinking telcos are already leading the way in this area, shifting from a mindset of selling products into an ecosystem of varied services that can widen the value proposition. Canada's Telus now gets a quarter of its revenues from B2B digital verticals for healthcare, agriculture, digital process and back office solutions, part of a diversification strategy that has helped the company's TSR outperform the country's broader telco sector.

[8] McKinsey Gen AI Telco CxO Survey, December 2025 (n = 49).

Several European telcos are adding additional services to their portfolio. Spanish operator MasOrange, for instance, has focused on expanding into adjacent categories such as energy, insurance, security, and health/telemedicine to serve its customer base with value-added services and to help increase ARPU and reduce churn, all while reinforcing its core connectivity value proposition. By leveraging a dynamic, TechCo operating model of independent units and agile platforms and scaling through partnerships and other flexible collaboration models, MasOrange has built a number of growing new businesses. These have included an app-controlled electricity-and-gas service, which it sold last year for 90 million euros; an insurance business (mobile, home, health, and payment protection) with a cumulative ambition of more than 7.5 million policies and over €1.5 billion in premium volume; a cobranded home security business (with ADT) that has reduced churn; and B2B focused offerings, including a cybersecurity and IoT offering that are supporting ARPU uplift and profitability through higher-value enterprise services. Italy's WindTre has enjoyed similar initial success by rolling out insurance and energy products, using upskilling and realigned incentives to enable its store clerks to help drive interest and sales of a significant volume of policies or agreements.

— **ICT products and services.** As more enterprises seek out third-party support in optimizing their rapidly evolving patchwork of technologies and systems, telcos have a lucrative opportunity to play a bigger role in digital transformation. By expanding from core B2B offerings into end-to-end solutions that span cloud, cybersecurity, and IT managed services, the industry can grow B2B ICT into a dynamic business that accounts for as much as 20 percent of total telco revenue. One European telco used that model to grow its share of the domestic cloud market, upgrading its infrastructure, upskilling employees, and striking key partnership on the way to more than doubling its cloud-related revenues and earnings in four years.

— **Network API.** The effective monetization of 5G has been a significant challenge for telcos since the wireless technology's debut several years ago. By working more closely together to roll out standardized, global API platforms, the industry can spur the creation of new 5G-related products and services that leverage features such as speed on demand, low-latency connections, speed tiering, and edge compute discovery. A trio of Brazilian operators recently partnered on an API initiative to help financial institutions improve digital security and antifinancial fraud efforts. Working with the industry's GSMA Open Gateway Initiative, TIM, Claro, and Vivo developed three customized APIs focused on continuous mobile number verification, SIM card swap checking (to prevent account takeover attacks), and instant device location and validation. The services are already reaching close to 150 million customers and early results have been promising, not only generating $9 million in incremental revenue from the SimSwap API alone but also delivering tangible consumer benefits.

— **Private 5G networks.** Another vehicle for increasing telcos' return on 5G is customized, vertical-specific enterprise architecture to enable advanced use cases. Many operators (including T-Mobile, O2, Vodafone, and KPN) are already launching such solutions, from smart hospitals or connected transportation to augmented retail experiences. By targeting areas such as logistics, manufacturing, and the public sector, such specialized networks could fuel incremental 5G monetization opportunities over time.

Achieving true commercial excellence requires a significant appetite for change as well as a number of specific enabling factors, including the following:

— embracing a new type of partner ecosystem and striking deals with vertical leaders to help drive the development of new offerings beyond the core; create an internal integration platform to ease the process of bundling, activation, and management of non-core services

— laying the operational and technical foundation to support ICT-driven partnerships, especially alliances with hyperscalers and provides of cybersecurity and software; that means developing a modular platform to manage full-stack solutions as well as a specialized salesforce skilled in catering to demanding enterprises and the integrated delivery and support systems that those customers expect

— establishing common application layer standards to minimize developer friction, and integrating network APIs with hyperscaler platforms to help expand the reach of new offerings

— obtaining reliable, enterprise-grade wireless spectrum access to back up private network service guarantees and design scalable solutions, often in partnership with industrial players or hyperscalers, and which combine hardware, software, and services targeting use cases in specific sectors or verticals

Infrastructure innovation

As AI reshapes economic competitiveness, digital infrastructure is becoming an even more valuable strategic asset. For telecom operators, this shift opens a path to move beyond the role of capital-intensive utility toward a more differentiated position in the digital economy. The combination of rising AI workloads, growing concerns around network resilience and security, and renewed emphasis on performance at the network edge is creating a set of concrete, near-term opportunities for telcos, particularly in areas where their assets, capabilities, and trusted role intersect with emerging demand.

Governments and enterprises are increasingly looking for secure, local, and reliable infrastructure to support AI-driven use cases, while at the same time expecting networks to deliver higher performance, greater reliability, and lower unit costs. Capturing this opportunity will require operators to think holistically about infrastructure, both their own internal systems and the wider external environment in which they compete. They'll have to focus not just on what they build, but how they operate it, modernize it, and turn it into a true source of differentiation and returns. The key areas include:

— **AI Infrastructure:** With AI fast becoming a private- and public-sector imperative, telcos have an opportunity to play a valuable role as a reliable provider of the enabling infrastructure, from cloud and edge to secure connectivity. This is particularly the case in the growing number of markets where governments (and the businesses they regulate) are looking to increase their technological resilience, or sovereignty, often backed by national AI or data center strategies that include mandates and sizeable public funding. Industry leaders recognize this opportunity as one of their most promising, with 55 percent of those we surveyed identifying AI infrastructure as a top strategic bet.[9] Numerous telcos are already actively exploring this potentially fertile ground, lured by the prospect of increasing their return on invested capital by as much as 50 percent and seizing a new chance to compete with established hyperscalers. In practice, telcos' advantage relative to hyperscalers lies less in scale of compute or model development and more in trusted local presence, regulatory alignment, and control of national infrastructure assets such as edge locations, secure connectivity, and data-center footprint. As a result, sovereign AI initiatives are most often pursued through partnership-based models, with telcos providing the infrastructure and compliance layer, and hyperscalers contributing cloud platforms, GPU-based compute, and foundation models. From Spain (Telefonica) and the UK (BT) to Saudi Arabia (STC) and Southeast Asia (Singtel), telcos are planning or launching AI-ready networks, data center platforms, and other cloud or edge-based products to serve government and enterprise customers seeking out secure, locally-housed (aka sovereign) AI infrastructure with models that are trained by local teams on local datasets.

— **Network excellence:** As AI workloads and digital services move closer to the edge, network quality is reemerging as a major differentiator. By incorporating AI and automation across the entire life cycle—from planning, construction, and deployment to fault management and inventory—operators can provide intelligent, high-performing, and resilient networks that be a key driver of value creation rather than just a technical asset. In certain instances, such an approach has shown the potential to reduce network operating expenses by up to 30 percent and total capital expenses by as much as 15 percent. For instance,

[9] McKinsey Telco CxO Survey, January 2026 (n = 125).

one North American telco used AI tools to improve the safety, efficiency, and customer satisfaction of its frontline technician force by unlocking 10 to 20 percent additional capacity and enabling up to a 30 percent reduction in overall field-force requirements.

— **Infrastructure and IT modernization:** Monetizing the past decade-plus of capital investments into 5G, fiber, and other technologies is as much about updating operators' legacy network infrastructure and IT systems as it is about devising new products and services. If operators can't successfully upgrade their own internal technology capabilities, they are much less likely to be able to effectively take advantage of cutting-edge innovations like AI, or to market new offerings tied to the most advanced network solutions. By embracing fiber rollouts, automation, and cloud-ready architecture that is both modular and API driven, forward-looking telcos can lay a foundation that could reduce IT costs by as much as 30 percent and boost ARPU by 10 to 15 percent; previous McKinsey research has shown that operators in the top quartile of IT maturity, on average, grow revenues three percentage points more year over year compared with their peers, and they do so at a higher rate of profitability (roughly 15.5 percent net operating profit after tax versus 14 percent).[10] Several European operators (such as Norway's Telenor and Spain's Telefonica) have already seen encouraging results from their decommissioning of copper networks, including double-digit decreases in energy consumption and fault rates, and 7 to 14 percent increases in customer satisfaction.

Successfully leveraging infrastructure as a value creation opportunity requires a number of key capabilities and approaches, including the following:

— **Compute power.** To seize the AI infrastructure opportunity, telcos will need to ensure they have sufficient access to compute capacity, either through GPUaaS or necloud providers, or through partnerships with hyperscalers. Their infrastructure will also have to be both AI- and edge-ready, able to satisfy power, cooling, latency, and location requirements of such advanced workloads. Maximizing the shared hardware will also be critical, such that it can simultaneously run more traditional network functions along with the AI inference demands via AI-RAN technology.

— **Unified data flow.** Optimizing the telco network can't happen without a comprehensive, unified data flow to fuel automated analytics. To achieve that level of transparency and understanding, operators must aggregate real-time telemetry from all the different network domains (RAN, core, transport). An autonomous, AI-based platform is equally important, enabling operations including planning, fault detection, incident resolution, and energy optimization with little to no human involvement. At the same time, a fully AI-literate network workforce of engineers and planners is essential for integrating AI tools into operations at the outset and laying the foundation for increased automation.

— **Planning and oversight.** A successful infrastructure and IT modernization campaign takes more than just an embrace of more advanced technologies. On both internal and external fronts, careful planning and oversight is needed for such significant changes. That means developing thorough, end-to-end technical and operational plans for the organization's journey from legacy to modern systems, as well as proactive, comprehensive management of customer migrations, including targeted outreach, guided onboarding, and incentive models.

Market shaping

Over the past two decades, telecom operators have faced a steadily tightening set of constraints. Capital intensity remains high, pricing pressure is persistent, and returns have struggled to keep pace with those in adjacent digital sectors. Investors, meanwhile, have become less bought in to complex, vertically integrated models that can

[10] "Transforming telecom tech: How IT excellence drives innovation and cost efficiency," McKinsey, February 28, 2025.

be difficult to manage, slow to adapt, and heavily influenced by evolving regulatory environments. All of this is playing out as customers' expectations for reliability, speed, and seamless digital experiences continue to rise.

Against this backdrop, many operators are recognizing that incremental improvement alone is unlikely to change the trajectory of the increasingly commoditized industry.

Traditional levers—cost discipline, selective network investment, and product refreshes—remain necessary but insufficient to the task at hand. More leadership teams are stepping back to ask more fundamental questions about structure, focus, and scale: which parts of the business truly benefit from being integrated? where can specialization unlock value? and how can capital be deployed more effectively?

Two market-oriented strategic responses are emerging most clearly from this reassessment: delayering and strategic M&A. Each represents a distinct way of addressing structural challenges in the telco value chain, and each has the potential to materially improve performance. At the same time, both require significant organizational change and management attention, making clear strategic intent and disciplined execution essential. Two actions in particular are required:

— **Delayering.** The challenges telcos have faced in the past decade or two, particularly with investors, have convinced some to make a dramatic break from the traditional integrated ownership model. To attempt to improve focus, innovation, and valuations, a small but growing number of operators is choosing to delayer or separate into two or more independent corporate entities, typically one built around network and infrastructure ("NetCo" or "InfraCo") and the other(s) centered on products and services ("ServCo"). While the "NetCo" can take advantage of its new status as a carrier-neutral wholesale access provider, the "ServCo" can further specialize into distinct units based on type of customer, offering, business model, investment horizon, or skills. In many cases, the establishment of a separate NetCo helps to spur consolidation in the telecom infrastructure market and greater economies of scale.

  The rewards of such a radical step can be significant—some operators that have taken it have realized as much as a ten percentage point uplift in EBITDA margins and a ten times increase in enterprise value multiple, but such positive impact is far from guaranteed. Delayering can take as much as two or three years and sizeable amounts of capital to pull off, and some industry observers think such a split can hamper each separate entity's ability to present a seamless, omnichannel customer experience, and hence requires significant change management. Still, several carriers have already experienced positive returns through delayering: one of the earliest such moves about a decade ago, by Czech carrier O2/CETIN, produced a 27 percent total market cap increase just one year after the split. More recently, carriers such as Etisalat (based in the UAE), and MTN (based in South Africa) have found success delayering into four or five different separate units or platforms.

— **Strategic M&A.** The more traditional, market-driven maneuver that telcos have looked to as a possible solution to their recent woes is consolidation. Such an approach can be successful in harnessing economies of scale, value proposition differentiation, and access to new markets, with the potential to achieve synergies amounting to more than 15 percent of combined capital expenditure and 30 percent of operating expenditure (and it can be especially useful in in-country consolidation plays). The 2024 combination of Orange Spain and MasMovil is expected to realize synergies of €490 million after four years, with the new company serving nearly 50 percent of the country's fixed and mobile subscribers and close to a third of the industry revenues; likewise, the 2025 pairing of Vodafone UK and Three UK will create the country's largest mobile operator, serving nearly 40 percent of subscribers as well as telco revenues, with an anticipated £700 million in synergies by the fifth year. More industry leaders are contemplating such moves; close to 60 percent of respondents to our most recent survey in January identified M&A as one of their top five strategic priorities going forward, a level of consensus that was only matched by AI-driven

cost efficiency.[11] But combining forces is by no means a panacea for value creation. Unless it's carried out thoughtfully, with significant cost transformation and disciplined, post-merger execution, it can be a damaging waste of time and resources.

Market-driven organizational and strategic moves that generate sustainable, incremental value are easier said than done; reaching that ambitious goal takes many key factors, such as the following:

— **Operational clarity.** Delayering a traditionally integrated operator in such a way that the separated parts are greater than the whole requires a number of key enablers: the clear delineation of the different assets and activities, including the decoupling of underlying IT systems, taking into account value pools, synergies, interdependencies and compliance; as part of this, it's important to carefully draft commercial and service-level agreements between the newly distinct entities, the NetCo and ServCo, to maintain adequate continuity; and finally, the holistic redesign of each entity's operating models and processes to spur focused execution.

— **Clear assets structure.** To increase the odds that strategic M&A will ultimately be viewed positively in hindsight, operators will want to have their assets in question structured with clear financial, legal, and operational separation between business units, which help to boost valuations and simplify deal-making. It's equally important to be organizationally prepared for the aftermath of any transaction; appropriate teams and tools should already be in place to speed post-deal integration and reduce potential disruption.

Finding a place within an AI-driven economy

As telecommunications leaders contemplate their industry's place in the emerging AI-driven economy, they have valid reasons to feel cautiously optimistic about their ability to successfully reinvent themselves for this new era. Recent improvements in financial performance, more disciplined capital allocation, the near completion of costly network buildouts, and customers' evolving expectations have created a valuable opportunity for the sector to move beyond the challenging past two decades into a healthy era of sustainable growth.

There is, however, no single playbook or prescriptive formula for turning that opportunity into significant enterprise value. Market structures, competitive dynamics, starting positions, and policy regimes, among other factors, mean that successful telcos will invariably end up taking distinct paths, and the scale and pace of impact will vary significantly.

What the strongest performing operators have already shown is that they share a common mindset, if not a common strategy. Rather than betting on one narrow initiative, they have made progress across several of the critical broad strategic areas outlined in this article. Some have focused more on embedding AI at the core of their operating models or redefining commercial engagement, others on innovating with infrastructure or actively shaping markets rather than reacting to them. Companies such as MasOrange, Singtel, Indosat, Telus, and Deutsche Telekom illustrate this diversity: each has emphasized a combination of strategic moves and approaches, yet all have shaped transformation with a portfolio of actions rather than a single lever.

[11] McKinsey Telco CxO Survey, January 2026 (n = 125).

Regardless of the specific path, any brand of large-scale telco reinvention depends on a small number of critical enablers. This starts with visible, sustained leadership from the top—setting a clear strategic direction and aligning the organization behind it. Equally important is disciplined investment in data and modular IT, which form the foundation for multiple strategic bets. And finally, it requires an execution-driven culture that favors experimentation and learning by doing, with leaders willing to move from plans to action quickly and make course corrections as warranted.

For telco leaders, the overarching lesson is straightforward. Value creation may not require following a uniform template, but it does demand focus, coherence, and the willingness to act across multiple fronts. Telcos that launch a handful of bold, context specific moves have a chance to materially improve returns and regain investor confidence; in the process, they can capture more of the value their networks enable and reestablish themselves as growth platforms at the heart of the AI economy. Those that hesitate, or pursue siloed initiatives without a broader strategy, risk repeating their recent, growth-challenged history, relegated to low-return, utility-like roles as the pace of dynamic, AI-enabled change accelerates across every industry.

Andrea Travasoni is a senior partner in McKinsey's Milan office, where **Marco Labianca** is an associate partner. **Dev Patel** is a senior partner in the Chicago office, **Theodor Vendrig** is a partner in the Oslo office, and **Tiago Palma** is a partner in the Madrid office. **Alima Ziman** is an associate partner in the Almaty office.

The authors wish to thank Azamat Bakytbekov and Blanca Alcala for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

Telcos' AI inflection point: What leaders do to capture value

Agentic AI offers telcos a potential opportunity to reorient their recent trajectory. The question is whether they can redesign their organizations enough to make the most of it.

This article is a collaborative effort by Benjamim Vieira, Guilherme Cruz, Ignacio Ferrero, and Tomás Lajous, with Borja Belda, Daniela Mendoza, and Manuel Palacios, representing views from McKinsey's Technology, Media & Telecommunications Practice and QuantumBlack, AI by McKinsey.

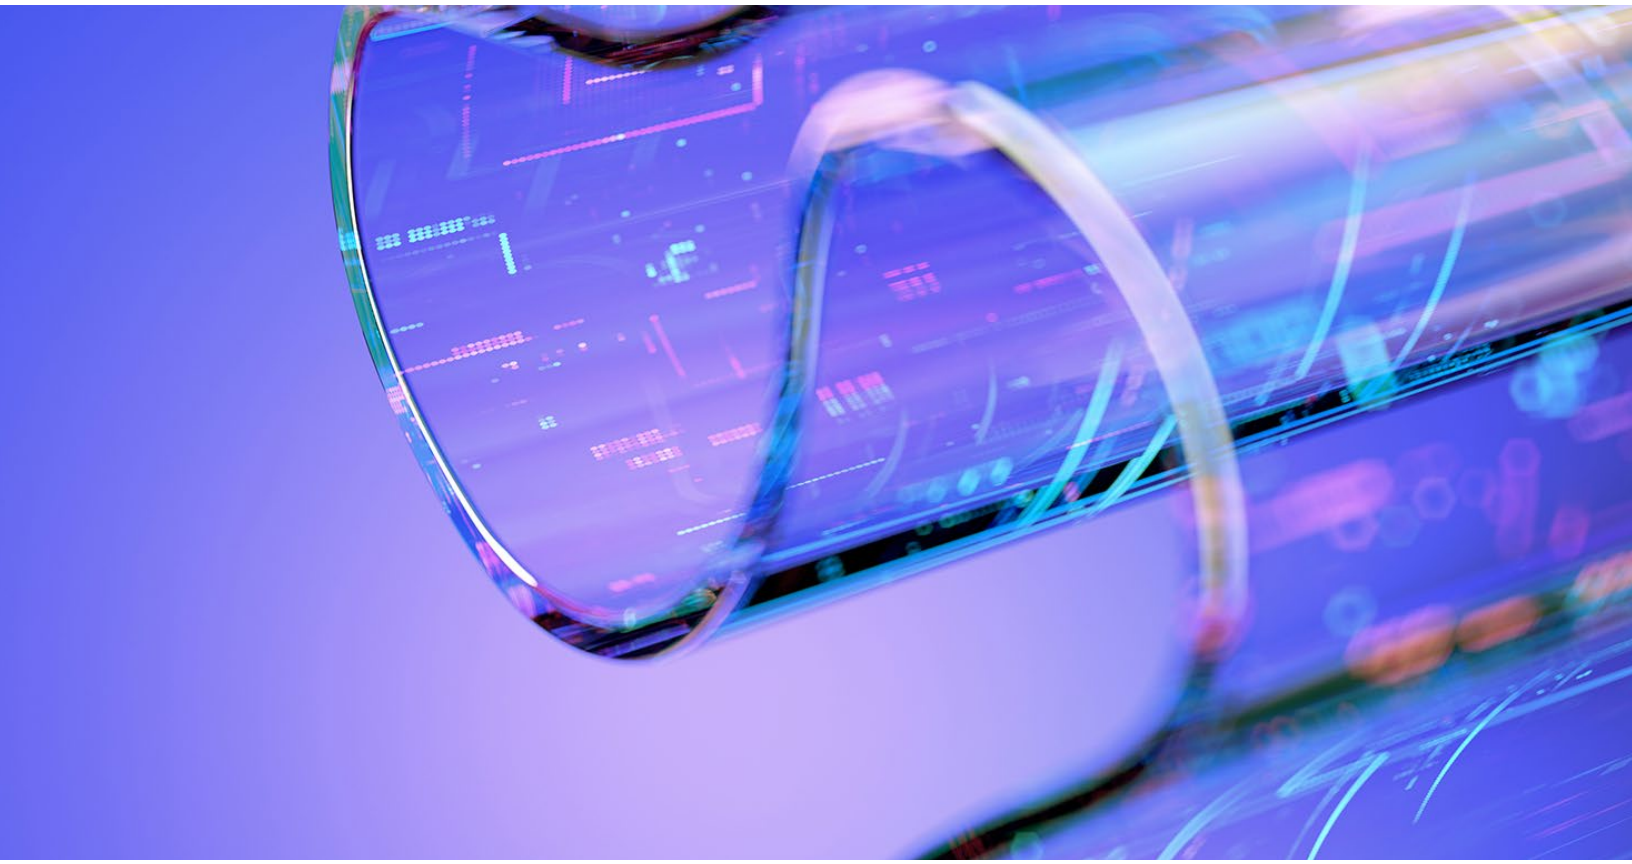

March 2026

For more than a decade, telecom operators (telcos) have pursued successive waves of digital transformation to offset slowing growth, rising capital intensity, and limited product differentiation. Most delivered incremental gains, but not the type of overarching structural change that the industry requires at a time when the pace of innovation, competition, and challenging secular issues show no signs of slowing down.

Agentic AI may be the first technology capable of altering that trajectory. Unlike earlier automation tools that improved individual tasks, AI agents have the potential to reshape entire workflows, even helping make operational decisions and coordinate work across functions. This shifts AI from being a powerful productivity tool to a full-scale execution layer that can fundamentally change how telcos design operations, deploy capital, and create value.

The question for telco leaders is no longer so much about where exactly to apply AI, but how to redesign the enterprise to operate with and optimize it.

Over the past few years, much of the sector has learned the hard way that there is no quick way to make that happen. But even as most telcos have failed to gain real value from their growing AI investments, rolling out too many fragmented or duplicative AI use cases and pilots, a smaller number of forward-thinking players are showing how the rapid emergence of agentic AI can be a paradigm shift, offering a rare opportunity to reorient the industry's prospects.

The telco leaders in harnessing AI to generate significant impact have already demonstrated it requires a sustained, long-term approach that includes CEO-led sponsorship, disciplined organizational transformation and change management, and a clear focus on end-to-end processes rather than individual tasks. This article, based on a recent survey[1] of top telco executives and our experience working with clients on AI deployments, provides industry peers with a road map to similar success, with the possibility to increase both ROIC and EBITDA margins by as much as ten percentage points within five years.[2] It lays out several key elements of a coherent agentic AI operating model that is necessary for telcos to capture that scale of new value and growth over the coming years and highlights the current state of telco AI adoption and impact for both leading, early-moving telcos and the rest of the industry.

The state of AI for telcos

While telco executives increasingly recognize the importance of AI to their future, relatively few have seen sizable impact from their embrace of the technology. And that embrace has grown significantly in recent years.

Just over half of the telcos we surveyed have at least 50 full time equivalents (FTEs) dedicated to AI, and fully two-thirds expect to increase the AI portion of their IT budgets this year; roughly half of that segment anticipates devoting more than 10 percent of that spending to AI, with the average industry allocation growing to 9 percent (Exhibit 1).

[1] The McKinsey Telco CxO Gen AI Survey was conducted in December 2025, with 49 respondents representing telecom operators from across North America, South America, Europe, Africa, the Middle East, and Asia.
[2] Depending on baseline performance and implementation scope.

Exhibit 1

Telcos are allocating more personnel and boosting IT budgets for AI to speed up deployment and begin to capture significant impact.

