

Agile Telco

OPTIMISED OPERATIONS FOR MODERN CONNECTIVITY PROVIDERS

Why telcos are integral to a converged AI infrastructure

Nokia's CTAIO Pallavi Mahajan explains the programmable infrastructure, datacentre fabrics, optical transport, autonomous network operations and AI-RAN innovations that telcos are adopting to succeed in the AI supercycle

- Automated NOCs - Into the heart of darkness?
- The unstoppable rise of the autonomous telco
- How multi-agent AI is replacing the telco BSS middle layer




WWW.THEAGILETELCO.COM



Change the future of global connectivity

Find out more



Become a TM Forum Member



INNOVATE AMERICAS

October 6-7, 2026
Dallas, USA

INNOVATE ASIA

17-19 November 2026
Bangkok, Thailand

DTW ignite

29 June - 1 July 2027
Copenhagen, Denmark

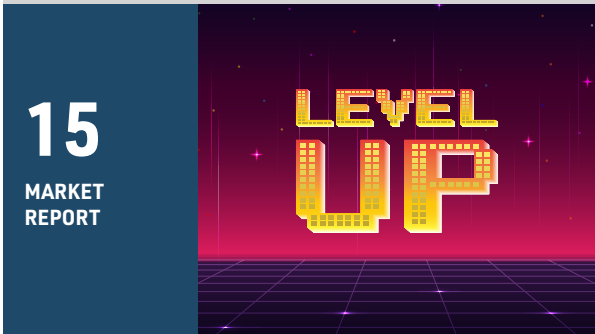
Future events

Join us at an upcoming event





06 COVER INTERVIEW
Pallavi Mahajan, Nokia



15
MARKET REPORT



21
AUTONOMOUS TELCOS



26
MULTI-AGENT AI

Agile Telco

TECHLED

IN THIS ISSUE

05 COMMENT
George Malim shines a light on the dark NOC

06 COVER INTERVIEW
Nokia's CTAIO, Pallavi Mahajan, explains why telcos are integral to a converged AI infrastructure

10 AUTONOMOUS NETWORKS
Vivek Jaiswal details how telcos will thrive in the AI supercycle

15 MARKET REPORT
The step from here to autonomy looks to be mainstream adoption of TM Forum's Level 4 autonomous networks maturity architecture. What will it take for telcos to rise up the levels, asks George Malim?

21 AUTONOMOUS TELCOS
Chantel Cary says the telecoms industry is at a decisive moment and reports on the rise of the autonomous telco

26 MULTI-AGENT AI
Martin Rueckert explains how multi-agent AI is replacing the telco BSS middle layer

30 AGENTIC AI
Tony Savvas assesses the latest developments involving agentic AI for telecoms

36 LEVEL 4 AUTONOMY
Peng Zheng sets out what's next for telcos at autonomous networks L4

38 EVENT PREVIEW
As the industry once again heads to Copenhagen in Denmark for DTW Ignite, don't miss our preview of the event

43 EVENT PREVIEW
GSMA's Sihan Bo Chen showcases the next phase of mobile innovations lined up for MWC26 Shanghai

45 AFRICAN CONNECTIVITY
GSMA reports on how Africa is rewriting its connectivity story

TECHLED Agile Telco

OPTIMISED OPERATIONS FOR
MODERN CONNECTIVITY PROVIDERS

WWW.THEAGILETELCO.COM



Telcos step fearlessly into the heart of darkness

The dark NOC sounds menacing and laced with evil intent. A solitary robot vacuum cleaner glides across the carpet, picking its path between abandoned operator chairs. Lily pads of mould linger on remains of hastily abandoned coffee cups. Discarded headsets gather dust and the space is lit only by the emergency exit sign, the displays are black and the desk lighting is off. The nuclear holocaust hasn't happened, it's only the adoption of agentic AI, digital twins and autonomous networks that has made the change.



George Malim,
managing editor

Control-alt-delete has been pressed for the last time as the users of these chairs logged-off permanently and the dark NOC is now the autonomous, beating heart of network operations led in continuous, automated processes by an artificial intelligence that is proving to be very real.

The decision-making, data analysis, modelling and outcomes previously steered by humans with all their inefficient personal, social and financial needs has been replaced with frictionless systems that predict, automatically redimension and keep the network protected.

The dark NOC is an all-orchestra symphony of systems, process and functions integrating optimised operations across bandwidth utilisation, infrastructure repairs and renewal, as well as real-time configuration and reconfiguration to meet traffic peaks and troughs. Just because it's dark and no one can hear it working, that doesn't mean the dark NOC isn't making beautiful music. In fact, it's replacing the noise and tensions of traditional network operations centres with a smooth, uninterrupted continuous movement to deliver the best experiences in the most efficient ways.

The reality today is that we're a lot of extinguished lightbulbs away from a truly dark NOC. The NOC relies on a complex web of systems, not all of which are ready to integrate efficiently with AI. Organisationally, few telcos are ready to trust their telco-grade operating systems exclusively to AI agents and the skills of humans are still required to deal with non-routine processes.

We are at the start of a new era and, while it's clear that significant spaces in the NOC can at least be dimmed as workers move to other tasks in the telco, humans remain essential controllers of agents and exclusively have the ability to manage many functions and operations. However, it's fair to say some desks are already dark and further NOCs will grow silent as humans move into higher level operations centres with smaller, localised NOCs operating humanless in the dark.

Stay agile – and carry a torch!

George Malim

MANAGING EDITOR
George Malim
Tel: +44 (0) 7930 301 841
g.malim@wkm-global.com

MANAGING DIRECTOR
Nathalie Millar
Tel: +44 (0) 1732 808690
n.millar@wkm-global.com

SALES CONSULTANT
Cherisse Jameson
Tel: +44 (0) 1732 807410
c.jameson@wkm-global.com

DESIGN
Jason Appleby
Ark Design
Tel: +44 (0) 7801 817 139

PUBLISHED BY
We Know Media Ltd, Suite 138, 80
Churchill Square, Kings Hill, West
Malling, Kent ME19 4YU, UK
Tel: +44 (0) 1732 807410

All rights reserved. No part of this publication may be copied, stored, published or in any way reproduced without the prior written consent of the Publisher

weknow

Agile Telco magazine is published as a digital edition at least 4 times a year. It is available free of charge to all readers worldwide, at the publisher's discretion. To subscribe **free of charge** go to: www.techledworld.com/agiletelco, click the "Subscribe" tab on the Home Page and answer the questions shown. The Publisher reserves the right to alter or end this free offer at any time without notice. No guarantee is stated or implied. You can unsubscribe at any time by emailing subs@techledworld.com with UNSUBSCRIBE in the Subject line.

AI-native networks combine programmable infrastructure, datacentre fabrics, optical transport, autonomous network operations and AI-RAN from the cloud to the datacentre and at the edge

As telcos prepare to support AI workloads at hyperscale, their role is shifting from being a connectivity supplier to becoming integral to a converged AI infrastructure that supports critical workloads from the edge to the datacentre to the cloud. Pallavi Mahajan, the chief technology and AI officer at Nokia, tells George Malim, the managing editor of Agile Telco, how she sees telcos enabling and thriving from the AI supercycle as they build on their trusted networking capabilities to play their part in the complex models and opportunities that the various iterations of AI demand.

