comviva
A TECH MAHINDRA COMPANY

DIGITAL RUSH 3.0

Designing for Autonomy



WHY ENTERPRISE PROCESSES, NOT AI MODELS,
ARE THE BOTTLENECK TO AGENTIC SCALE

Executive Summary

News stories, analyst reports, and boardroom reviews all point to the same issue: while AI pilot projects often look promising, many larger rollouts fall short. What's been missing is clear evidence explaining why this happens and how to fix it. That's what this study set out to uncover.

We looked closely at the AI value gap and found that

ONLY

13%

of companies say more than three-quarters of their AI projects delivered the expected ROI on time.

Half of those surveyed say that process architecture, workflows built for people, not AI, is a main barrier to scaling up.

At the same time, 65% admit their companies mostly use AI to automate current tasks instead of rethinking how work could be done. Our analysis shows the value of changing this approach: **companies that redesigned their processes specifically for AI are almost three times more likely than others to achieve expected ROI on over 75% of their AI projects.** The business challenge is clear: while three out of four leaders want AI to create offers on the fly, only 12% have agents that can finish a transaction.

To chart a clear path forward, we looked at four key areas needed to scale AI: Process, Architecture, Composition, and Enablement, together known as PACE. Process means redesigning core BSS workflows to work natively with agents, rather than just improving steps meant for people. Architecture is about building an API-first, standards-based, real-time data foundation that agents can use safely. Composition involves moving away from fixed catalogues and scripted channels to offers and transactions that are put together dynamically and handled by autonomous systems. Enablement covers governance, guardrails, skills, and new roles. It is not just a final step, but a layer that supports the other three and keeps autonomy safe as it grows.

If your AI program is stuck between a successful pilot and a rollout that has lost momentum, the following findings will help explain the reasons and show what leading operators are doing differently. This research makes one thing clear: scaling agentic AI is mainly a process-design challenge, and only after that a technology challenge.



Key Findings

Six findings frame the analysis that follows.

01

Agent-to-agent commerce is emerging as the next big step, but most operators are not prepared.

While 75% of leaders think AI should create personalized offers instead of relying on a set catalogue (with 38% strongly agreeing), only 12% have deployed autonomous agents that can understand intent, create offers, and complete transactions. Although 58% are starting to plan for customer AI agents to interact with their systems, only 11% are actively preparing for this change.

02

Redesigning processes pays off, but partial efforts do not.

Only 13% of all respondents see the expected return on investment from more than 75% of their AI projects. For organizations that fully redesigned their processes for agent-native operations, 28% saw these returns, compared to just 10% of those who only optimized existing processes. Making only partial changes does not provide any ROI advantage over simple optimization.

03

The main challenge is outdated processes,yet budgets do not reflect this.

About 28% say that legacy process architecture is the biggest barrier to scaling AI, even more than data quality (23%), and half see it as a main constraint. Still, most planned investments go to AI capabilities and platforms (27% on average), while only 21% is set aside for process redesign.

04

Operators are stuck in the automation trap.

About 65% agree that their company mostly automates current processes instead of rethinking how work is done. Nearly half (49%) focus on optimizing what already exists, and only 16% are fully redesigning processes for agent-native operations.

05

Ambition is outpacing actual adoption.

While 74% expect agentic AI to manage over half of routine operations in the next five years, and 83% see multi-agent orchestration as very important or critical, 66% are still in the early stages of exploring or testing AI. Currently, agentic AI is used in only 15% of operations.

06

There is growing agreement on what makes systems ready for agents, but a gap remains in standards.

API-first architecture (22%) and real-time data access (19%) are seen as the most important traits. About 66% have adopted or are considering TM Forum Open APIs, and 48% are doing the same with MCP-class agent-to-tool interfaces. Still, only 23% say standards are a key part of their agentic strategy, and just 22% consider their data foundation to be highly mature.

The Agentic Paradox

Rising Investment, Thin Returns

AI investment in the telecom industry covers many areas, but few companies have gone deep. Only 14% have made AI a core part of their operations, and 20% are using it in several areas. Most organizations are still testing AI in one or two areas (31%) or are just exploring it (35%) (see Figure 1). This pattern also appears in the types of AI used. Generative AI is the most common at 56%, followed by traditional and discriminative AI at 39%, and RPA at 31%. Fully agentic AI, which means autonomous agents that can plan, decide, and carry out tasks, is used in only 15% of operations. Most deployments focus on customer-facing areas: customer channels (62%), customer management (57%), and market and sales (55%) are the top uses. Other areas, such as business decision intelligence (42%), product management (40%), revenue management (34%), digital services (33%), business assurance (21%), and B2B2X (9%), are less common.

FIGURE 1

AI Investment Trajectory:
where is your organization in its AI investment journey for core BSS operations?



Figure 1. n=201; Source: Comviva Research Centre.

Returns are falling behind expectations

When asked how many of their AI projects met the expected ROI on time, 41% of respondents said that between a quarter and half did, while 34% said between half and three-quarters did (see Figure 2). Only 13% saw expected returns on more than 75% of their projects, and 10% saw them on less than 25%. Most companies are seeing planned returns from less than half of their AI investments.

FIGURE 2

AI ROI Reality:
approximately what percentage of your AI initiatives have delivered the expected ROI within the originally planned timeframe?

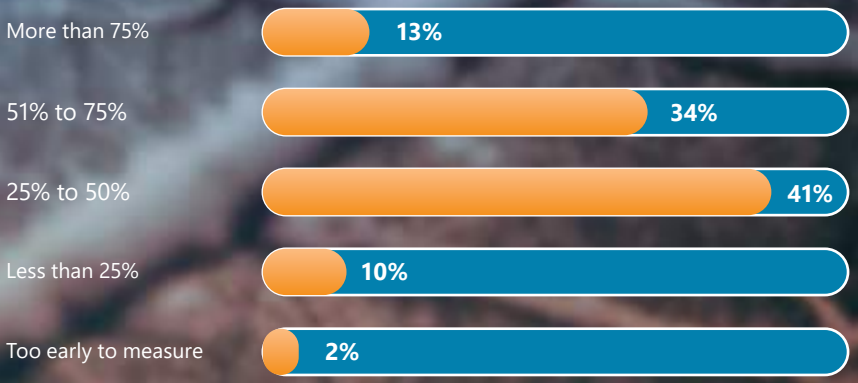


Figure 2. n=201; Source: Comviva Research Centre.

The Agentic Paradox

Rising Investment, Thin Returns

AI has delivered real, but limited, value so far. Respondents said the biggest benefit has been reducing operational costs (22%), followed by automating repetitive tasks (18%) and speeding up the launch of new products (15%). Only 5% said AI helped increase revenue through personalization or upselling. Right now, AI mainly helps companies work more efficiently rather than drive growth.



The Scaling Wall:
what is the single biggest barrier preventing your organization from scaling AI across operations?

