RCR Wireless News
INTELLIGENCE ON ALL THINGS WIRELESS

JULY 2025

MARKET PULSE REPORT

TELCO AI

BY SEAN KINNEY, PRINCIPAL ANALYST

Contents

Introduction

The telecoms industry is at a pivotal juncture. The adoption and scaling of artificial intelligence (AI) across networks and operations is no longer a question of if, but a question of how quickly. The journey from experimental trials to fully-realized production environments represents a technological transformation but, more fundamentally, it represents a reorientation of business models, operational paradigms, and strategic thinking.

RCR Wireless News recently convened a broad panel of industry experts to discuss these pressing issues at the Telco AI Forum. This report is meant to encapsulate those discussions and explore how predictive, generative, and agentic AI can unlock substantial operational efficiencies and significant new revenue streams, provided they are built on scalable cloud architectures, coherent data strategies, and strong ecosystem partnerships.

As AI moves into all network domains, particularly the radio access network (RAN), and becomes foundational to emerging 6G standards, success will depend on an organizations' ability to balance rapid innovation with sustainable scalability. Striking the appropriate balances requires technological capabilities and a thoughtful approach to organizational culture and trust.

Top *RCR* takeaways

AI is central, not supplemental. Executives must prioritize holistic AI integration across business models, infrastructure, and operations, moving quickly from proofs-of-concept into scalable production deployments.

Balance ambition with sustainability and scalability. Scaling AI intensifies energy demands and architectural complexity. Executives must prioritize efficient operations (AIOps) and services {edge inference), built on a standardized modular design.

Culture and trust underpin AI success. Successful telco AI adoption and scaling hinges on organizational culture, transparency, and trust as much as it hinges on technical capability.

Interoperability and standardization will determine winners. Fragmented, proprietary AI solutions echo past industry missteps predicated on vendor lock-in. To maximize ROI, leaders should champion interoperable, open, and standardized AI frameworks across cloud, core, RAN, and edge domains.

AI-native thinking sets the stage for AI-native 6G. If 6G is going to truly embed AI across every layer of networks and services, executives must proactively participate in and shape standards, positioning AI as a foundational design principle.

1

AI ambitions are predicated on cloud, data, and ecosystem imperatives

AI, in its various forms, is not just another tool in the telco transformation playbook. It is the playbook. Whether it's optimizing network performance, enriching customer interactions, creating new monetization pathways, or powering autonomous workflows at the edge, AI is now central to every aspect of communication service provider (CSP) strategic planning. While telco AI adoption and experimentation is booming, widespread activity does not yet equal impact. For all the momentum, the telecoms industry now faces a more difficult challenge: moving from proofs of concept to production-scale platforms, then using those platforms for systemic automation and value creation.

The distinction between AI for networks and networks for AI is increasingly critical. In the early phases of telco AI adoption, the emphasis was rightly placed on internal efficiency: leveraging predictive analytics and automation to reduce cost and complexity across operations. That continues, but the frontier has shifted.

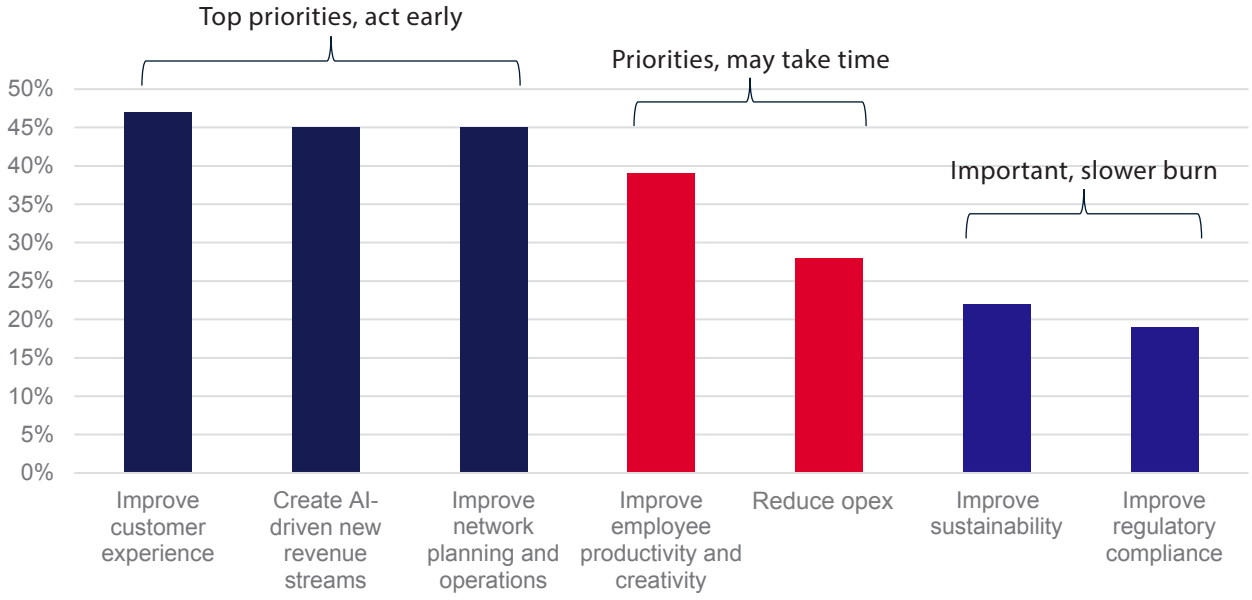

Data courtesy of GSMAi.

Today, the network itself is becoming an AI delivery platform capable of hosting inference workloads, enabling real-time decision-making at the edge, and supporting sovereign AI services. To unlock this next phase of value creation, CSPs must ground their ambitions in modern cloud architectures, coherent data strategies, and partner ecosystems that can drive the industry toward industrialized AI.

GSMA Intelligence (GSMAi) Head of Research and Consulting Tim Hatt cited research that more than 90% of operators are active in AI, and working to prioritize use cases with an eye on operational efficiency and revenue generation. The top priorities are customer experience, new revenue streams, and improved network planning and operations. Priorities that may take more time include employee productivity and automation-enabled opex reduction. Longer-term telco AI goals are improving sustainability and regulatory compliance. See the graphic on this page for a detailed look from GSMAi.

Given the gravity of how AI strategies implemented today will inform future success, vendors are working to help CSPs keep up with the pace of change in the world of AI while delivering turnkey solutions — developed using deep ecosystem-led collaboration — that are easy to consume.

Dell Technologies Director of Telecom Marketing Sandro Tavares talked through the role of vendor ecosystems in helping CSPs scale up AI-enabled solutions. "Our industry is known for being really good in experimenting with new technologies, not as good as actually going from experimentation and science fair projects into real implementations," he said. "And that is really a call for action…Our job here as provider is to make sure that… they have the solutions available, and most importantly, that they are easily consumable. So we have to make this process as simple as possible."

"It's not what is happening, but what isn't...It's important to not try to boil the ocean but rather to understand where the priorities lie."