Mahajan has shaped industry-defining platforms and innovative business models across AI, infrastructure and networking. Her experience spans the full compute spectrum – from accelerated computing across the datacentre, cloud and edge to leading the first exascale supercomputer deployment, Frontier, and architecting software-defined networks. Today, as chief technology and AI officer at Nokia, she is driving the shift to AI-native, distributed systems where compute, networks and intelligence converge. Her remit spans technology, AI and platform strategy, steering Nokia's roadmap and scaling its role in next-generation infrastructure.

Before Nokia, Mahajan held senior leadership roles at Intel, Hewlett Packard Enterprise and Juniper Networks. In each role, she led at the frontier across AI datacentres, high performance computing and autonomous networks. Her work has consistently focused on platformisation at scale – building systems designed for performance, adaptability and impact. The result: a track record defined by turning complexity into platforms that power the future of AI infrastructure.

SPONSORED INTERVIEW

George Malim: The AI supercycle has largely been framed as a compute story composed of GPUs, frontier models and hyperscaler capex. From your vantage point leading AI and technology at Nokia, what is missing from that framing?

Pallavi Mahajan: The compute-only framing isn't wrong, it's just incomplete. In frontier AI, the unit of compute has evolved from chips to entire AI factories. The real breakthrough for the AI supercycle isn't bigger silicon, it's the network, and that is what turns many machines into systems.

As the unit of compute changes, AI workloads themselves are changing and have become far more distributed. They are no longer monolithic, they are fragmented and distributed by design. We are starting to disaggregate inference across multiple dimensions at once – prefill and decode are being separated because they have very different compute profiles. Distributed inference is now a co-ordinated system rather than a single process and that co-ordination is the new complexity.

The Mixture-of-Experts (MoE) models route tokens dynamically across different experts that may sit on entirely different GPUs while agentic systems chain multiple models together in coordinated workflows which have to address perception, memory and control all in real-time.

Think of a coordinated swarm of robots carrying a fragile object together. You need compute coordination because if one participant falls behind, the system itself becomes unstable. In an AI grid, the network is not just carrying workloads, it's part of the infrastructure itself.

Over the next decade, the industry will increasingly realise that AI infrastructure is not only a compute story, it is fundamentally a distributed systems story. And, in distributed systems, the network becomes part of the infrastructure itself.

GM: As AI workloads become increasingly distributed, how does the relationship between workload orchestration and the network itself change?

PM: The workloads are increasingly getting distributed across racks and across pods. Today workload orchestration and network orchestration are starting to happen but the old model was that the compute is going to decide and the network is going to follow. That whole model is breaking and AI is fundamentally changing it because two control planes that have to work together

continuously are emerging. One is the workload control plane, which decides where models, experts, prefill, decode and inference tasks should run. The other is the network control plane, which exposes the real-time fabric state and adapts the connectivity dynamically to support those placement decisions.

Now the challenge is that both these control planes operate at different layers and often at different timescales. The workload scheduler may say it can see GPU capacity but the network sees congestion. It's seeing degraded optical paths or changing traffic conditions and, fundamentally, you need these two brains to act as one.

The future architecture is not just about faster networking, it is about converged orchestration. In this case, the two brains, the two control planes, are coming together and we're building up autonomous networks because compute orchestration without network awareness is just impractical.

Pallavi Mahajan
Nokia



The critical interface becomes the boundary where workload intent and network state exchange information continuously. This becomes even more important as infrastructure extends beyond the datacentre, especially when you start looking at edge compute, the operator's network and the AI radio access network (RAN) environment.

Historically, infrastructure was optimised around centralised pools of compute but AI is creating a far more dynamic environment with compute placement, workload movement and network behaviour increasingly having to adapt together in real-time.

And as distributed inference expands outward, the network stops being passive infrastructure and starts participating directly in workload execution behaviour.

For Nokia, we see this as a major opportunity. We operate across datacentre fabrics, optical transport, autonomous network operations and AI-RAN so, for us, it's not just an opportunity to build high performance connectivity. It's about creating the programmable control architecture that allows distributed systems to operate coherently and at scale.

GM: Most people experience AI as a chat window or an image generator. Underneath, the workloads driving those experiences look nothing like traditional internet traffic. What is actually different and what does that mean for how networks are built?

PM: When it comes to the traffic pattern, networks were first built to carry voice and then came data, next came rich media content and now it's about AI. And even within AI, there is foundational AI, vision-based AI, generative AI, agentic AI and physical AI to consider. With generative AI, what we started to see was that the communication is more like downlink-heavy but now there's so much traffic getting uploaded. Think about if you have your AR/VR glasses, there is so much content that's actually being uploaded versus what is getting downloaded from those.

When it comes to agentic AI, networks were always designed with a busy hour traffic paradigm. Agentic AI changes that because agents never sleep, they work 24/7 and the traffic is very bursty and it's very, very sensitive to latency. When it comes to physical AI, going back to my example of swarm of robots, a 10 millisecond latency is not like a pixelated movie experience that you can just forget. It is actually a critical safety issue so there are a lot of changes that happen from a network traffic point of view and what is demanded from the network.

With a chat window, every prompt can trigger coordination across thousands of GPUs and specialised models. Now, the AI workload behaves very differently from web traffic. The first difference is disaggregation, because modern AI inference pipelines are increasingly fragmented across different compute domains, such as racks, pods and datacentres and different stages of inference, whether it's reasoning, retrieving, retrieval or agentic coordination, they all now operate across this distributed infrastructure.

Traditional internet traffic tolerates unpredictability surprisingly well but AI workloads do not. Once the network becomes unstable, the entire stack above it has to compensate with larger safety margins, more retries and more context handling.

Earlier we talked about the coordinated swarm of robots carrying a fragile object together. That kind of system only works if all participating models, agents, inference components and physical actors stay coordinated in real-time. The moment one participant falls behind because of network instability, the entire system has to compensate, slow down or fail safely.

We are already seeing this in frontier AI infrastructure. OpenAI's recent MRC work is a good example where the goal is not simply

With a chat window, every prompt can trigger coordination across thousands of GPUs and specialised models

more bandwidth but better predictable performance under congestion and failures. At large training scale, even a late transfer can ripple through the entire job and leave GPUs idle.

What does this mean from a network design point of view? So far, I've just painted a picture of the problems that we have right now. We need to change how infrastructure gets built. Inside the datacentres, the fabrics are increasingly optimised for the traffic patterns that we just talked about. We have optical interconnects that will support 800 gigabit to 1.6 terabit class architectures. We are increasingly coming up with architectures where lossless transport becomes increasingly important.

In addition, the inter-datacentre connectivity increasingly behaves like part of the same coordinated execution environment rather than a traditional WAN. AI workloads do not stop at the datacentre boundary and, as distributed inference expands outward towards operator networks and edge infrastructure, the same architectural principles increasingly extend across the entire stack.

GM: Inside the AI datacentre itself, there is a quieter revolution happening at the physical layer: optical circuit switching, co-packaged optics, the shift away from copper at almost every scale. What is driving this, and where does it lead?

PM: The biggest constraint in AI infrastructure is power. It's not like floor space or compute availability, it is just power. Today, there's a lot of electrical to optical to electrical conversion that happens in the data path. Every time you do this conversion, it consumes energy, which we cannot afford to waste so to me, that is a forcing function pushing optics deeper into the datacentre.