Telco full-time equivalents dedicated to AI development and implementation, % of respondents

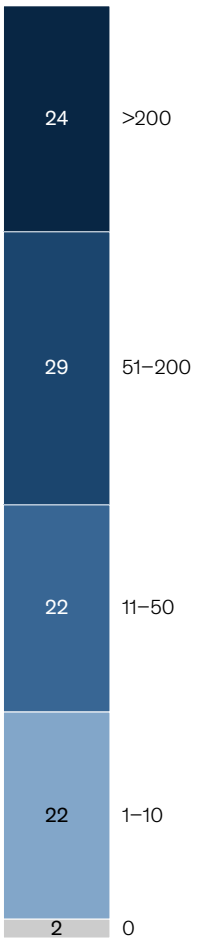

Share of telcos' IT budgets allocated to AI, % of respondents

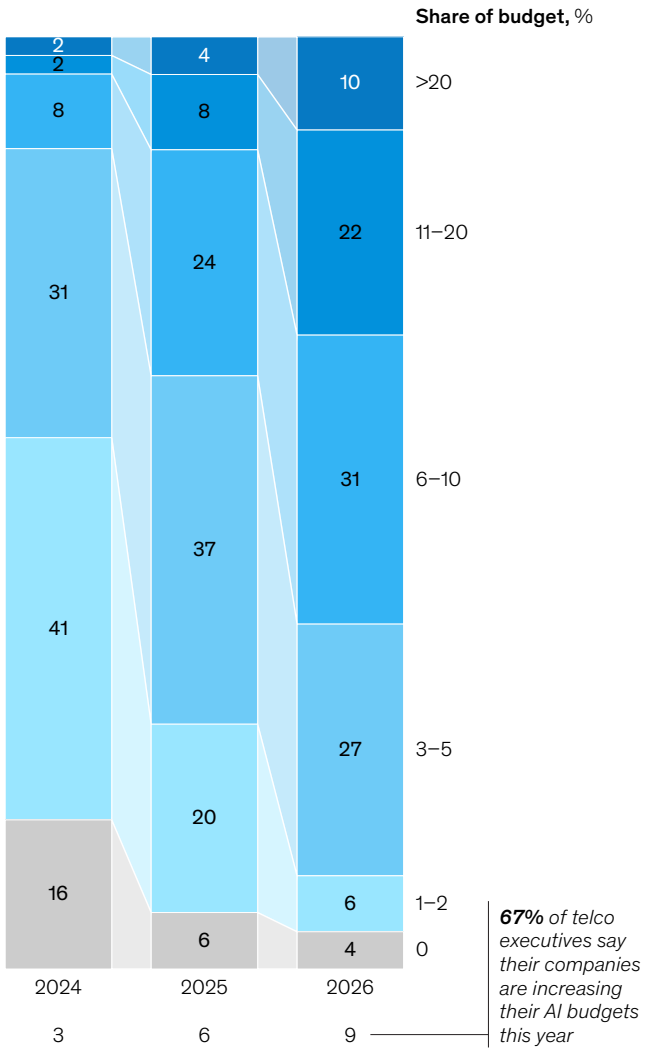

Note: Figures may not sum to 100%, because of rounding.
[1]To calculate these weighted averages, we took the midpoint of the budget % ranges (0, 1–2, 3–5, 6–10, 11–20, and >20), multiplied each by the corresponding % of respondents in the category, and summed the multiplied numbers.
Source: McKinsey Telco CxO Gen AI Survey, Dec 2025, n = 49

McKinsey & Company

Though virtually all telcos we surveyed are piloting AI, only 57 percent report scaling use cases across multiple domains, led by customer service and networks, virtually unchanged from almost a year ago. Far fewer—just 16 percent—characterize AI as the "new normal" across their organization. At the same time, only 51 percent view AI as a "blockbuster technology" that will fully transform the industry, compared to 61 percent in the prior survey.[3] Most of that shift in sentiment was toward a more sober view that the technology is "something relevant that will generate impact in the industry," which is now shared by 47 percent of respondents, up from 36 percent the previous year.

This more realistic, but still bullish, mindset, which is focused on tangible results rather than the technology's seemingly limitless capabilities, reflects the industry's initial experiences with incorporating AI. While telcos have seen meaningful cost savings in certain areas, most have yet to realize widespread productivity gains. Only 12 percent of respondents report having already captured sizable impact, with most others caught in what we call "the money step," where investments have yet to yield any real balance sheet benefits despite spurring other impacts. A majority of industry leaders, however, remain cautiously optimistic that those benefits will materialize toward the end of the decade; they view the most fertile areas as customer support and network functions, followed closely by IT, where most telcos expect to enjoy at least 10 percent cost savings over the next one to two years and close to 30 percent by 2030 (Exhibit 2).

Despite that relatively rosy long-term view, telco executives acknowledge the substantial challenges they face in scaling the impact of AI, such as moving beyond limited productivity gains from isolated use cases. More than three-quarters of those surveyed agree that immature operating models, data limitations, and lagging adoption due to ineffective change management are the top issues. When it comes to translating impact into measurable, substantial value, there is less consensus (Exhibit 3). A little over half of respondents identify employee or team adoption as the biggest hurdles, while a third point to the difficulty of implementing process changes at scale, and a quarter noted inflexible budgeting as the key obstacle.

Telcos' AI inflection point

A critical reason for telcos' relative bullishness about overcoming those challenges is the recent emergence of agentic AI. Unlike previous AI tools, which have largely enabled impressive but still incremental improvements, agents represent the possibility of an automation-driven inflection point for industry. By 2030, according to the McKinsey Global Institute, AI automation could create as much $16 billion in new economic value for the telecom sector in the United States alone; and redesigning workflows around human-agent collaboration, whether for sales reps, equipment installers, or electronic engineers, may be key to capturing it. Roughly half of the industry's workforce may have the potential to be remade with collaborative human-agent automation, more than double most other sectors. That prospect of using agents to redefine (or re-create) workflows means the industry may be able to move beyond the same outdated operational processes and models that have proved to be stumbling blocks to unlocking new value.

While much of the industry discussion around agents a year ago was largely theoretical, more than half of telco survey respondents now say they are deploying agentic use cases across at least one function. Several early movers are already showing how the technology can be leveraged to capture sizable value, as the following examples across different functions illustrate.

Customer service

Across European telecom operators, rising customer service volumes and high call center costs have made improving efficiency in customer care a strategic priority.

[3] McKinsey Telco CxO Gen AI Survey, February 2025 (n = 52).

Exhibit 2

Telcos expect AI to help generate much greater cost savings by 2030, particularly in customer service, network, and IT domains.

Telco estimated cost reduction attributed to AI, by domain,[1] % of respondents

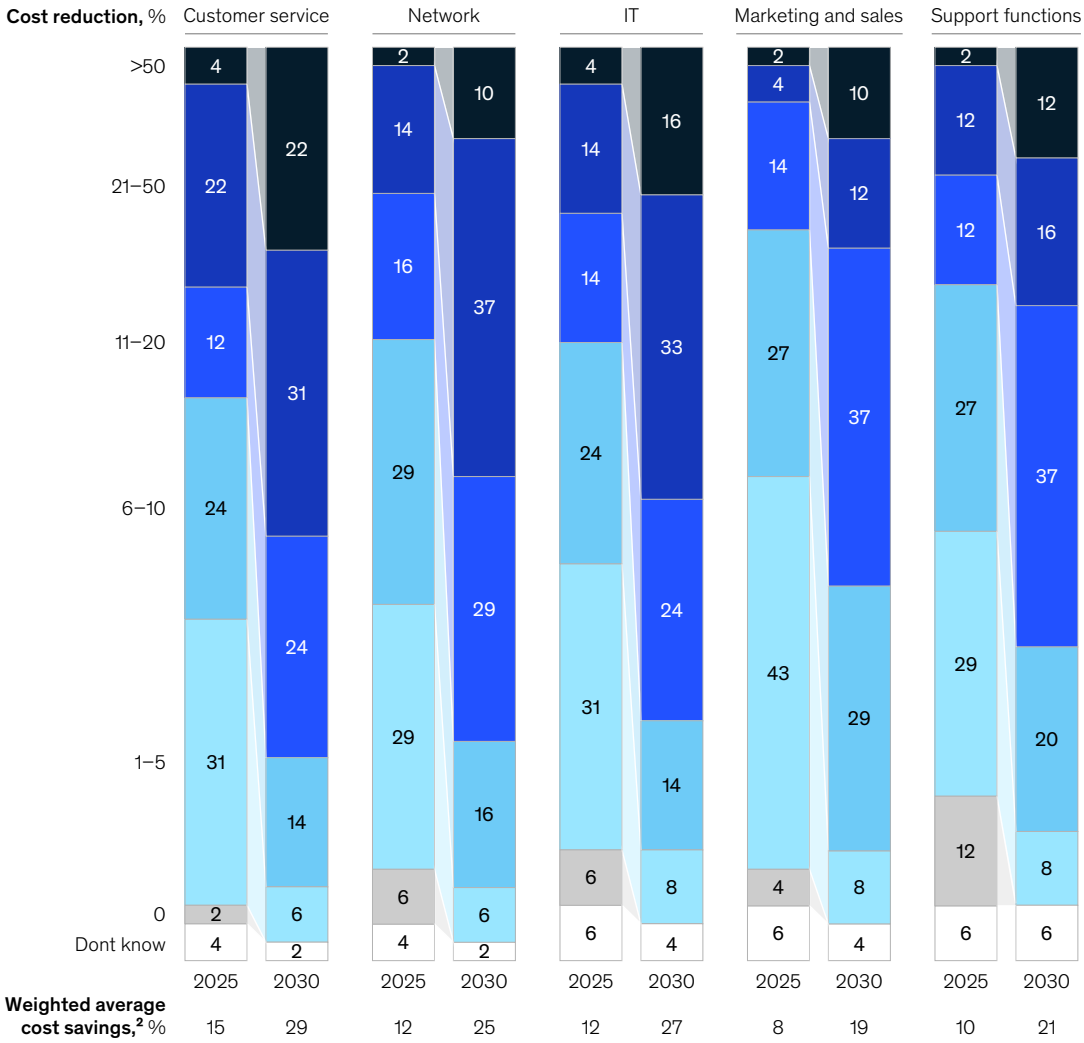

Note: Figures may not sum to 100%, because of rounding.
[1]Question: What is the expected impact (% cost reduction) attributed to AI (including gen AI and agentic AI) in different domains in [year]?
[2]To calculate these weighted averages, we took the midpoint of the cost % ranges (0, 1–5, 6–10, 11–20, 21–50, and >50), multiplied each by the corresponding percentage of respondents in the category, and summed the multiplied numbers.
Source: McKinsey Telco CxO Gen AI Survey, Dec 2025, n = 49

McKinsey & Company

Exhibit 3

Employee adoption and deploying process changes at scale are the biggest obstacles telcos say they face in generating value from AI.

Biggest challenges telcos face in generating value from AI, % of respondents

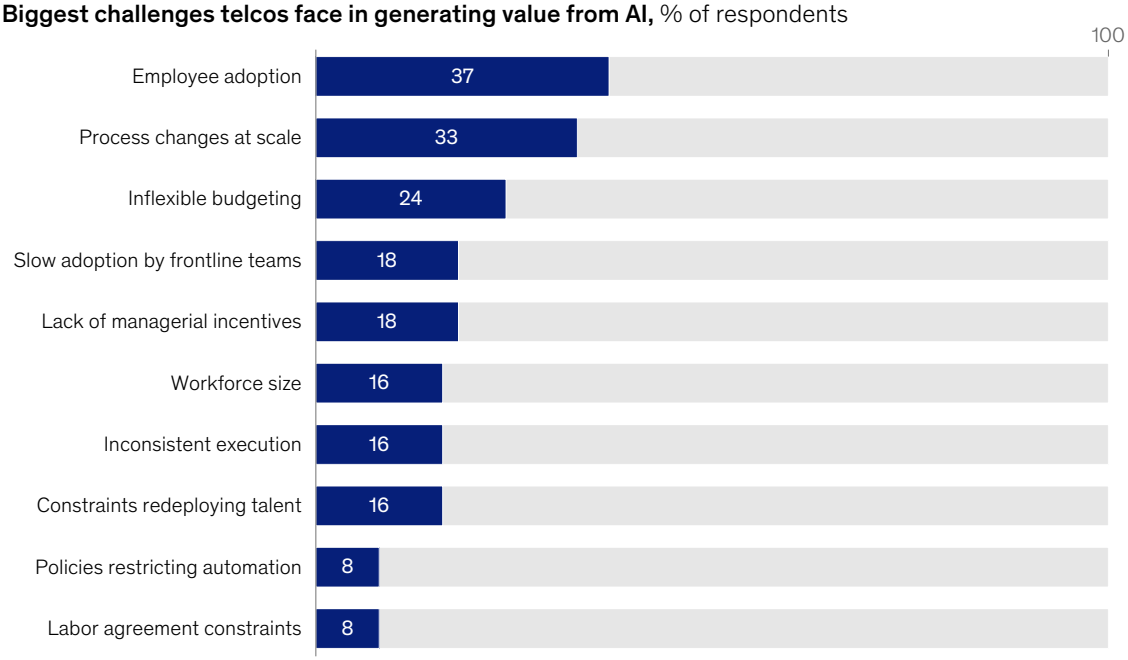

[1]Question: What are the biggest challenges your company faces in materializing value from AI? Please identify a maximum of 3.
Source: McKinsey Telco CxO Gen AI Survey, Dec 2025, n = 49

McKinsey & Company

At KPN, this has translated into developing voice-to-voice agentic AI to handle customer care interactions end to end. To build this voice-to-voice agent, the team took a three-step approach. They first used call volumes and transcripts to identify and analyze interactive voice response end points; next, they drilled down into "subintents" (such as fixed line no dial tone, defective modem or router replacement, or mobile voice/SMS roaming issues) and customer journeys to assess feasibility and agentification complexity; and then they built use cases based on impact value and development complexity and feasibility.

Leveraging this analysis, the operator prioritized six of the 18 use cases, including invoice agent, generic Q&A, and authentication for initial piloting. As a result, the operator is aiming to reduce the average handling time of calls requiring human experts, significantly reducing the run rate in overall call center spending within a year.

Network

For two leading network operators serving savvy, demanding customers in very different regions—MASORANGE in Europe and NTT DOCOMO in Asia—a deep understanding of network customer experience (CX) recently surfaced as a key competitive differentiator. This happened at a time when capital expenditure (capex) levels heightened the need for more-efficient investment decisions.

In response, each operator has used an AI for CX index, creating a daily metric that correlates the operator's network performance with customer satisfaction and risk. As a result, the operators were able to identify customers unsatisfied with the network, determine the primary cause of their subpar experience, and link

those insights to capital interventions (such as site upgrades) or other new customer offerings to help improve customer retention. For example, sites with low CX index and high revenue were identified as priority sites for the operators to invest capex.

To optimize the opportunity, both operators redesigned capex planning workflows and led efforts to upskill network teams on this AI application and ensure sustained adoption of the new decisioning approach across areas and branches. With this approach in place, the operators either assessed CX along the most heavily used commuting lines (affecting an estimated 5 million to 10 million users, for example) and reprioritized their 2025–26 capex plan, or identified areas where customers were the most sensitive to network performance fluctuations and drove actions to improve the network and help enhance the loyalty and experience of those customers. (For more on NTT DOCOMO's use of the AI for CX metric, see sidebar, "NTT DOCOMO: Using AI to keep communities better connected.")

Case study

NTT DOCOMO: Using AI to keep communities better connected

In Japan's hyper-connected cities, network quality isn't an abstract metric. It's felt in fleeting moments on packed commuter trains, in dense urban corridors, during short, repeated interactions like checking a map between stations, refreshing a news app on a crowded platform, or resending a text message that lags just long enough to be noticed.

For NTT DOCOMO, those moments were becoming a warning sign.

Even as the company continued to invest heavily in network infrastructure, customer experience in some high-traffic areas was deteriorating. Traditional engineering metrics could confirm that the network was busy—but they couldn't explain why customers were increasingly frustrated.

NTT DOCOMO recognized that closing this gap required a fundamentally new way of approaching the network: one that connected technical performance to lived customer experience—and linked both directly to investment decisions. The goal wasn't incremental optimization. It was a systematic, experience-led approach to how the network is planned, operated, and improved.

Experiencing the network as a customer

To make that shift, NTT DOCOMO worked with McKinsey as its thought partner to co-develop the Customer Network Experience (CNX) index—an AI-driven metric designed to quantify network customer experience in a way that is actionable across the organization.

Built using more than 400 terabytes of network-side data, the CNX index applies advanced analytics and machine-learning models to capture how users really experience the network across locations and time periods, rather than how the network performs in theory.

The methodology was tailored to Japan's unique usage patterns. In particular, "commuting-line" features were developed to reflect the impact of repeated, short-duration connections in dense urban transit corridors—a critical driver of customer perception that conventional metrics, such as download speed, had failed to reflect.

"As networks become more complex, improving customer experience requires more than adding infrastructure," says Yoshio Umezawa, VP of R&D Innovation at NTT DOCOMO. "The CNX Index allows us to connect network performance directly to how the network is actually experienced, giving us a common, objective foundation for decision-making across the organization. AI enables us to move from reactive responses to a more systematic, experience-led approach to network management."

Beyond the analytical solution: driving real change

Crucially, CNX was never treated as a standalone analytical output. NTT DOCOMO embedded the index directly into core workflows, including network planning, preventive maintenance, and capital expenditure (CAPEX) decision-making. CNX insights were explicitly linked to site revenue, allowing teams to identify where poor experience overlapped with high economic importance. Planners could then compare interventions against experience and site revenue in a consistent, objective way.

To ensure long-term ownership, McKinsey also supported NTT DOCOMO in building internal capabilities through an AI academy that trained engineers and planners on both the fundamentals of analytics and the practical use of CNX. The two organizations also co-developed a comprehensive change-management program to drive adoption across headquarters and regions—aligning teams around shared definitions, dashboards, and decision routines.

Turning insight into better moments and better decisions

The new perspective quickly surfaced insights that had previously been hard to identify. In Tokyo and the surrounding regions, CNX analysis revealed that central urban areas often delivered lower customer experience than surrounding zones, despite having a high number of 5G sites. The issue wasn't simply coverage but had multiple causes identified by CNX.

By linking the experience index to site revenue, NTT DOCOMO could pinpoint critical locations where targeted intervention would deliver disproportionate customer impact.

The result: Approximately 10 to 30 percent of planned CAPEX interventions have the potential to be reprioritized, allowing capital to be redirected toward higher-impact uses without increasing overall investment levels. Today, CNX is used by teams across headquarters and regions to support more consistent prioritization, tighter alignment between central planning and local execution, and more disciplined monitoring of outcomes.

Beyond efficiency gains, CNX gives NTT DOCOMO something more durable: a foundation for managing the network around how customers actually experience it. By embedding AI-driven insight into everyday decision making, the company is showing how advanced analytics can move from analysis to action, even as network complexity continues to rise.

"What makes this transformation powerful isn't the technology alone—it's the shift in mindset," says Takamichi Watase, a McKinsey partner based in Tokyo. "By anchoring network decisions in real customer experience, NTT DOCOMO has created a common language that aligns engineers, planners, and executives around what truly matters: how the network feels to the people who use it, every day."

Marketing and sales

Most telecom operators now use AI-augmented agencies and in-house tools to accelerate core marketing activities such as copy generation, promotion design, and campaign execution. At the same time, some operators are extending AI into frontline sales and contact center operations to unlock additional value from existing customer interactions.

The business process outsourcing (BPO) unit of leading South American telco Entel, for instance, faced a common challenge in its contact center operations: Changing customer expectations and engagement patterns were driving a decline in outbound sales, while the company continued to handle hundreds of thousands of inbound service calls each month. Yet fewer than 1 percent of these calls was systematically analyzed, limiting

visibility into sales opportunities, agent performance, and evolving customer needs. This gap represented a significant value opportunity.

To unlock this value, Entel Connect, in partnership with a hyperscaler, implemented an agentic AI tool that processes and analyzes 100 percent of inbound calls daily. The solution decomposes each conversation to identify leads, flag missed opportunities, and track customer sentiment while also generating personalized coaching insights for every agent and actionable dashboards for supervisors. The tool was designed to operate at scale, processing large volumes of calls overnight at very low unit cost, while continuously refining recommendations based on real operational feedback. Over time, the contact center evolved from a reactive service operation into a proactive, data-driven sales and learning engine. Within ten weeks, inbound sales increased by 40 percent, driven by a doubling of sales attempts during service calls, with no negative impact on customer satisfaction.

Support functions

A South American operator whose finance team was spending a disproportionate amount of time on repetitive, manual tasks decided to build an agentic solution to tackle the long turnarounds of basic analyses and manual data preparation and free up capacity for higher-value work.

The team defined more than 200 recurring question archetypes across operating expenditures, capex, and profit and loss, had the solution ingest more than 15 structured and unstructured data sources, and deployed a multiagent architecture (including intent recognition, entity extraction, query generation, visualization, and guardrails) to deliver accurate, explainable answers to questions of what, why, and how. This solution automated roughly 90 percent of management report findings and produced more than 95 percent accuracy on core use cases, allowing the teams to turn their attention from manual analysis toward faster, informed decision-making.

Building a winning agentic telco

A year ago, in "Scaling the AI-native telco," we outlined the key elements for telcos to scale AI efficiently.[4] These include targeting domain- and workflow-specific transformation opportunities, building scalable, modular AI platforms, implementing adequate data foundations, and fueling adoption with change management best practices. While those overarching pillars remain valid, our recent experience in helping telcos deploy AI has highlighted some critical, more granular lessons about what it takes for telcos to build a coherent agentic operating model capable of unlocking sustainable new value and growth.

1. If AI isn't in the budget, it isn't real (like its impact)

One of the primary pitfalls of early AI initiatives across both telcos and other industries has been the failure to explicitly tie them to, and track their progress against, clearly defined, benchmarked goals. Several leading organizations from the telecom sector and beyond are already taking this approach, rigorously reviewing their budgets to pinpoint where AI can deliver measurable impact and tying it to a clear business case. That impact can be significant (Exhibit 4). In this way, companies can prioritize what drives value on both the top and bottom lines rather than pain points or challenges that they think AI could help solve. This approach gives business leaders tangible, measurable targets to achieve and be judged by. As part of this, AI should be a key component of the entire budgeting process; for instance, if a division expects a certain amount of impact from its AI initiatives in the coming year, then the company may consider adjusting budgets accordingly to reinforce accountability for delivery on the goal.

[4] "Scaling the AI-native telco," McKinsey, February 27, 2025.

Exhibit 4

Telcos have significant top- and bottom-line opportunities to use AI to improve EBITDA by up to 20 to 30 percent.

Telco today vs future telco powered by AI/gen AI (~5 year horizon), illustrative

Profit and loss (P&L) domain	What is included	Share of revenue, %		Impact on P&L, pp[1]	Examples of how AI will drive impact
		Telco today	Future telco		
Revenue		100	103	+2–3	Hyper-personalized digital communications, refined CVM,[2] revamped S2S[3] (gen AI transcripts)
COGS[4] and bad debt	Cost associated with interconnect, roaming, content	32–33	31–32	–2 to –3	Optimize connectivity footprint, predictive bill shock to manage collections up front
Operations Margin[5]		67–68	71–72		
Marketing and sales	Cost from sales channels, marketing, product development for both B2C, B2B	9–10	7–8	–2 to –3	Automated sales channel calls, e-commerce reduction, store footprint optimization
Customer service	Cost from care (ie, sales to support, inbound, digital), customer value management, loyalty, and field operations	2–3	1–2	–1 to –2	Automated care channels, optimized CVM[2] (eg, predictive issue detection and proactive resolution)
SG&A[6]	Cost from HR, finance, legal, and cross-function (eg, supply chain)	6–7	5–6	–1 to –2	Automated contract management, invoicing, supplier searches, refined AI-driven HR and workforce
Network	Cost from network planning, engineering, and run (eg, network operations center)	12	10–11	–1 to –2	Self-healing agentic-driven networks, self-optimized networks (focused on network run)
IT	Cost from operation, maintenance, and development of IT systems	2–3	3–4	+0–1	Accelerated software development (automated testing, debugging, refactoring, and documentation)
EBITDA/ Total		33–34[7]	41–43	+8–10	
		Capital expenditures, %		**Change,** pp[1]	
Capital expenditures	Network expansion, real estate	100	92–96	–4 to –8	Optimization to increase capital expenditures ROI with impact on customer experience, real estate fine-tuning

Note: Numbers based on wireless operators with comparable characteristics (eg, ~$500 million to $1 billion in annual revenue, more than 5 million in active users). Figures may not sum, because of rounding.
[1]Percentage point. [2]Customer value management. [3]Service to sales. [4]Cost of goods sold. [5]Impact on volume of full-time equivalents (FTE) calculated based on the expected impact on P&L cost related to workforce × a correction factor to adjust for different ranges (eg, expected that agentic AI automates first low-wage workers, therefore the FTE impact is bigger in pp versus that on P&L). This impact excludes third-party spend (eg, external call centers). [6]Selling, general, and administrative. [7]EU average = 33; US average = 36.
Source: Investor reports; McKinsey Global Institute; McKinsey analysis

McKinsey & Company

2. Isolating key tasks and jobs to be done are essential for redesigning workflows

Much of AI agents' transformative power stems from how they can be leveraged to fundamentally reinvent end-to-end workflows across business functions. Yet most telcos continue to capture only a fraction of this value. Nearly three-quarters of the surveyed leaders report using AI primarily to support minimal or moderate redesigns of existing processes rather than fundamentally rethinking how work should be done in an AI-first environment. As a result, AI often accelerates legacy ways of working instead of delivering notable improvements in cost, speed, and effectiveness.