- Tim Hatt, Head of Research and Consulting, GSMA Intelligence

From science project to strategy

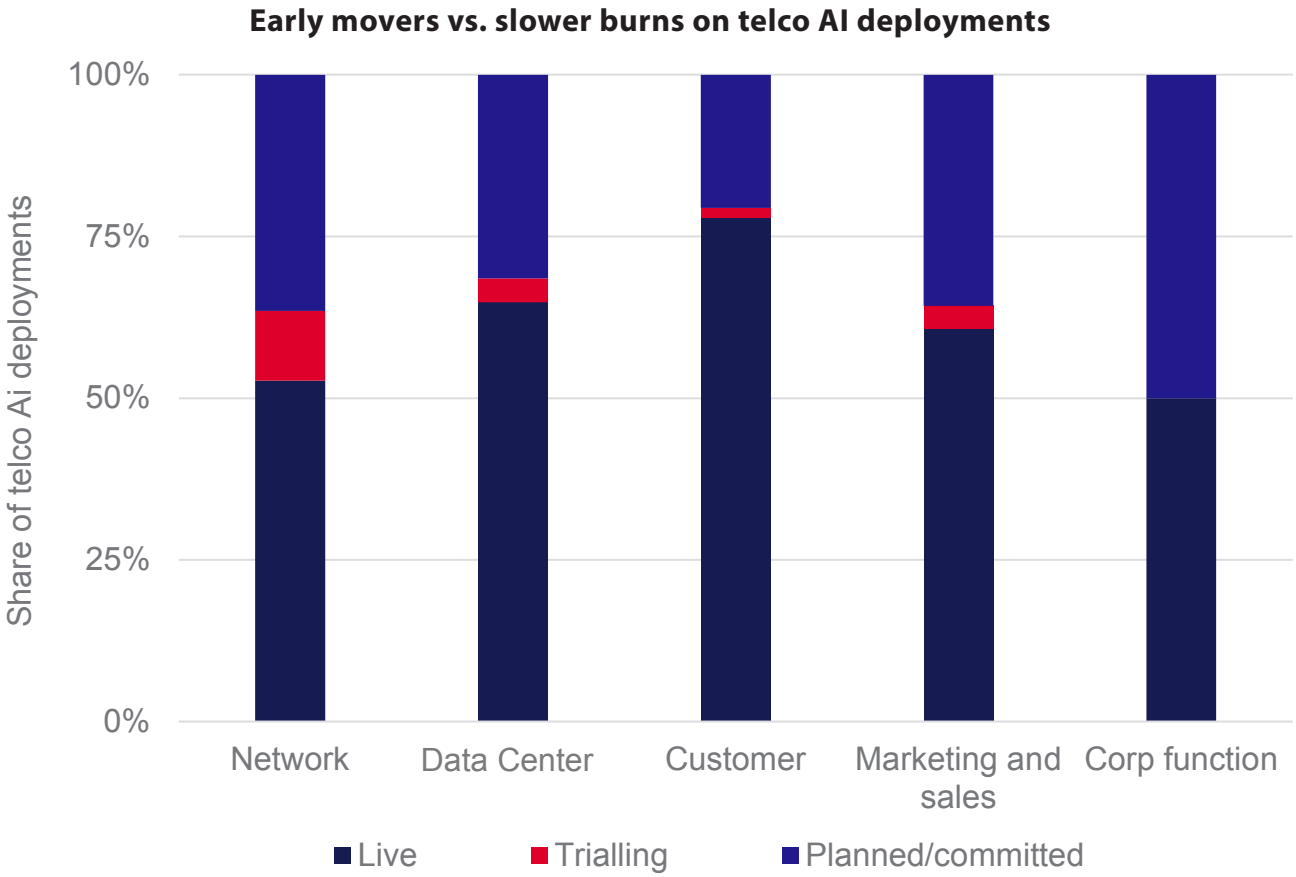

Figures based on data collected from 250 operators as of May 2025. Data courtesy of GSMAi.

A valuable portion of GSMAi's research is a "pipeline view" of telco AI adoption tracking use cases as they move from planned/committed to trial stage to live. This cuts across network, data center, customer, marketing and sales, and corporation functions. For a detailed view, see the above graphic. "A lot of the early deployments have focused on customer care and the sales touchpoints," Hatt said. But, adoption is now accelerating across domains and functions, indicating a tipping point toward holistic integration rather than isolated initiatives. "I think as we start to see a network effect play out as more operators demonstrate proof points and ROI, we'll see" more (and faster) adoption.

Underlying this analysis is that CSPs are focused today on customer experience, making money, and saving money. Longer-term objectives include regulatory compliance and environmental sustainability. But the through line of this transformation is the idea of evolving from a classic telco — a company that sells connectivity against metered rates — into a "tech-co," which is marked by providing a sticky services platform built on top of its core assets and competencies.

For telecoms and the other verticals it services (all of them, essentially) NVIDIA is pushing the idea of an "AI factory," GPU-accelerated computing infrastructure deployed in whatever manner serves the customer's objective. The general approach, as described by NVIDIA Global Head of Business Development for AI Operations Lilac Ilan, is "understanding the value that they can extract with it, building the skillset that is needed, understand how to talk about it, and then looking at it externally as [a] monetization opportunity."

The last two years, Ilan said, were all about proofs of concept. The idea was, "Let me understand if it's beneficial for me…It wasn't about scale. It was about understanding the technology. And I think that understanding of the technology then enables [CSPs] first to then scale it this year…and then take it to enterprises and governments, start to take it outside."

Networks for AI and AI for networks – a two-way value chain

The arc of telco AI adoption started, naturally, with easy wins like customer care and network optimizations that build on proven predictive AI solutions. The next step is simultaneously about incorporating gen AI and agentic AI, but also about understanding how AI can benefit networks and how networks can benefit AI — that is to say, how can the bundling of network connectivity and AI-enabled services lead to new products, new services, and new revenue.

"Are we talking about, in the telco context, AI helping networks or networks helping AI?" Hatt said. "This is not a new kind of way of thinking about it, but I think it's a helpful one. A lot of this sort of early happenings in 2023, 2024, concentrated on how AI can help. [And] understandably because you can automate, you can take costs out of the business, and make things more efficient. Now a lot of the focus is really moving to leveraging network capabilities with AI, and with them new commercial models. That's, I think, where we're seeing a lot of interesting stuff, especially at the edge."

Hatt discussed edge-based inference, particularly in situations where compute efficiency, latency, and data sovereignty are involved, as an opportunity for operators to leverage distributed compute infrastructure and connectivity for new types of services. As industry attention turns to how the introduction of GPU-accelerated computing to CSP's distributed infrastructure can be monetized for GPU-as-a-Service (GPUaaS) and sovereign AI, Amdocs Head of R&D Doron Sterlicht looked at the three challenges that need to be addressed.

"The first, of course, achieving multi-tenancy and security — the ability to truly slice the GPU cluster to support multiple tenants that might have varying needs from startups to enterprises to government bodies, all this while ensuring strong isolation, security and compliance. This is critical to allow CSPs to serve a broader customer base and maximize the utilization of the infrastructure."

Sterlicht continued: "The second thing is to go beyond just the naked GPU-as-a-Service and to move up the value chain into AI-as-a-Service, meaning models as a service and application and agent-as-a-service, leveraging on a platform that will allow internal and external partners to build and deploy gen AI apps and agents across different verticals in the region. And third is monetization and automation to leverage a full-feature storefront to onboard customers to automate the provisioning of either the infrastructure or the models that they're interested in, and to manage flexible billing models at a time-based token consumption and so on."

"I believe this is also kind of a time-boxed opportunity and CSPs have to move fast and, therefore, having a turnkey solution is critical in terms of time-to-market."

- Doron Sterlicht, Head of R&D, Amdocs

"And most importantly, to have all of these components in one unified turnkey stack covering the infrastructure orchestration, the platform services, monetization. So CSPs can rapidly deploy and scale GPUs and any AI service as service with confidence."

To play that back, the business opportunity is real, but capturing it requires the nuts-and-bolts work of sharing network and AI resources, building horizontal infrastructure, and making it all easily consumable.

Industrializing intelligence – it takes a village

Telco AI sits at the intersection of business model transformation, capital investment, cloud-native design, disaggregated infrastructure, organizational change, and rapid technological evolution. For the CSPs tasked with providing always-on connectivity and services, this complexity amplifies both operational risks and strategic stakes.

As explored earlier, Amdocs, Dell Technologies, and NVIDIA have co-developed a joint solution that ties these elements together. The combined AI factory scalably combines Dell servers, NVIDIA GPUs, and verticalized AI-enabled applications meant to help operators "address soaring AI demand while delivering tailored solutions that drive industry transformation and business growth."

The value proposition centers on embedding AI across a cloud-native, horizontal infrastructure — spanning core, RAN, edge, and far edge — to run networks more efficiently while delivering new predictive, generative, and agentic AI services. Building on Sterlicht's insights on GPUaaS, Ilan re-emphasized the opportunity for CSPs to sell sovereign AI.

"Telcos are positioned very well to be the provider of sovereign AI to the market," Ilan said. "Why? Because… they have data centers. Data centers are a very constrained environment…Telco[s] have access to data centers, telco have access to government and education and enterprises in that market. They're in that journey of moving from a telco to a tech-co and kind of bringing that to that enterprise."

Returning to the complexity of aligning underlying cloud and data strategies, Tavares said, "It's a huge undertaking." CSPs "have to build this AI infrastructure while they're operating a mission-critical network that has to be on pretty much 100% of the time."

Culture and change management

If the core narrative is the challenge of scaling AI to deliver both internal and external value, the subtext is clearly that technology alone is not enough. Any investment in technological transformation requires a

> "It's almost like changing the wheels of the car with it moving. So that's our job here...We have to go together...to present the solution and define with the customers the plan to adopt and industrialize solutions."
>
> **- Sandro Tavares, Director of Telecom Marketing, Dell Technologies**

proportional investment in change management to support new ways of working like AIOps, CI/CD, and agile delivery models.