The easiest way to think about this transition is across three scales. Scale-up is inside the rack with GPUs communicating across very short distances. Copper has dominated this layer traditionally but it is approaching its physical and power limits. Scale out is where 800 gigabit and eventually 1.6 terabit optical interconnects become increasingly important. Scale across is where the inter-datacentre communication improvements are needed.

Coherent pluggables are now extending high-capacity optical connectivity across campuses and metro-scale AI infrastructure and we are seeing several important transitions happen simultaneously. Co-packaged optics bring optical engines close to the switch application-specific integrated circuit (ASIC) itself, reducing the electrical path where power is heavily consumed. At the same time, optical circuit switching allows optical paths to be reconfigured dynamically without repeated optical-electrical-optical conversion. Google has already demonstrated this in production environments.



We are also seeing transitional approaches like linear-drive optics and coherent-lite architectures focused on reducing power and latency while remaining deployable at hyperscaler volume.

Now, for me, the long-term trajectory is clear, optics is moving inwards. For decades, when we thought of optics, we primarily thought of the technology living between buildings. Now optics is moving into the rack, into the package and eventually closer to the silicon itself. The deeper architectural shift is that the optical layer is becoming programmable so it is no longer just passive infrastructure sitting underneath the workload. Increasingly, the workload and the orchestration layers interact directly with the optical fabric itself so the fabric is becoming software driven.

For us, the opportunity is not just to build components, it's about building up a layer of autonomous network which is designed from the start to participate in the broader AI infrastructure control architecture.

GM: The boundary between the datacentre and the radio access network is starting to dissolve. AI-RAN is one of the most watched architectures in the industry right now. What is AI-RAN actually solving and why is it more than just having GPUs in a base station?

PM: I think the GPUs in a base station framing is right but it misses the deeper architectural shift. AI-RAN is not simply adding AI to RAN, it is about programmable RAN, not GPUs. AI traffic is changing the shape of the network demand. AI traffic just breaks the old models. More than 50% of ChatGPT and Gemini's traffic is originating from mobile devices and these workloads are more bursty, more uplink-intensive and more latency-sensitive than traditional mobile traffic.

Agentic systems involve continuous background interactions and physical AI brings in the whole concept of deterministic control requirements. Think about the swarm of robots example and how it's so different from traditional consumer workloads. Now, looking at all of that, it's clear the RAN is where much of the pressure first lands. If you look into it historically, RAN performance has depended on tightly coupled algorithms with custom silicon. That delivered very high efficiency but it also tied the innovation cycles to hardware refresh cycles.

AI is evolving differently. Generative AI was 3.5 years back, agentic AI happened in 2025 so AI is moving at software speed and networks must go out and follow the same speed. You cannot just be innovating at the speed of hardware so, for us, AI-RAN starts with hardware-software decoupling. Nokia's approach is to use a hardware abstraction layer that allows software innovation to evolve independently from the underlying compute platform. The hardware

provides the stable execution capability but the intelligence sits with the software and it evolves at the software cadence.

In Nokia's architecture the intelligence is injected rather than fused permanently into the platform. Once intelligence enters the execution path, the RAN can no longer behave like a fixed-function pipeline and becomes a programmable platform. Models enter through governed interfaces and can be validated, replaced, refined or removed independently without destabilising the underlying system. This is an important balance because operators need both rapid innovation and carrier-grade stability.

To me, AI-RAN is larger than radio optimisation. The edge now becomes a runtime environment, not just an access layer. AI-RAN is not something that we are waiting for 6G to happen to provide. We are already on the path and when 6G happens, it's just a software upgrade for us. I feel the industry is already moving towards a software-led architecture where intelligence evolves continuously rather than waiting for the next G to happen to make that innovation come in.

GM: As AI takes on more operational decision-making across hyperscaler fabrics, operator networks, and increasingly across both, the question of trust becomes critical. How is Nokia approaching this, and where do you see the company playing a defining role?

PM: The trust problem is actually remarkably similar across hyperscalers and operator environments. Autonomy without governance is chaos so, whether it is an autonomous controller sitting inside a datacentre fabric or an autonomous controller managing slicing and policy decisions inside a telecoms network, the underlying challenge is the same. If you have AI systems that cannot explain themselves, they cannot be bounded safely, you do not have the ability to roll back reliably, then you know they're not going to be deployable at infrastructure scale.

The simple truth is if you can't see what the network is doing, you simply cannot trust it. Autonomous infrastructure has to operate as a glass box, not as a black box and it means three things. First of all, it should be inspectable, which means that operators can understand what the system is doing and why. The second is that it is bounded, which means that the system operates within explicit safety limits. Finally, the third is about reversibility, which means that it must always be possible to roll back safely when needed. Glass box autonomy makes networks trustworthy, explainable and reversible.

This is where digital twins also are very important in this model. Before any autonomous actions touch production infrastructure, they need to be validated against digital representations of the resource and the service layers. Given that we are the only trusted vendor that has a portfolio that ranges all the way from wireless networks to datacentres, this is where, across the multiple layers of the infrastructure stack, we are building up the autonomous network with a very consistent trust and governance model.

The human role also changes from approving every action to defining policies, trust boundaries, escalation rules and rollback conditions rather than approving every action individually. It has moved from human-in-the-loop to human-on-the-loop.

Of course, we have a lot of partnerships going on. We're working with Google Cloud on agentic automation and digital twins, Microsoft on data integration and cybersecurity automation, and AWS on intent-based network slicing. For us, we see that the future is converging around programmable infrastructure, optical innovation, autonomous operations and distributed AI execution extended all the way to edge. And across all these layers, trust, governance and observability are foundational.

www.nokia.com

Telcos need autonomous networks to thrive in the AI supercycle

The AI supercycle offers a transformative growth opportunity for telecoms providers. To seize it, networks need to provide deterministic connectivity. A critical part of this shift will be investing in highly adaptive autonomous networks that can anticipate changing traffic mixes and operate at machine speed to enable rapid service creation and seamless execution, writes Vivek Jaiswal, the senior vice president of Autonomous Networks at Nokia.



SPONSORED ARTICLE

Vivek Jaiswal
Nokia

Nokia's AI-native automation framework was built for the AI era, combining data, orchestration and resource automation to deliver essential and trusted networks in the age of AI. Importantly, AI isn't just more traffic, it's different traffic and the AI supercycle is unfolding in waves.

Generative AI is already mainstream. Agentic AI is making it possible for goal-oriented software agents with 'human-like' reasoning capabilities to drive business and operational outcomes, enabling them to interact and coordinate with each other. Physical AI will make latency and reliability non-negotiable, merging the digital, physical and human worlds.

Nokia's global wide area network (WAN) forecast projects AI traffic will reach 921 exabytes per month by 2034, growing at an annual rate of 23% and accounting for around 30% of total global WAN traffic.

But the bigger disruption is qualitative. AI workloads are throughput-hungry, latency-sensitive and, crucially, bursty and distributed across datacentres, metro networks, edge clouds and enterprise sites.

Network behaviour determines which AI architectures are best suited, impacting inference placement, model synchronisation and control loop speed for industrial automation. For CTOs, connectivity increasingly needs to become a programmable element of distributed AI computing.

Why best-effort operations fall short in the AI era

If traffic changes in character, not just volume, the network can't remain a system that is only optimised at its time of design. It must evolve continuously, at software speed, without sacrificing predictability. AI-native networks therefore need to be both adaptive and deterministic.