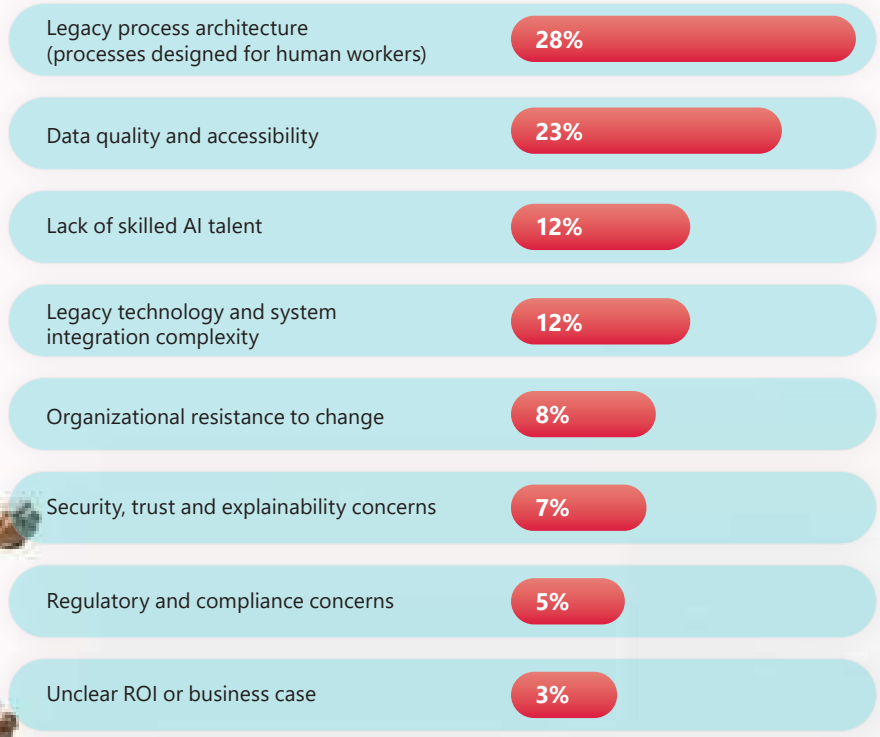


Figure 3. n=201; Source: Comviva Research Centre.

The cause is process, not the model

When leaders talk about the biggest barrier to scaling AI, most do not mention budgets or algorithms. Instead, 28% say the main issue is legacy process architecture, systems, and workflows designed for people, not agents. Data quality and accessibility come next at 23% (see Figure 3). Legacy technology, system integration complexity, and a lack of skilled AI talent tie for third at 12% each.

The urgency comes from the gap between where we are now and where we expect to be. **Agentic AI is used in only 15% of operations today, but 83% of leaders think multi-agent orchestration will be very important or critical in the next three to five years.**

Also, 74% expect agentic AI to handle more than half of routine operations within five years. This gap cannot be closed just by buying more models. Telcos need to redesign the processes, data, and controls that agents will use. The rest of this report looks at the four PACE dimensions, Process, Architecture, Composition, and Enablement, to show where value is stuck and how to unlock it.

Process

Redesign the Work, Don't Just Automate It

The survey shows that many telecom companies use AI in processes that were not built for it. Respondents say that process architecture not designed for AI is the main reason for the AI value gap, making up 31% on average. Half of those surveyed see process architecture as the main barrier to scaling AI, while another 39% see it as a secondary issue (see Figure 4). When asked to pick just one barrier, 28% say it is the biggest challenge.

Process Architecture Awareness:
To what extent does your organization recognize process architecture as a distinct constraint on AI scaling?

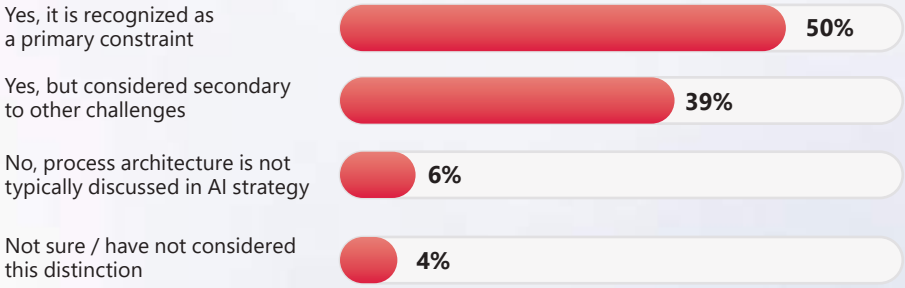


Figure 4. n=201; Source: Comviva Research Centre.

Optimization, not reinvention

Most operators have used AI to improve their current processes instead of starting from scratch. About 49% say their main approach is to optimize what already exists, automate steps, and add intelligence. Only 22% have partially redesigned processes for AI, 13% use a mix of methods, and just 16% have fully redesigned processes for AI-first, agent-native operations (see Figure 5). Many recognize this as a problem: 65% say their company mostly uses AI to automate existing processes instead of rethinking how work should be done. When companies automate workflows that still rely on human handoffs, approvals, and unwritten knowledge, they keep the same limits that block true autonomy.

Redesign vs. Optimize:
for processes where AI has been deployed, what approach has your organization primarily taken?

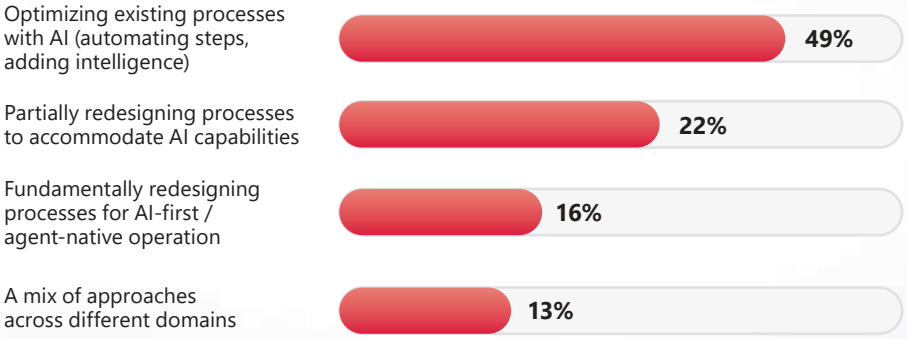


Figure 5. n=201; Source: Comviva Research Centre.

Where the friction concentrates

The areas that could benefit most from autonomy are often the ones where processes designed for people are hardest to change. In these places, old workflows create the biggest challenges for AI, even though the potential is there. Survey respondents ranked customer service and case resolution highest (52%), followed by billing, invoicing, and settlement (43%), and then fraud detection, revenue assurance, and risk (40%) (see Figure 6). Next come order management and fulfillment (33%), lead-to-order sales execution (30%), and partner or supplier onboarding (30%). These are all high-volume, rules-based processes where autonomy should bring quick results. The real barrier is the existing workflow, not the lack of opportunity.



Process-Level Bottlenecks:
Which of the following processes represent the greatest bottlenecks to scaling AI in your organization?