Hatt acknowledged the technological investment required but emphasized, "A lot of this comes down to sales strategy and culture in terms of how you do business in the AI era — be much more of an IT consultant rather than a pure-play connectivity seller."

Ilan echoed this sentiment: "In order to get to the scale, in order to get to the full potential of it, we have to lift up the skills."

"Everything goes well until it gets to operations," Tavares said. "Of course there is a strong bias for conservatism in the operation of the network because it can fail." In terms of what this means for an ecosystem of vendors who benefit from supporting telco AI adoption, "We have to make sure that the solutions are there, they're easily consumable, and we have to work with our customers to basically crack the code of how do we implement this whole new array of things without breaking the telecom operations…This is the key for the success in the industry. I have no doubt whatsoever it is going to happen. But, yeah, there is a lot of work ahead of all of us here."

In other words, the road to telco AI isn't blocked by innovation; it's blocked by inertia.

AI is the playbook

All this to say: AI isn't a bolt-on or an overlay. It's the playbook. Operators provide the capabilities that power global commerce and shape how consumers live their lives. The strategic addition of AI requires infrastructure, operational, and organizational overhauls — all while operators examine, or re-examine, underlying cloud and data strategies.

Network automation is here today, and it's only going to gain more efficacy. We'll drill into operationalizing AI, and what that looks like in the real world, in further chapters. Spoiler alert: success depends on a synchronized overhaul of infrastructure, operations, and organizational models; and decisive, fast action could lead to a durable competitive advantage.

"Agentic AI opens the door and accelerates the path to autonomous more than anything else."

- Lilac Ilan, Global Head of Business Development for AI Operations, NVIDIA

2

Predictive, generative, and agentic – building a holistic telco AI strategy

As mobile networks become more cloudified, distributed, disaggregated, and software-defined, AI is of course an efficiency play and a lever for new monetization. But it's also a scalability imperative. Analysys Mason projects that from now through 2027, operators will invest $9.2 billion on AI infrastructure for internal use cases (efficiency) and $21.6 billion on AI infrastructure for external (monetization) use cases.

Spirent Communications Vice President of Product Management Anil Kollipara put it this way: "It's becoming unscalable. And this is an area where agentic AI can be really, really helpful."

While Kollipara focused his commentary on the outlook for agentic AI, he also detailed the important role of more classical predictive AI and generative AI as operators invest in network automation and service differentiation. In general terms of adoption: predictive AI is widely deployed, gen AI is entering early adoption, and agentic AI is in the proof-of-concept phase.

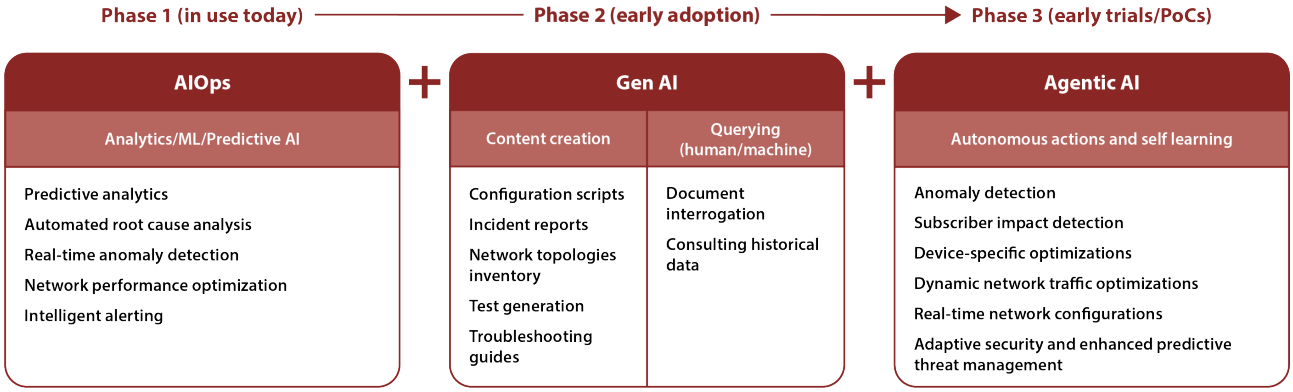

Data courtesy of Spirent Communications.

For a detailed look at the broad use cases associated with these types of AI, courtesy of Kollipara, check out the above graphic. Directionally, predictive AI uses data to accelerate decision-making; gen AI augments human intelligence to enhance productivity and insight; and agentic AI fully automates discrete and systemic processes.

AIOps comes of age – operational intelligence at scale

The most mature of the three approaches is predictive AI which is institutionalized via an AIOps methodology. Impactful use cases include predictive analytics, anomaly detection and root cause analysis. Kollipara's colleague Stephen Douglas, head of market strategy, said many operators are moving hard on AIOps which "is not massively new within the [operation] support space."

He continued: "What is probably newer is just more about productionization of the predictive AI…for things like root cause or automated root cause analysis."

Multinational operator VEON — which has taken a unique approach to delivering connectivity-enabled digital services in its various markets (something we'll explore later in this report) – uses predictive AI for a range of use cases. VEON Chief Information Officer Inanç Çakiroğlu said it's all about making the best use of the "data ocean" available to operators.

"On the revenue side, we have been developing upsell, cross-sell, churn models, best offers, recommendation engines for our…digital products. And all those predictive models we have been working maybe more than 20 years…We had credit scoring, fraud management systems, network optimization, scenarios, fault prediction, self-organizing network use cases, and the contact center efficiency and the effectiveness use cases."

"It's becoming unscalable. And this is an area where agentic AI can be really, really helpful."

- Anil Kollipara, Vice President, Product Management, Spirent Communications

With this high degree of maturity, clear relationship to business value, and scope of adoption, all of these predictive AI-based use cases are no longer novel — they're foundational capabilities essential for unlocking more sophisticated, higher-value transformations.

The rise of gen AI makes data interactive and insightful

Enter gen AI. Now that "data ocean" Çakiroğlu described can serve as an assistant; natural language processing and other modalities allow human workers to interact with data in ways that enhance decision-making and productivity; this is particularly important as the telecom workforce ages and needed skillsets become harder to find.

Tupl's Hector Montes, COO and co-founder, described how the company builds on the predictive AI foundation with gen AI tools. "We do it in a way similar as how the human brain works," he said. "So we have two hemispheres in the brain, each one representing a crucial layer capability. On the left side, we have predictive AI, the analytical powerhouse. This is like the brain's logical hemisphere designed to process vast amounts of data, just like an engineering expert would do. On the right side of the brain we have the generative AI that is the emotional, creative counterpart. And that's enabled a more natural human-like interaction and communication. Think of a customer service agent powered by gen AI responding accurately with empathy, clarity and context awareness. That's improving the customer satisfaction while reducing the human workload."

Kollipara gave a real world example of a network failure. Traditionally, network engineers would collect relevant KPIs, log files, and the like, then conduct an analysis to identify the potential cause or causes of the failure. That information would then be relayed to another functional group and be formatted into a trouble ticket which would then be sent to another stakeholder for resolution. "These are all manual steps today." With gen AI, those steps can become easier and faster; and with agentic AI, that workflow could evolve from human-in-the-loop to fully-autonomous closed-loop execution.

Toward autonomy – getting to closed-loop network automation

Agentic AI — systems of agents and sub-agents that can translate an intent into an outcome by interacting directly with network hardware and software elements — will drive cost-reducing automation and revenue-enabling capabilities. "The goal with agentic AI," Kollipara said, "is you can completely automate this where the AI is executing that workflow on your behalf."

Focusing on real world impact, Montes said Tupl worked with T-Mobile US to deliver more than 600,000 proactive customer fixes per month. "So quite significant," he said. The general evolution is to start with a minimum viable product and expand from there.

"You're seeing a lot of focus on anomaly detection, trying to identify these very, very small issues, very, very hard to pinpoint in the network but can cause a lot of issues…so you can prioritize potential fixes or a human interaction where it's still required."

- Stephen Douglas, Head of Market Strategy, Spirent Communications

Montes described Tupl's strategy for introducing autonomy. "We automate the complex tasks that otherwise are handled manually by human experts." The approach is to develop a minimum viable product "and expand. You start small and you expand."