This is where legacy network operation models break down. AI services don't experience the network one domain at a time. They experience it continuously.

As a result, the network cannot respond fast enough when AI workloads shift traffic patterns, compute demand or service paths in real-time. That makes legacy best-effort networking unsuitable for services such as edge AI

Network behaviour determines which AI architectures are best suited, impacting inference placement, model synchronisation and control loop speed for industrial automation



inferencing, real-time physical AI-powered industrial automation, enterprise AI assistants, video analytics and distributed model synchronisation.

Autonomy is the new prerequisite

Autonomous networks are becoming essential for telecoms providers that want to turn AI demand into profitable growth. The core idea is simple: the network must maintain continuous awareness through deep observability, apply predictive analytics and policy-based optimisation to decide the best response. Closed-loop automation ensures that the network acts as one – before the customer notices anything has drifted.

Done right, this is what makes services-on-demand real. The operator can spin up an assured service, adjust resources as the workload moves and keep outcomes inside service level agreement (SLA) boundaries. This is why the industry is so focused on Level 4 autonomy. Translating into high automation with minimal human intervention, but still bounded by business intent, policy constraints and governance guardrails.

A complete network acting as one

Autonomy doesn't emerge from a collection of isolated AI features. It requires complete automation architecture that connects strategy to execution, business, service and resource layers. Practically, that means intent models that translate objectives into measurable targets, such as orchestration that coordinates actions across domains, automation that works in multi-vendor environments, and an observability fabric that correlates telemetry, service health, experience signals,

flows and security events in near real-time.

Underneath, you need the operating system for AI at scale: data products, machine learning and large language model operations (MLOps/LLMOps) pipelines, reusable workflow libraries and governance. Nokia's approach is to unify these elements into a complete network, acting as one, so operators can industrialise cross-domain use cases.

Turn intent into real network behaviour

If there is one factor that separates automation from real autonomy, it's intent: intent captures the essence of what the customer wants – abstracting and taking care of the details of how this happens. Intent translates customer needs into concrete requirements for the network, using closed-loop automation to coordinate action across domains – and most importantly, ensuring that customer needs are fulfilled.

In practice, intent means targets like latency ceilings, throughput guarantees, reliability bounds, or priority treatment for enterprise and AI workloads – and Nokia has already proven this in TM Forum catalyst projects with customers.

Agentic AI matters because it doesn't stop at dashboards or recommendations. It can monitor goals, detect drift, trigger remediation and rebalance resources in near real-time – so contracts promise operational reality rather than post-mortem reporting.

None of this starts from zero. Agentic systems build on today's AIOps – analytics, observability, assurance models and existing workflows – but shift them from human as doer to human as observer. Safety becomes

Autonomy doesn't emerge from a collection of isolated AI features



Zero-touch is about keeping complexity from exploding opex as services become more dynamic

the differentiator. Digital twins provide a controlled way to validate intent translation, simulate policy changes and test agent actions before they run live. And, as ever, the fuel is trusted data: real-time inferencing is only as good as the data pipelines feeding it.

The glass box requirement

As autonomy rises, so does the blast radius of mistakes. That's why telecoms can't accept black box AI in operations – especially when the goal is deterministic connectivity for high-value workloads.

Nokia describes the answer as a glass-box approach to governance. As our chief technology and AI officer Pallavi Mahajan tells Agile Telco in the interview that precedes this article, glass box autonomy makes networks trustworthy, explainable and reversible.

This type of governance extends beyond code to model, policy and data lineage – so autonomy can scale without sacrificing control, compliance or operational confidence.

Zero-wait, zero-touch, zero-trouble

Zero-wait is about compressing time-to-revenue for assured enterprise services. For example, Lightstorm India's automation initiatives have resulted in close to 90% faster time-to-revenue with automated service provisioning in fixed networks, and 500 times faster order processing and delivery. Similarly, Saudi Arabia's STC reports that its autonomous networks initiative resulted in significant reduction in time to fulfil new service requests –

from six days to seven minutes.

Zero-touch is about keeping complexity from exploding opex as services become more dynamic. This has enabled UAE telco du to benefit from nearly 80% improvement in triage and troubleshooting across its IP network, while STC has been able to perform 10,000 autonomous actions an hour in its RAN. And for multi-domain services like 5G slicing, the direction is clear: automation must coordinate resources across RAN, transport, core and cloud so slices can be created – and adjusted – at the pace enterprises expect.

Zero-trouble is the payoff. Nokia has reported that applying AIOps into fixed broadband assurance scenarios can improve assurance effectiveness by more than 40% for customers – helping detect issues earlier, reduce noise and focus remediation where it matters. In an AI economy, this is not just better operations. It is what protects premium revenue by keeping outcomes inside deterministic bounds.

Together, these three zeros turn autonomy into a monetisation engine: faster productisation, lower cost-to-serve and higher confidence in premium contracts.

They also get operators ready for the upcoming stage of the AI supercycle, where connectivity and computing work closely together, and services are assembled in real-time across cloud, edge and network areas.

All told, the AI supercycle is a once-in-a-generation opportunity for telecoms providers to move beyond connectivity and become foundational enablers of the AI economy. And that requires building the autonomous networks foundation to deliver assured, adaptive and deterministic connectivity.

www.nokia.com



June 23rd | 3PM BST

Networks for the AI Era

Join this free webinar to learn how AI is transforming IoT connectivity, and what operators and enterprises must do to stay ahead

- ✓ Understand how AI is reshaping the traditional IoT network model as inference and training workloads distribute across edge, on-device, & cloud
- ✓ Learn regulation & the threat of non-compliance have disrupted traditional connectivity strategies
- ✓ Discover the key capabilities that enterprises need from their connectivity service providers
- ✓ Best-practice learnings from floLIVE

Register Now →




1P



HIGH SCORE

1234567890

2P



CREDIT 01

Autonomous networks go up another level

TM Forum's Level 4 target architecture for autonomous networks to achieve self-healing, self-optimising networks with zero touch presents a mountain to climb for telcos. Most right now are working hard to achieve Level 2 and Level 3 autonomous operations in parts of their infrastructure and face substantial challenges in moving their siloed architectures into the autonomous era. Key challenges include avoiding replication of those siloes in the new era and enabling autonomous operations to co-exist with partially-automated or manual processes, writes George Malim, the managing editor of Agile Telco.

TM Forum defines autonomous network maturity Level 4 as describing a highly autonomous network. The Forum says that "in a more complicated cross-domain environment, the system enables decision-making based on predictive analysis or closed-loop management of service-driven and customer experience-driven networks via AI modelling and continuous learning". That's one step below Level 5 – the fully autonomous network – and significant steps past partial, Level 2, and conditional, Level 3 autonomous networks. Arguably, Level 4 presents a sweetspot of achievable technological transformation balanced with operational benefits. Even so, the complexity and challenges involved in achieving Level 4 are far from trivial.

This is illustrated in TM Forum's own research published in its regional guide to autonomous networks in June 2025. The research uncovered that just 4% of respondents had achieved Level 4 last year, while mainstream adoption will happen at some point between now and 2030. This demonstrates growing appetite for reaching Level 4 but suggests that telcos are taking a longer-term trajectory to get there.