Leading operators shift the focus of redesign from processes to tasks and the accompanying skills that they require. By systematically decomposing priority value chains into their underlying activities and isolating the essential capabilities, they create a clear, detailed view of the work that is performed and what it takes to get it done. Each activity is then reassessed based on its suitability for automation, human-AI collaboration, or continued human ownership (Exhibit 5). This granular approach allows operators to redesign workflows deliberately, rather than layering AI onto inherited structures.

Exhibit 5

Reimagining telco processes from scratch by mapping existing workflows is key to capturing impact from agentic AI.

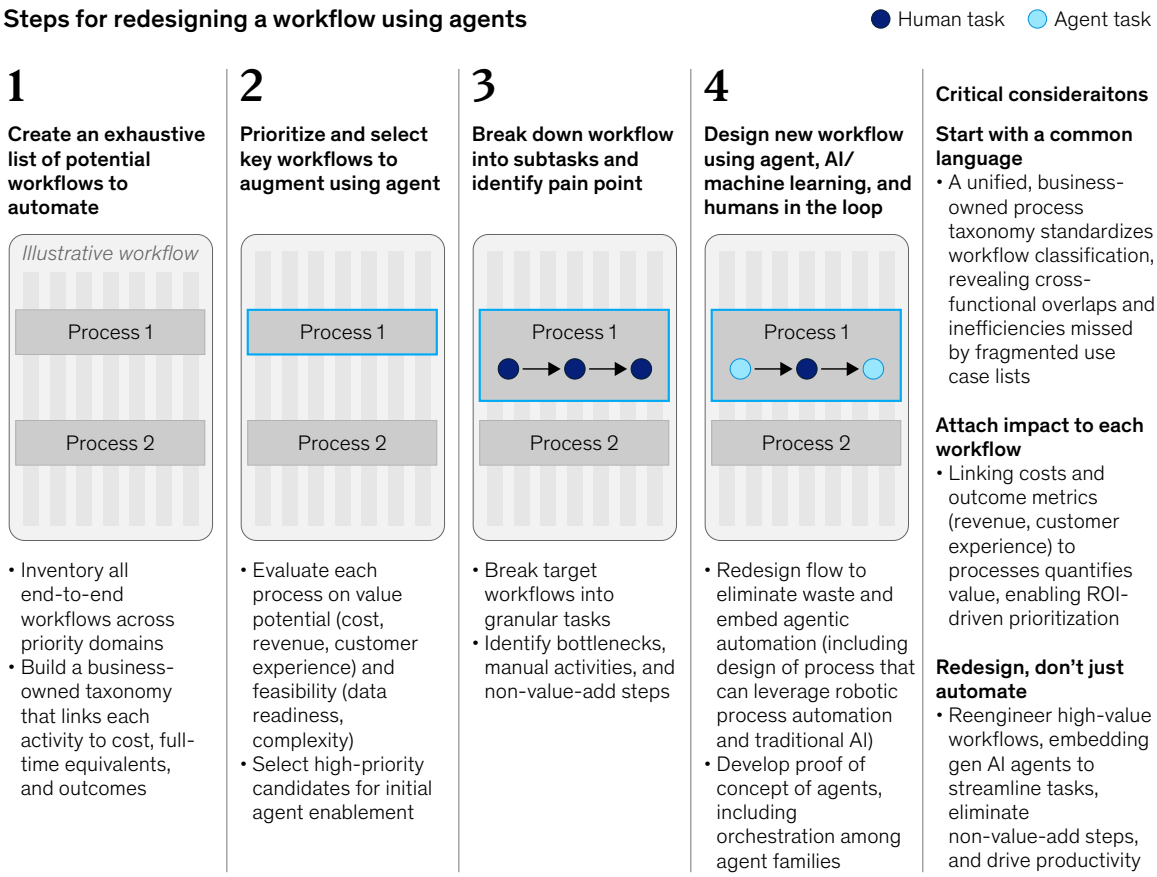

McKinsey & Company

One North American operator applied this approach to its marketing and commercial operations. It started by mapping various teams' day-to-day activities to ground the transformation in existing opportunities. It then fully redesigned all major workflows (such as evergreen, always-on marketing) with AI at their core. Only after the redesigned workflows were clear were technology choices and organizational changes finalized.

This approach is critical because it can fundamentally change outcomes. By redesigning workflows from scratch (looking only at the "necessary steps"), entire stages, handoffs, and coordination layers are eliminated, with additional value unlocked. By contrast, embedding AI into existing processes mainly accelerates the status quo, just as applying AI to isolated tasks often shifts bottlenecks rather than removing them, resulting in limited productivity gains. Telcos that redesign end to end, anchored in task-level clarity and AI-first assumptions, can unlock sustained improvements in cost, cycle time, and output quality that incremental automation cannot achieve.

3. Give agents and humans clarity about their new roles

Redesigning workflows is just the first step of the organizational overhaul telcos will need to generate substantial impact from AI. It's equally important to transform their operating structures and both agent and human roles to fuel lasting gains in cost efficiency, cycle time, and overall capacity.

A critical aspect of this is explicitly defining how AI agents are embedded within teams—what work they perform end to end, how they collaborate with humans, how exceptions are escalated, and how performance is tracked and improved. Once telcos have identified the opportunities for embedding AI agents into workflows, they can turn to redefining human roles and responsibilities. Research by the McKinsey Global Institute shows that while 30 percent of all telecom workforce hours could realistically be automated by 2030, fully 91 percent of all sector jobs will require fundamental redesign and rethinking to make the most of automation. Even as traditional roles in areas such as DevOps (software development and IT operations), product marketing, customer support, and field engineering are likely to experience declines in numbers, positions like solutions sales specialists, forward-deployed engineers, product builders, and network performance analysts are just as likely to grow; entirely new roles, including AI workflow architect and AI risk manager, will also emerge and gain prominence.

Today, many responsibilities that could be automated or orchestrated by AI agents are fragmented across numerous roles, teams, and functions. Without rethinking the structure to consolidate those responsibilities into clear, accountable roles, the impact of AI is likely to remain limited. Organizing these evolving roles around employees' skills and expertise, based on a chapters model, helps create belonging and clarity regarding role expectations. As employees rely less on cross-functional coordination, team interactions become faster and more agile. Individual employee capabilities, however, will shift to include designing, supervising, and continuously improving agents.

The benefits of these moves are manifold. Clearly defining how agents are integrated reduces structural overhead by removing layers of supervision, coordination, and approvals. Shifting roles, team structures, and agent ownership allows AI to take on end-to-end work while freeing up human capacity for higher-value tasks, rather than just making those individuals marginally more productive. And transforming operating structures and roles along with workflows helps prevent humans from serving as bottlenecks in these new, AI-driven processes.

4. Change management can't be a one-off exercise

Telco leaders often assume that technical upgrades and AI pilots will eventually translate into commercial impact. Our research shows otherwise. Slow adoption, rooted in weak change management, is the primary barrier to scaling AI value. Converting AI capability into sustained productivity and margin improvement requires an operational discipline that hardwires new behaviors into daily work.

Leading operators pair a CEO-led mandate with frontline execution teams, including change managers, analytics translators, and product owners, that are embedded directly into sales and service teams. Large language models (LLMs) make it possible to measure AI usage and adoption at the individual level, enabling much finer-grained visibility into how work is done. This data foundation allows operators to track behavior with greater precision, hold supervisors accountable for reinforcing usage, and shift incentives from outcomes to inputs to drive company effectiveness and efficiency. Such proximity ensures that adoption is measured daily at the individual level, frontline supervisors are accountable for reinforcing usage through coaching and structured performance conversations, and behavioral levers such as incentives and gamification are used to normalize AI usage. Finally, change teams cannot be treated as a short-term tactic; they must remain embedded with frontline operations. Together, these practices ensure that increased AI usage translates directly into measurable improvements in sales and service performance.

At Indosat Ooredoo Hutchison (IOH), CEO Vikram Sinha has framed the company's AI transformation approach as "70 percent people, 20 percent process, and 10 percent technology," underscoring that sustained adoption depends primarily on organizational change.[5] IOH paired this vision with a scaled capability-building and change program: the full C-suite team convenes quarterly for immersive AI sessions, AI translators have been trained across every function, and employees participate in gamified recognition programs that encourage experimentation and reward impact. These efforts embed AI directly into day-to-day work and build a broad base of AI champions, creating the momentum needed to drive adoption across the organization.

Operationalizing these lessons requires concrete program design. Capability building must be role specific and reinforced through coaching incorporated into daily routines. Core workflows should be rebuilt so that AI outputs are always factored into decision-making, with clear ownership and transparent tracking of adoption. Senior leaders must model AI use, while progress is communicated frequently across the organization. Sustained engagement comes from running short, focused campaigns that combine incentives, targeted learning, and rapid iteration based on frontline feedback.

Ultimately, to turn AI into a durable source of competitive advantage and value creation, telcos need to convert change management from a one-off program into a repeatable operating rhythm.

5. AI factories model success and accelerate value capture

Telecom leaders eager to move their AI initiatives beyond a handful of "pilot purgatory" or isolated use cases often run into the same problems: fragmented efforts, high costs, slow scaling, and insufficient risk management. In practice, this often means duplicative model builds, inconsistent foundation model choices, and limited reuse of data, prompts, or agent architectures. This is even more relevant in the current market, where users increasingly leverage AI solutions in their day-to-day work (such as for code development) without clear guardrails or training, thus posing a potential risk for their organizations.

An AI factory model addresses these issues by introducing a common playbook and set of capabilities to guide how the organization delivers and scales AI responsibly. By bringing together the right process, talent, and technology, it accelerates time to value and makes it easier to redesign workflows with greater automation and consistency (Exhibit 6). Rather than creating one-off solutions for every new project, the factory model treats AI assets—platforms, tools, and reference architectures—as standardized, core capabilities that can be built once and reused. These assets can include reusable agent frameworks, shared model environments, and standardized governance controls.

[5] "McKinsey and IOH discuss cracking the AI paradox at Innovate Asia," McKinsey, November 26, 2025; Khoo Tee Tan, "Vikram Sinha on Indosat's role in shaping Indonesia's AI transformation," McKinsey, December 3, 2025.

Exhibit 6

An AI factory model gives telcos a common playbook and capabilities to successfully scale AI across the organization.

Critical elements for building an AI factory

Strategy

- Structured ideation
- AI opportunity identification
- Process redesign
- Value realization
- Portfolio strategy and governance

Tech foundations

Technology

- Monitoring
- Agentic AI platform
- Modular AI agents
- AI sandbox
- AI productivity tools
- On-premises capability

Data

- Central AI platform
- Data strategy and protocols
- AI FinOps
- AI-driven data quality
- AI orchestration
- AI data discovery
- Operational monitoring
- Unstructured data handling
- CI/CD[1] starter kit
- AI data remediation
- Experiment tracking
- AI data products
- AI LiveOps tooling

Rewiring capabilities

Operating model

- Use case intake
- AI standards
- AI control tower
- AI architecture review
- Ways of working/ pod structure

Talent

- Capability assessment
- Strategic workforce planning
- Career path
- AI-enabled HR engine
- Partner ecosystem
- Managed service provider partnerships

Scaling and activation

Change management

- AI marketing assets
- AI champion network
- Executive AI training
- Adoption model

Responsible AI

- AI principle and mandates
- Risk log tool
- Risk control
- Responsible AI training

Key attributes

- Pre-provisioned infrastructure
- End-to-end life cycle support
- Flexible model integration
- Reusable components and frameworks
- Scalable resource management
- Unified orchestration and monitoring
- Enterprise-grade security and governance

[1]Continuous integration/continuous delivery.

McKinsey & Company

Without such standardization, scaling AI across the enterprise can become very difficult. Teams tend to develop solutions independently, adopting different tools, design patterns, and operational practices, which often lead to fragmented or siloed implementations that are hard to govern and even harder to scale. Over time, this fragmentation can increase technical debt, raise inference and infrastructure costs, and create uneven performance across similar AI use cases. This inconsistency also increases complexity and introduces significant risk, particularly around security, compliance, and data governance.

The AI factory also acts as a connective layer between an organization's varied AI efforts and senior leadership, increasing transparency and helping prioritize the most valuable, cost-effective opportunities while making trade-offs explicit. By centralizing visibility into use case value and model performance, it strengthens accountability for AI impact.

When standing up an AI factory, organizations must be focused on the highest-priority use cases. This discipline determines which capabilities are established first and defines the road map for scale. The AI factory is built on a modular set of capabilities, including a structured intake process to assess and prioritize use cases, and shared, reusable components that enable teams to leverage off-the-shelf solutions.

— **Intake** provides a standardized process to review and prioritize AI initiatives, ensuring that resources focus on the highest-impact opportunities. Some telcos now deploy AI agents within the intake workflow to assess readiness, recommend development stage (proof of concept, scale, or deprioritize), and define a clear action path.

— **Shared, reusable components** reduce development time while maintaining quality and consistency. With proven assets such as propensity model blueprints and enablement tools, teams avoid reinventing solutions for each use case and accelerate delivery at scale.

In addition to these two core capabilities, an AI factory can take on some or all aspects of talent and change management. It is not a static construct but a living capability that must be built and sustained within the organization, providing a repeatable model for moving from isolated pilots to enterprise-level results.

6. Moving forward: Agents must understand the business and the data to truly unlock its power

As telcos transition toward AI-native, self-directing enterprises, the industry is approaching a structural inflection point. The next generation of agent-led operating models will likely not be limited by algorithmic capability, but by whether AI systems can truly understand the business and the network they are meant to help run. Even the most capable LLMs struggle when key concepts, policies, and dependencies are scattered across systems or left implicit.

As they work to address this, leading operators are starting to invest in knowledge frameworks (semantic data layers) to become a foundational layer of AI-native telco architecture. Those conceptual models represent the institutional knowledge "locked" in the organization. Sharing this knowledge with agents can unlock new capabilities, cut down on duplicate efforts, reduce technical complexity, and enable agents and workflows to build on one another, turning fragmented business and network data into trusted, reusable intelligence.

This approach can have multiple applications. For example, in network operations, when service quality drops, agents can analyze the network layers to narrow down root causes, identify affected customers and services, and understand effects on service-level agreements. Instead of speeding up ticket handling, agents can determine whether a single network fault is driving multiple incidents, weigh different remediation options, and guide engineers toward the best intervention. For example, a large Australian telco has accelerated service automation 30 percent by encoding engineer's knowledge on the network into agents.

One implementation approach to enable LLMs to understand enterprise data has been the use of ontologies and knowledge graphs. These explicit conceptual models of core business entities—and the relationships between them—aim to provide a shared, consistent representation of the organization's data landscape for both humans and AI agents. While powerful in theory, ontology-driven approaches can be costly and complex to design, govern, and continuously update at enterprise scale. In practice, many organizations struggle to operationalize them quickly enough to keep pace with evolving business needs.

An emerging alternative is the use of modular, context-specific AI "skills"—reusable workflows or capability layers that encapsulate business logic, decision rules, and data access patterns in a structured but lighter-weight way. Rather than requiring a fully harmonized enterprise data model up front, AI skills allow organizations to operationalize AI around well-defined tasks (for example, pricing analysis, customer segmentation, and root-cause analysis) and progressively build capability through repeatable, composable units. Tools such as Claude Skills illustrate how this approach could enable faster deployment, clearer governance, and more pragmatic codification of institutional knowledge across AI agents.

While telcos' initial efforts to incorporate AI have shown genuine promise across pilots and fragmented use cases in areas such as customer service, network, and marketing and sales, only a small number of operators have been able to capture significant value. But as those operators are showing and their peers are learning, the emergence of agentic AI may present telcos with an automation-driven inflection point, in which agents' ability to redefine entire workflows is critical to leveraging AI's ability to radically improve customer experience and operational efficiency.

The development of a coherent agentic operating model is essential to take advantage of the new agentic opportunity. This involves prioritizing the agentic redesign of some end-to-end processes, tying AI initiatives to clear financial goals and budget outcomes, defining and clarifying new roles and responsibilities for both agents and humans, building ontologies to give agents the necessary business context, and implementing comprehensive change management to drive adoption. Finally, AI factory models help bring together all these elements, creating reusable trusted data assets and processes that can be harnessed broadly across the organization to help lower unit costs and improve economics as scale increases.

Telcos that treat agentic AI as simply another layer of tooling may see marginal productivity gains and growing complexity. Those that commit to building a coherent agentic operating model will fundamentally reshape how work gets done—and, in doing so, unlock material improvements in ROIC, margins, and customer experience. As agentic AI advances, the gap between these two groups could widen. For telecom leaders, the question is increasingly less about whether AI can create value, but more focused on whether their organizations are prepared to operate differently enough to capture it.

Benjamim Vieira is a senior partner in McKinsey's Lisbon office; **Guilherme Cruz** is a partner in the New York office, where **Tomás Lajous** is a senior partner, **Borja Belda** is an associate partner, and **Manuel Palacios** is a consultant; **Ignacio Ferrero** is a partner in the Miami office; and **Daniela Mendoza** is an associate partner in the Madrid office.

The authors wish to thank Bruce Philp, Cindy Xiao, Leire Merino, Max Shackelford, and Takamichi Watase for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

Winning B2B customers in technology and telecommunications

Our latest B2B Pulse Survey of 3,000 decision-makers shows how agentic AI is creating value beyond connectivity, raising the bar on execution, integration, and trust.

This article is a collaborative effort by Andrea Travasoni, Dev Patel, Giacomo Dolci, and Naveed Niwaz, with Bulat Gaifeev, Luca Furlani, and Olimpia Gascó, representing views from McKinsey's Technology, Media, & Telecommunications Practice.

March 2026

B2B technology and telecom growth is entering a new phase—one defined less by disruption than by rebalancing. Core connectivity remains a critical foundation, but it is no longer sufficient to define where value is created or where growth accrues. Instead, growth is increasingly shifting to higher layers of the stack, where security, intelligence, seamless engagement, and trust shape customer outcomes.

Insights from McKinsey's latest Global Technology and Telecommunications B2B Pulse Survey, based on more than 3,000 enterprise decision-makers across 11 industries and 18 countries, show that several fundamentals are holding. Overall investment intent remains positive, driven primarily by large enterprises, even as growth in core connectivity continues to lag behind higher-value domains such as security and agentic AI. Customers increasingly recognize operators' relevance beyond the network, and omnichannel engagement has become a baseline expectation across the buying journey.

At the same time, new openings for growth are emerging. Investment momentum is broadening as small and medium-size enterprises (SMEs) spend more decisively. Customer stickiness toward telecom operators is stronger than expected, though loyalty is becoming more conditional and increasingly anchored in security and trust rather than connectivity alone. Quality of experience has emerged as a decisive swing factor, shaping how customers engage and transact across channels.

Agentic AI is beginning to reset the growth equation. Adoption is accelerating across customer segments and regions, reshaping how enterprises buy, interact, and operate. As customers themselves become more AI-enabled, digital channels are moving from a supporting role to a primary growth engine—enabling larger, more complex transactions to be executed end to end and raising expectations for frictionless, outcome-driven engagement.

For telecom operators, these shifts create a dual imperative. Agentic AI offers a tangible lever to improve internal efficiency and service performance today, supported by strong customer openness to AI-enabled assistance. This makes operational deployment an immediate priority, with clear potential to reduce cost to serve and scale service quality. Over time, these capabilities may also provide a foundation to extend offerings beyond core connectivity into trusted, AI-enabled solutions. The ability to materialize these opportunities will depend less on technology access than on disciplined execution—specifically, operators' ability to deploy reliably at scale and maintain trust in an increasingly regulated, AI-driven environment.

1. Growth fundamentals are holding

While the competitive landscape is evolving, several structural dynamics from prior years remain firmly in place.

Investment outlook is positive, with core connectivity lagging behind

Survey results show that more than 60 percent of respondents plan to increase technology and telecommunications spending, driven primarily by large enterprises. This momentum reflects a broad recognition of the need to invest in innovation and remain competitive at the forefront of technology.

The trend holds across geographies, though optimism varies significantly. Africa (74 percent expecting growth) and Asia–Pacific (70 percent) show markedly stronger investment intent than Europe (54 percent), where macroeconomic pressures appear to be dampening spending appetite.

Growth expectations also diverge sharply by product domain. While core connectivity remains foundational, the strongest momentum has shifted to higher-value areas centered on protection, intelligence, and automation. Security and agentic AI are the fastest-growing domains (Exhibit 1), each expected to expand by more than 9 percent over the next 12 months. Core connectivity, by contrast, is projected to grow at less than 5 percent, making it the slowest-growing major category in the portfolio.

Exhibit 1

Security and agentic AI lead growth beyond the core.

Expected budget or spend evolution over next 12 months (net of inflation), by product category,[1]
% of respondents across business segments

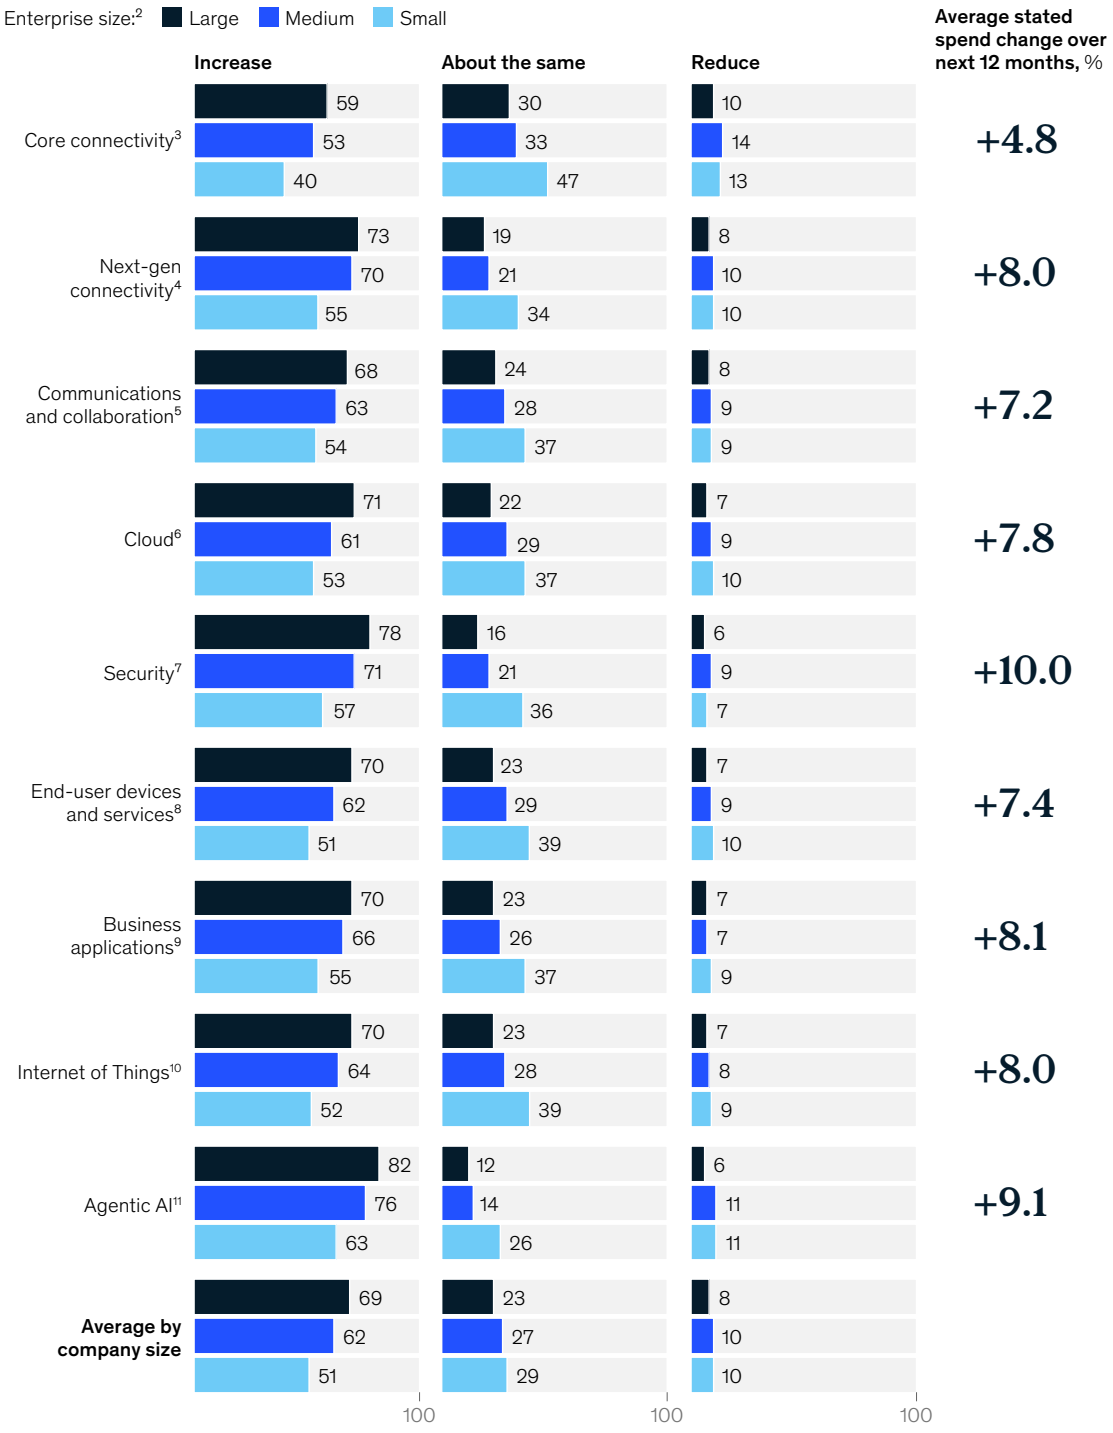

[1]Question: How do you expect your budget or spend by product category to evolve in the next 12 months net of inflation? [2]Large = >500 employees; medium = 100–500 employees; small = <100 employees. [3]Mobile and fixed core connectivity. [4]SD-WAN solutions, 5G, private networks. [5]Collaboration software including share and sync; unified communication tools including audio and video. [6]Public, private, and hybrid cloud models, legacy on-premise, IaaS (infrastructure as a service), PaaS (platform as a service). [7]Data security, endpoint security, network security, security applications, web and email security. [8]Mobile/PC, other hardware (eg, tablets, Wi-Fi), device management, support contract/contact center. [9]On-premise/licensed software, SaaS (software as a service), internally developed software. [10]Hardware embedded with electronics, software, sensors, and connectivity, which enables these objects to connect and exchange data including all related services. [11]AI systems that can plan and execute multistep tasks and take actions.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Operators have a right to play beyond core connectivity

B2B customers now expect operators to act as integration partners and end-to-end solution providers, placing greater emphasis on the ability to orchestrate complex, multidomain solutions rather than on network assets alone (Exhibit 2). As expectations rise, execution has become the primary constraint.