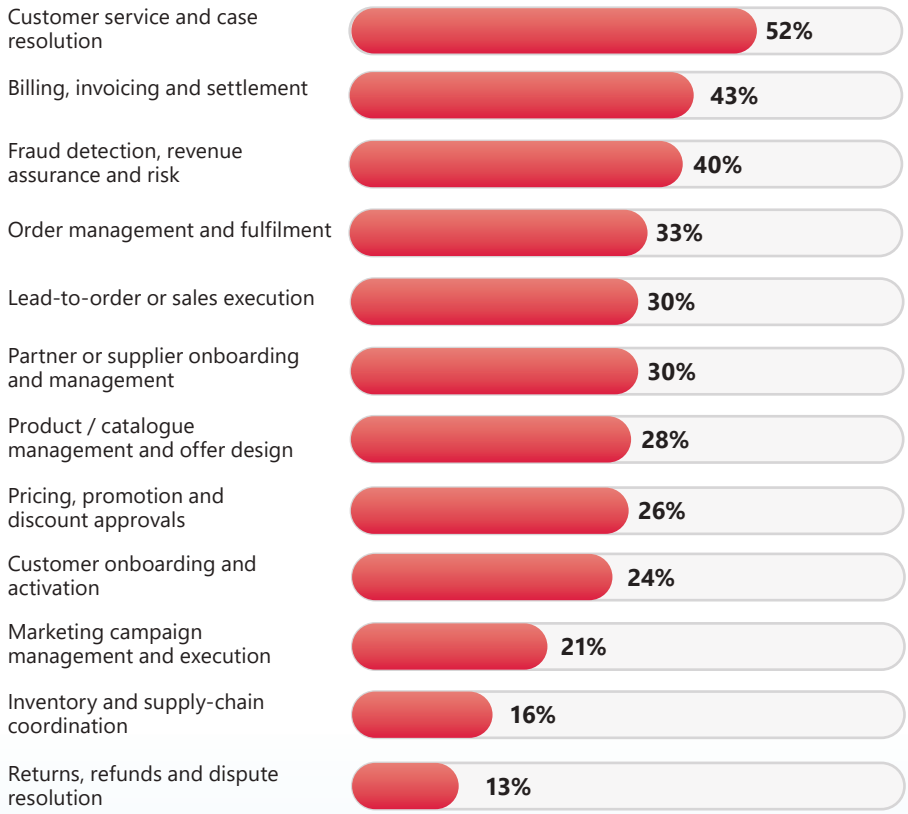
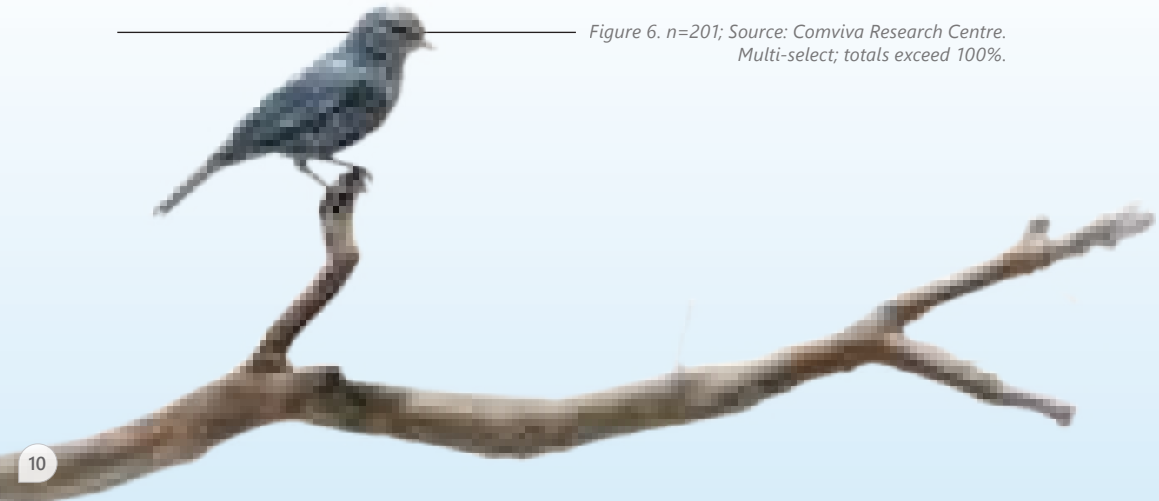


Figure 6. n=201; Source: Comviva Research Centre. Multi-select; totals exceed 100%.



The redesign dividend

Organizations that fully redesign their processes for agent-native operation are almost three times more likely than those making smaller changes to see expected returns on over 75% of their AI projects (see Figure 7).

The benefits of redesign are not gradual, as they only show up with full commitment. Partial changes keep the old process's limits and add the hassle of change, but do not unlock the real value of autonomy.

The Redesign Dividend:
share of AI initiatives that met expected ROI, by transformation approach

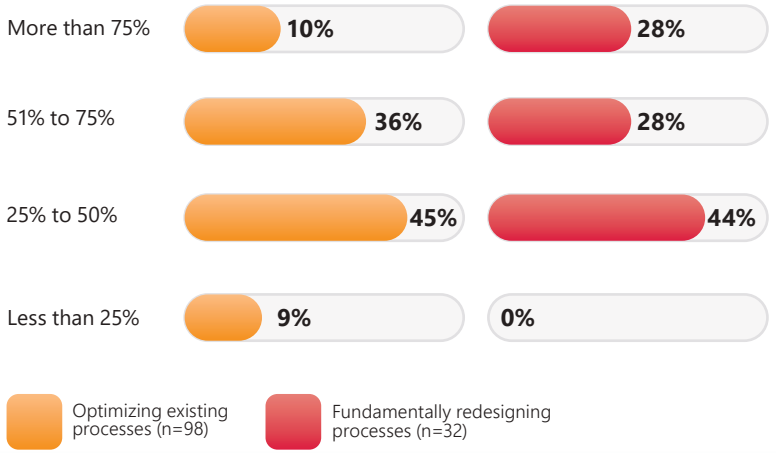


Figure 7. n=201; Source: Comviva Research Centre. Bases: fundamentally redesigning n=32, optimizing n=98

In practice

Vodafone Italy and Fastweb redesigned customer support around an agent-based system rather than just adding a chatbot to the existing setup. Using LangChain's agent framework (LangGraph), the agent manages complex support tasks end to end, with 90% response accuracy and 82% of handled issues resolved. The takeaway is the sequence: the process was decomposed so an agent could own it first; the model came second.

Source: LangChain, "How Vodafone Italy Built an Agentic RAG System for Customer Support", https://www.langchain.com/blog/customers-vodafone-italy

A Architecture

Build the Foundation Agents Can Act On

Agents carry out processes, but the architecture enables them to do so. Even with a redesigned process, things can still go wrong if agents cannot access data in real time or depend on reliable system interfaces. The survey shows that operators recognize this and are moving toward a standards-based, API-first approach.