"The industry is making real progress, but we should be realistic about where most CSPs are today," explains John Harrington, the executive vice president and head of Europe at Nokia. "In 2024, our Bell Labs division noted that the majority of CSPs are still understood to be between Level 1 and Level 2, with a more recent report from the TM Forum stating that most CSPs are aiming for Level 4 autonomy by 2030."

No path to nirvana

Ahmed Abdelaziz, the vice president of automation and transformation at Rakuten Symphony, sees most telcos currently focused on Level 3 autonomous networks. "Most companies are now trying to crack Level 3 and move towards Level 4. Getting to Level 4 isn't like an easy guided navigation where you reach the end of the road. You need to do a lot of housekeeping. You need a lot of systems talking to each other and a foundational layer that in many telcos has been either forgotten or people haven't found worth the investment to unify – for example, your data or your inventory layer."

Ahmed Abdelaziz
Rakuten Symphony

“The good thing about reaching Level 4 is that you don’t reach a nirvana”



Real world Level 4 adoption

Dell’Oro Group reports that while some advanced domains, notably radio access network (RAN) optimisation are approaching Level 3, the overall industry remains in the Level 1-2 range. Automation is largely domain-specific and rule-based. However, Level 4 adoption is no longer theoretical and progress is starting to be seen at a few innovative telcos. These include:

- Rakuten Mobile has proven Level 4 at scale in a live RAN with TM Forum validation. Based on production traffic, the telco is able to realize 20% RAN energy savings using AI-driven closed-loop control with no impact on the customer experience.
- TDC Net and Ericsson have achieved TM Forum Level 4 autonomy certifications for a live RAN deployment. The validation focused on Ericsson’s PCEM software which reduced the energy required to transmit 1GB of data by approximately 5% under live network conditions.
- Nokia and STC have also demonstrated Level 4 autonomy in live RAN operations. During the Hajj period, in spite of traffic increasing by 40%, the network executed 10,000 autonomous operations per hour, helping to improve downlink throughput by approximately 10%.
- In China, China Mobile has reported Level 4 progress, receiving certifications across domains including service assurance, wireless energy optimisation and IP fault management. China Telecom has applied automation in high value domain such as optimisation, assurance and energy savings and is an emerging Level 4 RAN adopter, reports Dell’Oro Group. China Unicom is also demonstrating targeted L4 RAN automation in traffic optimisation, automated assurance and energy efficiency.

“The good thing about reaching Level 4 is that you don’t reach a nirvana,” he adds. “It’s not a goal you reach one time and that’s it. You reach it in specific use cases, in high value scenarios. The reason people like the TM Forum way is because it lists high value scenarios where you reach Level 4, and that makes you able to judge the benefit of each use case and its monetary value.”

Based on TM Forum’s definition, which AN level has your company achieved and what are your goals (all respondents)?

2025					
LO	L1	L2	L3	L4	
12%	36%	31%	17%	4%	
2026					
LO	L1	L2	L3	L4	
0%	6%	25%	46%	23%	
2030					
LO	L1	L2	L3	L4	
0%	0%	5%	10%	85%	

TM Forum

Figure 1: Telcos’ views on their achieved autonomous network levels

Recognising that moves to Level 4 will be piecemeal is essential and demands mechanisms to manage all networks simultaneously. That also involves taking in technological innovations that are emerging along the way, affecting the scope of existing projects. What is clear is that most telcos will need autonomous systems that form part of a wider network fabric that combines traditional infrastructure, Level 2 and Level 3 systems and a growing estate of autonomous Level 4, and later higher, systems that make agentic AI for network operations a reality.

Built-in structural integrity

Ultimately, previous generation systems will retire but today these need to operate efficiently alongside automation so it’s imperative that autonomous networks can see inside non-autonomous or partly-autonomous infrastructure and that the legacy can understand what the autonomous infrastructure is doing. That visibility is essential to enable holistic, optimised networks as the pace of transformation accelerates over the next decade.



John Harrington
Nokia

“We believe that autonomous networks must evolve into AI-native, cross-domain architectures that can support AI as a structural workload”

“We believe that autonomous networks must evolve into AI-native, cross-domain architectures that can support AI as a structural workload, not simply as another application running on top of the network,” says Harrington. “As AI creates more machine-to-machine traffic and more real-time demands, networks will need to work together across RAN, core, IP, optical and cloud. This means moving away from separate automation tools in different parts of the network to an integrated approach that is trusted, secure and able to support AI at scale. Transparency, traceability, security and safety controls must be built in from the beginning, so networks can deliver intelligence securely, predictably and with the right performance.”

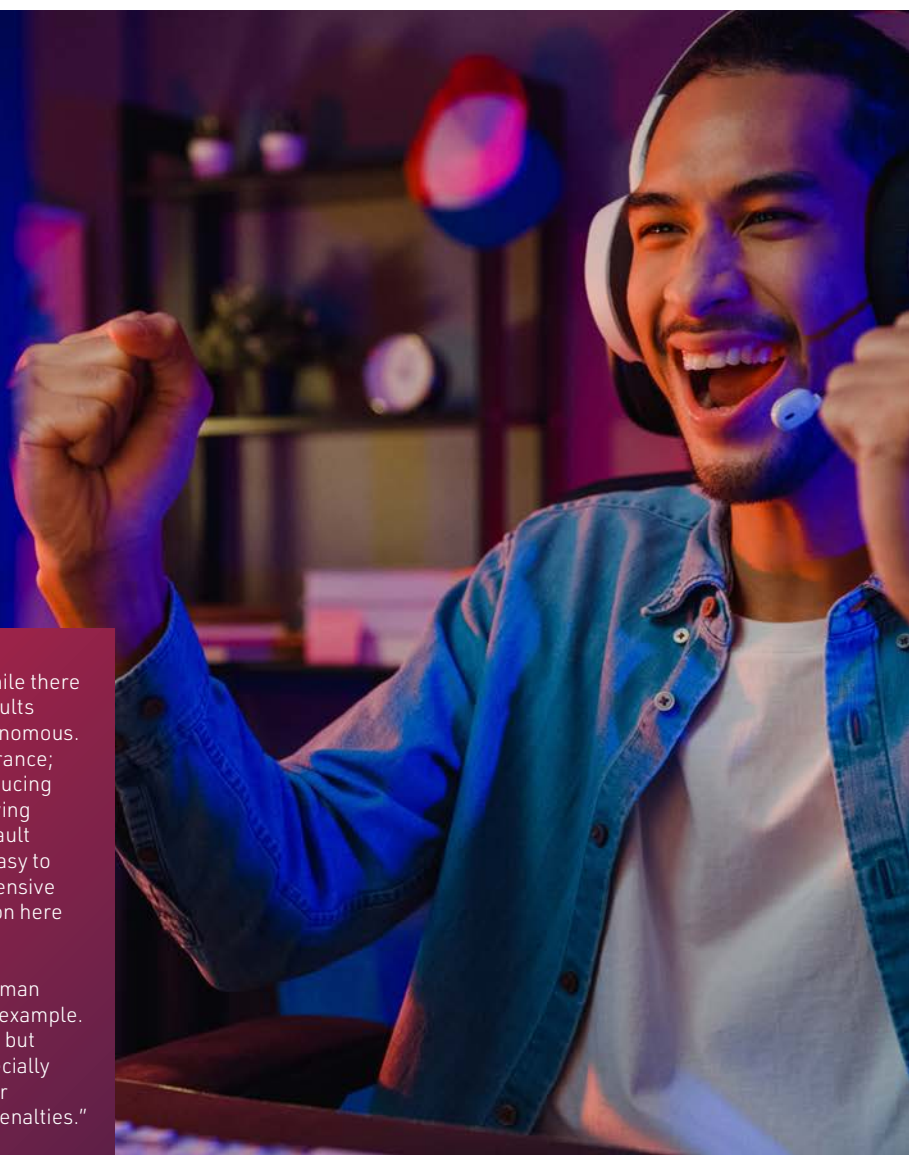
Harrington identifies three essential ingredients for telcos to accelerate their journeys to Level 4. “Three things need to happen,” he confirms. “First, CSPs need trusted, high-quality, cross-domain data rather than duplicated data trapped in operational silos. Second, they need intelligent automation that can correlate incidents, test corrective actions and learn from outcomes. Third, they need open, cloud-native, multi-vendor architectures so Level 4 capabilities can co-exist with Level 2/3 and manual processes during the transition.”