Nearly 80 percent of survey respondents recognize operators as relevant players beyond mobile and fixed connectivity, particularly in areas such as network orchestration, security, cloud provision and hosting, as well as the delivery of AI capabilities. Customer preferences, however, vary by segment. Small enterprises show the strongest interest in one-stop-shop offerings. Medium enterprises are more selective, favoring operators that concentrate on a defined set of domains. Large enterprises sit between these positions but show growing interest in advanced, integrated solutions as agentic AI adoption accelerates.

As B2B customers increasingly expect operators to act as integration partners and end-to-end solution providers, success hinges less on network assets and more on the ability to deliver complex, multidomain solutions at scale.

Exhibit 2

Operators have a right to play beyond connectivity.

Role telecom operators should play in the industry in the future,[1] % of respondents

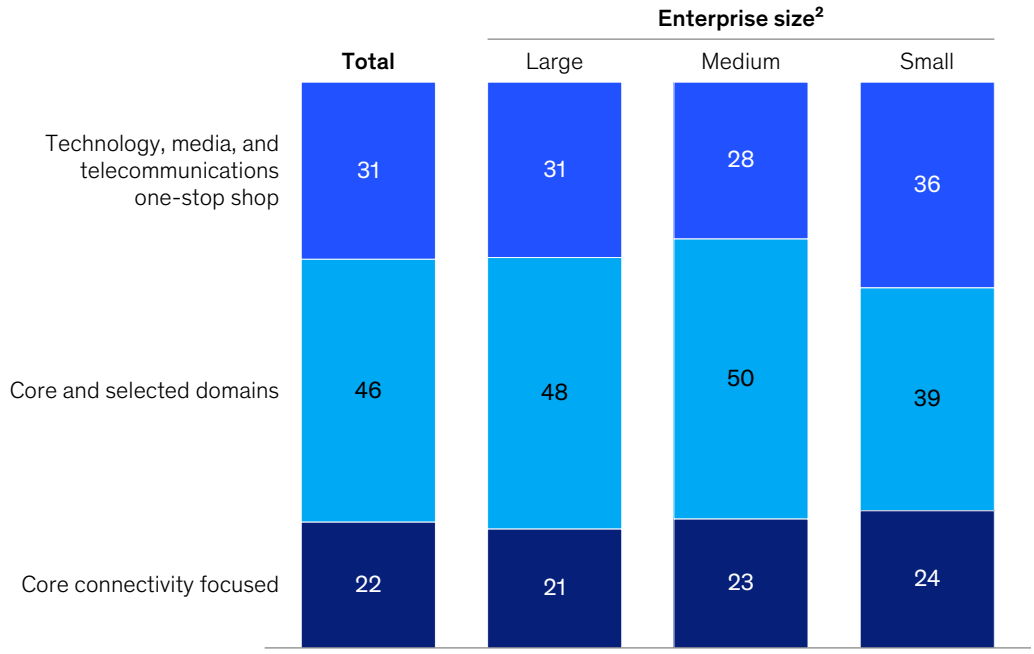

Note: Figures may not sum to 100%, because of rounding.
[1]Question: What role should telecom operators play in the industry in the future?
[2]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Customers want omnichannel

Delivering on the promise beyond core connectivity increasingly depends not only on what operators offer but on how they engage customers across channels—particularly digital ones. After a period of restraint, appetite for fully digital engagement is rebounding. Survey respondents indicate that approximately one-third of B2B customers expect to rely on digital-first interactions over the next 12 months (Exhibit 3). As a result, technology and telecom providers are being assessed less on channel availability and more on their ability to deliver seamless, consistent experiences across digital interfaces and human touchpoints.

This orchestration matters most where relationships are strongest. Nearly 60 percent of B2B survey respondents prefer hybrid engagement models that combine digital and remote human interaction, particularly in reordering and servicing journeys. Customers want the efficiency of digital channels without sacrificing continuity or trust built through existing relationships.

Exhibit 3

Digital intent outpaces usage.

Preferred communication channel with provider across the buying journey,[1] % of respondents

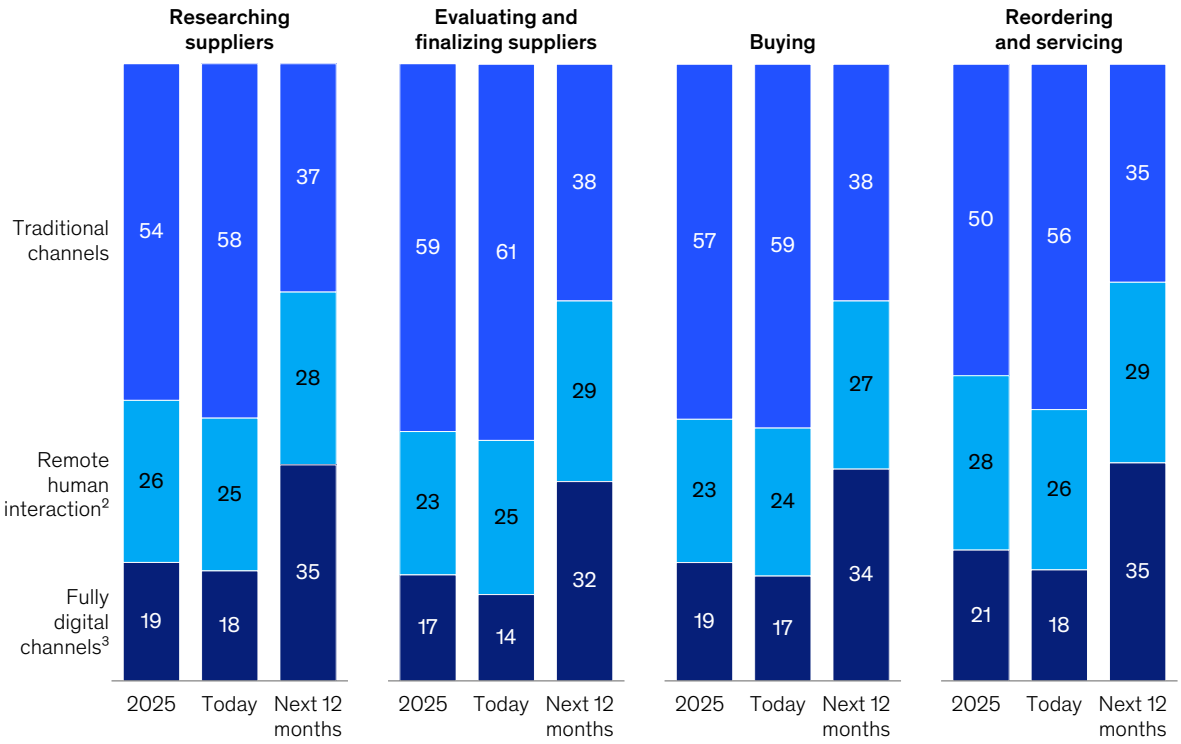

Note: Figures may not sum to 100%, because of rounding.
[1]Question: Please select the preferred channel of communication with your provider across the buying journey over last/next 12 months.
[2]Remote includes phone calls, video conference calls, emails, fax, etc.
[3]Fully digital includes company websites, e-commerce, chatbots, internet searches, mobile apps, etc.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

2. New trends are creating renewed openings for growth

At the same time, shifts in demand, retention, and digital behavior are opening new paths to growth.

Spending intent is rising, led by SMEs

Enterprise technology budgets point to a solid growth outlook for the year ahead. Survey results indicate that more than 60 percent of companies plan to increase technology spending, driving a 7 percent year-over-year increase in overall budget-growth intent (Exhibit 4).

This renewed growth expectation is no longer driven solely by large enterprises. SMEs, often more cautious, are showing renewed confidence, with budget-growth intent among small and medium-size enterprises rising by eight percentage points year over year.

Exhibit 4

Small and medium-size enterprises show the strongest year-over-year increase.

Expected budget or spend evolution over next 12 months (net of inflation), by product category,[1]
% of respondents across business segments

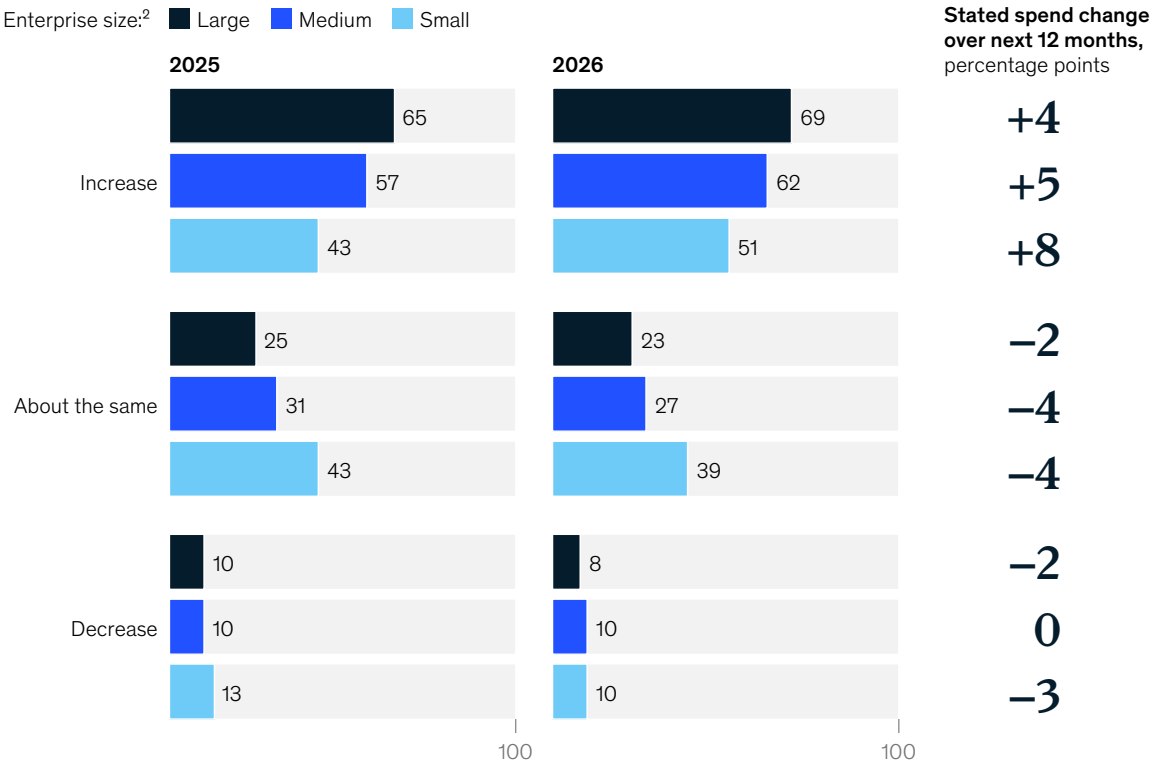

[1]Question: How do you expect your budget or spend by product category to evolve in the next 12 months net of inflation?
[2]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

While large enterprises continue to account for the majority of absolute spend, the broadening of confidence marks a structural change from prior years. As SMEs reenter the market more decisively, the addressable growth base is expanding and competition for incremental demand is intensifying, particularly in high-growth domains beyond core connectivity.

Customer stickiness is stabilizing, but loyalty is increasingly conditional

After a period of elevated churn risk, B2B customer stickiness toward telecom operators has begun to stabilize. Nearly three-quarters of surveyed customers report no intention to switch provider types over the next 12 months, up from roughly two-thirds in prior years (Exhibit 5). This shift points to a more resilient customer base and suggests that competitive pressure is translating into more selective switching rather than broad-based defection.

This resilience, however, is uneven across the portfolio. Stickiness is strongest, and improving fastest, in business applications and end-user devices and services, where customers report a seven- to nine-percentage-point increase in their likelihood to remain with telecom operators between 2024 and 2026. In these domains, operators appear to be consolidating their role as trusted delivery partners.

Exhibit 5

Telecom operator retention has improved across most product categories.

Change in share of respondents choosing telecom operators as their primary vendor,[1] 2026 vs 2024, percentage-point difference in % of respondents

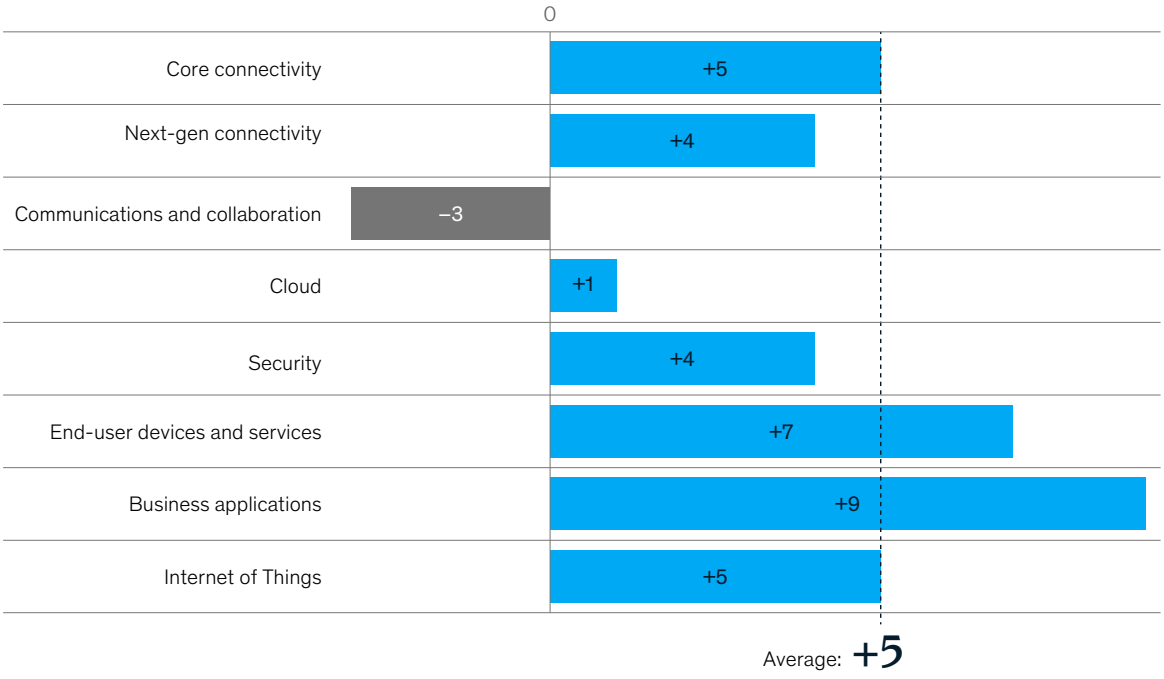

[1]Question: When sourcing technology and telecommunications products, who is your primary vendor today and most likely to be in 12 months?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

By contrast, other domains remain structurally contested. Communications and collaboration, as well as cloud services, show flat or declining stickiness, reflecting intense competition and a crowded supplier landscape. In these areas, incumbency alone offers limited protection, and customer relationships remain more fluid.

When customers do consider switching, the underlying drivers are also shifting. Among the roughly 30 percent of surveyed B2B customers that report having switched providers in the past year, cybersecurity has emerged as the leading trigger, overtaking traditional factors such as price, coverage, or service reliability (Exhibit 6). Even in relationships that are primarily connectivity-led, security considerations now rank as the number one reason for defection.

Taken together, these patterns point to a fundamental change in loyalty. While switching behavior has not yet increased materially, stable churn rates should not be mistaken for durable loyalty. Customers are no longer anchored to operators by connectivity alone; instead, loyalty is becoming increasingly conditional and must be actively earned through credible, outcome-driven security and trust capabilities. For operators that meet this expectation, cybersecurity can serve as a powerful lever to defend and expand share of wallet. For those that do not, incremental demand—and eventual erosion—is likely to accrue, especially in favor of cloud and technology providers rather than remain within the telco ecosystem.

Exhibit 6

Security now ranks as the number one reason for switching.

Top 3 reasons to switch providers per category,[1] % of times ranked as most important by switchers

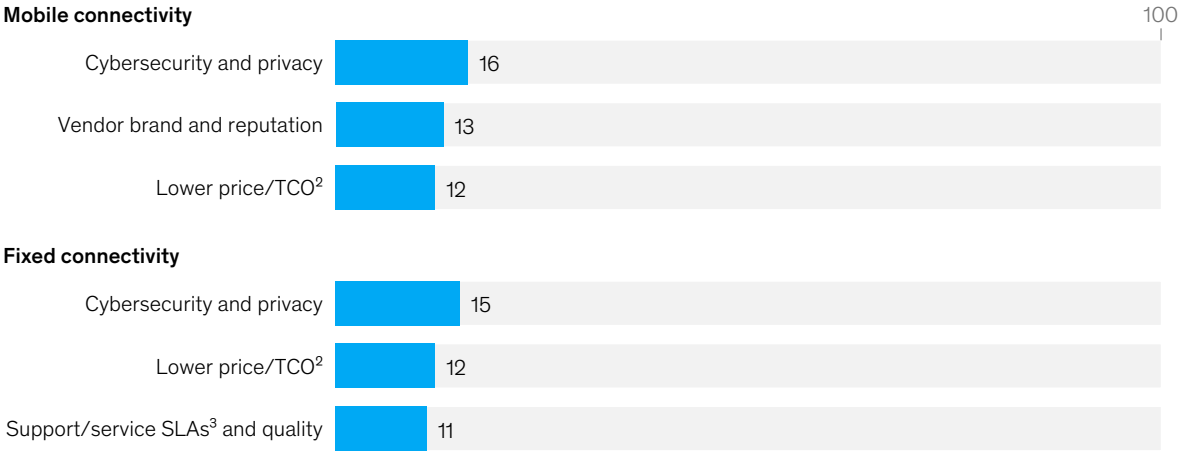

[1]Question: Which of the following reasons prompted you to switch the [product category] provider? Provider switch reasons respondents selected from: lower price/TCO; product/service features and capabilities including redundancy/reliability; support/service-level agreement and quality; vendor brand and reputation; existing relationship with the provider; product innovation/industry and vertical-specific solutions; end-to-end solution and/or ease of integration with your systems; sales rep knowledge and ability to shape solutions; cybersecurity and privacy; experience in engaging across different physical and digital channels.
[2]Total cost of ownership.
[3]Service-level agreements.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Customer experience is a swing factor for channel choice with AI-forward customers driving digital channel growth

According to the survey, B2B customers continue to prefer hybrid engagement models that combine digital channels with human interaction. At the same time, new evidence suggests that quality of experience has become the decisive swing factor in channel choice. Rather than channel availability, improved quality and reliability of digital journeys is supporting adoption growth (Exhibit 7). As digital experiences improve across the buying cycle, customers are increasingly willing to rely on digital-first interactions, particularly where journeys are intuitive, consistent, and frictionless.

Exhibit 7

Reliable digital channels drive 1.4 times higher fully digital adoption.

Share of respondents preferring fully digital channel across the customer journey, %

[1]Question: Please select the preferred channel of communication with your provider across the buying journey over last/next 12 months. (Share of respondents choosing fully digital channel.)
[2]Question: What are the challenges you encountered with digital channels across the buying journey? (Respondents selecting "No challenges" for the corresponding step of the buying journey.)
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

While these improvements have supported steady growth in digital usage, customer experience alone explains only part of the shift. The rapid advancement of agentic AI among B2B customers seems to be an even more powerful "accelerator," pushing digital channel adoption, as well as the size of digital transactions.

Enterprises that have fully implemented agentic AI at scale in at least one use case, or that are actively scaling live use cases across functions, show a materially higher willingness to transact through fully digital, end-to-end channels. For these AI-mature customers, the average "ticket size" for digital transactions exceeds $1 million—roughly twice the level reported by peers still in early pilots or without operationalized AI. As customers become more comfortable delegating decisions and execution to AI-enabled workflows, digital self-service seems to increasingly extend to complex, high-value purchases (Exhibit 8).

Exhibit 8

AI-forward buyers are driving digital spend.