The stack is converging on open standards

When looking at how organizations enable AI agents to work with core systems, 66% of respondents mention using open, industry-standard APIs (TM Forum). This is higher than the 53% who use internal proprietary APIs and microservices (see Figure 8). Almost half (48%) have adopted or are considering Model Context Protocol (MCP) or similar agent-to-tool interfaces. Other approaches include event-driven architecture or streaming platforms (42%), RPA bridges to legacy systems (36%), and Agent-to-Agent (A2A) communication protocols (30%). While many are exploring these options, fewer are fully committed. Only 23% say industry standards are central to their strategy, and 46% see them as important but have adopted them only to a limited extent so far.

Agent-System Integration Approach:
which technical approaches has your organization adopted or is actively evaluating to enable AI agents to interact with core systems?

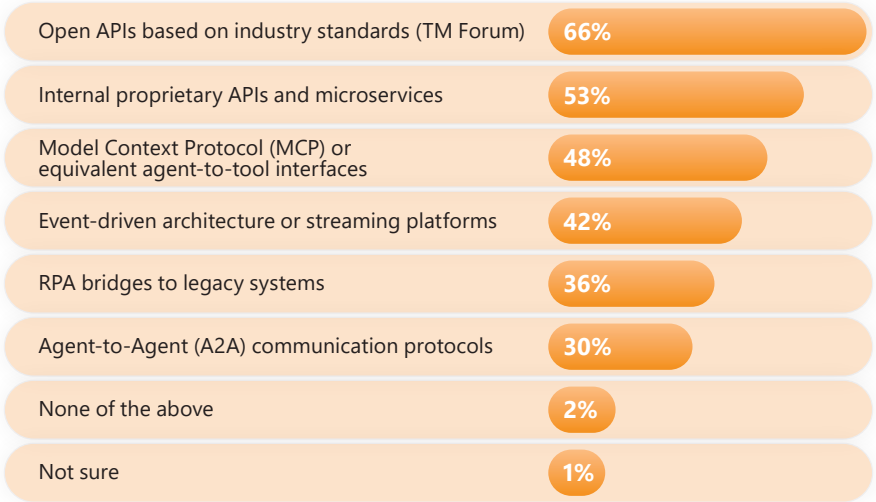


Figure 8. n=201; Source: Comviva Research Centre. Multi-select; totals exceed 100%.

Budgets point at the foundation

Planned investments mostly match the identified needs, but not completely. Over the next two years, respondents expect to spend an average of 27% on AI capabilities, models, and platforms; 22% on data quality, accessibility, and governance; 21% on process redesign and re-architecture; 16% on talent development and organizational change; and 14% on governance, compliance, and security.

About 60% of the average budget goes to the basics: data, processes, and the people needed to improve them, rather than just models. However, when compared to where respondents see the biggest value gaps, there is a clear mismatch. Models get 27% of planned spending but are blamed for only 13% of the gap, while process redesign gets 21% of spending but is blamed for 31% (see Figure 9).

The Diagnosis-Budget Gap.
Where leaders attribute the AI value gap vs. where they plan to invest (average % allocation)

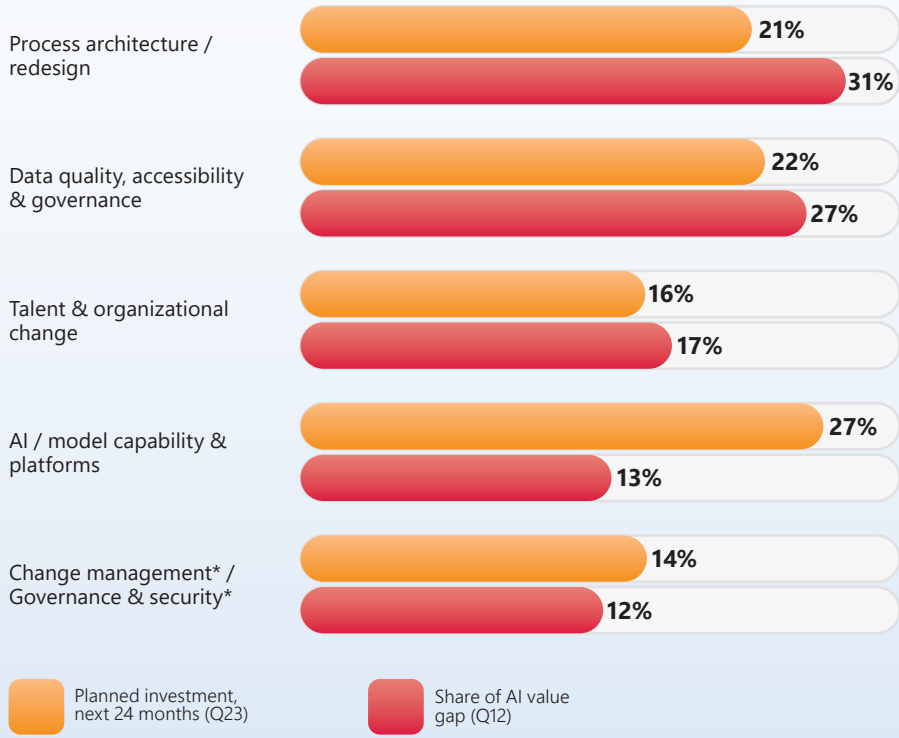


Figure 9. n=201; Source: Comviva Research Centre.

Data foundation is the weakest layer

Only 22% of organizations consider their data foundation highly mature, meaning their data is real-time, unified, and consistent across different areas. Another 36% say their data is mostly mature but still has some gaps, 24% report a mix of strengths and weaknesses, and 18% are at an early stage with major fragmentation and quality problems (see Figure 10). Agents tend to make data issues worse instead of working around them. While a human analyst might find ways to handle outdated fields or inconsistent IDs, an autonomous agent can spread these problems into every decision it makes.

Since 42% describe their data as mixed or early stage, most operators are close to using data they cannot fully trust. To address this, the architecture plan should focus on three steps: open up the core systems using standard Open APIs, fix real-time access and consistency before increasing autonomy, and intentionally adopt protocols for agent-to-tool and agent-to-agent communication.

Data Foundation Readiness:
how would you rate your organization's data foundation for supporting autonomous agents (real-time access, semantic consistency, lineage)?