Take it to the next level

Growing maturity means an increased need to tie the value Level 4 autonomous networks can deliver to the cost of doing so. “Everybody talks about being autonomous and about AI until they hit the token maxing,” explains Abdelaziz. “That’s a new term in the market – where you’re so hyped about doing something with autonomy and AI that you overdo it and max out your tokens and cost your company. There’s an enterprise company that allowed its employees to run AI unlimited, and they had a bill of US\$500 million in one month without anywhere near that level of benefit.”

“Running AI and autonomy without a clear goal and endpoint is mostly going to be non-beneficial in terms of cost and outcome,” he says. “The industry is largely ready for Level 3 in many areas and aspiring for Level 4, but not uniformly – some high value scenarios are more important than others.”





Harrington at Nokia goes further, suggesting that while there are clearly valuable use cases in finding and fixing faults more quickly, other areas should never become autonomous. "The main use cases include improving service assurance; optimising the RAN; automating the mobile core; reducing energy use; speeding up service delivery and improving network planning," he says. "However, of all these, fault management is the top automation priority and it's easy to see why. Domains like the RAN are complex and expensive to maintain. Automating fault detection and resolution here brings clear operational benefits."

"By contrast, some areas may always benefit from human oversight," he adds. "Security, for instance, is a good example. Autonomous capabilities in security can be powerful, but judgement, context and governance still matter, especially when potential threats risk exposing customer data or breaching regulations with corresponding financial penalties."

How to pick the winners

Even though benefits are accruing from Level 3 activities and achieving Level 4 autonomy is undoubtedly more complex, the effort should be worthwhile. "You can already see value at Level 3 but Level 4 is important because in operational use cases, the actual benefit is reducing your opex, reducing your time to resolve and having a tangible outcome of reducing the number of tickets and issues in the network," explains Abdelaziz. "Level 4 is about the ability for the machine to handle the trigger, the diagnosis, the remediation, the execution and the feedback loop. If you're able to do that, you can practically handle about 90% or more of your issues in a very short time without requiring manual analysis, and you can save money on opex."

"There are other use cases - around customer support and new customer services - where Level 4 allows you to add services with a shorter time to market and handle customer issues much faster," he adds. "Every one of these has its own monetary value and tangible business outcome."

Harrington shares the view that Level 4 is not a destination and not the only solution for achieving autonomous operations. "Level 4 is not about making every part of the network fully autonomous straight away," he emphasises. "The more realistic approach is to use Level 4 where it will make the biggest difference. For example, in service assurance, Level 4 can help the network spot issues, understand the cause and take action much more quickly than a manual process but in some other areas, Level 2 or Level 3 may be good enough."

"That is why Level 4 can be seen as a practical sweet spot," Harrington says. "It is advanced enough to deliver real benefits, such as faster fault resolution, lower operating costs, better customer experience and less manual work, but it is still achievable in the medium term because telcos can apply it first to the most valuable parts of the network rather than trying to automate everything at once."

What comes first - autonomous networks or AI workloads?

Adoption of autonomous networks to at least Level 4 appears to be a foundational enabler for profitable AI-oriented telecoms services. Yet AI itself is an enabler of Level 4 autonomy so, to an extent, the two are parallel innovations.

"AI and autonomous networks support each other," explains Harrington. "AI helps the network become more autonomous, because it gives the network the ability to spot problems, make decisions and improve performance without waiting for manual intervention. AI is therefore not just another service running on the network. It is one of the main technologies that helps the network configure itself, optimise itself and fix itself. This becomes especially important at higher levels of autonomy, where the network has to deal with more complexity and continuously make changes."

"At the same time, many new services will need more autonomous networks to work properly," he adds. "Services such as 5G slicing, private networks, network-as-a-service,



AI-RAN and edge/cloud services all need a quick set-up, strong service assurance and flexible use of network resources. Manual processes are often too slow for this."

Abdelaziz also sees the need to zero-in on specific AI capabilities and needs from the network. "AI is a huge factor for Level 4," he acknowledges. "Reaching Level 4 without AI is literally like having a car go to work without a tyre. AI is the moving engine of the decision, the diagnosis and the recovery. There are different parts of AI though. If we're talking about agentic frameworks, using agents becomes an architectural choice. Some companies might want to use inference models and training close to their data, running ML and AI models close to their data. Other companies will prefer to just let agents handle every step of the process."

"Either way, AI is the wheel that moves the decision," he says. "When we talk about autonomous, it means the system can self-diagnose and self-execute decisions - and you need that intelligence to mimic human behaviour. That is what AI does best."

Does AI depend on Level 4 autonomous networks?

AI and Level 4 autonomous networks are not inextricably linked. "Level 4 is not needed for every AI use case," points out Harrington. "For example, a simple AI assistant that helps a CSP's operations team answer questions could work at a lower level of autonomy but more advanced agentic AI use cases, such as where AI recommends an action, checks it, carries it out and monitors whether it worked, need the kind of data, visibility, controls and closed-loop automation that come with Level 4. Level 4 is not the starting point for AI, but it becomes much more important as AI use cases become real-time, cross-domain and customer-facing."

That importance is morphing into pressure on telcos growing. "Yes, moving towards Level 4 is becoming urgent for telcos,

but that does not mean every part of the network has to reach Level 4 at the same time," says Harrington. "The pressure is mainly coming from rising costs, the growing complexity in how networks are managed, changing customer expectations, energy use and competition. As networks become more complex, manual operations are becoming harder to manage. Telcos are introducing new services, new technologies and more cloud-based systems, but they still need to keep costs under control and deliver a good customer experience. More automation is needed to do that."

Harrington cites research from STL Partners which estimates that AI, automation and analytics could save an average CSP around US\$650 million a year in capex and opex and create another US\$144 million a year in extra revenue from faster and more profitable new services. "Pressure also comes from scale," he adds. "For example, creating a 5G network slice manually across 10,000 gNodeBs could involve more than 100,000 managed objects and updates to several hundred thousand radio parameters. That could take hundreds of hours of work. With automation, it can be reduced to minutes or hours."

Economic pressure drives adoption

For Abdelaziz, the biggest pressure is the economics of the market. "The telco industry is a rare example of an industry where average revenue per user (ARPU) year over year has always been decreasing rather than increasing - you're providing better service and newer technology at lower cost, because connectivity is now a commodity," he explains. "It's no longer something only a wealthy person can access. The business pressure is always there: reducing opex, reducing capex, dealing with fierce competition and churn, having shorter time to market for features, and providing something substantially different to the customer at lower cost. There is already business pressure driving you towards autonomy."