Maximum order value respondent would purchase through end-to-end digital self-service,[1]
% of respondents

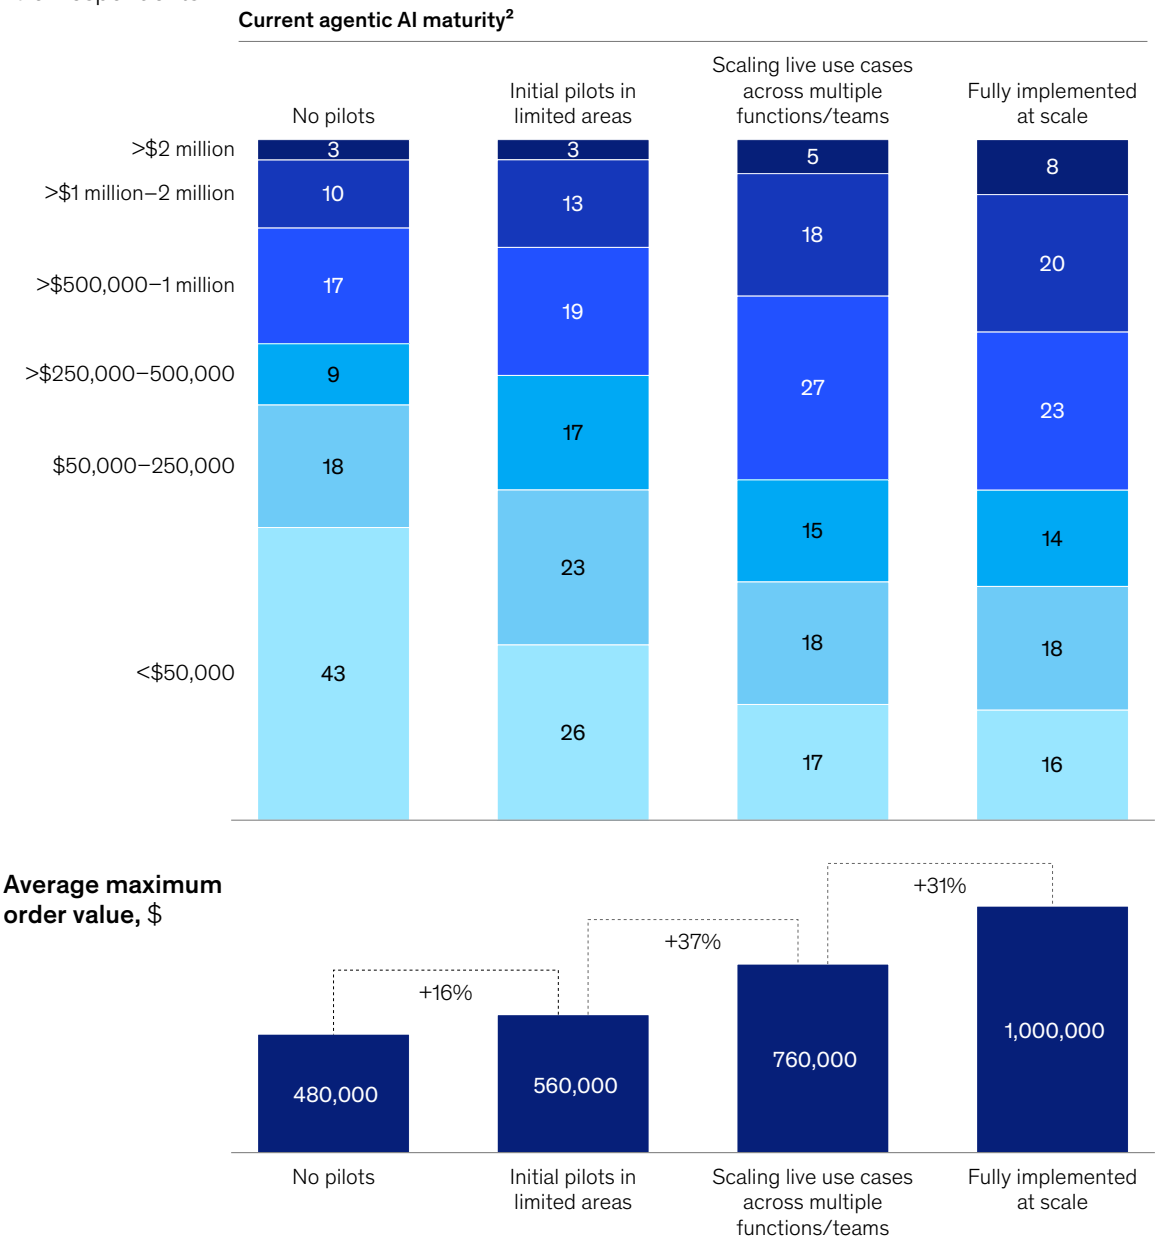

Note: Figures may not sum to 100%, because of rounding.
[1]Question: What is the maximum order value you would purchase through end-to-end digital self-service for both a new product or service?
[2]Question: With regard to agentic AI adoption, how would you best describe your company's current maturity?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

This relationship appears systematic rather than episodic. According to the survey, as customers progress along the agentic AI maturity curve—from experimentation to scaled deployment—the value they report being willing to transact digitally increases in lockstep (Exhibit 9). Each step up in AI maturity is associated with a self-reported 20 to 45 percent increase in the absolute dollar value customers are prepared to execute through operators' end-to-end digital channels.

This pattern holds across enterprise sizes but is most pronounced among large organizations. According to the survey, large enterprises with agentic AI fully implemented at scale report a willingness to transact up to one-third more value digitally than peers using AI in more limited or selective ways.

Exhibit 9

AI maturity boosts digital spend across business segments.

Maximum order value respondent would purchase through end-to-end digital self-service,[1] $

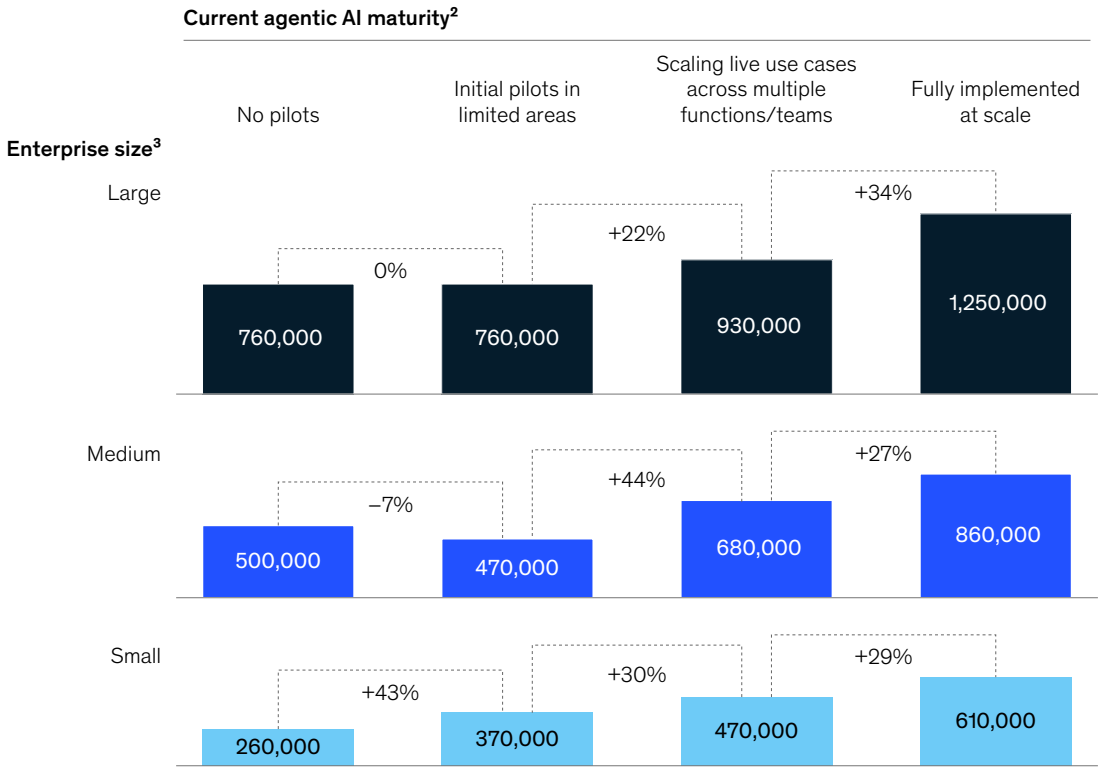

[1]Question: What is the maximum order value you would purchase through end-to-end digital self-service for both a new product or service?
[2]Question: With regard to agentic AI adoption, how would you best describe your company's current maturity?
[3]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Taken together, these dynamics point to an impending acceleration in digital channel adoption. Improvements in customer experience have laid the foundation, but agentic AI is now changing how enterprises buy—reducing reliance on human-mediated interactions and increasing comfort with automated, end-to-end digital journeys. As agentic AI adoption continues to expand rapidly across customer segments, digital channels are likely to shift decisively from a complementary interface to a primary growth engine, enabling providers to scale sales productivity and support larger, more complex transactions online.

3. Agentic AI is resetting the growth equation

Agentic AI now sits at the center of these shifts, accelerating change across buying, operations, and delivery models.

Agentic AI adoption is already widespread but not at the same pace everywhere

Roughly 80 percent of surveyed large and medium-size enterprises report having agentic AI in place in some form, ranging from early pilots to fully scaled deployments (Exhibit 10). Adoption among small enterprises is lower but still material, with approximately 65 percent of surveyed B2B customers reporting some level of implementation.

Exhibit 10

More than 80 percent of companies are expected to be active in agentic AI in the next 12 to 18 months.

Current level of agentic AI adoption maturity vs future ambition (next 12–18 months),[1]
% of respondents

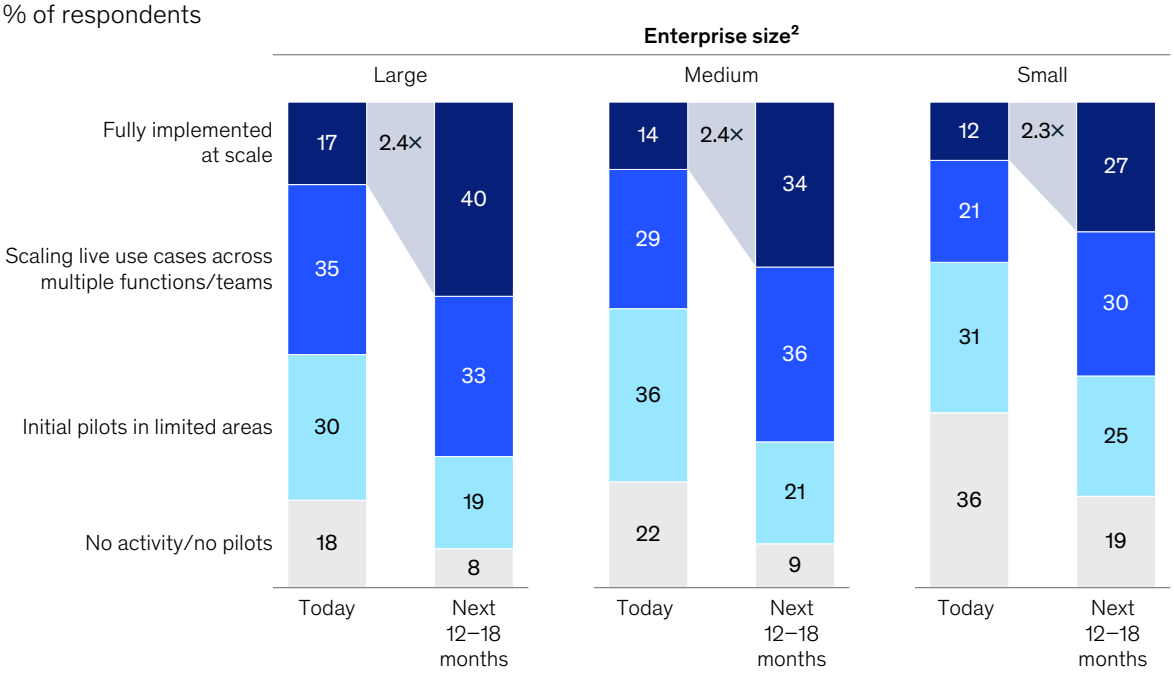

Note: Figures may not sum to 100%, because of rounding.
[1]Question: With regard to agentic AI adoption, how would you best describe your company's current level of maturity and future ambition (next 12–18 months)?
[2]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

The more meaningful distinction lies in maturity. Among large and medium-size enterprises, about 40 to 50 percent report to have already moved beyond experimentation, with agentic AI either fully implemented at scale or actively scaling across multiple domains. In contrast, only about 30 percent of SMEs report to have reached a similar level of maturity.

Over the next 12 months, companies across all size segments expect the share of fully implemented agentic AI solutions to increase by a factor of 2.3x to 2.4x. This signals a decisive shift from isolated use cases to enterprise-wide execution, bringing fundamental changes to operating models, decision-making, and system architectures. For B2B operators, the customer base is becoming more automated, more digital, and more capable of engaging through AI-enabled, end-to-end journeys. Portfolios, sales motions, and service models will need to evolve accordingly to remain relevant.

Across regions, momentum in agentic AI adoption remains very strong, with roughly 70 to 80 percent of surveyed large enterprises actively engaged, ranging from initial pilots to fully implemented, at-scale use cases (Exhibit 11). While this high level of activity is evident globally, meaningful regional differences are beginning to emerge.

Exhibit 11

There is strong momentum in agentic AI, with 70 to 80 percent of large enterprises actively engaged across regions.

Current level of agentic AI adoption maturity among large enterprises,[1] % of respondents

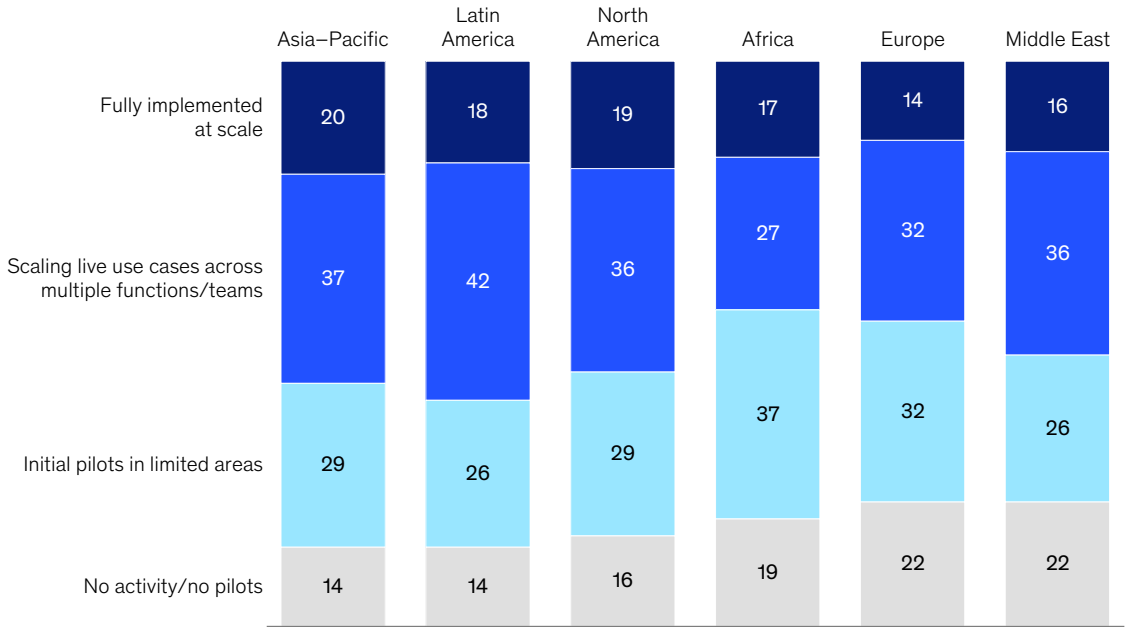

[1]Question: With regard to agentic AI adoption, how would you best describe your company's current level of maturity?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Asia–Pacific and Latin America stand out as the most advanced regions, with more than 80 percent of large B2B companies with already initiated or scaled agentic AI. This reflects faster progression from experimentation to enterprise-wide deployment.

By contrast, Europe and the Middle East have a lower share of enterprises having implemented agentic AI fully at scale or at scale-up stage. More than 20 percent of large B2B companies report no active involvement in agentic AI, indicating a slower transition from pilots to scaled implementation.

While these gaps may appear moderate at first glance, they are strategically meaningful. Geographies further along the adoption curve are likely to experience earlier shifts in buying behavior, greater reliance on digital self-service, and rising expectations for AI-enabled engagement. As these expectations evolve, competitive dynamics will increasingly be shaped by how quickly providers adapt their capabilities to customer readiness and local market conditions, rather than by technology availability alone.

Data privacy and sovereignty are becoming non-negotiable

Rising geopolitical fragmentation and expanding regulatory scrutiny are elevating data privacy and sovereignty from compliance considerations to decisive buying criteria. Across regions and industries, surveyed B2B customers increasingly view sovereign compliance as a prerequisite for engaging with technology and telecom providers, an expectation that is intensifying as agentic AI becomes more deeply embedded in enterprise operations.

Data privacy and regulatory compliance now rank as the single most important customer concern, cited by more than half of respondents (Exhibit 12). Sensitivity is particularly acute as AI agents gain broader access to enterprise data and take on more autonomous decisioning and execution roles. At the same time, regulatory frameworks (most notably around AI governance in the European Union) are becoming more explicit and consequential, raising the stakes for customers seeking to deploy AI at scale without regulatory risk.

As a result, providers that fail to meet sovereignty, security, and compliance expectations are increasingly excluded from consideration, regardless of price or functionality. Trust has become a gating factor rather than a differentiator. Telecom operators enter this environment with a structural advantage, grounded in decades of experience operating under stringent regulatory regimes and managing critical infrastructure. Operators that can simplify compliance, provide transparency, and absorb regulatory complexity on behalf of customers have an opportunity to convert a baseline requirement into a source of differentiation, particularly in the context of agentic AI-enabled solutions.

Concerns are most pronounced among small enterprises, where nearly half of respondents cite data privacy and security as their top consideration, significantly higher than among large and medium-size enterprises. This gap reflects structural differences in internal capabilities—smaller organizations typically lack dedicated security and compliance resources and therefore place greater reliance on external providers to manage risk on their behalf.

Across all segments, customer preferences converge on a clear deployment mandate. Nearly two-thirds of respondents require their data to run in private cloud environments they directly control or within on-premise infrastructure (Exhibit 13). Appetite for less restrictive models remains limited, with only a small minority expressing comfort with public cloud or provider-hosted environments. Innovation and scalability matter, but they do not outweigh the need for control.

Exhibit 12

Data security and compliance dominate agentic AI concerns.

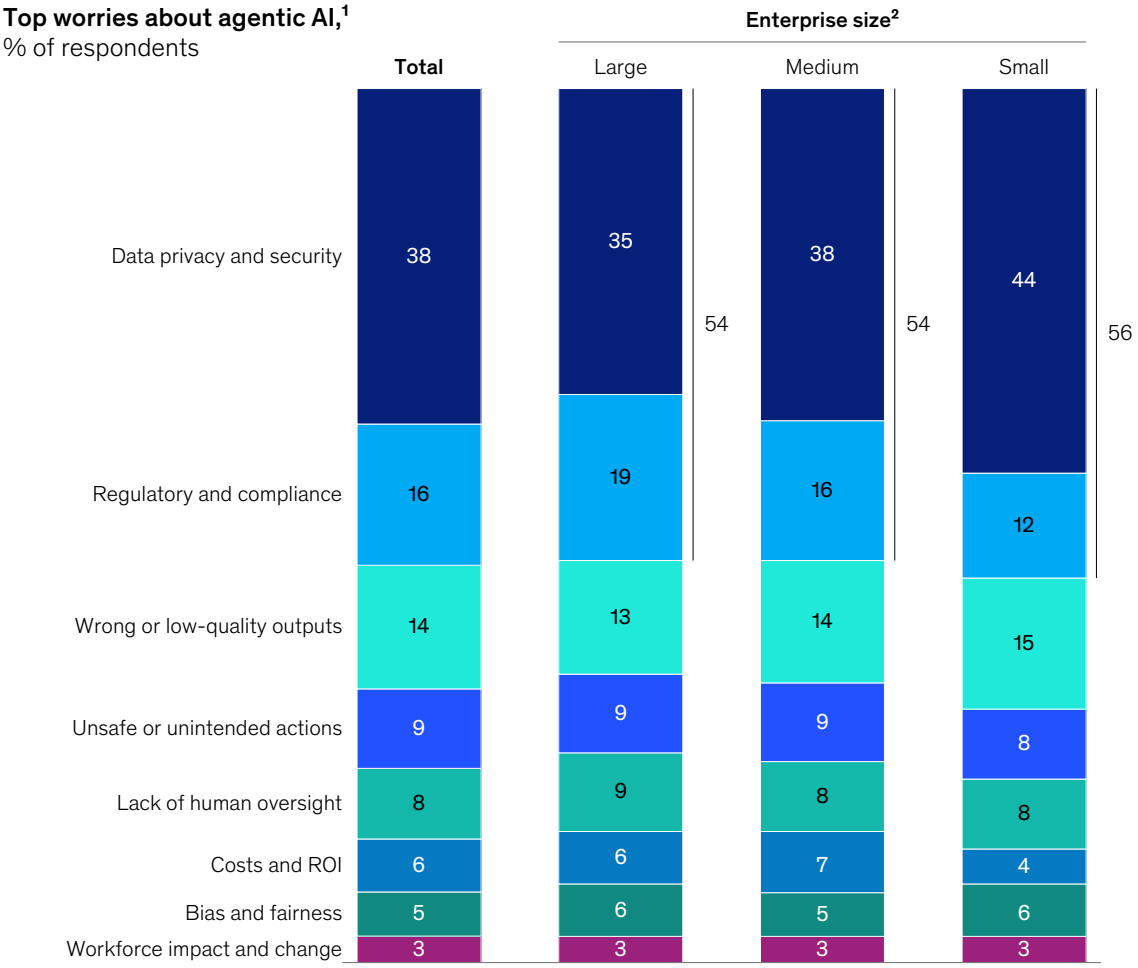

Note: Figures may not sum to 100%, because of rounding.
[1]Question: What are your biggest worries when thinking of agentic AI?
[2]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Looking ahead, these dynamics are likely to intensify. As agentic AI adoption accelerates and enterprises delegate more data access and operational authority to autonomous systems, expectations around privacy by design, sovereign deployment, and regulatory assurance will rise in parallel. Providers that can natively support flexible deployment models and embed trust into AI-enabled offerings will be best positioned to capture growing demand. Those that cannot, risk being structurally sidelined as AI-driven buying behavior reshapes the market.

Exhibit 13

A preference exists for private cloud and on-premise data management.

Setup considered acceptable for where data is stored/processed,[1] % of respondents

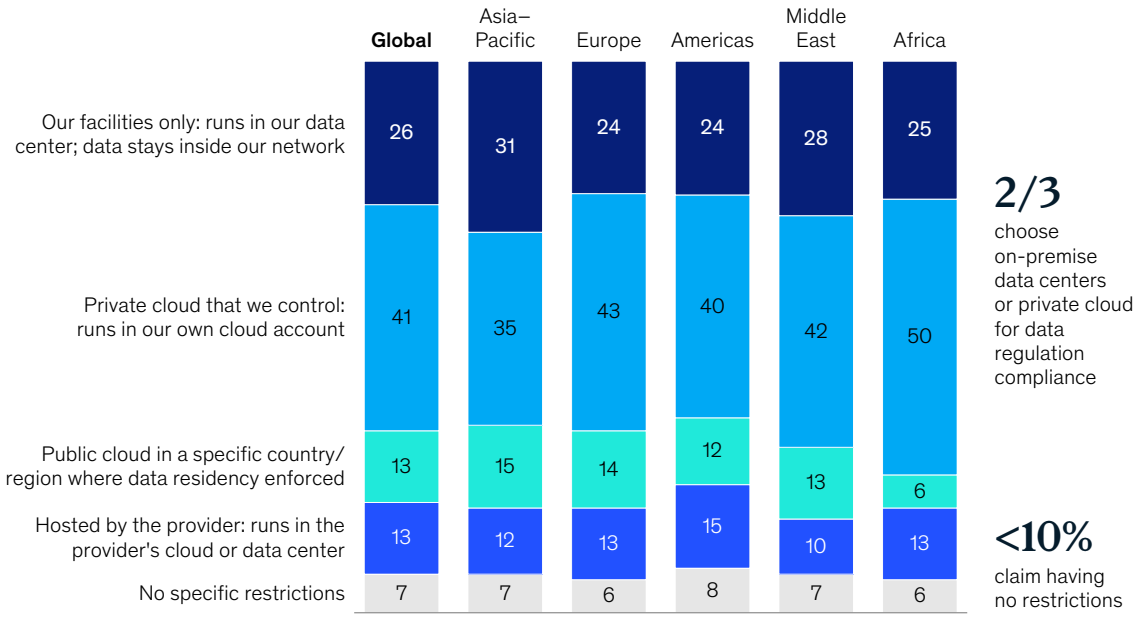

Note: Figures may not sum to 100%, because of rounding.
[1]Question: Thinking about your company's data regulation and compliance policies, which setup is acceptable for where your data is stored/processed?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Agentic AI adoption is growing, especially within customer care use cases

As agentic AI adoption moves from experimentation toward broader deployment, B2B customers are converging on a clear set of priority use cases. Customer care has emerged as the primary entry point: 80 to 90 percent of respondents report they have already implemented, or plan to implement, agentic AI in customer service processes (Exhibit 14). This consistency across industries reflects both the maturity of available solutions and the immediate, measurable impact on service efficiency and experience.

Beyond customer care, adoption is expanding into more operationally complex domains. IT and development functions, along with fraud and security, represent the next wave of high-prevalence use cases, with more than 70 percent of surveyed customers indicating active or planned implementation. Revenue-facing applications—particularly in sales and marketing—are more uneven but feature prominently in advanced industries such as technology, media, and telecommunications; retail; financial services; and energy and utilities, where competitive intensity and growth pressure increase the value of AI-enabled engagement and decision support. Taken together, these patterns suggest that agentic AI adoption is evolving toward more differentiated, industry-specific deployment paths rather than a one-size-fits-all trajectory. For B2B operators, understanding which use cases customers prioritize for optimization through agentic AI also provides a clear signal on where to focus internal efficiency efforts. By mirroring customer priorities internally, operators can accelerate capability development, prove impact at scale, and subsequently monetize these capabilities through deployment at customer sites.