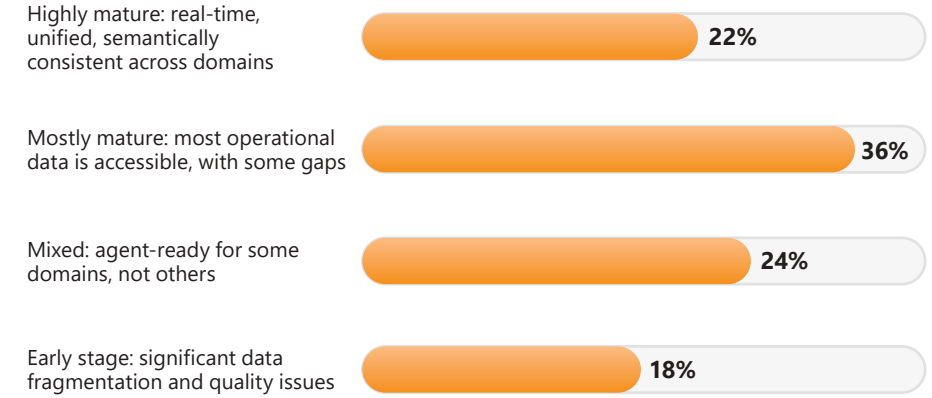


Figure 10. n=201; Source: Comviva Research Centre.



In practice

Vodafone consolidated fragmented customer data on the Tealium Customer Data Platform to feed real-time propensity-to-buy models, foundation first, autonomy second. The companies report 10-30% higher conversions from personalized web experiences, and more than 50% for high-intent segments prioritized by the models.

Source: Tealium, "How Vodafone is Using Tealium for More Personalization", https://tealium.com/blog/customer-data-platform/vodafone-use-cases-key-lessons-from-omr-2025/

Composition

From Catalogue to Autonomous Commerce

Composition is where agentic AI meets customers, and it's also where this study found the most distinctive results. The traditional fixed catalogue, long central to telecom sales, is losing ground. Most respondents (75%) think AI should build personalized offers and bundles from modular parts in real time, rather than using a set catalogue. This challenges current processes for catalogues, pricing approvals, and bundle sign-offs.

Intent is running ahead of deployment

Today, 36% of organizations use rule-based chatbots that guide customers through set steps, and 32% use NLP-powered assistants that understand natural language but still follow scripts (see Figure 11). Another 15% use AI co-pilots to help human agents, 5% have not used AI in customer interactions, and only 12% have autonomous agents that can understand intent, create offers, and complete transactions. The gap between the 75% who want dynamic composition and the 12% who can actually do it highlights both the opportunity and the execution gap.

From Chatbot to Autonomous Commerce:
which statement best describes your organization's current AI-powered customer interaction approach?

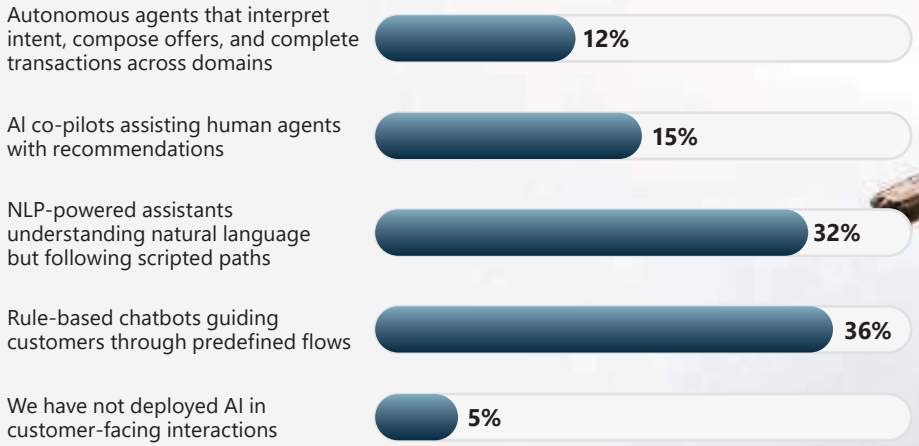


Figure 11. n=201; Source: Comviva Research Centre.

Priority follows proof, and misses the catalogue

Where does redesign attention go? Comparing current deployments with future redesign priorities shows that organizations focus on areas where AI is already proven, not on unexplored opportunities (see Figure 12). Respondents using AI in a domain are much more likely to prioritize it for agent-native redesign: 1.4 times more for customer management, 1.7 times for revenue management, and 1.8 times for market and sales. Less-used domains like B2B2X (9% deployed) and business assurance (21%) get little redesign focus, even from their own users.

Growth happens where results are already visible, which matches how respondents say they plan transformation. But there is an exception: the catalogue is often overlooked. Only 7% of respondents put product management first for agent-native redesign, and just 5% of those already using AI there do, even though 75% support dynamic offers. Ambition for composition has not yet influenced the decisions needed to make it happen; dynamic offer assembly is still low on the list.



Deployment vs. Priority:
Where AI runs today vs. the domain ranked first for agent-native redesign