While fully closed-loop automation remains aspirational, the trajectory is clear. Douglas noted, "We're not there yet," but the industry is making progress from conversational gen AI-based agents to "agent-to-agent, post-loop types of environments." As these early experiments go forward, the agents are speaking to each other…And if I go back to that sort of AIOps idea of fault prediction, you can start to imagine here an agent talking to a test agent saying, 'Hey, we predict there's a potential fault. Trigger a suite of tests to verify is it a real fault. If it is, tell the orchestrator to maybe reroute the traffic, run the test to validate did it actually resolve it, did the traffic have no degradations,' all without human interaction."

Complexity, explainability, trust

As it stands, the combination of predictive, gen, and agentic AI, is seemingly how operators will be able to manage increasing complexity while controlling costs, evolving workforce dynamics, and enabling net-new growth.

Kollipara pointed out that scaling technology has to be aligned with scaling the larger business toward specific goals. For operators, those goals include monetizing new types of network-enabled services. In order to sell those new services, a fundamental blocker is the ability to continuously deliver against various service level agreements and otherwise assure end user quality of experience and service. "That scaling is very hard to achieve with a manual, human-based approach," he said. Throwing people at the problem is cost prohibitive, so another fundamental blocker. Therefore, in order to grow, in order for "our customers to scale their business and their operations, they have to leverage the latest and greatest in the AI…and the automation."

And in order to scale AI to realize proportionate benefit, Douglas explained that operators need to be able to trust and validate how and why autonomous systems are doing what they're doing. "I think this is one of the interesting challenges at the moment…is how do you guarantee the models being used, they're not hallucinating or gradually deteriorating in their performance? And even the outcomes, if you're going to automate them, how do you guarantee that the actual decision that's being put in place is actually a good one for the operational live network?"

This trust is a function of ongoing testing and validation, to Douglas's point, but Montes also called out explainability as another key to confidently embracing autonomy. "Whenever you have any autonomous decision-making engine, then you need to get explainable AI. We're talking about complex processes… Explainable AI is a key [to] success for adoption…It needs to be completely opposite to a black box. And in a way AI is perceived as a black box. So this is a key that companies like Tupl need to provide to the telcos."

"We automate the complex tasks that otherwise are handled manually by human experts."

- Hector Montes, COO and Co-founder, Tupl

Mapping for impact – redefining KPIs for the AI era

"AI is a fast-moving object in the world of technology, so it is a big challenge," said David Kypuros, Principal AI Architect - Global Telco Partner Solutions at Red Hat. "It's a very difficult pathway for CSPs." As service providers navigate this path, they're gaining clarity on where AI makes a real difference both technically and organizationally. "CSPs are getting smarter," he said. "CSPs are definitely more intelligent" about where and how to apply predictive and gen AI for business value. The key, he explained, is in "mapping for impact" — identifying where AI can be embedded into the business to create tangible outcomes, rather than getting stuck in proof-of-concept limbo.

Kypuros laid out a framework in three categories. First, predictive and generative AI can be applied internally to help employees be more productive. Second, AI can enhance lifecycle management efficiency for operational teams. And third, it can support the monetization of new services and capabilities. For example, at a Tier 1 operator with an established portfolio of enterprise APIs and services, gen AI can add value by improving how those services are discovered and consumed. "Using gen AI to make those API services easier for enterprises and governments to consume is like a fast track to take what they already have, don't disrupt it, but add the power of gen AI to make those services move quicker into production for enterprise."

As gen AI becomes more widely accessible, operators face a balancing act. When does gen AI offer unique value? When is predictive AI good enough? "The lines will be blurry for a while," Kypuros acknowledged. He advised operators to revisit the "mapping for impact" principle to target "low-hanging fruit" — like customer care and call centers — where benefits are proven and adoption is relatively simple. In these areas, leveraging hyperscalers can enable agentic call flows. Still, hybrid deployment is often preferred given operator's security needs, he said. Red Hat has supported on-prem fine-tuning efforts that match or beat the performance of large foundation models. "It's faster and it's cheaper. Everything doesn't need to be connected to a frontier model."

That hybrid model isn't just about cloud or on-prem; it's also about aligning AI across a complex ecosystem of hardware, software, and partnerships. "Telcos have quite a burden to integrate a lot of capabilities," Kypuros said. This includes coordination with network equipment providers and accelerator vendors like Intel and NVIDIA. "With Red Hat, that's one of our key capabilities — having this platform that's already partner-enabled and AI-enabled so a telco can tell us, 'Oh, we're working with this set of accelerators, or this technology'... and what Red Hat can do is… show you how to integrate this into a cohesive solution." That integration is a prerequisite to delivering AI-powered services that work in production.

The final measure of success, Kypuros said, is not just infrastructure performance but whether operators can create services their customers truly value. That means "bringing value to the end customer so that if the telco develops a new service… their end customer in the B2B, that last B, finds value in the service from their KPI perspective." ROI comes from alignment: "The telco needs to be able to drive their value... aligning these new

"It's not just performance, it's not just low level KPIs, but it's KPIs, ROI and ultimately end value to making a difference in their end customers."

- David Kypuros, Principal AI Architect - Global Telco Partner Solutions, Red Hat

services to the telco's own services that come from their 5G core or RAN or edge." That alignment makes AI services more consumable and integrated with a telco's existing offerings.

The common thread running through these approaches is strategic alignment of AI with operations and of AI with business evolution.

AI is key to transitioning from communications service provider to digital service provider

Back to AI as a customer-facing enabler of new monetization, Çakiroğlu laid out VEON's philosophy of using AI, "augmented intelligence" in his definition and the company's application, to improve customer experience. "We provide digital services, digital experiences, to customers. We aim to enrich our customer life with our digital offering and the digital services. And in that sense, AI is a great differentiator.

With the combination of AI, customer data, and digital services, "It is possible to create personalized interaction in real-time with all your customer base…automatically and in real time. This is what we lacked before. We couldn't interact in real-time with our customer base before because it was basically not possible and we didn't have enough tools, but now we have."

As the industry pushes toward a more automated, adaptive, and service-centric future, operators must take a layered approach to AI adoption. But it's not a linear progression, rather an integrated strategy. Predictive AI enables scalable operational efficiency. Generative AI enhances the interface between people and data. And agentic AI introduces the possibility of full workflow autonomy. Together, these capabilities are emerging as a response to rising network complexity and the widening skills gap, and as strategic enablers of monetization and differentiation. With trust and explainability as critical prerequisites, AI must be deployed thoughtfully — augmenting rather than replacing human expertise and aligning with long-term business objectives.

But as operators graduate from narrow automations to holistic, AI-native operations, a new frontier comes into view: the radio access network. Often the costliest and most resource-intensive domain, the RAN is also where some of the biggest opportunities lie. AI-infused RAN promises new efficiencies and new revenue streams if operators can modernize their compute architectures and deployment strategies.

"We think that AI is the only way [of] making digital services better digital services, making our users better."

- Inanç Çakiroğlu, Chief Information Officer, VEON

3

The present and future of AI RAN – AIOps today, GPUaaS tomorrow

At Mobile World Congress Barcelona 2024, the AI-RAN Alliance unveiled its foundational vision for AI for RAN, AI and RAN, and AI on RAN. Those respectively speak to "advancing RAN capabilities through AI to improve spectral efficiency, integrating AI and RAN processes to utilize infrastructure more effectively and generate new AI-driven revenue opportunities, [and] deploying AI services at the network edge through RAN to increase operational efficiency and offer new services to mobile users."

It's an ambitious job but AI-RAN Alliance Chairman Alex Choi, also a research fellow with SoftBank Group International, said, "We've made tremendous progress since our debut." He said membership has passed 90 companies from 18 countries, and the alliance has demonstrated proof-of-concept breakthroughs in areas like AI-native air interface design, deep learning for uplink channel estimation, mobility-aware power saving, advanced spectrum sensing, and dynamic AI model partitioning. The takeaway, Choi said, is "all these demos

are not just hypothetical." They validate the organization's mission and signal movement "from concept to early real world experimentation."

As it pursues its mission and moves from lab to live, Choi discussed some of the key infrastructure decisions operators will face in the AI era, particularly the role of expensive, power-hungry GPUs (like those provided by AI-RAN Alliance anchor member NVIDIA). Still, Choi said discourse and thinking around the right approach to RAN compute is evolving and is not one-size-fits-all.