That pressure should not allow the benefits of autonomous networks to be skipped over in haste. "The key is to make sure today's investments do not create a new set of silos," says Harrington, pointing out Nokia's Autonomous Network Fabric creates one shared intelligence layer for the network. "Telcos should avoid buying automation tools that only work in one part of the network and cannot easily connect with other systems. Instead, their investments should be open, cloud-native, multi-vendor in approach and API-based. They should also use shared data and intelligence across the network, so different domains, such as RAN, transport, core and cloud, can work together rather than separately."

"Futureproofing also means accepting that full autonomy will take time," he adds. "For many years, Level 4 systems will need to work alongside Level 2, Level 3 and even manual processes. That means telcos should choose technologies that can be introduced step by step, starting with the areas where automation delivers the biggest benefit. We also believe that telcos should design for intent-based operations from the start. In simple terms, this means the operator tells the network the outcome it wants, for example better performance, lower energy use or faster service delivery, and the network decides the best way to achieve it within agreed rules. We believe autonomous networks need to be designed for AI from the start, not treated as an extra layer added to existing networks. The best investments are the ones that will deliver benefits today while also creating the foundations for Level 5 autonomy further down the line."



Real Intelligence.

Understanding how telecom is being redefined by AI, cloud-native software, automation and autonomous networks takes decades of accumulated knowledge.

Turning that into your advantage takes intelligence. The real kind.

Appledore provides the insights needed to make the right choices in times of transformational change and unprecedented potential.

appledoreresearch.com



ANALYSIS | RESEARCH | CONSULTING

The rise of the autonomous telco



The telecoms industry is entering a decisive moment. AI is no longer just improving individual processes, it is exposing the limits of the way telecoms operations have been organised for years, writes Chantel Cary, the head of product marketing at Oracle Communications.

For decades, operators have managed complexity through human coordination, process discipline and incremental automation. That model delivered progress, but it was designed for a slower, more predictable environment. It is no longer enough.

Networks are becoming more software-driven and distributed. Services span cloud, edge and core environments. Customer expectations continue to rise. AI is accelerating how quickly telecoms organisations are expected to detect issues, understand impact, make decisions and act.

The gap is becoming clear. The operating structures many communications service providers still rely on were not built for the speed, complexity and cross-domain coordination the industry now requires.

Operations teams still monitor environments, analyse alarms, escalate issues, correlate data across systems and coordinate remediation manually across network, service, customer and business domains. OSS and BSS modernisation has improved parts of this process but the core structure remains fragmented.

Data still sits in silos. Workflows still stop at system boundaries. Decisions still depend on too much manual intervention.

That fragmentation has become a strategic liability. It slows response times, limits visibility, increases cost-to-serve and makes it harder to deliver consistent customer outcomes. As AI becomes more embedded across operations, customer engagement and network systems, it is pushing the industry toward more autonomous ways of operating.

This is not simply about adding more automation to existing operations, and it is not another network automation initiative under a new name.

The idea of an autonomous telco represents a broader shift in how the service provider operates. It points to a model where AI, operational systems, workflows, network intelligence and business processes work together continuously to improve outcomes with progressively lower levels of manual intervention.

That shift extends well beyond the network. It changes how service providers think about OSS and BSS architectures, customer engagement, service assurance, workforce coordination, revenue management and the systems that connect them.

SPONSORED ARTICLE

The layers of a more autonomous telco


A more autonomous telco is not a single capability or a single domain. It is a connected operating model built across four layers: the network, operations, business functions and customer experience. Each layer matters, but none can deliver the full autonomous telco vision on its own.

The first layer is the autonomous network. This is where much of the industry's work has already started, through closed-loop control, intent-driven operations and more self-optimising network domains. Networks need to adapt to traffic patterns, demand fluctuations, service requirements and performance conditions with less manual intervention. The network is also becoming more than infrastructure. It is becoming a source of real-time intelligence. Voice and video interactions, for example, carry intent, sentiment, identity signals and decision context that can inform customer experience, service assurance and operational workflows when they are governed and connected properly. Yet, many network and communications environments remain fragmented across domains, vendors, legacy and cloud-native technologies, and operational systems. Autonomy in the network matters, but it cannot stop at the network.

The second layer is autonomous operations. This is where service providers move from reactive monitoring and manual remediation towards operations that are more predictive, proactive and increasingly self-managing. The goal is not simply to detect issues faster. It is to identify, diagnose and resolve issues before they impact customers, while automating more of the service lifecycle end-to-end. AI can help detect abnormal conditions earlier, understand the likely root cause, assess customer and service impact, and recommend or trigger the next best action. That requires assurance, orchestration, inventory, field service and customer care to operate as connected closed-loop processes that can detect, predict and act. Without that connection, AI becomes another insight layer rather than a mechanism for execution.

The third layer is autonomous business. This layer is often underemphasised in telecoms conversations, but it is critical to whether the operating model can scale. Workforce management, supply chain, finance and revenue management all shape how efficiently a service provider can run, invest and serve customers. AI creates the opportunity to make these functions more adaptive: aligning capital plans to network demand, anticipating supply constraints before they affect fulfillment, improving workforce allocation for installation and repair, and making charging, billing and revenue operations more responsive to new pricing and partner models. The benefit is not just back-office efficiency. It is the ability to connect commercial, financial and operational decisions to what is happening in the network and with the customer in real-time.

The fourth layer is autonomous customer experience, spanning marketing, sales and customer service. This is where engagement shifts from reactive, campaign-based and siloed interactions to a more continuous model across the customer lifecycle. Marketing shifts from periodic campaigns to always-on engagement that adapts as customer behaviour changes. Sales moves beyond static segmentation and manual offer construction towards more dynamic, guided journeys that can assemble the right offer, context and next



action in the moment. Customer service moves from resolving issues after they happen to anticipating needs, guiding resolution and closing the loop back into the next interaction. The experience may feel seamless to the customer, but making it work requires shared context across customer data, product catalogue, configure-price-quote, service qualification, network and product inventory, field service, provisioning, billing and care.

This is why the autonomous telco concept matters now. These layers cannot develop as separate transformation tracks. A service issue, sales order, network constraint, billing event or customer interaction can cut across all of them. Automating each layer independently will not be enough. Service providers need to connect intelligence across them so the business can operate with greater speed, adaptability and coordination.