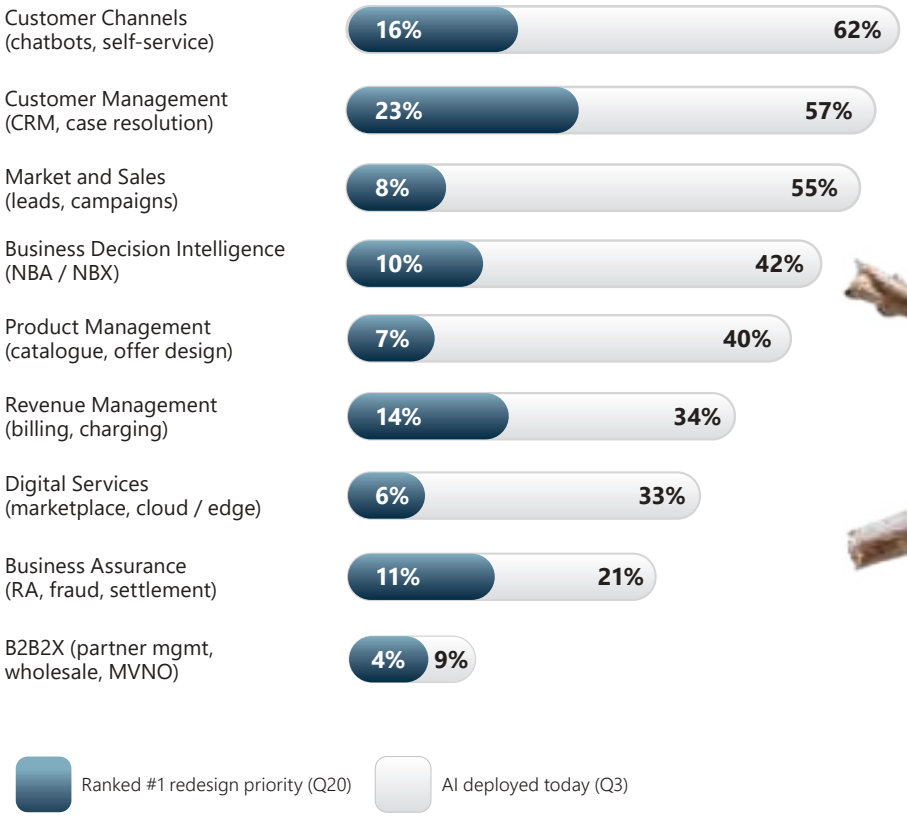
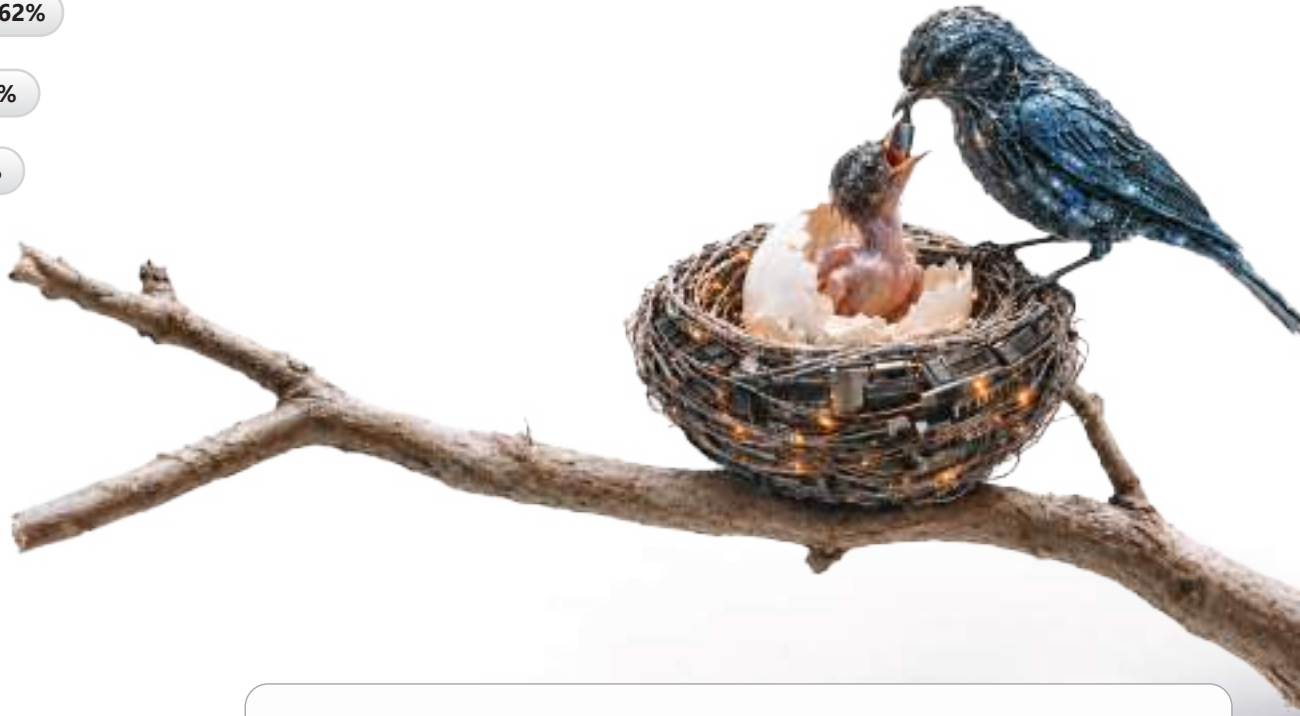


Figure 12. n=201; Source: Comviva Research Centre.



The way forward is to treat offers as modular parts that agents can put together based on real-time context and eligibility. This approach is supported by the Process and Architecture agendas, and Enablement must ensure it is done safely. Agents should be able to set prices and create bundles within clear risk limits, instead of sending every decision to a human.



In practice

In September 2024, T-Mobile announced IntentCX, an intent-driven AI decisioning platform custom-built with OpenAI under a multi-year partnership, with launch planned for 2025. The stated ambition goes beyond chatbots: comprehend customer intent and sentiment in real time, apply knowledge of the customer to every interaction, propose resolutions and take proactive actions on the customer's behalf.

Sources: T-Mobile Newsroom, "T-Mobile and OpenAI Join Forces to Revolutionize the Customer Experience" (18 Sept 2024); Light Reading, "T-Mobile taps OpenAI to reduce customer churn", https://www.lightreading.com/ai-machine-learning/t-mobile-taps-openai-to-reduce-customer-churn

E Enablement

Govern the Autonomy, Staff the Transition

Enablement is not the final step in PACE. Instead, it supports the other three steps at every stage. Each redesigned process, new API, or autonomous offer increases the area that governance, guardrails, and staff must oversee. Organizations are most intentional here, but resources are often limited.

Guardrails keep humans accountable

Respondents focus on practical guardrails. The top priorities are real-time policy and compliance enforcement (49%), data privacy and consent controls (47%), human override and escalation mechanisms (46%), and risk-based authority limits like monetary thresholds (43%) (see Figure 13). These are followed by explainability and decision audit trails (36%), continuous monitoring of agent behavior and drift (31%), bias detection and fairness checks (21%), and cybersecurity and identity controls for non-human actors (19%). These measures allow operators to approve autonomy while staying accountable for results.

Trust, Risk and Guardrails:
Which guardrails do you consider most critical to enable autonomous AI agents in your operations?

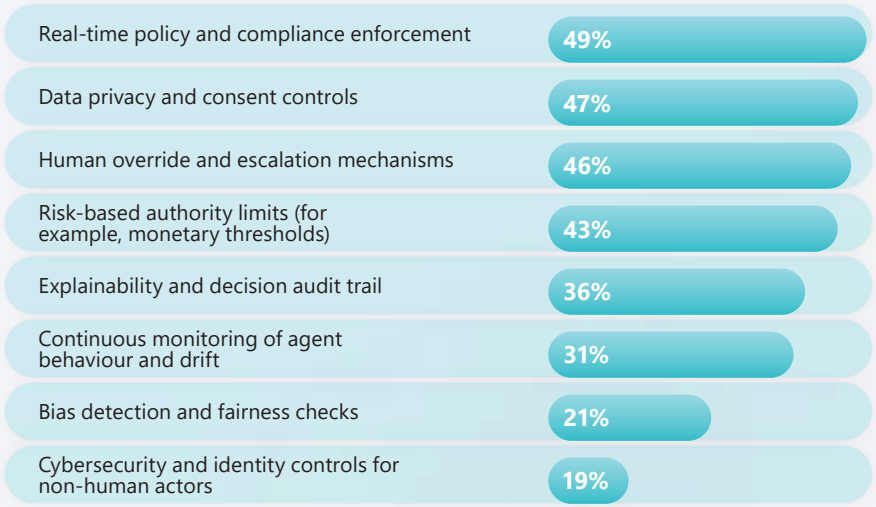


Figure 13. n=201; Source: Comviva Research Centre. Multi-select; totals exceed 100%.

Governance preferences show a similar trend. About 37% prefer human-on-the-loop models, where agents act but humans monitor and step in as needed. Another 34% support human-in-the-loop, where agents suggest actions and humans make the final decisions (see Figure 14).

Only 11% support human-out-of-the-loop with exception escalation. Meanwhile, 10% choose graduated autonomy, reducing supervision as trust grows, and 8% would use different models for different risks and processes. Most organizations favor supervised autonomy, relaxing control gradually instead of giving it up completely.