Understanding the hybrid AI RAN computing architecture

Traditionally, he said, RAN processing is anchored in digital signal processors contained in an ASIC-based architecture. "We are seeing clear indications it's evolving to a hybrid computing platform meaning combining CPUs and GPUs" which bring to the table "distinct strengths." CPUs, for instance, are well-suited for Layer 2 and Layer 3 control plane tasks. But, "When you get down to the Layer 1 task…the physical layer like 5G NR signal processing or massive MIMO beamforming where the massive parallel computation is required, so GPUs now start to make a lot of sense. It delivers clear cost/performance advantages."

Bottomline, Choi said, "Realistically we will very likely see hybrid CPU and GPU platforms emerging as the dominant architecture for AI RAN in the coming years. It's not a one-size-fits-all solution." In rural or less traffic-intensive environments, CPU-only RAN computing will work; however, in dense urban environments with a higher capacity demand, "It's increasingly likely that accelerators like GPUs will be needed." Combine that with AI-enabled use cases for RAN optimization and new services, "the case for GPU adoption becomes even stronger in the long term."

For operators, this means "having a very clear forward-looking roadmap is essential because RAN computing choices are no longer about hardware. They are about enabling the next generation of services and capabilities."

The counterpoint to that comes from Qualcomm Vice President of Product Management Ofir Zemer who looks after RAN automation and attendant AIOps. Regarding deployment of GPUs throughout a radio network, "I struggle to see who foots the bill for that," Zemer said. His customers are currently focused on cutting costs which embedding GPUs does not support. "The next thing they want to do…is using AI in order to do intelligent new services," like roll out fixed wireless access or deliver network slices. "You don't need GPUs for that," he said. "Only when carriers will be able to cut costs very significantly and generate new revenue streams, then I see the possibility that they'll start investing in turning their RAN into a hyperscaler with GPUs. I don't see that happening anytime soon."

Of course like with most polarizing questions, perhaps there's a middle path. Fujitsu Head of Wireless Marketing Rob Hughes recommended starting small and scaling what works. "It doesn't mean you necessarily have to do something and cover the entire network day one." He threw out the idea of using GPUs to deliver

"Realistically we will very likely see hybrid CPU and GPU platforms emerging as the dominant architecture for AI RAN in the coming years."

- Alex Choi, Chairman, AI-RAN Alliance and Research Fellow, SoftBank Group International

enterprise applications that have a clear business case and which operators can apply a clear value proposition to based on their assets and capabilities.

Regardless of the evolutionary path operators take related to AI RAN, Red Hat's Guy Turgeon, senior principal industry specialist for telco, said more foundational technology decisions will inform how AI is (and can be) put into the RAN. "I think it's going to be very specific to the operator," he said. Those that have invested in cloud-native RAN architectures are "in a good position to be able to deploy new workloads." COTS hardware can accommodate the strategic addition of GPUs because the infrastructure is "already…very configurable, software-defined." For operators who take a more appliance-based approach, expect a more forklift-type upgrade process, not to mention the skills needed to lifecycle manage these platforms. In that case, "It's going to be a very difficult transition."

AIOps and picking the low-hanging fruit

To Turgeon's point about the foundational role of cloud-native architectures as a precursor for AI, IT methodologies and ways of working, DevOps and CI/CD particularly, have to be adopted by operators to make the most of those investments in cloud-native. An extension of that is AIOps, AI for IT Operations; in the context of AI for RAN, he said, "It's all about AIOps," particularly as operators adopt gen AI and work toward closed-loop automation in disaggregated, open networks. "With 5G, with O-RAN…the RAN has become more software defined, so that's opened up the opportunity to have more automation which will eventually bring us to the dark NOC concept."

Hughes agreed, calling AIOps "the area that's moving the fastest at the moment." Why? Because it streamlines operations, enables more efficient resource usage, and cuts opex. And, to boot, it's "fairly easy" to implement, comparatively low cost, and doesn't require wholesale changes to the network.

Zemer reiterated the demand for AIOps among Qualcomm's customers. "The way that we've been working with our key customers is actually analyzing what the engineers are doing." That analysis can be used to improve performance and, again, reduce opex. Instead of armies of network engineers crunching KPIs, AIOps can "take that load away." He said if you can 10% in opex, that's generally billions of dollars for Tier 1 carriers. "You need to really do a bottom up analysis of what the engineers are actually doing."

Is GPUaaS the next big revenue stream or a MEC redux?

So AIOps is what is being done today with apparent success, saving operational costs and making engineering teams more productive, but whether GPUaaS can translate technical capability into real revenue at scale is uncertain. Fortune Business Insights reckons the GPUaaS market will grow at a 35.8% CAGR during the forecast period of 2024 to 2032 to a value in the next decade of nearly $50 billion. The questions here are two-fold with

"The way that we've been working with our key customers is actually analyzing what the engineers are doing."

- Ofir Zemer, Vice President, Product Management, Qualcomm Technologies

"These GPUs are valuable assets that people really want to take advantage of, and there's an opportunity to monetize the excess capacity that is in the network."

- Rob Hughes, Head of Wireless Marketing, Fujitsu

maybe a nesting question under the second part: how much of that market can operators capture and will the idea of sharing RAN compute resources pass organizational muster, and if it does pass organizational muster, will sales teams actually be able to sell the thing and avoid a mobile edge computing (MEC) redux?

GPUaaS for operators is fairly nascent; trial activity to date is largely concentrated in dense urban areas in the APAC region, including work by PLDT, Singtel, SK Telecom, and SoftBank. SoftBank, where Choi is a research fellow, has a number of trials ongoing with NVIDIA using the GPU giant's hardware and software for a variety of AI RAN use cases. Choi didn't provide any specific updates on how that work is progressing but he did address the idea that network engineers may be reluctant to share RAN compute resources with third parties.

"Network engineering teams understandably worry about sharing critical RAN resources with external workloads," he explained. "Their KPIs are based on uptime, quality of service, and meeting SLAs, so any proposal to mix workloads can trigger strong pushback. But here's the key point: We're not talking about jeopardizing network performance and reliability." He said that smart orchestration enables workload isolation, meaning "RAN traffic always is a top priority no matter what. GPU resources allocated to AI applications are flexibly pulled back and throttled during peak telecom demand…RAN traffic shows highly predictable patterns whether we look at daily, weekly, monthly cycles. This predictability,combined with advanced AI-driven pull testing models, allows us to make extremely accurate resource allocation predictions."

Hughes described a "chicken and egg kind of situation" that, again, is highly reminiscent of early (and still largely fruitless) operator-led to sell MEC-based services. Essentially, he said, no one wants to develop applications until infrastructure is there, and no one wants to deploy infrastructure until the applications are there. Regardless, "These GPUs are valuable assets that people really want to take advantage of, and there's an opportunity to monetize the excess capacity that is in the network…That really derisks the investment for the operator because they can get started now, they can use those GPU resources to do some of the AI things to improve RF performance…get those advantages, and then as new services become available they can take advantage of those. But in the meantime, you're not sitting on an asset that's just collecting dust."

Whatever the strategy, ROI will guide the way

While the technical case for AI RAN continues to strengthen, operator strategies for implementation remain highly variable—shaped by infrastructure maturity, vendor relationships, cost structures, and internal skill sets. Some operators are moving aggressively toward cloud-native, software-defined architectures that can support AI workloads alongside traditional RAN functions. Others remain anchored in appliance-based, vertically integrated networks that make incremental upgrades more challenging.

This divergence is also geographic and economic. Operators in dense urban markets or enterprise-heavy regions may find a clearer near-term ROI from GPU-based compute or AI-enabled edge services, while rural and cost-sensitive operators may prioritize lightweight AIOps use cases like power savings and log analysis.

A phased approach helps operators avoid the trap of overbuilding without service readiness while creating a path to eventual monetization.

Ultimately, each operator must weigh its own assets, challenges, and business priorities when mapping out an AI RAN evolution. Some will lead with automation and operational efficiency; others will chase new revenue streams. But across the board, there's a shared recognition that investment decisions must be rooted in business logic rather than hype. As UScellular CTO Mike Irizarry put it: "Whatever we do will be guided by the return."

"Whatever we do will be guided by the return."