Exhibit 14

Customer care leads agentic AI adoption, with more than 80 percent project adoption.

Existing and future use-case plans across B2B,[1] % of use cases selected among top 3 per industry either as already implemented or planned to be implemented within next 12–18 months

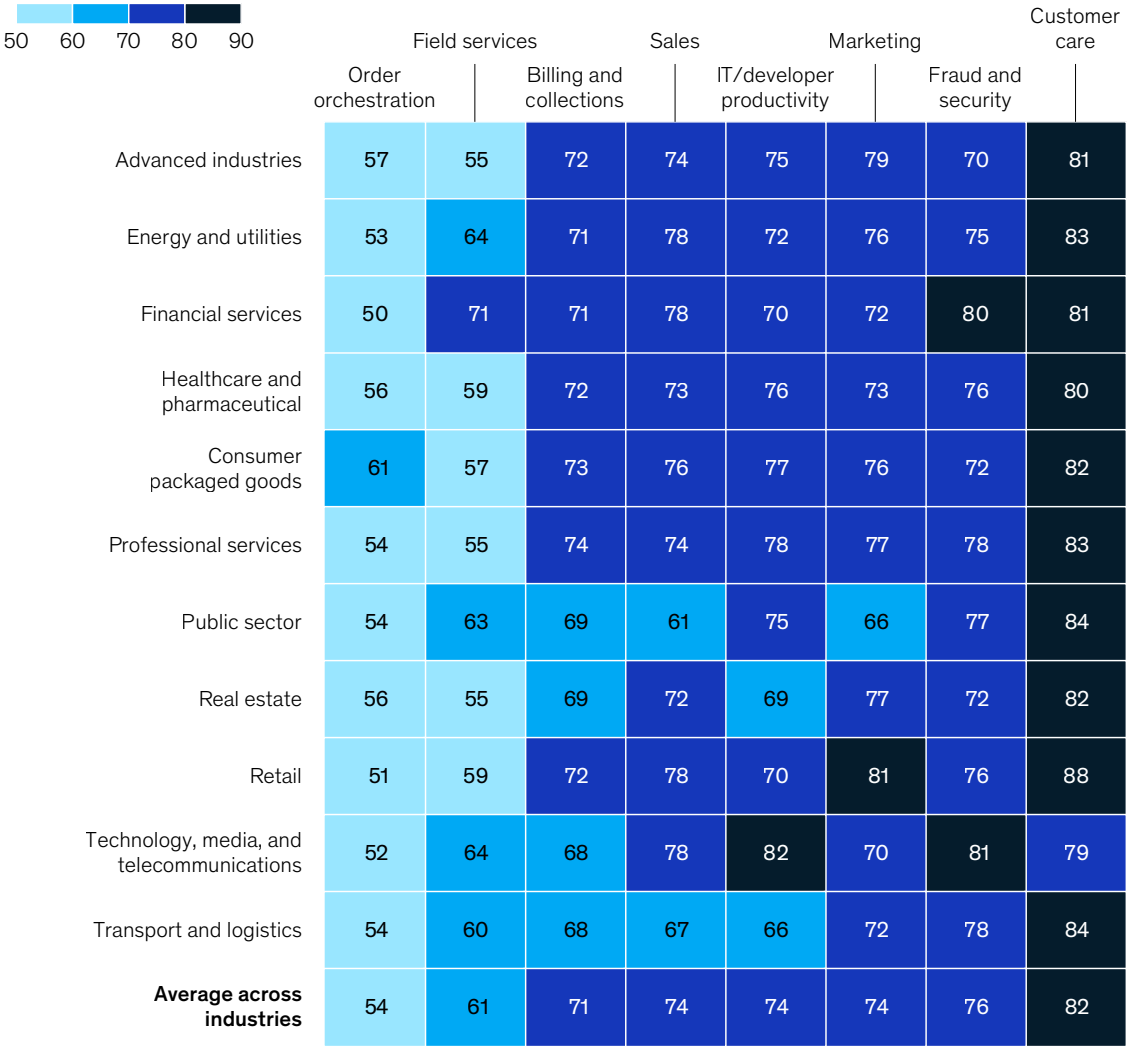

[1]Questions: In which areas has your company already implemented agentic AI use cases? In which areas does your company plan to introduce agentic AI use cases in the next 12–18 months?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,038

McKinsey & Company

Despite strong intent and widespread experimentation, however, at-scale implementation remains the exception rather than the norm. Among surveyed customers that have initiated agentic AI adoption, only 3 to 11 percent report full-scale deployment for any given use case (Exhibit 15). Even when including organizations that are actively scaling AI across multiple domains, adoption reaches only 7 to 21 percent. Across use cases, the largest share of respondents remains in the pilot phase, highlighting a persistent gap between ambition and execution.

Exhibit 15

Less than 11 percent of companies have fully implemented agentic AI solutions.

Current level of agentic AI adoption maturity by use case,[1]
% of respondents

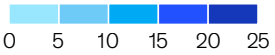

	Current agentic AI maturity[2]		
	Fully implemented at scale	Scaling live use cases across multiple functions/teams	Initial pilots in limited areas
Customer care	11	21	23
Marketing	7	15	18
Fraud and security	9	14	15
Sales	7	14	15
IT/developer productivity	7	15	15
Billing and collections	6	14	14
Field services	6	11	11
Order orchestration	3	7	7

[1]Question: In which areas has your company already implemented agentic AI use cases?
[2]Question: With regard to agentic AI adoption, how would you best describe your company's current level of maturity?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,308

McKinsey & Company

This execution gap is evident even in customer care, the most advanced and mature use case. While the majority of organizations are either piloting or scaling agentic AI in this domain, only around one in ten surveyed customers report having achieved full-scale deployment. Nearly half remain in early stages, with a substantial share still working to move beyond pilots and isolated team-level implementations.

Taken together, these findings point to a critical inflection point. Customer demand for agentic AI is increasingly well defined, but the ability to execute reliably at scale remains constrained. Providers that can bridge this gap—by integrating agentic AI into core processes, operating models, and governance frameworks—will be best positioned to capture value as adoption accelerates from pilots to enterprise-wide deployment.

4. Agentic AI creates immediate efficiency and a longer-term growth option for operators

For operators, this creates a dual opportunity: capture near-term operational gains while selectively building toward differentiated AI-enabled offerings.

Agentic AI improves efficiency while preserving customer trust

Agentic AI presents telecom operators with a near-term opportunity to materially improve internal efficiency. Applied to high-volume, operationally intensive interactions, agentic AI can reduce cost to serve, extend service availability, and improve responsiveness, while preserving human-led engagement where commercial stakes

and relationship value are highest. Rather than replacing human interaction, agentic AI enables a more effective allocation of human effort.

Customer acceptance provides a strong foundation for this shift. Across surveyed B2B segments, 98 percent of customers indicate openness to receiving AI-enabled assistance in at least one interaction (Exhibit 16). Acceptance is highest in service and support contexts, where speed, availability, and consistency are most critical. Between roughly half and two-thirds of respondents are open to AI-assisted 24/7 phone or chat support, and close to half would accept AI guidance when resolving technical issues consistently across enterprise sizes.

By contrast, tolerance for AI involvement is lower in commercially sensitive interactions. Only around one-third of customers are comfortable with AI support in sales and proposal activities, underscoring that trust, judgment, and negotiation remain distinctly human domains. This boundary is instructive. It suggests that the value of agentic AI lies not in full automation, but in intelligent orchestration.

Exhibit 16

Ninety-eight percent of B2B customers are open to AI-assisted service.

Share of respondents open to be serviced with AI in at least one use case,[1] % of respondents

Areas where respondents are open to AI-enabled servicing from their connectivity provider,[1] % of respondents

30 40 50 60 70

	Enterprise size[2]		
	Large	Medium	Small
24/7 answers by phone or chat: AI that replies to common questions without waiting for a human	62	56	52
Fix technical issues: AI that guides me step by step to resolve service or equipment problems	55	50	48
Self-service tasks online: AI that helps me change plans, add services, or check status	50	47	42
Prevent problems proactively: AI that monitors services and fixes issues before we notice them	50	46	41
Manage support cases: AI that opens, routes, and keeps me updated on support tickets	50	44	39
Billing and accounts: AI that explains charges and helps correct or change items	40	41	40
Sales and proposals: AI that learns my needs and drafts tailored offers/proposals	39	36	38

[1]Question: Thinking about how your connectivity provider supports you today and in the future, where would you be open to be serviced with the use of AI?
[2]Large = >500 employees; medium = 100–500 employees; small = <100 employees.
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

For operators, the winning model is therefore hybrid by design. Agentic AI should be deployed to absorb volume, complexity, and variability at scale, freeing human experts to focus on high-value interactions that drive growth, differentiation, and long-term relationships. Operators that successfully rebalance this mix can unlock meaningful productivity gains in the near term, while strengthening service quality and customer trust rather than eroding it.

The provider market remains highly contested

As agentic AI adoption accelerates, enterprises are making a clear build-vs.-buy choice and are increasingly opting to buy rather than develop capabilities in-house. Across use cases, the survey shows that roughly 80 percent of B2B customers prefer to source agentic AI through off-the-shelf solutions, managed services, or third-party-delivered custom builds (Exhibit 17). This preference holds consistently across major use cases, from customer care and sales to billing, marketing, and field services.

Despite this strong demand signal, the provider landscape remains fragmented. No use case requires a fundamentally distinct delivery approach, and no provider type has yet emerged as a clear winner. Customer preferences vary by application. Surveyed foundation model and AI-native providers currently lead stated preferences, particularly in customer care and sales, reflecting their early internal adoption and proven functionality. System integrators follow as a credible alternative.

Exhibit 17

Over 80 percent prefer agentic AI from external providers.

Preferred delivery approach for agentic AI use cases implemented/planned within next 12–18 months,[1] % of respondents

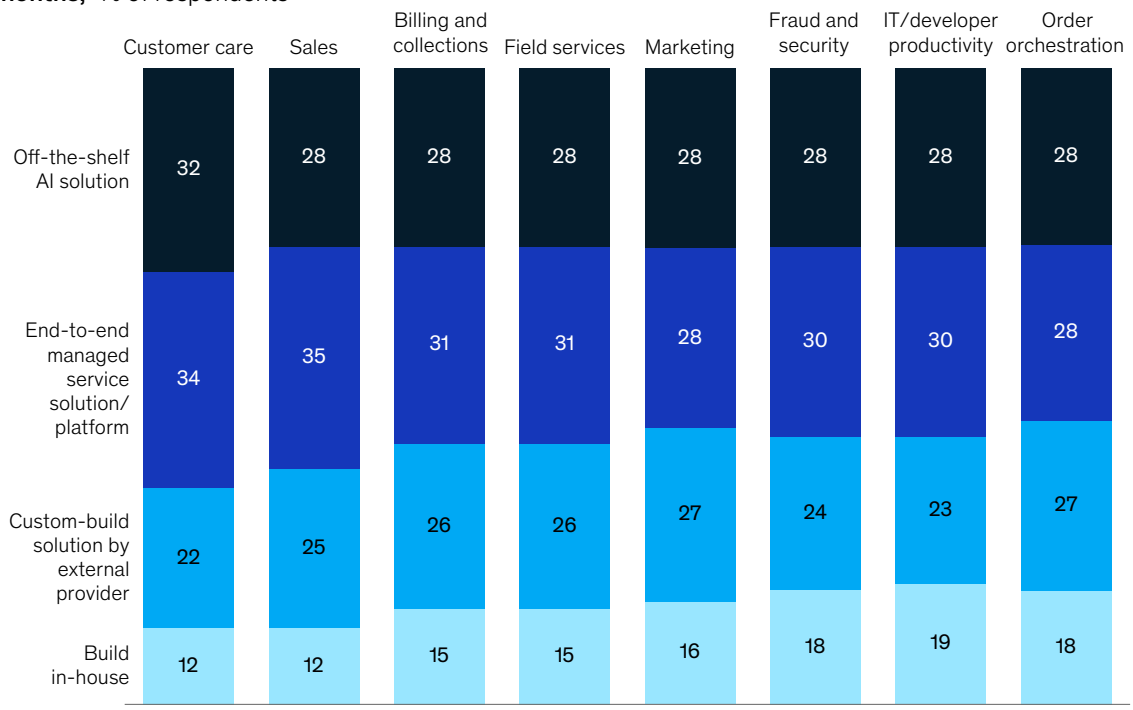

Note: Figures may not sum to 100%, because of rounding.
[1]Question: For each agentic AI use-case area you said your company has implemented or plans to introduce in the next 12–18 months, how would you prefer it to be provided?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

More unexpectedly, telecom operators also feature among the top three preferred provider types (for the next 12 to 18 months) across several use cases. This signal should not be read as a reflection of today's market reality; few operators currently offer a clearly defined, scaled agentic AI proposition. Rather, it appears to reflect a latent customer expectation—or even aspiration—around the role operators could play, anchored in trust, operational reliability, and deep integration into existing processes.

This lack of consolidation could be strategically significant. While customer openness to buying agentic AI solutions is clear, no operator has yet established a clearly defined, scalable agentic AI proposition with a proven go-to-market model (Exhibit 18). As a result, the opportunity for operators exists in principle, but remains early, uncertain, and largely untested. Realizing it would require rethinking value propositions, delivery models, and commercial packaging rather than extending existing offerings.

Exhibit 18

No provider type consistently dominates across agentic AI domains yet.

Preferred provider for agentic AI use cases implemented or planned within next 12–18 months,[1]
% selecting telecom as provider

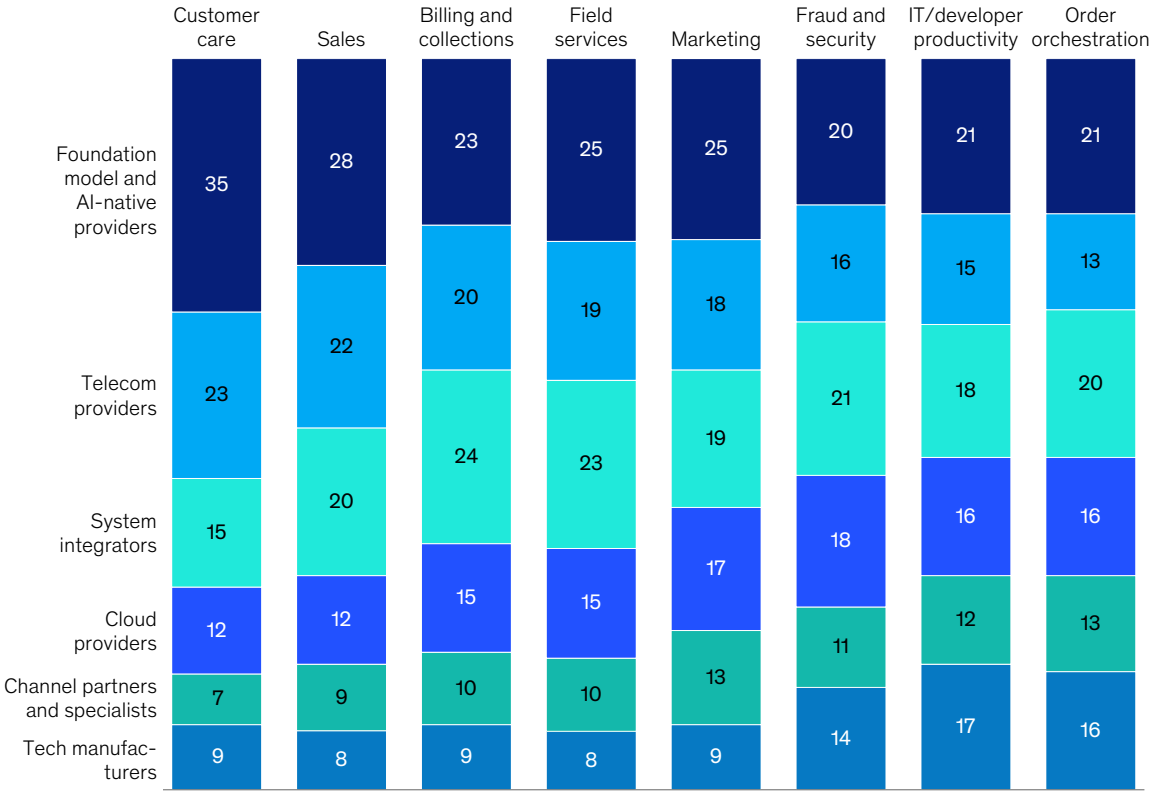

Note: Figures may not sum to 100%, because of rounding.
[1]Question: For each agentic AI use-case area you said your company has implemented or plans to introduce in the next 12–18 months, who would you prefer to provide it?
Source: McKinsey analysis; McKinsey Global Technology and Telecommunications B2B Pulse Survey, Q1 2026, n = 3,214

McKinsey & Company

Customer behavior further clarifies where this opportunity is, and where it is not. When speed, standardization, and ready-to-deploy functionality are the primary decision criteria, enterprises gravitate toward AI-native, off-the-shelf offerings, which are preferred roughly twice as often as custom-built or fully managed alternatives. By contrast, integrated, managed, and domain-embedded deployments (those requiring deep process knowledge, trusted operations, privacy controls, and long-term accountability) remain less clearly served.

It is in these environments that telecom operators could selectively explore a differentiated role, particularly in domains such as customer care, sales, and billing, where they already own customer relationships and process context. Whether this opening can be translated into a scalable and repeatable business will depend less on technology access and more on disciplined execution, clear value articulation, and the ability to operate agentic AI reliably at scale in a fast-moving and still-maturing market.

Connectivity still underpins B2B telco value, but it no longer defines where growth is created. As value moves up the stack, agentic AI is accelerating both efficiency gains and new ways for operators to engage customers beyond the network. Capturing this opportunity will depend less on technology access than on execution, combining security, digital experience, and thoughtful AI–human orchestration.

Andrea Travasoni is a senior partner in McKinsey's Milan office, where **Giacomo Dolci** is a partner and **Luca Furlani** is an associate partner; **Dev Patel** is a senior partner in the Chicago office; **Naveed Niwaz** is a partner in the Copenhagen office; **Bulat Gaifeev** is a capabilities and insights specialist in the Lisbon office; and **Olimpia Gascó** is a consultant in the Geneva office.

Issue briefs

McKinsey & Company

Technology, Media & Telecommunications Practice

AI-driven telecom networks

Structural and cost pressures have leading telecom operators looking to AI to reset network economics, enable disciplined growth, and unlock new revenue streams.

This article is a collaborative effort by Gustav Grundin, Sebastian Cubela, and Tomás Lajous, with Borja Belda, Matyas Zetek, and Sebastián González, representing views from McKinsey's Technology, Media & Telecommunications Practice.

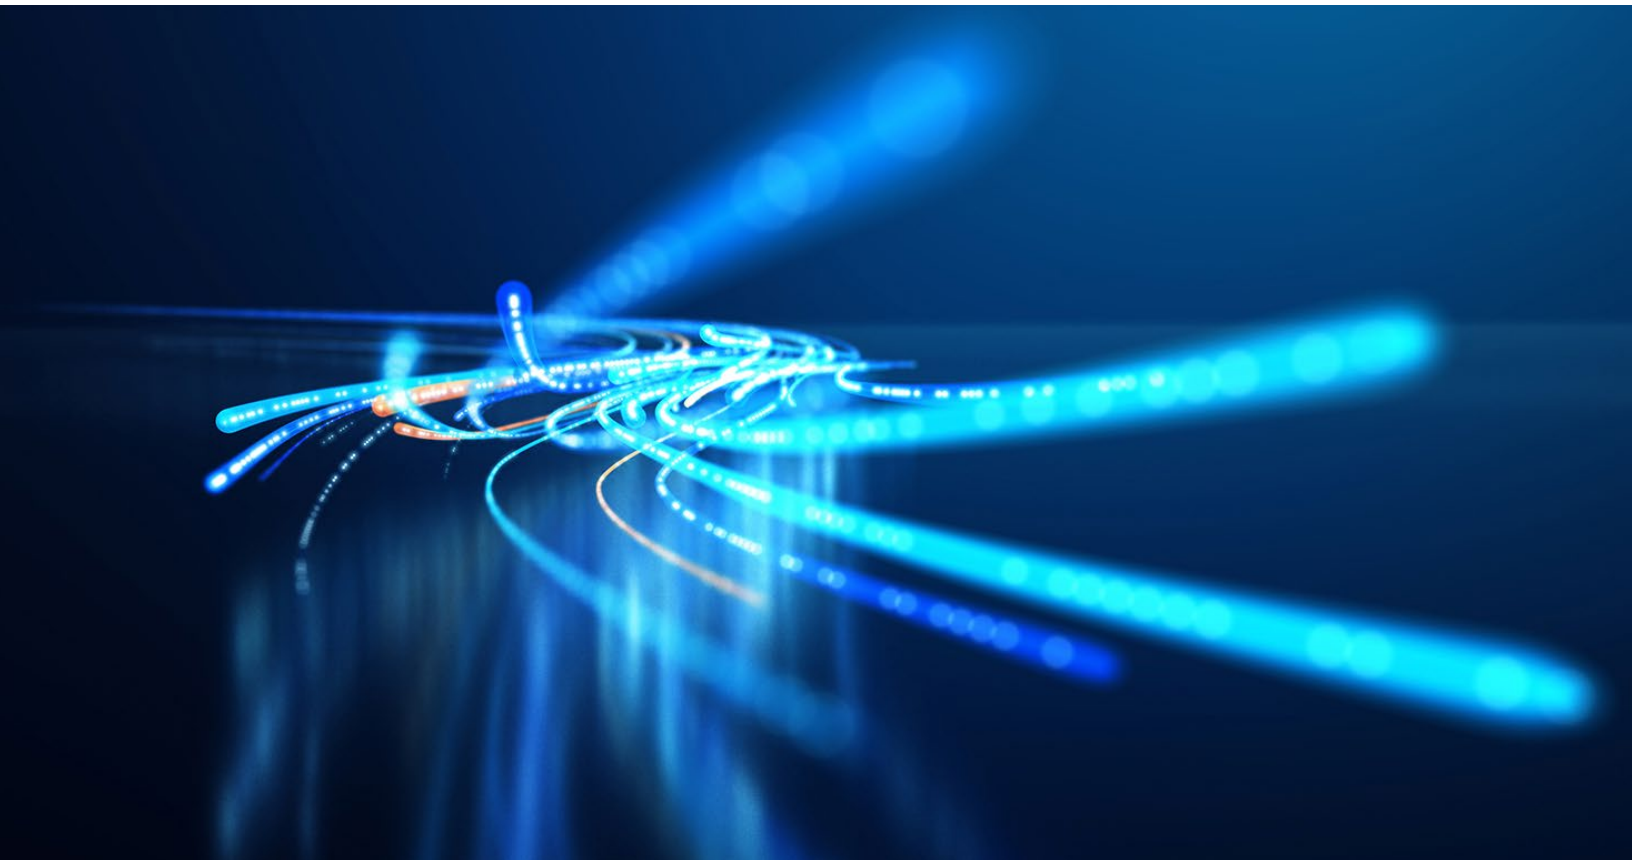

March 2026

Telecom network economics are under structural long-term pressure. Network complexity continues to rise, customer expectations for performance remain high, and cost pressures arise in the context of an expected data traffic growth deceleration in developed markets.

The rise of AI provides multiple opportunities for telcos to turn the tide. For starters, AI can be embedded in telco value propositions offering differentiated services. Operators can also participate in the emerging AI value chain by becoming the AI infrastructure backbone (as explored in an accompanying issue brief in this compendium). And telcos can leverage AI to transform their operations from the ground up.

Most directly, in an industry where improved operational efficiency has become central to the networks' value equation, AI represents a rare opportunity to reimagine the domain and reset its economics, while delivering exceptional performance. According to our latest survey of global telco top executives, the network domain will be one of the primary focus areas for AI deployments during the next two years, alongside customer care.

However, the difference between marginal gains and structural impact lies not in the technology itself but in how operators redesign processes, roles, governance, and budgeting around it. To capture these opportunities, operators need to raise their ambition and apply the same type of discipline to the network domain.

Recent progress in this field is encouraging. One of the loftiest and potentially most consequential ambitions many operators have articulated—a fully autonomous, self-optimizing, self-healing network—is no longer a distant vision, but an achievable goal in the coming years.

What's at stake?

AI's impact on networks spans both capital expenditure (capex) and operating expenditure (opex), across the following three major domains:

1) Value-led planning and simulation:
Network planning is shifting from static engineering thresholds to AI-driven, value-based optimization. To complement traditional network planning based on capacity and coverage estimates, advanced machine learning models and digital twins now simulate thousands of rollout and upgrade scenarios before capital is deployed. These simulations consider and estimate impacts on customer experience (CX), traffic evolution, customer churn and average revenue per user (ARPU), and potential competitive moves from other operators.