Governance for Human-Agent Workflows:
Which governance model is most effective for hybrid human-agent workflows in your context?

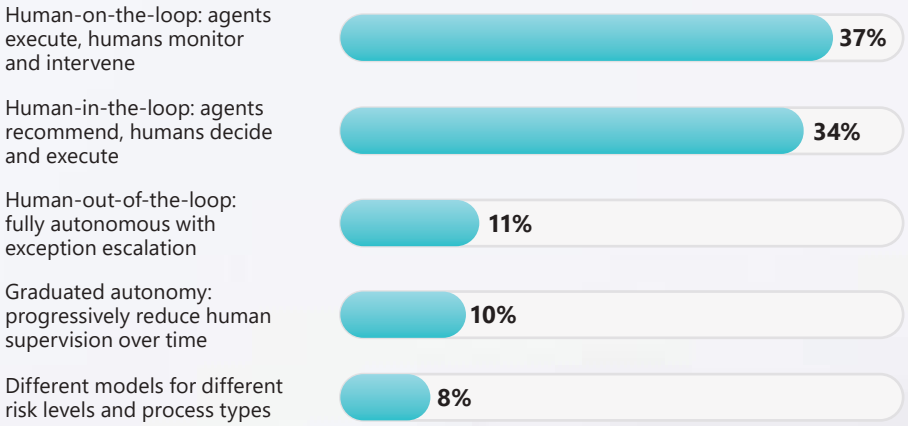


Figure 14. n=201; Source: Comviva Research Centre.



Staffing the transition

Organizations are changing their structures to adapt. They are hiring new roles like AI process architects and agent workflow designers (48%), setting up AI governance boards or centers of excellence (46%), and building cross-functional teams that bring together domain experts and AI engineers (43%). Other steps include creating dedicated AI or automation teams within operations (40%), working with external specialists (40%), shifting roles from process executors to process designers and AI supervisors (37%), and reducing middle management as AI handles more coordination (22%) (see Figure 15).

Organizational Redesign Actions:
How is your organization adapting its structure to support agentic AI? (Multi-select)

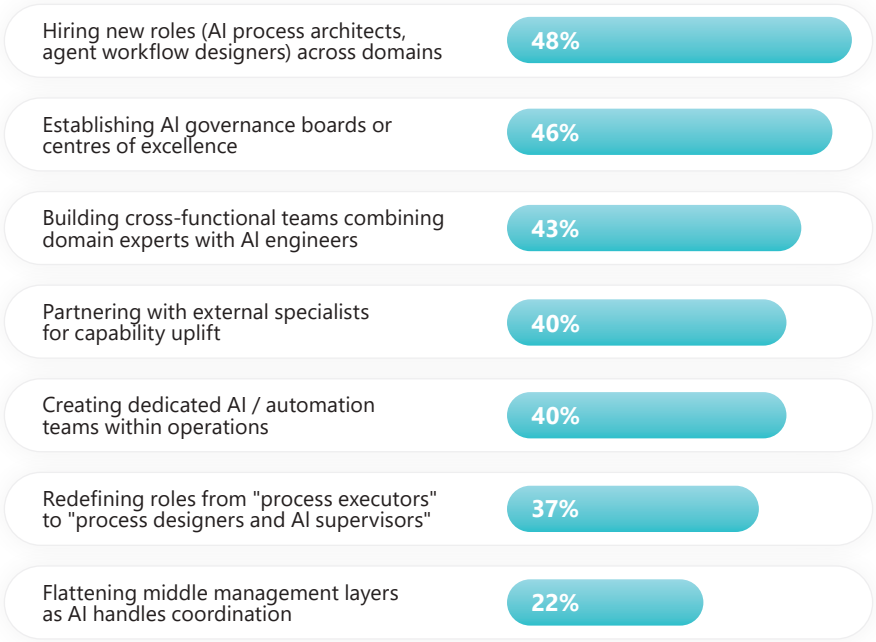


Figure 15. n=201; Source: Comviva Research Centre.

The hiring trend is clear. Over the next three to five years, the most in-demand new role will be the AI Process Architect, who designs agent-native workflows. This role is cited by 59% of respondents, followed by Agent Supervisors (50%), AI Ethics and Governance Officers (46%), AI Business Analysts who connect business outcomes to agent goals (43%), Data Curators (39%), Human-Agent Interaction Designers (37%), and Prompt and Policy Engineers (29%).

The build shift is real, and bigger than one line

Ninety percent of organizations surveyed plan to use AI development frameworks to build or expand their software, instead of relying only on commercial off-the-shelf products (see Figure 16). Of these, 30% will adapt or extend COTS solutions with AI frameworks, 22% will build complete solutions in-house, 21% will create client- and user-facing components, and 15% will develop add-on functions and extensions. Only 10% will stick with COTS-based sourcing. AI enables changes to both the software supply chain and the organization's structure.

COTS to "Vibe-Coding" Shift.
Is your organization planning to shift from purchasing COTS (Commercial Off-the-Shelf) software components toward building solutions in-house using AI development frameworks?

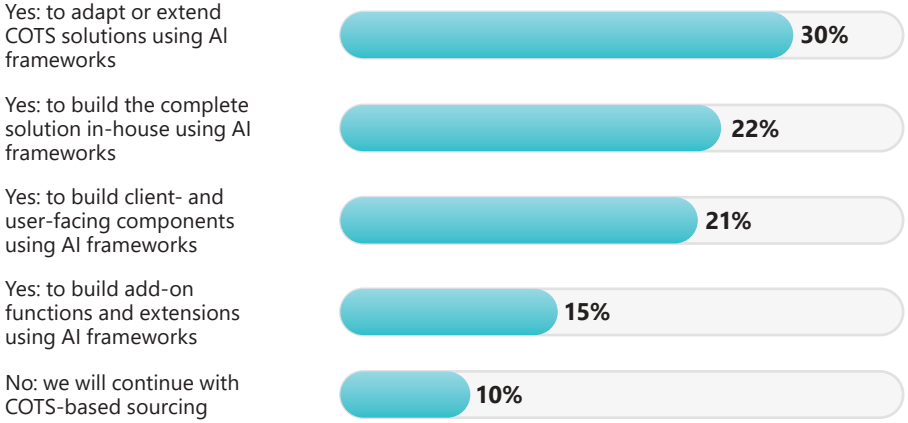


Figure 16. n=201; Source: Comviva Research Centre.

In practice

Vodafone and ServiceNow launched an AI-first service management platform that uses agentic AI to handle queries and resolve issues inside an enterprise-grade governance envelope. In an Ireland pilot, the companies report digital channel usage up 45% and a fourfold improvement in customer satisfaction; governance and autonomy working as complements, not trade-offs.