- Mike Irizarry, CTO, UScellular

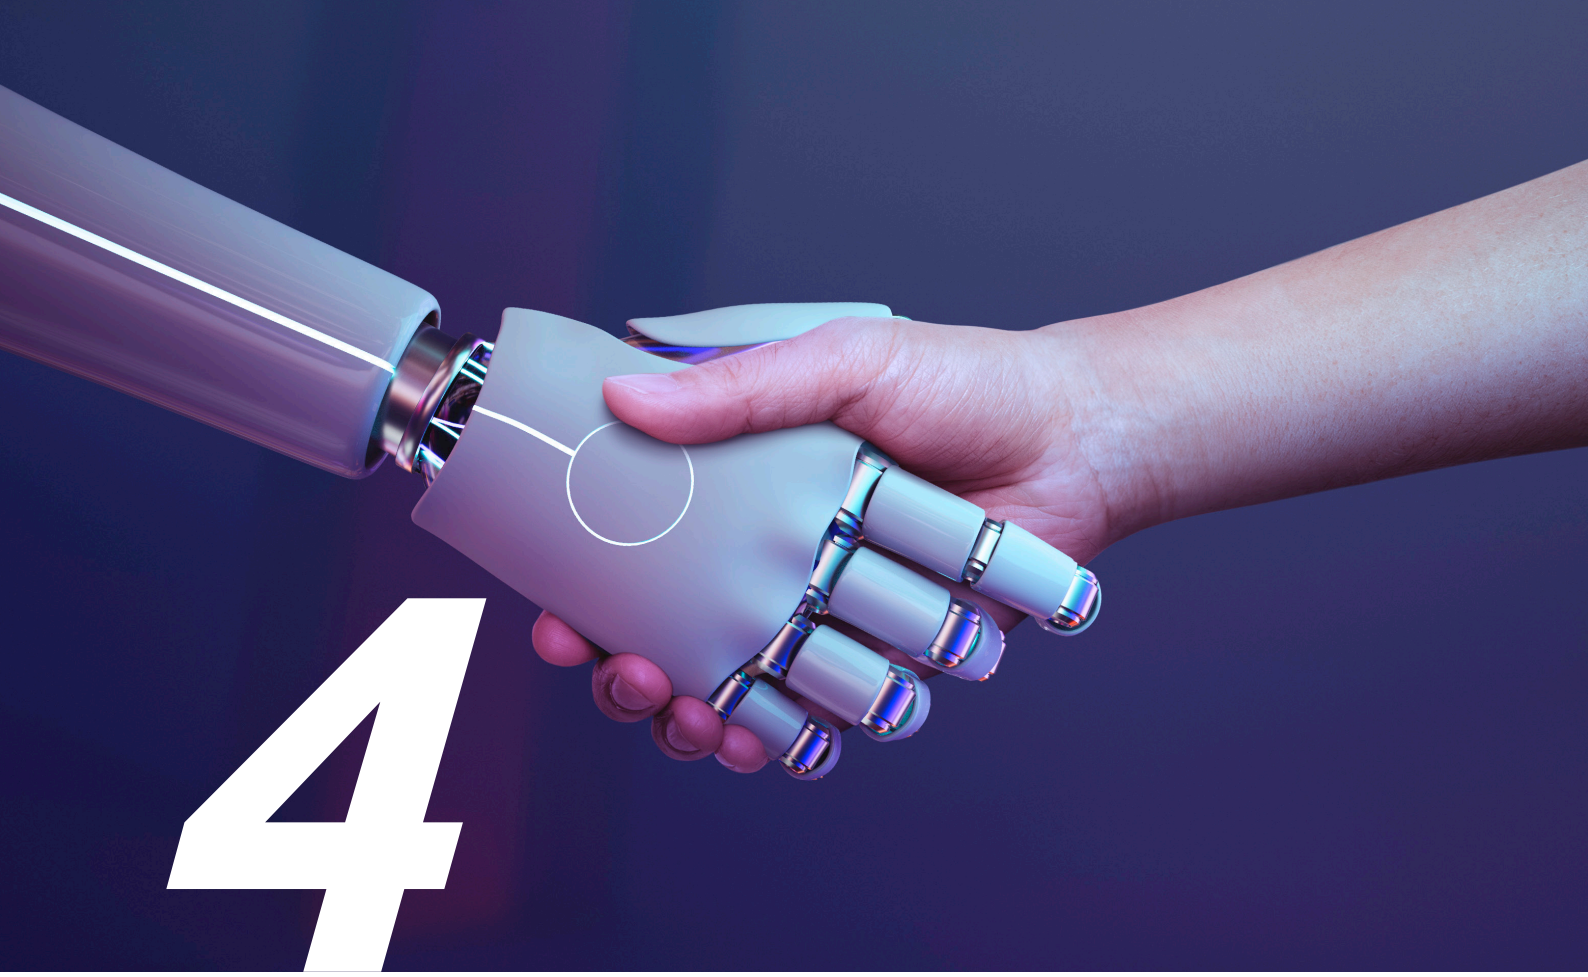

4

From hype to human — telco AI success starts with culture and trust

To realize the full value of telco AI, operators must localize at the edge, humanize at the core

Whether they're pursuing internal efficiencies or launching new services, operators are learning that AI strategy is intrinsically linked to human strategy. Conversations with VEON Group and Telus point to a common truth: successful AI strategies are of course technical but very much cultural — and that applies to consumer culture and organizational culture.

VEON, operating across six emerging markets — Bangladesh, Kazakhstan, Kyrgyzstan, Pakistan, Ukraine, and Uzbekistan— has evolved from a connectivity provider to a digital services company. Group Chief Digital Operations Officer Lasha Tabidze shared that while VEON serves 160 million connectivity users, its super apps now reach 125 million people — and 32.5 million of them aren't even SIM customers. In Q1 2025, those apps accounted for 14.5% of the company's revenue.

VEON provides connectivity and digital services in seven culturally and geographically diverse markets.

"We truly believe that digital services are not just an extension of telecom," he explained. "Creating digital services for the communities, for the customers, it's very important. Especially for the frontier and emerging markets, it's the way how they connect to the world — I mean financially, personally, socially."

He laid out VEON's DO [digital operator] 1440 strategy: "We wanted to be really relevant to all the customers every single minute of the day. That was the foundation and the cornerstone of how we started to change our strategy." Whether via serving educational, financial, healthcare, or other services, Tabidze said success is based on taking a bottom-up look on a market-by-market basis, and understanding what's "valuable for them." For 1,440 minutes per day, the "customer is at the core…We see this as our commitment and our responsibility."

Hyper-local by design: local languages, super apps, and AI for everyday relevance

Further to this point, VEON delivers its digital service offerings via super apps — another white whale many operators have been unsuccessfully chasing for time immeorial— that are intentionally branded distinctly from the local opco, eg its Pakistani operators is Jazz, the Pakistani super app is SIMOSA. "We opened up completely

"It's not just about running the race. It's about running the race differently – to do what is necessary for the market."

- Lasha Tabidze,
Group Chief Digital Operations Officer, VEON

for everyone," Tabidze said. He stressed the important combination of local data, local ecosystems, and local languages. "It's not about playing a game alone. It's about really strategic partnerships."

An extension of the DO 1440 strategy is the AI 1440 strategy where AI stands for "augmented intelligence… You have to make sure that AI, when you are embedding this into the applications or platforms or any kind of services that you deliver, it's really for the market specifically." For instance, VEON worked with the GSMA Foundry to develop a local language model for Kazakhstan which is leveraged for its consumer- and enterprise-facing services. This is meant to address systemic bias in AI and further VEON's larger business objectives.

Practical deployments to make AI a "daily business"

Specific to VEON's own use of AI, Tabidze said the company has a range of network optimization use cases, but also use cases for every operational module of the business, including customer service, human resources, marketing planning, and sales, among others. He called AI "a daily business." His advice to other operators: "There is no one recipe that everyone will use but the sooner you start, the sooner you will understand which use cases work, where you need to train more, and go forward."

Of course we use AI for our networks, sales planning, marketing activities, but it's efficiency which for VEON is "a daily business." addressing market specific nuances is very important. HR agents for local markets, for example… large scale companies in the countries need these products. Bias that AI speaks only English is deep in the mindset. You have to invest in local languages and cultures, then you create a digital product but really recreate the digital life to make it useful for each and every person in that market.

Big picture, he described a careful balance between global progression at group level and hyper-local focus at the market level. Rather than create services at the center then bring them to individual markets, VEON employs "knowledge sharing [and] practice sharing" to understand what works in specific markets that's applicable to others. "That's in the DNA."

Telus on building trust before building tools

For Telus, Chief Information Officer Hesham Fahmy explained how the company has worked to foster AI-enabled innovation by taking "a human-centric approach." For employees, that means Telus created a "walled garden to make it safe for employees." The idea was to get people comfortable and familiar so "magic can happen…really the adoption and literacy has gone across the whole enterprise."