Operators deploying AI-based planning engines are already reporting success metrics:

— 10 to 20 percent lower greenfield rollout capex

— 15 to 35 percent lower upgrade capex

— Reallocation of 20 to 30 percent of annual capex away from low-ROI interventions

However, AI alone will not deliver results. Operators need to embed those simulation engines into their capital process and governance. Budget decisions need to be explicitly tied to the outputs, planning teams need to be trained to act on and improve model outputs, and engineering teams must provide feedback, so models consider new or unexpected restrictions.

2) AI-driven operations:
AI is also reengineering and optimizing network operations across multiple areas: In energy management, AI dynamically optimizes energy consumption by managing sleep features and detects anomalies without affecting service quality. In field operations, route optimization and automated scheduling reduce idle time and unnecessary dispatches. In maintenance, predictive models shift operators from reactive repairs to proactive interventions on critical assets.

Combined, AI-driven operational use cases can reduce total network opex by 15 to 30 percent.

To maximize value, operators need to redesign current workflows around human–AI collaboration, automating workflow steps, eliminating redundant handoffs, coordination efforts, and inefficiencies.

3) Automated repair and self-healing networks
Issue resolution and self-healing capabilities are emerging as one of the most widespread AI applications in network operations. Operators are deploying AI across the entire "issue management journey." For example:

Advanced anomaly detection and CX monitoring models identify potential faults early, even before customers are aware of a problem. Smart co-pilots and root-cause models analyze historical incidents and equipment documentation to recommend remediation steps, while dynamic matchmaking systems assign tickets to engineers with the most relevant expertise. Critical change agents automatically identify planned updates with disruption risks and design fallback or remediation plans to prevent major disruptions.

At scale, these capabilities have enabled operators to achieve 30 to 70 percent fewer troubleshooting tickets, leading to 55 to 80 percent reductions in network operations center costs, and 30 to 40 percent faster mean time to repair, alongside measurable improvements in customer experience.

What's ahead?

As operators move beyond pilots, four AI-enabled shifts are emerging:

— **Increased focus on "fixing the foundations:"** Network data has historically been fragmented and less mature than commercial data due to its complexity and volume. Scaling AI requires centralizing, governing, and standardizing network data, and defining a hybrid architecture (cloud and on-prem) to support AI use cases and workflow automations at different levels of the network.

— **Capturing institutional knowledge:** Much network know-how resides in engineers' experience. To enable network autonomy, operators are systematically codifying this expertise—through document ingestion, interviews, and structured modeling—so AI systems can infer network dependencies and remediation logic. This semantic layer becomes the foundation for self-directed decisioning.

— **Growth of network APIs and programmable connectivity:** AI-native and enterprise applications increasingly require identity, location, security, and quality-on-demand as programmable network functions rather than static connectivity. Network APIs allow operators to expose these capabilities directly to developers and enterprises, enabling differentiated, usage-based monetization. API coverage already spans around 80 percent of global mobile connections, positioning operators to participate in application-layer value creation rather than remaining transport providers.

— **Intelligent network services and cloud interconnect:** As AI workloads become more distributed, enterprises need software-defined network services to dynamically manage latency, routing, bandwidth, and regulatory constraints across cloud and hybrid environments.

What does it take to succeed?

Access to AI doesn't separate leaders from laggards; instead, it's the ability to embed AI as a continuous management capability across the network life cycle. Top performing operators take a business-led approach to AI, prioritizing use cases by value rather than technology, investing early in data foundations and digital twins, building internal capabilities beyond pilots, and partnering selectively with hyperscalers and AI specialists. Crucially, they recognize that AI is evolving at exceptional speed—models, tooling, and use cases are advancing every month.

AI is now the primary lever to reset network economics, protect margins in mature markets, enable disciplined growth in emerging markets, and unlock new, differentiated revenue streams. Operators that industrialize AI end-to-end can materially outperform peers on ROIC, EBITDA resilience, and long-term strategic optionality.

Gustav Grundin is a McKinsey partner in the Helsinki office; **Sebastian Cubela** is a partner in the Miami office; and **Tomás Lajous** is a senior partner in the New York office, where **Borja Belda** is an associate partner. **Matyas Zetek** is an associate partner in the Prague office, and **Sebastián González** is a consultant in the Connecticut office.

The authors wish to thank Abhimanyu Harlalka, Cindy Xiao, Eduardo Zamudio, and Manuel Palacios for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

AI infrastructure

Playing a major role in enabling the AI era could reignite growth for telcos along the AI infrastructure value chain, utilizing assets many telcos already control.

This article is a collaborative effort by Gustav Grundin, Miguel Frade, Sebastian Cubela, and Tomás Lajous, with Borja Belda and Lorraine Salazar, representing views from McKinsey's Technology, Media & Telecommunications Practice and QuantumBlack, AI by McKinsey.

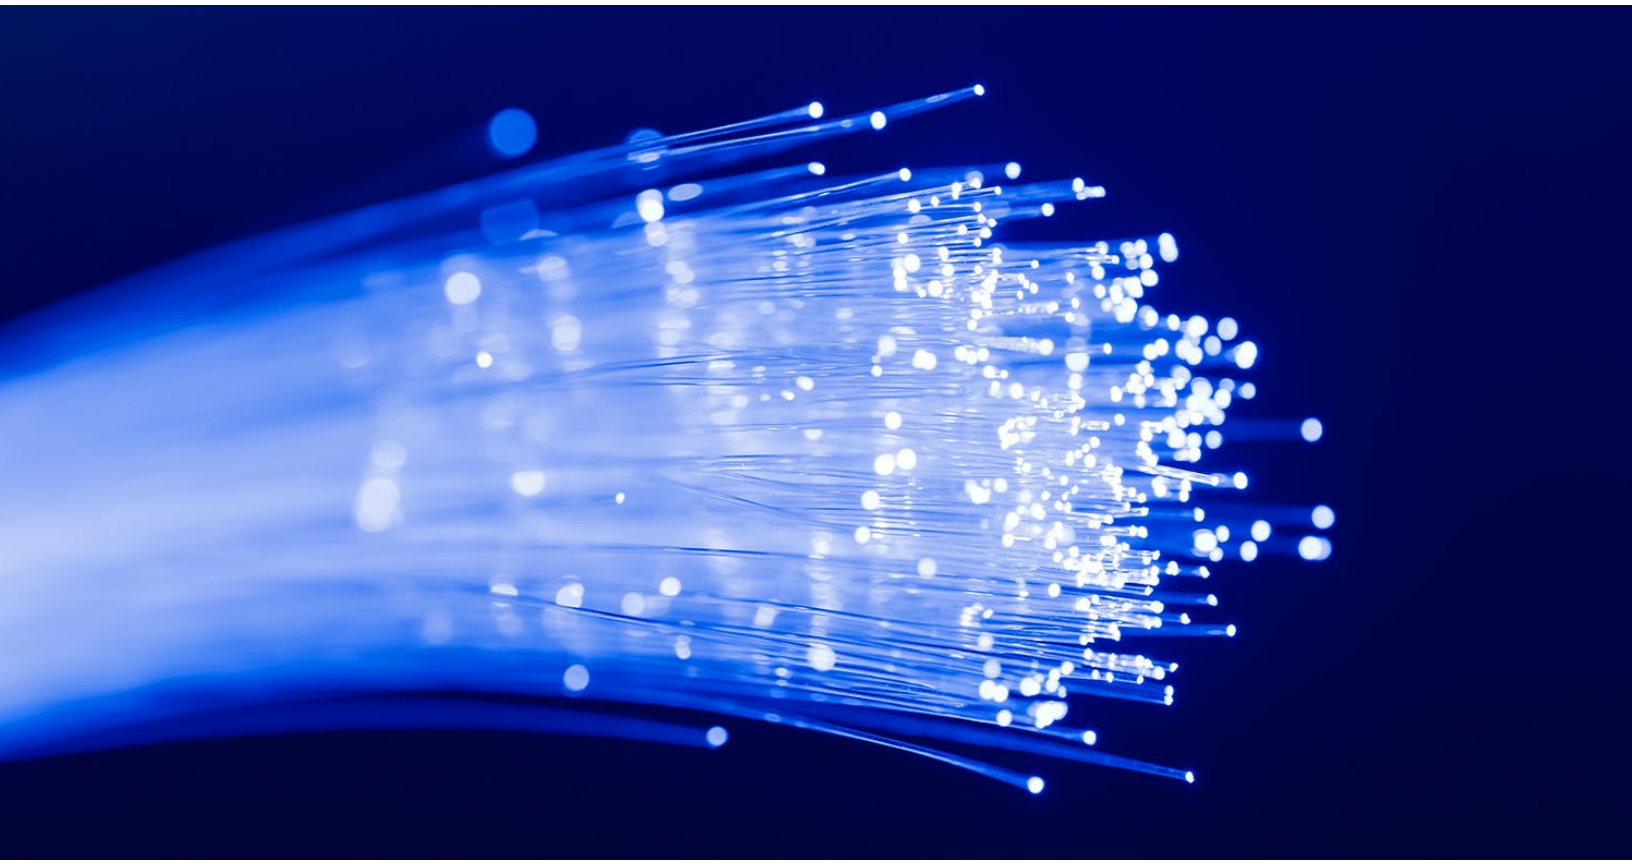

March 2026

AI's continued unprecedented growth depends heavily on accelerating investment in the physical and digital infrastructure required to train, deploy, and scale AI workloads—from data centers and fiber connectivity to intelligent networks, power, real estate, and accelerated, GPU-based compute. Historically, telcos have not captured a fair share of growth from tech disruptions that substantially increased data traffic such as the rise of video and social media consumption.

Today that dynamic is beginning to shift, giving telcos a new opportunity to play a central role in enabling the AI era. In the process, this could help reignite the industry's growth.[1]

The rapid adoption of generative and agentic AI is driving demand not only for centralized compute but also for distributed infrastructure closer to end users. Building and operating this infrastructure require assets that many telcos already have: extensive fiber networks, national footprints, space and presence at the edge, access to power, and experience managing complex, high-availability networks. The question is no longer whether telcos are relevant to AI infrastructure, but where along the value chain they can compete effectively and profitably.

What's at stake?

Between 2012 and 2025, global mobile data traffic grew by over 50 percent per year, while telecom service revenues barely increased by 1 percent annually, with hyperscalers capturing most of the economic value.

The AI wave offers telecom operators an opportunity to address that long-standing value capture challenge. Global data center demand could more than triple by 2030, reaching at least 170 gigawatts, driven largely by AI workloads. Importantly, this is not a single "bet-the-company" opportunity. The AI infrastructure value chain offers multiple entry points with different risk, capital intensity, and return profiles, allowing operators to align ambition with balance-sheet capacity and strategic intent.

Within this space, three value pools are particularly relevant for telcos:

— Fiber connectivity for new data centers in core and Tier 2 cities represents a global revenue opportunity of approximately $30 billion to $50 billion by 2030, as new facilities require high-capacity and often dark-fiber[2] connections.

— Beyond connectivity, opportunities are present in offering intelligent, software-defined network services to help enterprise customers manage AI workloads, control annual cloud data transfer fees, and meet latency and regulatory requirements. Commonly referred to as egress costs, cloud data transfer fees are estimated at $70 billion to $80 billion annually, creating scope for differentiated telco offerings focused on performance, cost efficiency, and control.

— Telcos may unlock new value by utilizing existing space, power, and cooling for distributed compute needs, including participation in the rapidly emerging GPU-as-a-service (GPUaaS) market, which is estimated to be worth $35 billion to $70 billion by 2030 (excluding hyperscalers). This opportunity is supported by accelerated compute workloads that are growing at more than 30 percent CAGR and expected to represent more than two-thirds of data center demand within five years.

[1] "AI infrastructure: A new growth avenue for telco operators," McKinsey, February 28, 2025.
[2] Dark fiber is unused or "unlighted" fiber optic cable.

Who are the key players and stakeholders?

The ecosystem is broad and increasingly competitive. Key stakeholders include hyperscalers and colocation providers, cloud operators, and GPUaaS players—many of which prefer dark fiber. Telcos enter this landscape as infrastructure owners and operators, often competing with—and simultaneously partnering alongside—hyperscalers. Other important players include the public sector; regulators that play a central role by shaping fiber access, data sovereignty, and licensing requirements; and key regulated industries that drive demand for sovereign AI.

Additional value chain participants include chip makers, software providers, systems integrators, and partners supporting data center retrofits and cooling solutions. Finally, energy providers and investors are increasingly influential, given power constraints and the capital intensity of AI-ready infrastructure.

What are the recent important developments?

As AI infrastructure scales and decentralizes, several developments are worth watching.

— **AI data center capacity is expanding and spreading to new geographies.** Hyperscalers and colocation providers have announced plans for more than 2,600 new data centers. Roughly one-quarter are to be located in cities with no existing data center footprint, pushing demand into new markets. Procurement preferences are also becoming clearer, with growing emphasis on dark fiber. Overall, the global footprint is expected to approach 11,000 facilities by the early 2030s.

— **Power and permitting constraints are reshaping where value is accruing.** Grid power limitations and long permitting timelines are slowing new builds, increasing the value of existing space and available power. This dynamic favors players that control "ready" assets—particularly sites with assured, contracted power in locations that can support distributed compute for inference.

— **Telco participation is shifting from pilots to execution.** Early movers among operators are signaling a transition from experimentation toward scaled delivery. Some are partnering with hyperscalers to provide dark fiber and edge compute; others are launching GPUaaS offerings or repurposing central offices or spaces with power and cooling to support AI workloads.

What are the biggest challenges?

Despite the opportunity, risks are substantial. There are still uncertainties around demand, particularly across use cases and geographies, as well as the technology and connectivity needs to enable computing at the scale required by AI providers. Competition from hyperscalers and cloud-native players is intense, with superior scale, capital, and developer ecosystems. And it's still not a given that investments can generate enough compute capacity to meet AI ambitions, while still delivering strong returns.

Commercial and competitive uncertainty also remains. Intelligent network services do not yet have a clearly established market definition, and telecom operators increasingly compete with hyperscalers and network service providers across connectivity, space and power, and GPUaaS offerings.

Capital requirements are high, especially for GPUaaS, dark fiber, and AI-ready data centers, while technology cycles are short and pricing pressure could intensify as supply expands. Many telcos lack the software, sales, and partnership capabilities required to compete beyond traditional connectivity.

What does it take to succeed?

Capturing value from AI infrastructure requires several shifts for telcos:

— Commercial capabilities including dedicated hyperscaler-focused sales teams, faster design and contracting cycles, digitized routing and asset data, and productized connectivity solutions are essential. GPU-as-a-service offerings require advisory-led selling and customer success capabilities. Close partnerships with hyperscalers are key to understand their evolving needs and new emerging technologies in this space.

— Partnering with and leveraging infrastructure investors, software providers, systems integrators, and retrofit specialists is key to lower capital needs, de-risk projects, and accelerate service delivery and facility readiness.

— Disciplined underwriting is important, including evaluation of fiber routes for multi-tenant upside and careful management of cost structures across compute, power, cooling, networking, and memory.

— Risk management for space and power, including prioritizing sites requiring limited retrofit or committing capital only after securing firm tenant demand, is critical; for edge inference, at least 500 kilowatts of power are typically required to achieve economic viability.

What should players be watching?

Telco leaders should monitor where hyperscalers and colocation providers expand into new markets and how power constraints reshape infrastructure economics. Operators have an opportunity to be early movers in the build-out of distributed compute or edge infrastructure, but much will depend on the level of demand for GPUaaS offerings and whether demand for intelligent network services picks up and can be successfully monetized. Ultimately, the biggest risk facing telcos may not be choosing the wrong path—but failing to choose at all.

Gustav Grundin is a partner in McKinsey's Helsinki office, and **Miguel Frade** is a partner in the New York office, where **Tomás Lajous** is a senior partner, and **Borja Belda** is an associate partner. **Sebastian Cubela** is a partner in the Miami office, and **Lorraine Salazar** is a senior knowledge expert in the Singapore office.

The authors wish to thank Arjita Bhan, Eduardo Zamudio, Manuel Palacios, Matyas Zetek, and Nicholas Shaw for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Issue brief

B2C growth beyond the core

The beyond-the-core model offers telcos an attractive path to B2C growth as traditional growth engines are slowing. Leaders are learning how to capture the opportunity while addressing the challenges.

This article is a collaborative effort by Davide Schiavotto and Santiago Fernandez, with Louis Maus, Mate Deutsch, and Tamas Kabay, representing views from McKinsey's Technology, Media & Telecommunications Practice.

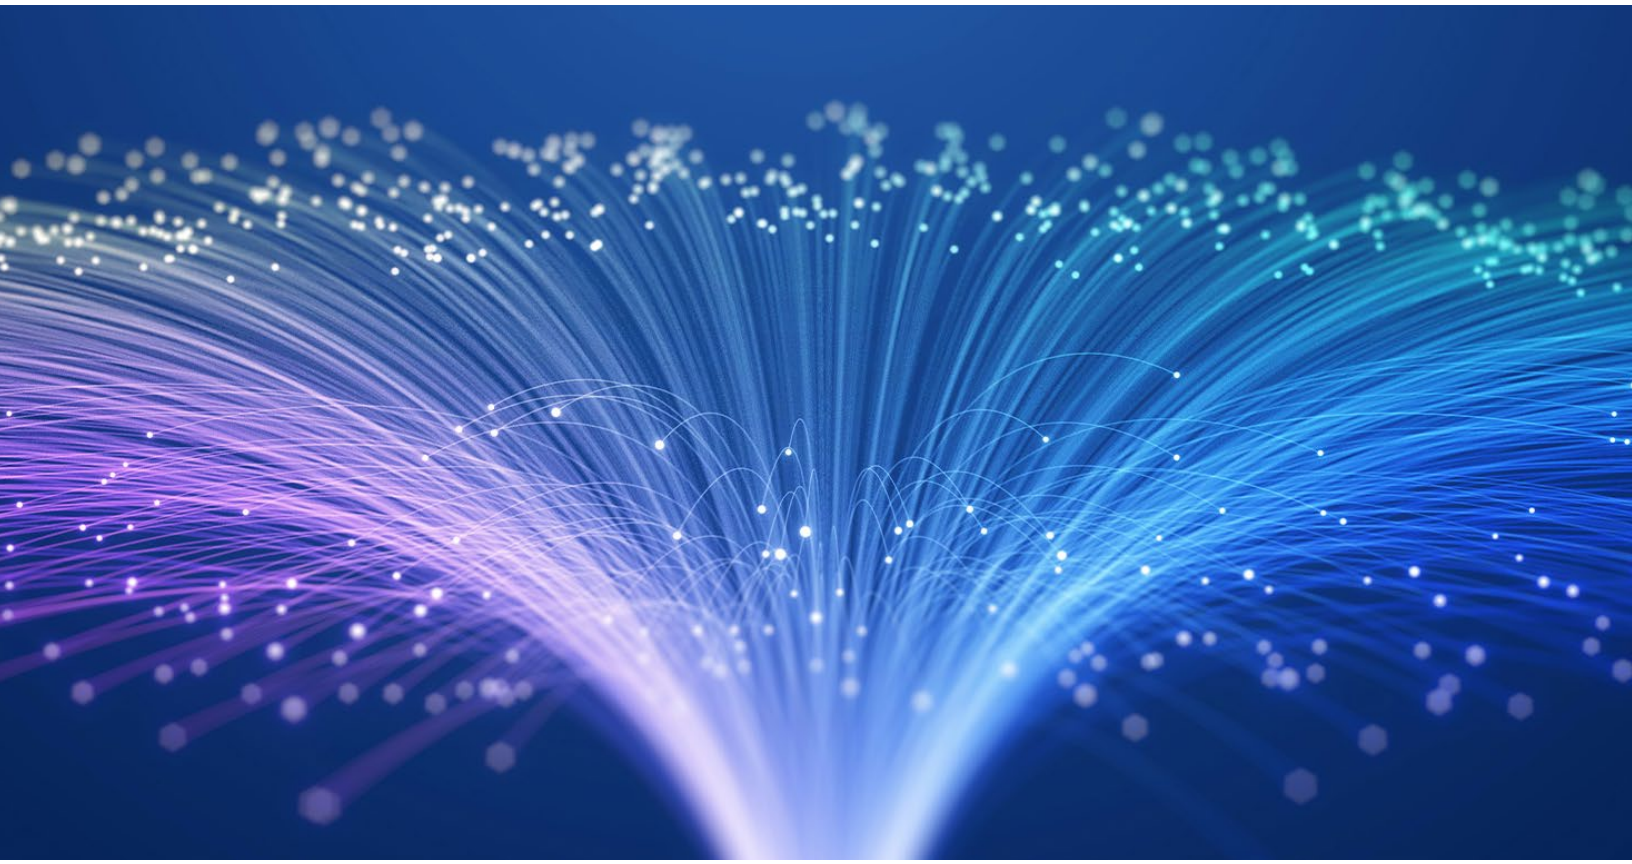

March 2026

Traditional B2C telco growth engines are increasingly constrained. Upsell opportunities face more challenges due to unlimited mobile voice and data plans, while fiber uptake is commoditizing previously premium, gigabit speed offerings. At the same time, many markets are running out of cross-sell opportunities due to fixed-mobile convergence (FMC) saturation; FMC has at least 50 percent penetration in major markets such as Belgium, Portugal, and Spain, with accompanying significant decline in linear TV viewership. Lastly, price increases are delivering decreasing returns, as less than 50 percent of headline price uplifts are typically retained once customers migrate to lower-priced, front-book offers.

As a result, telcos face a strategic question: whether to retain and maximize profitability via cost optimization levers or search for growth beyond the core. Our experience and research show that beyond-the-core value propositions can unlock new revenue streams, deepen customer engagement, and support long-term value creation; roughly 70 percent of consumers are willing to purchase or own at least one adjacent service via their telco provider.

Telcos are well positioned to capture this opportunity, given their trusted customer relationships, billing integration, extensive distribution and service footprint, and depth of proprietary assets such as data or existing partnerships. All these advantages are further accelerated by technological advancements (such as APIs and agentic AI) that make the expansion to new verticals easier than before. They are also reflected in rising telco investor expectations.

What's at stake?

The scale of the beyond-the-core opportunity extends well beyond incremental adjacency. By 2035, the integrated network economy (both B2B & B2C) is expected to represent a total global revenue pool of some $125 trillion. Within this, B2C-relevant connectivity adjacent domains such as housing (~$7.4 trillion), health (~$7.9 trillion), digital content (~$4.8 trillion), mobility (~$3.0 trillion), education (~$5.5 trillion), and B2C marketplaces (~$12.1 trillion), create meaningful headroom for telcos to expand their role in consumer value chains.

Beyond-the-core models also offer a structurally attractive financial profile. Distributor and B2B2C partnership models require limited incremental capital, leverage existing assets (retail, billing, customer care, field force), and can generate high-margin revenues while reinforcing the core business through higher engagement and reduced churn.

Importantly, investors increasingly recognize B2C and digital revenues as a driver of improved valuation multiples for integrated telcos relative to pure connectivity players.

What's happening?

Historically, beyond-the-core ventures have been difficult to execute. A previous McKinsey survey of top global telecom executives showed that a high share of new business builds failed or remained unprofitable in earlier waves of experimentation.[1] Recent evidence, however, shows that several leading telcos are now generating meaningful revenues from beyond-the-core activities. These operators derive between 5 and 40 percent of total revenues from adjacencies such as fintech, insurance, digital media, energy, and security, corresponding to annual revenues of ~€1 billion to ~€15 billion. These players demonstrate that beyond-the-core models can scale when pursued with focus, the right operating model, and strong ecosystem partnerships.

[1] "How telcos can succeed in launching new businesses beyond connectivity," McKinsey, February 25, 2022.

What does it take to succeed?

Despite growing traction, beyond-the-core growth remains challenging, with the primary barriers being organizational and operational rather than strategic. Traditional telco structures—designed for scale, reliability, and cost efficiency—often lack the speed, experimentation, and customer-centricity required to execute ecosystem plays effectively.

Leading players address these challenges through an innovative, platform-led approach with clear organizational focus. Once a telco has found the right vertical(s) relevant to its footprint, success depends on rigorous execution, including operating model (structure, process, people, technology) and link to core business P&L accountability and shared targets and incentives with core business.

If early results in these verticals are promising, longer-term, scalable success requires "bringing it together." This entails defining the overarching customer value proposition with strong ecosystem synergies. As part of this, it's critical to have the platform-first proposition design centered on orchestration, integration, and bundling, rather than as standalone products.