Source: Vodafone and ServiceNow press release, https://news.europawire.eu/vodafone-and-servicenow-launch-global-ai-platform-to-reinvent-enterprise-customer-experience/eu-press-release/2025/04/25/09/39/54/152981/

Setting the PACE:

The Next 24 Months Decide Who Scales

This research does not suggest that telecom is lagging in AI adoption. Instead, it shows that many companies are focusing AI on the wrong challenges. Agentic AI is often used for tasks meant for people, and the results reflect this. The real advantage goes to companies that rethink how work is done, not just those with the largest AI budgets. The benefits of redesigning processes are clear and measurable.

PACE sets the order for this work. Process comes first because it is the biggest barrier and the main source of untapped value. The highest returns come from fully redesigning processes. Next is Architecture, since new processes need real-time data and standard interfaces. Composition is where investments start to pay off, through dynamic offers, automated transactions, and preparing for agent-to-agent business.

Enablement supports all three steps by providing guardrails, governance, and new roles like the AI Process Architect to make sure autonomy can scale safely. When it comes to where to start, opinions differ (see Figure 17). About 28% would begin with back-office operations for lower risk and quicker returns. Another 25% would choose areas where AI has already shown results, which matches the findings of the deployment-versus-priority analysis. Eighteen percent would focus on domains with the most legacy issues, while 12% would start with customer-facing processes for greater visibility. Ten percent would tackle all areas at once. Another 6% are unsure and want a framework to help decide. While this is a small group, it matters because the industry lacks a common way to choose where to begin.

Transformation Sequencing.
What is your preferred approach to sequencing the agent-native transformation?

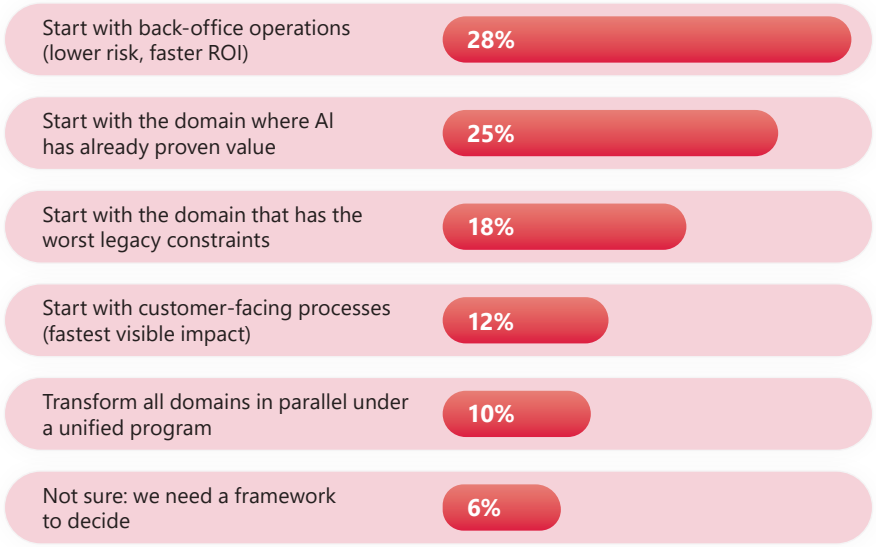


Figure 17. n=201; Source: Comviva Research Centre.



The window is short. About 74% of respondents think agentic AI will handle more than half of routine operations within five years, and 34% expect this to happen in just two years. Also, 83% say that multi-agent orchestration is very important or critical for their future operations.

When asked about the top priority for the next 24 months, they listed three main areas: redesigning core processes for agent-native work (26%), building trust, governance, and risk frameworks for autonomous AI (21%), and investing in strong data and integration systems (20%). In two years, the difference will be clear: some operators will have rebuilt their work for agents, while others will have only upgraded to newer models.

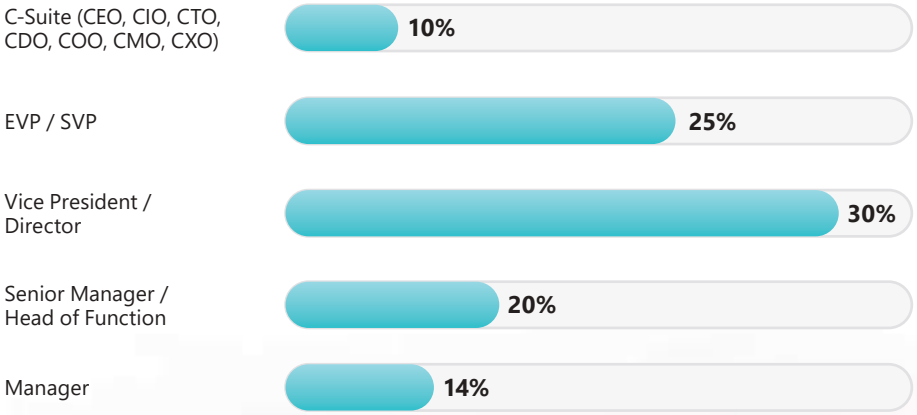
Methodology

Comviva's Research Centre commissioned Curious Insights, an independent global market research firm, to conduct this study as an online panel survey of 201 senior executives at telecommunications and adjacent connectivity organizations worldwide. Respondents were qualified through a firmographic screener covering organization type, role, function, decision authority, revenue, and headquarters region.

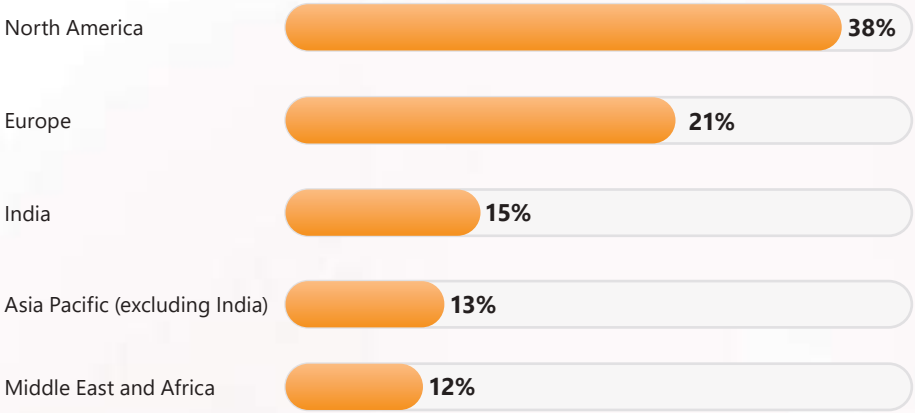
Which of the following best describes your organization?



S3: Which best describes your current role?



S7: Where is your organization primarily headquartered?







comviva

A TECH MAHINDRA COMPANY

EXPERIENCE **CREDIBLE IMPACT**

About Comviva Research Centre

Comviva's Research Centre examines current and future trends that will shape the future of the telecom industry. The Centre provides thought-provoking research and analysis of business and technology trends and dynamics, and collaborates with a wide range of business, technology, and academic thinkers regarding what the future of telecom will look like as technology changes every aspect of the industry.