Fahmy said around 60,000 Telus employees are using AI to augment their own workflows using a trusted, safe interface. "It's really about empowerment and adoption." He gave the example of customer support agents that are driving material KPI improvements around call containment, diversion, and customer experience; this AI assistant has responded to some 700,000 queries to date. "These are kind of behind-the-scenes, human-centric uses of AI that have also led to great outcomes for us."

Organizationally, he said, successfully implementing AI is heavily-reliant on creating "psychological safety." This includes safe sandboxes, communities for knowledge sharing, and stressing that AI isn't being deployed to displace human workers.

Sovereign and sustainable – AI as a platform for growth and responsibility

Fahmy also looked at the role of AI in helping Telus meet its own sustainability goals. Alongside the Vector Institute, Telus worked to address data center energy consumption by reducing cooling costs. He said 40% of data center energy consumption is for cooling; the two firms developed (and integrated) algorithms that have delivered a 12% reduction in annual energy consumption associated with data center cooling. They went on to open source that code. "That 12% should be for everybody to take advantage of," Fahmy said.

"Sustainability is core to our values," he continued, adding that some of Telus's financial bonds are tied to meeting sustainability targets. "We put our money where our mouth is…Consumption is increasing yet you want to drive down your environmental footprint." That's a problem suited to algorithms, math and science," Fahmy said. "We will look at every dimension we have and say, 'Is there an intelligent way…of finding creative ways to offset that growth and consumption across the board.'"

In terms of AI for revenue generation, Telus has established two sovereign AI factories in Canada. Fahmy said the goal was to address multiple dimensions of sovereignty — compute, control plane, and transport, underscoring sovereignty as a compliance measure and a strategic differentiator in selling AI services. He said sovereign AI is "most definitely" a multi-layer monetization vector ranging from compute- and GPU-as-a-Service, then services on top of compute and GPUaaS. "These are all service offerings that are revenue streams, and I think any global operator would benefit from that."

Bottomline, Fahmy said, operators need to make AI as accessible as possible so that the technological "tide rises all boats."

As VEON and Telus demonstrate, realizing the full potential of AI requires more than algorithms but further, it demands cultural fluency, a deep understanding of local context, and trust-building. With that human foundation in place, operators are better positioned to tackle the next challenge: deploying AI across distributed environments and edge networks without losing coherence or control.

"You create this really great environment and culture where people really want to lean in…I think that's been the formula Telus has had, and to really great success."

- Hesham Fahmy, Chief Information Officer, Telus

5

Getting to AI-native 6G – standards and sustainable scale

A new architecture for a new era

As operators adopt AI to solve clear and present challenges around optimizing resource utilization, gaining operational agility and efficiency, and creating new pathways to growth, AI also has a vital role to play as 6G standardization progresses. The convergence of AI for networks and networks for AI will perhaps unify to support the vision of AI-native 6G. In addition, the enormous power costs associated with scaling AI, particularly in hyperscale data centers where frontier large multimodal models are trained, should be considered alongside corporate, national, and global sustainability targets. As the space develops, delivering AI inference from the edge can help reduce the overall power budget while creating new service delivery pathways for operators. AI-RAN Alliance Chairman Choi called for a new operating system to "enable transformation" of operators and their infrastructure. "We anticipate that edge AI data centers will emerge in the near future because it's closer to the end users." But, to manage this distributed intelligence effectively, "We will need a new AI operating system for edge AI data centers specifically designed for managing fully-distributed geographically edge AI data center environments." The AI-RAN Alliance will work on this issue as part of its R&D process, Choi said.

AI for 6G spectrum sharing and orchestration

The current (and increasing) complexity of modern mobile networks already demands AI to make the machine run. With 6G targeting substantial system-level and capability enhancements as compared to 5G, the industry has already clocked another potential issue where, again, AI has a role to play: the availability of massive 400 megahertz channels which, in many markets, are not currently available in 6G candidate spectrum bands.

This, Choi said, means spectrum sharing becomes key. He described multi-radio access spectrum sharing (MRSS) as the mechanism to govern shared spectrum across 5G, 6G, Wi-Fi, satellite-based, and other radio access mediums. "As we move toward 6G, we are entering an era where spectrum sharing, massive bandwidth, and ultra-dense deployments make…conventional, manual network management completely impossible."

AI models can predict spectrum occupancy and demand patterns based on historical and real-time data to enable proactive spectrum allocation before congestion or interference occurs. Reinforcement learning agents can continuously calibrate the optimal way to allocate spectrum across different access technologies. And AI agents, working across radio and edge sites, can autonomously balance traffic loads.

"Without artificial intelligence technology, the complexity and speed required for the efficient MRSS operation in 6G would be practically unmanageable, making AI-native RAN architecture in 6G an essential foundation for the future," Choi said.

The 6G standardization outlook

The 3GPP recently hosted its 108th Plenary meeting in Prague, Czech Republic; the meeting marked the initial development of consensus on 3GPP Release 20 which will result in 6G-related study items and be followed by "normative work" in Release 21. This is aligned to meet the International Telecommunication Union's IMT-2030 timeline, meaning that technology proposals (what 3GPP does) need to be ready to go by early 2029. Keysight Technologies Chief Technologist for 6G Balaji Raghothaman framed some early priorities for 6G and looked at the role of AI in delivering on those priorities.

"In 5G, AI was an afterthought," he said. "It was an add-on…The nature of this process has brought in certain inefficiencies. So the vision for 6G is to really design from the ground up for AI/ML." Raghothaman explained that the explosion in computing power delivered by AI accelerators, as well as the optionality that exists within that technology category, will enable a distributed AI approach spanning centralized hyperscale data centers all the way out to the far edge — the devices we all use every day.

He called out emerging consensus around target 6G spectrum in the upper-mid band, use of sub-THz frequencies for things like wireless data center communications, integration of terrestrial and non-terrestrial networks for seamless, ubiquitous coverage, and the idea of integrated sensing and communication (ISAC). And AI/ML are relevant throughout.

"This is enabling us all to dream bigger and think about much larger use cases."

- Balaji Raghothaman, Chief Technologist, 6G, Keysight Technologies

Raghothaman said AI will be present in the air interface and the MAC layers, as well as in the network automation layer. "When we talk about radio access network optimization, we talk about radio resource management, traffic shaping, and then when we come down to the MAC layer we talk about scheduling and beam management… And of course the holy grail of AI for wireless has been the physical layer where there is a dream of being able to do coding modulation and receiver in a joint manner using machine learning which is completely tailored to the wireless channel at a particular instant in order to achieve the best possible capacity."

In terms of futuristic use cases 6G could enable, he discussed intelligent robots capable of environmental sensing being used for elder care, and using converged terrestrial/non-terrestrial networks to stream data from remote environmental sensors monitoring climate change-related metrics. In those and any other cases, Raghothaman said Keysight is focused on evolving end-to-end testing needs as AI permeates increasingly complex networks.

"There is a whole new aspect to the testing," he said, describing lifecycle testing of AI models, testing for data fidelity, and other externalities that come into play. "Digital twins are going to be key," he said.

Sustainable scaling – energy efficiency, heterogeneous operations, and avoiding vendor lock-in

Operators are integrating AI across distributed infrastructure — from the core to the edge — to improve network performance, customer experience, and operational efficiency. Gen AI, in particular, is gaining traction, offering opportunities for cost savings and automation, and also for monetization through emerging models like GPU-as-a-Service and sovereign AI.

"Maintaining a balance in terms of cost, maintaining a balance in terms of the real-time decision making point, and then the challenges it brings with the high power consumption – I think maintaining a balance in all these three points is going to be the hot topic for the industry to process."

- Ravi Sinha, Vice President, Technology and Solution Development, Reliance Jio

"When we on the ground want to hand over control to AI systems, often we meet some resistance with engineers…They used to understand everything and AI is sort of a black box that you sometimes can't really explain…I think it's very important to address." ."

- Johannes Bjelland, Program Direction, Intelligent Networks, Telenor

However, scaling AI in a distributed telco environment comes with steep challenges, including:

- Inconsistent infrastructure across cloud, RAN, edge, and customer premises.
- Limited compute and power capacity at the edge.
- Fragmented data pipelines and lack of interoperability.
- Vendor lock-in and overlapping standardization frameworks.
- And cultural barriers and operational conservatism.