Davide Schiavotto is a senior partner in McKinsey's Milan office; **Santiago Fernandez** is a partner in the Madrid office; **Lous Mas** is a fellow associate in the Munich office; and **Mate Deutsch** is an associate partner in the Budapest office, where **Tamas Kabay** is a research science expert.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

Satellite connectivity

As LEO consumer broadband and D2D connectivity push satellite capacity into the telecom mainstream, telcos must decide where to compete, partner, or ignore.

by Brooke Stokes and Javier Gil Gomez, with Karol Pyrzak

March 2026

Satellites are back in the telecom industry conversation, as the topic of former last resort connectivity is now landing on the desks and strategic agendas of telecom operators around the world. Two markets are moving from "niche" to more mainstream and relevant at scale: (1) Low-Earth-Orbit (LEO) consumer broadband and (2) direct-to-device (D2D) connectivity. Both shifts ride the same recent supply-side step change: Cheaper launches, higher cadence, mass-produced spacecraft, and falling terminal costs are turning satellite capacity into something that can be scaled and packaged like a consumer product.

Satellites won't replace terrestrial networks, but they will reset customer connectivity expectations at the "edges" of fixed and mobile, and those boundaries are expanding as LEO capacity scales. Because satellite players compete with a different economic model, they introduce new strategic uncertainty to the market and leave telcos with some key questions to grapple with: where to compete, where to partner, and where to ignore.

1. Consumer satellite broadband

As LEO capacity grows, consumer satellite broadband's target markets could expand significantly.

What's on the horizon?

Geostationary Earth orbit (GEO) broadband served remote households because terrestrial last-mile economics were poor, but latency and quality of experience limited adoption. LEO changes the experience profile. Latency is structurally far lower than GEO, enabling more demanding applications. Performance still varies with cell loading and handovers, but the addressable segment is no longer limited to the fully unconnected.

Two main adoption axes will affect take-up:

— **Strength of competing fixed alternatives:** Where fiber-to-the-home (FTTH) and modern cable are widespread, LEO will skew to rural coverage. Where fixed access is scarce, slow to build, or low quality (legacy xDSL, congested fixed wireless access [FWA]), LEO becomes a credible primary option.

— **Pricing sophistication and affordability thresholds in some markets:** LEO constellations are deployed globally, but pricing can be highly local. Providers can flex prices to maximize beam utilization, for instance, discounting in underfilled areas and selectively undercutting incumbents. That could pressure telco pricing norms. Upfront terminal costs, however, remain a material barrier and will likely cap adoption in some emerging markets.

What are the challenges?

So far, LEO broadband has been capacity-constrained above a certain household density (beams saturate and peak-time performance degrades). But sector roadmaps—including Starlink's V3 next-gen satellites—point to step-changes in capacity as larger spacecraft and higher launch cadence arrive; Amazon LEO and other entrants will add supply alongside Starlink's existing fleet. As "capacity in the sky" rises, the density constraint should ease, pulling competition from rural into suburban markets (and selected urban niches).

What could happen?

— **Capacity unlock plus local pricing.** With next-gen capacity (such as Starlink V3) and new constellations, LEO pushes into suburban markets, using so-called beam fill, reduced pricing to discount selectively and win share.

— **"Served-on-paper" broadband is vulnerable.** LEO targets households with weak real-world experience (congested FWA, copper-based infrastructure), not just rural areas.

— **ARPU compression to fill the capacity.** As satellite capacity scales and competition intensifies, increased satellite pricing erosion may anchor lower price expectations in competitive, non-fiber technologies.

2. Direct-to-device connectivity

The extent of D2D satellite connectivity market growth depends on which delivery models and commercial patterns become dominant and how the sector overcomes physical challenges.

What's on the horizon?

D2D satellite connectivity is gaining attention not because it can replace terrestrial mobile service, but because it can cost-effectively close coverage gaps where propagation limits and site economics make towers impractical. Unlike consumer satellite broadband, D2D positions satellites as another mobile access layer, embedded in smartphones and integrated into spectrum, roaming, and plan design.

Two delivery models are emerging:

a. **Mobile network operator (MNO)-spectrum (partnered) model:** Satellites use an operator's terrestrial mobile bands via commercial partnerships (such as Starlink–T-Mobile; AST SpaceMobile with different operators). While this has the ability to reach many existing phones with minimal user behavior change, scaling is limited by the number of partnerships and requires market-by-market deals with tight regulatory and interference coordination.

b. **Mobile-satellite-spectrum (MSS) model:** Uses dedicated mobile-satellite spectrum, typically enabled through OEM/chipset integration (such as Apple–Globalstar or Skylo). This has the advantage of clearer spectrum rights and more consistent cross-border footprint, but it depends on device penetration and multi-party ecosystem alignment (OEMs, chipsets, satellite operators).

What are the physical challenges?

Three physical constraints define near-term D2D: (1) Handset uplink: Unmodified smartphones are built for nearby towers and are capped by power, antenna gain, thermal limits, and battery; satellites can boost downlink, but not the uplink. (2) Indoor coverage remains weak because building penetration and clutter erode link margin and users often lack clear sky view. (3) Dense, urban capacity is structurally limited: satellite "cells" are large and shared, so capacity per km^2 is far lower than terrestrial dense reuse. As a result, early D2D fits low-rate messaging, safety, and intermittent voice in rural/outdoor settings; high-rate data degrades quickly as user density rises.

What could happen?

D2D go-to-market models are defined by who owns the customer/brand, who controls the user experience, and how integration is delivered (direct versus hub/platform). Four primary commercial patterns are forming:

— **MNO-led add-on/wholesale non terrestrial network (NTN):** The operator owns the customer, sells D2D as a plan add-on, and buys satellite capacity directly from the provider.

— **OEM-led feature:** The device maker owns activation and UX (free safety first, likely paid tiers later); the operator plays a limited bundling role.

— **Satellite-led retail/hybrid mobile virtual network operator (MVNO):** The LEO player owns the customer under its own mobile brand, uses terrestrial networks as MVNO/roaming for most traffic, and adds satellite as edge coverage.

— **Hub/platform-led enablement:** A third party provides a "one integration" layer (authentication/routing/charging/APIs) and can aggregate multiple satellite networks, accelerating multioperator and multicountry rollout.

There are three main scenarios for operators to plan for:

— **Safety becomes table stakes.** As emergency messaging/location standardizes, operators bundle and resell the service, with customers expecting the feature as part of premium plans.

— **NTN becomes "just another roaming layer."** Hubs normalize integration and accelerate cross-market launches.

— **Spectrum and regulation set the pace.** MNO model scales through partnerships; MSS scales with chipset penetration.

What can telcos do next?

1. **Map "underserved but paying" pockets** in fixed and mobile by micro-geography, where complaints, usage patterns, and churn signal a readiness to switch.

2. **Choose posture by segment and market.** Compete or partner (bundle/roaming) depending on market conditions and ignore altogether only where the market is genuinely saturated with high-quality fixed and mobile.

3. **Protect the experience layer.** Even when capacity is wholesale, telcos can win on packaging: plan design, billing, care, device programs, and predictable tiers.

Satellite connectivity won't replace terrestrial networks. But it will change what "good enough" looks like at the edges, and who owns the customer there.

Brooke Stokes is a partner in McKinsey's Southern California office, **Javier Gil Gomez** is a partner in the Madrid office, and **Karol Pyrzak** is a capabilities and insights specialist in the Wroclaw office.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

Sovereign AI

Despite huge waves of momentum and large spend commitments, many sovereign AI initiatives embody intent without readiness. Sequencing is key to making sovereignty repeatable and adoption scalable.

This article is a collaborative effort by Alex Panas, Ali Ustun, Arnaud Tournesac, and Luca Bennici, with Justin De Niese, Kaavini Takkar, and Newfel Drahmoune, representing views from McKinsey's Technology, Media & Telecommunications and Public Sector Practices and QuantumBlack, AI by McKinsey.

March 2026

Disclaimer: *This article is descriptive and analytical. It does not provide policy, regulatory, or national security advice, and it does not recommend specific national strategies. Examples are included for illustration of ecosystem patterns and should not be interpreted as endorsements or prescriptions.*

Over the last year or two, sovereign AI has moved from a niche policy concept to what many view as a strategic imperative. Business and political leaders increasingly see the capability to deploy and scale AI while retaining appropriate control across key dimensions of sovereignty as central to economic competitiveness, national security, and societal trust. But sovereignty is multi-dimensional (operational, technological, legal, and territorial). The practical goal is to apply sovereignty deliberately at the control points that drive trust, auditability, and accountability, rather than maximizing sovereignty everywhere. Yet many initiatives stall because they build "sovereign assets" (data centers, GPUs, sovereign regions, national model announcements) faster than they build the demand, governance, and trust needed for scaled adoption.

What's at stake?

Sovereign AI represents one of the largest opportunities in AI: Up to 30 to 40 percent of AI spending is influenced by sovereign requirements, comprising an estimated $500 billion to $600 billion market by 2030. But the strategic prize is broader than spend. The outcomes include (1) resilience and autonomy for critical services (reducing exposure to geopolitical shocks), (2) end-to-end risk management across data, models, and applications with controlled access and auditability, and (3) economic value capture through productivity and service-quality improvements, startup and scale-up formation, higher-value jobs, and new digital exports as ecosystems mature.

Who are the key players/stakeholders?

The sovereign AI ecosystem comprises four groups, each of which control a different constraint on progress:

— **Government (orchestrator, regulator, investor, anchor customer):** Governments set sovereignty goalposts and make trade-offs explicit (which workloads require strong sovereignty versus hybrid versus largely global). They build a "trust market" through standards, certification, and enforcement so adoption becomes repeatable, not bespoke. They catalyze supply through enabling policy and targeted investment (including permitting and energy planning). And they aggregate demand into multi-year frameworks and lighthouse programs to create an adoption flywheel and justify early investment.

— **Technology providers:** Providers span hyperscalers, regional cloud and data center operators, telecom operators, model developers, and integrators. Effective ecosystems rarely pick either hyperscalers or local providers. Instead, they design an architecture where players collaborate and compete at the layer where each has advantage, while localizing trust through sovereign controls.

— **Investors:** Investors bridge the "valley of uncertainty" before utilization is proven, funding the stack across energy, connectivity, data centers, cloud services, model development, applications, cybersecurity, and integration.

— **Enterprises:** Enterprises convert infrastructure into economic value. They decide where sovereignty truly matters, tier workloads accordingly, and redesign governance and operating models so AI can scale beyond pilots.

What are the recent important developments?

Three signals stand out in our research and survey results:

1. **Momentum and headlines.** Commitments are surging—billions toward GPUs, data centers, and sovereign regions—yet scaled outcomes lag.

2. **Demand is real but pragmatic.** More than 70 percent of enterprises see sovereign AI as strategically important, but decisions remain workload-specific and dominated by price, performance, and reliability.

3. **Execution is the constraint.** Many organizations have intent but lack the operating model and "trust fabric" to move regulated workloads quickly and repeatedly.

What are the biggest challenges?

Sovereign initiatives can fall victim to mis-sequencing: building assets before demand and governance are ready. Practically, this shows up as:

— **Intent without readiness:** While 72 percent of enterprises say sovereign AI is part of their 2026 roadmap, only 28 to 29 percent report having a sovereign AI strategy, a detailed action plan, or a budget set aside—and only 25 percent have workload tiering by sovereignty levels.

— **A perceived cost premium:** Sovereign solutions are widely perceived as 10 to 30 percent more expensive across the stack, making buyers selective about where sovereignty is worth paying for.

— **Slow transformation cycles:** Sovereign cloud and AI migrations typically take three to four years, driven less by technology constraints and more by governance redesign, risk/compliance alignment, operating model shifts, and workforce retraining.

— **Fragmentation and lock-in:** Without shared standards and interoperability, deployments become bespoke, slowing procurement and scaling.

What are the key success factors and critical enablers?

A set of key moves can make sovereignty repeatable and adoption scalable:

1. **Build a scalable trust fabric:** Codify sovereignty into standards, reference architectures, and certification regimes (including model governance expectations such as evaluation, monitoring, and documentation), so regulated industries can adopt quickly and repeatedly.

2. **Make sovereignty a portfolio decision:** Segment workloads into sovereign, hybrid, and global lanes to avoid all-or-nothing debates and focus controls where risk and exposure demand them.

3. **Design for interoperability:** Build portability and common control patterns so enterprises and startups can scale through interoperability rather than bespoke lock-in.

4. **Sequence execution in three waves:** A practical roadmap is: ***Wave 1:*** Baseline controls plus lighthouse demand; ***Wave 2:*** Scale shared infrastructure plus data ecosystems; ***Wave 3:*** Specialize in selected domains/languages and export solutions once the foundation is proven.

What should players be watching over the coming year?

Watch whether ecosystems shift from announcements to adoption mechanics: credible certification regimes and enforcement; governments bundling demand into multi-year frameworks; providers narrowing the perceived 10 to 30 percent premium via standardized, auditable controls and scale; and enterprises accelerating workload tiering plus operating model changes that shorten multi-year migrations. The leaders will be those that treat sovereign AI as an ecosystem build—sequencing demand, governance, and trust in step with infrastructure investment.

Alex Panas is a senior partner in McKinsey's Boston office; **Ali Ustun** is a senior partner in the Doha office; **Arnaud Tournesac** is a partner in the Paris office; **Luca Bennici** is a partner in the Dubai office, where **Justin De Niese** is an associate partner, and **Newfel Drahmoune** is a consultant; and **Kaavini Takkar** is an associate partner in the Seattle office.

This article was edited by Larry Kanter, a senior editor in the New York office.

McKinsey & Company

Technology, Media & Telecommunications Practice

Telecom infrastructure

Telecom infrastructure is at an inflection point due to slower rollout and other growth challenges. Success creating value depends on whether players can transition toward a more integrated platform.

by Luca Fiandro and Nemanja Vucevic, with Lorraine Salazar and Shu Chern Lim

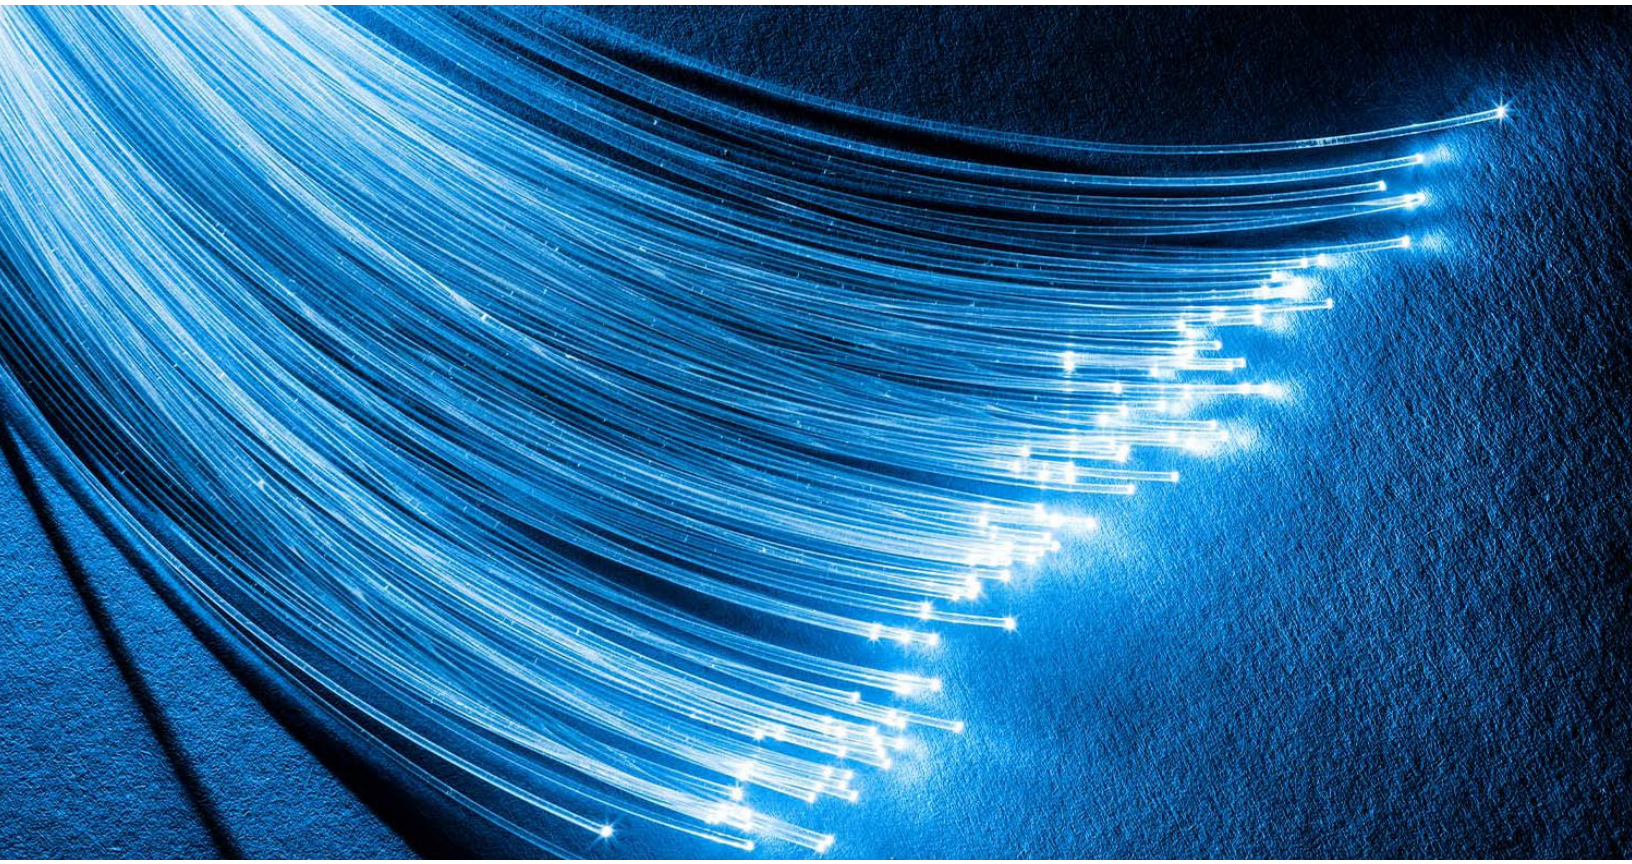

March 2026

Telecommunications infrastructure—anchored by towers and fiber networks—has long been a cornerstone of digital connectivity and a resilient investment class. Today both asset types face a shared inflection point: Traditional growth engines are slowing, even as demand for connectivity accelerates and diversifies, driven by data-intensive applications and new compute architectures.

Over the past two decades, telecommunications towers (which provide the physical locations for radio equipment that enables mobile connectivity) have benefited from rapid mobile adoption, network densification, and growing tenancy. Fiber networks, which form the backbone of fixed broadband, mobile backhaul, enterprise connectivity, and long-haul connectivity, have replaced copper as the future-proof wireline technology, delivering stable, infrastructure-like returns.

As coverage has expanded and markets mature, however, both towers and fiber face slower rollout and challenges to maintain utilization growth. These dynamics are pushing operators and investors to rethink how value is created, sustained, and differentiated across the infrastructure stack.

Though both remain foundational assets, the next phase of value creation will favor players that move beyond passive ownership and reposition their infrastructure as dynamic platforms at the center of the digital economy.

What's at stake?

Despite slowing deployment, infrastructure executives remain optimistic. In towers, 98 percent of executives expect long-run profitability to increase, reflecting confidence in the sector's ability to generate a second wave of value. For towers, the opportunity lies less in adding new sites and more in extracting additional value from existing assets. Future growth will depend on tenancy optimization, selective densification, and expansion into adjacent services including small cells, private networks, and edge data centers, rather than footprint expansion.

Fiber coverage has similarly reached scale—with approximately 80 percent of Europe, 70 percent of Asia–Pacific, and 60 percent of North America. Executives nevertheless expect rollout to continue, adding roughly 15 percentage points of additional coverage globally over the next five years. While rural and subsidized builds remain relevant, a new stimulus for growth is emerging: the rapid expansion of data centers, particularly in Tier 2 cities, where grid constraints and compute demand are reshaping fiber routing and investment priorities. Beyond coverage, the largest value pool lies in improving take-up, minimizing future duplication, and expanding the range of services delivered over fiber infrastructure.

Who are the key players and stakeholders?

In towers, key stakeholders include independent and telco-owned tower companies; mobile network operators as anchor tenants; infrastructure investors; and regulators governing site permitting and spectrum usage. Public authorities and enterprise customers are increasingly relevant as nontraditional tenants. More recently, there is growing interest in towers as potential locations for distributed compute supporting AI inference workloads, although this remains at an early stage.

The fiber ecosystem is broader and increasingly interconnected. It includes fiber operators, fixed and mobile service providers, infrastructure investors, hyperscalers, enterprise customers, public-sector stakeholders, and data center operators. Telecom operators remain the primary deployers of fiber infrastructure, but ownership and influence from specialized infrastructure players and investors continue to grow. Regulators play a critical role, shaping wholesale access models, fiber migration policies, and rural deployment subsidies. Adjacent participants—such as energy providers, private-network users, and edge data center customers—are becoming more important as fiber players extend their scope.

What are the recent important developments?

First, tower valuations have declined, with average transaction multiples falling from approximately 22 times EBITDA in 2022 to around 17 times in 2025, reflecting market maturity and slower organic growth. In response, attention is shifting away from pure site expansion toward active equipment, energy management, and operational efficiency.

Fiber deal multiples have likewise declined, albeit at a more modest pace, from 20 times EBITDA in 2022 to around 19 times in 2025. Fiber deal activity has also slowed, dropping at an annual rate of 20 percent between 2022 and 2025. Owners appear reluctant to sell at lower prices, leading to longer holding periods.

Third, massive investment in data centers to support AI workloads—both training and inference—is reshaping fiber demand. Constraints on grid power in major hubs are driving new subsea, long-haul, and metro fiber builds, particularly toward secondary markets and Tier-2 cities that can support incremental compute capacity.

Finally, low-Earth-orbit (LEO) satellites are emerging as a complementary technology. Executives expect LEO to capture close to 20 percent of existing fiber's addressable market, particularly in hard-to-reach areas, reinforcing a hybrid connectivity model rather than direct substitution.

What are the biggest challenges?

Both asset classes face structural constraints. Towers confront expiring or renewing anchor-tenant contracts and limited greenfield rollout opportunities, putting pressure on revenue growth. Fiber networks—particularly consumer-focused builds—struggle with low take-up rates, often near or below the roughly 30 percent threshold required for sustainable profitability. This challenge is compounded by overlapping networks and slow migration from legacy copper infrastructure.

Across both sectors, expansion into new services—such as active equipment, energy management, or enterprise connectivity—requires new capabilities and presents new risks, altering the traditionally infrastructure-like investment profile. Regulatory complexity, overbuilding, and operational inefficiencies further complicate value capture.

What does it take to succeed?

Six levers underpin successful value creation:

1. **Consolidation and optimization** to unlock scale efficiencies and avoid duplicative networks

2. **Selective new rollouts,** using more efficient designs such as small cells, "street work," or targeted expansion into Tier 2 cities to support AI infrastructure growth

3. **Extending the scope of services,** including radio access networks (RANs), energy management, and operations and maintenance for tower operators; and active services, enterprise connectivity, and cloud interconnection for fiber players

4. **Growth in adjacent areas,** such as small cells, distributed antenna systems, private networks, and edge data centers

5. **Boosting tenancy and take-up** through wholesale models, faster fiber migration, and nontraditional tenants

6. **Operational efficiency,** leveraging AI, analytics, standardized processes, and energy optimization, to reduce costs

What should leaders be watching over the coming year?

Key developments include the pace of tower and fiber consolidations; adoption of active-equipment and energy-management models; regulatory progress on fiber migration; the real impact of LEO satellites on capital allocation; and the accelerating race to build fiber infrastructure that supports AI-driven data center growth.

More broadly, success will hinge on whether infrastructure players can transition from narrowly defined, real-estate-like models toward integrated, engineering-led platforms—while preserving the stability investors expect.

Luca Fiandro is a partner in McKinsey's Milan office, **Nemanja Vucevic** is a partner in the Madrid office, **Lorraine Salazar** is a senior knowledge expert in the Singapore office, and **Shu Chern Lim** is a capabilities and insights expert in the Boston office.

The authors wish to thank Aleksander Matlok, Alex Mestre, Asif Khan, Boris Gusev, Federico Mentesana, Gerardo de Geest, Gustav Grundin, Jiaqi Zhang, Kah Man Leo, Lieven Verboven, Margherita Piccini, and Ming Shen Oon for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

Designed by Darby

McKinsey.com

@McKinsey
@McKinsey
@McKinsey