As established throughout this report, the technological demands of AI adoption must be matched by organizational transformation. For the telecoms workforce to use AI, it will take more than just performant solutions; it will also take cultural change and trust.

The vendor lock-in question is also a bit layered. Taking Open RAN as an example, operators are very much invested in hardware/software disaggregation in the RAN, mirroring earlier transformations of the core network. Yet increased interoperability comes with integration, testing, and assurance burdens. Whether the lower upfront capex offsets higher opex remains unresolved. These tradeoffs extend beyond RAN to cloud platforms and AI systems, both additional areas where lock-in risks are material.

Given these complications, the question arises: how can operators sustainably scale AI, especially compute-intensive gen AI, across a highly distributed, energy-constrained infrastructure while addressing complexity, cost, and IT sprawl?

The answer isn't clear or simple but it involves model evolution toward lightweight, fine-tuned models, silicon enhancements for energy-efficient inference, federated edge platforms, standardized APIs and service management and orchestration architectures to reduce integration complexity, organizational buy-in on AI as core to future operating models, and agentic AI systems that support dynamic, efficient workflows.

Sinha described an architecture where small language models (SLMs) do 80-90% of the work in distributed agentic systems, with LLMs reserved for complex, global coordination. "In that way you can also save a lot of

"We don't want to repeat the former SON experience where you had a lot of vendor-proprietary applications and algorithms in place...We want to open that up and make sure that we have interoperability and portability across the apps."

- Olli Andersson, CTO, Telecom Infra Project

"Gen AI is not just a use case, but it's a strategy to reshape the infrastructure. And without a continuum between cloud and edge, it becomes unsustainable at scale."

- Luca La Brocca, Business Development Manager, Tiscali

energy and a lot of value you can add, and reduce the cost as well." He said infrastructure for agentic systems should target 80-90% SLMs.

Discussing standardization approaches to speed up cross-domain AI and data systems, Bjelland said the problem Telenor faces today is quickly scaling which impacts time-to-value. "That's really hard because now we spend so much time setting up an AI agent and making the data streams work…The time it takes to get that up and running, it needs to be reduced." He pointed out that operators' cloud strategies, which also speaks to the vendor lock-in point, can materially change AIOps and operation and management of agentic systems. "So getting kind of standardized and getting those platforms…in place, I think it's a must before we really can see efficient scaling."

Andersson also emphasized the need for well-defined, modular blueprints and reference architectures as a key piece of sustainable scaling. Tsicali Business Development Manager Luca La Brocca highlighted the need for composable AI that can be deployed across the entire infrastructure in a heterogeneous manner that will drive scale and, in turn, impact. This is particularly true for resource-intensive gen AI use cases.

Building the AI-native telco

As 6G research moves into early standardization, AI is emerging as both a foundational capability and a design principle. Alexandros Kaloxylos, executive director of the 6G Industry Association, said that across more than 80 R&D projects under the European Smart Networks and Services Joint Undertaking, over 200 AI solutions are being trialed. These span network management, control, security, and AI-as-a-service — with particular emphasis on energy efficiency.

DOCOMO Euro-labs Director Riccardo Guerzoni added that AI is already influencing 5G-Advanced via 3GPP Release 18, and will take on a central role in 6G system architecture beginning with new study items now under discussion.

One major area of focus is building the data frameworks needed to train and validate AI models in a trustworthy, regulation-compliant, and efficient manner. Kaloxylos stressed the need for realistic, benchmarkable datasets — ideally synthetic but operator-approved — to support reproducibility and cross-regional collaboration. Guerzoni described ongoing 3GPP work on AI lifecycle management and data governance, with the likely introduction of a dedicated data plane in 6G to handle massive volumes of data from core, RAN, and user devices. Both speakers noted the tension between innovation and regulation, particularly around user consent, ownership, and privacy.

Another challenge lies in aligning AI performance with energy efficiency. AI can optimize network energy usage, but it also consumes power, creating a double-edged sword. Kaloxylos advocated for sensible design choices and federated learning architectures that can distribute training loads. He cautioned against applying complex AI models where simpler logic would suffice — a theme echoed in other sections of this report.

"The projects have demonstrated that indeed AI is... the cornerstone of the 6G networks."

- Alexandros Kaloxylos, Executive Director, 6G Industry Association

"AI can be used…to reduce the energy consumption in the network," Kaloxylos said. "That's one thing. But…this does not come for free. And AI has to consume energy in order to provide the solutions."

When asked about 6G's "killer app," Kaloxylos offered a grounded answer: energy optimization will likely be the most immediately valuable internal application, while AI-powered personal assistants could emerge as a consumer-facing service opportunity.

Guerzoni expanded on the role of Open RAN, confirming that 3GPP and the O-RAN Alliance are now coordinating 6G-era efforts. 3GPP will define the overarching AI architecture and lifecycle management, while O-RAN will build complementary services such as AI models for RICs and energy-aware optimization agents. This convergence underscores the importance of architectural interoperability and energy-aware design principles. This coordination crystallized in a joint meeting between the two groups in March.

He also highlighted how user equipment itself could contribute data for training air interface optimization models, creating both opportunity and complexity around edge inference and upstream data transport.

Taken together, the discussion reflects how the ambitions of 6G — spectrum sharing, AI-native design, hyper-flexibility — cannot be realized without revisiting core architectural assumptions. They also reinforce a central truth from this report: AI in telecoms is not an overlay, but a reorientation. Designing AI-native systems will require robust foundations in data governance, open standardization, and cross-domain collaboration. And while no one can fully predict 6G's killer app, one thing is clear: AI will be at the heart of both the problem and the solution.

"O-RAN is part of the picture and we'll work together with 3GPP for the 6G standardization."

- Riccardo Guerzoni, Director, DOCOMO Euro-Labs

"AI will be part of 6G from day one, from optimizing RAN performance to addressing the complexities of radio beam management, optimizing circuit design, facilitating more efficient traffic flows, and reducing power consumption."

- Nancy Friedrich,
6G Solutions Expert, Keysight

Conclusion

Telco AI represents both extraordinary opportunity and unprecedented complexity. As operators evolve from proofs-of-concept to production-ready AI deployments, the path to sustainable success demands cultural transformation, strategic clarity, and technical acumen. Key lessons that emerged during the Telco AI Forum include: holistic integration across infrastructure, operational workflows, and customer services is essential; interoperability, modularity, and standardization will mitigate complexity and vendor lock-in; and cultural readiness, underpinned by trust, will determine the pace and effectiveness of AI adoption.

Looking toward the AI-native 6G era, executives must actively engage in shaping emerging standards and infrastructure paradigms, ensuring they are built from the outset to embed and leverage AI capabilities fully. In doing so, telecoms leaders will not only future-proof their organizations but position themselves to capitalize on AI-driven opportunities that redefine the competitive landscape. The strategic imperative is clear — embrace AI comprehensively, scale sustainably, and position for a future where AI is the central axis of operational and market differentiation.

About the sponsors

Dell Technologies

Dell Technologies offers a portfolio of use-cases with Dell AI for Telecom, combining AI Factories, edge computing, and 5G solutions to transform telecom networks. With open, scalable architectures and strong industry partnerships, Dell empowers operators to optimize networks, enhance efficiency, monetize in the enterprise and deliver smarter, more connected experiences.

Learn more: www.dell.com

Keysight

Keysight inspires and empowers innovators to bring world-changing technologies to life. As an S&P 500 company, Keysight is delivering market-leading design, emulation, and test solutions to help engineers develop and deploy faster, with less risk, throughout the entire product life cycle.

Learn more www.keysight.com

NVIDIA

NVIDIA and Dell Technologies collaborate through the Dell AI Factory, combining Dell's scalable infrastructure with NVIDIA's powerful GPUs and AI platforms. This partnership empowers telecom operators to harness AI for optimizing networks, accelerating 5G deployments, and delivering smarter, more connected experiences.

Learn more www.nvidia.com

Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions and services, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT, and automate, secure, and manage complex environments.

Learn more: www.redhat.com

Spirent

Spirent provides innovative products, services, and managed solutions that address the test, assurance, and automation challenges of a new generation of technologies from the lab to the real world, including 5G, SD-WAN, Wi-Fi, O-RAN, cloud, autonomous vehicles, and beyond.

Learn more www.spirent.com

MARKET PULSE REPORT

Telco AI

RCR Wireless News
INTELLIGENCE ON ALL THINGS WIRELESS