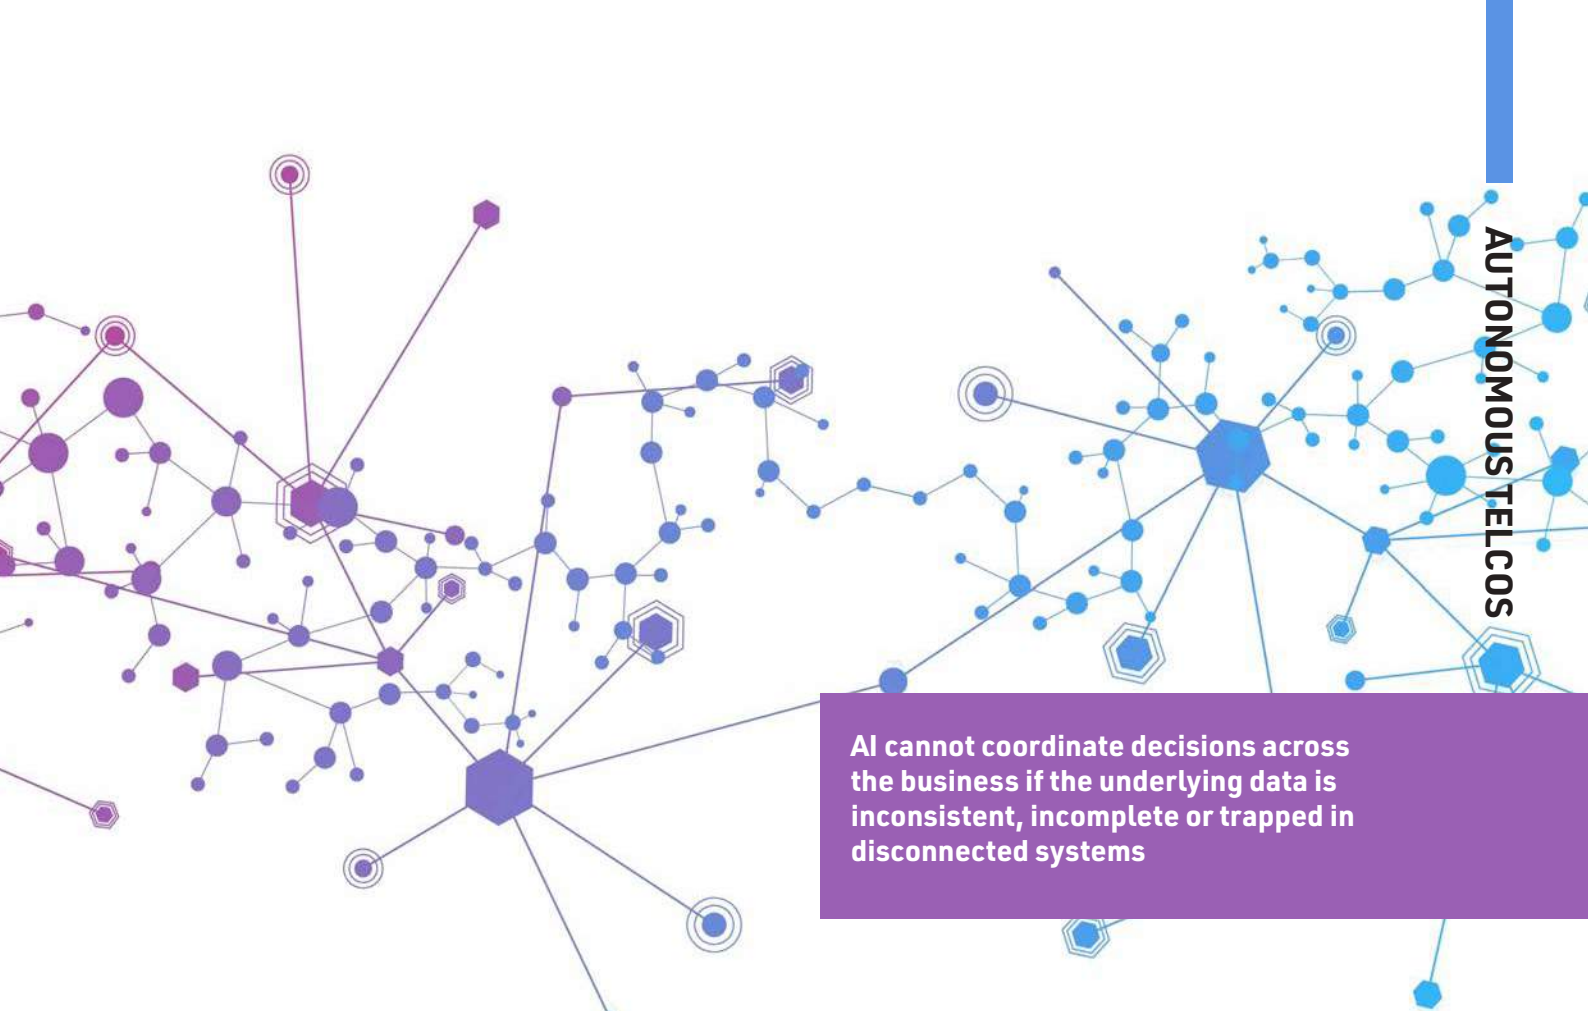
What has to change

AI is often discussed in terms of productivity and automation. For telecoms, the more important shift is operational.

AI exposes how disconnected telecoms environments still are. A model that detects an issue is useful, but it is not enough. It also needs to understand customer impact, service dependencies, operational policy, available resources and the actions that can be taken across the business.

That changes the architecture underneath. A more autonomous operating model requires environments that are more modular, interoperable and event-driven than most operators run today. AI agents, orchestration systems, OSS, BSS and operational workflows all need to exchange context and coordinate actions across domains.

This increases the importance of standardised interfaces and open operational frameworks. TM Forum Open APIs provide a common way for systems and workflows to exchange operational context across the stack, while Open Digital Architecture creates a more composable foundation for connecting network, customer and business systems.



AI cannot coordinate decisions across the business if the underlying data is inconsistent, incomplete or trapped in disconnected systems

It also increases interest in emerging coordination models such as model context protocol (MCP) servers and agent-to-agent (A2A) communication frameworks. As AI agents become more deeply embedded into telecoms operations, operators will need architectures that can securely share context, coordinate actions and govern interactions across increasingly distributed environments.

This changes the role of OSS and BSS.

For years, modernisation focused on integration, cloud migration and replacing legacy infrastructure. A more autonomous operating model changes what OSS and BSS need to do. They become coordination layers that connect network operations, customer experience, service orchestration and business functions such as revenue management.

That matters because telecoms operations no longer break cleanly between technical and business systems. AI makes those connections harder to ignore. OSS and BSS need to help connect the data, workflows and decisions required to move from order capture through to fulfillment and service delivery.

The future operating model is not about adding AI on top of existing systems. It is about building environments where intelligence can move continuously across networks, services, customer operations and business systems.

How service providers can move towards it

The path to a more autonomous operating model will not come from a single platform decision, a single AI initiative or a single network automation programme. It requires a phased, pragmatic

approach that helps service providers modernise while maintaining control over critical data, models and regulated operational environments.

The first step is establishing a unified data foundation. AI cannot coordinate decisions across the business if the underlying data is inconsistent, incomplete or trapped in disconnected systems. This is one area where Oracle is focused: helping CSPs make data across network, OSS, BSS and IT systems consistent, governed, accessible and available to AI in real-time, so intelligence can move from analysis into action. Without that foundation, AI remains limited to isolated recommendations rather than operational execution.

The second step is deploying AI against targeted, high-impact use cases. Operators do not need to automate everything at once. They should start where AI can deliver immediate value and where the business case is clear, such as customer engagement, network optimisation, operational efficiency or sales order fulfillment. These are the places where fragmented workflows create visible cost, delay and customer impact.

The third step is scaling through AI agents embedded directly into workflows. This is where many initiatives fall short. AI models may identify a risk, recommend an action or generate an insight, but the value is limited if that output still depends on manual interpretation and handoff. From there, intelligence has to be embedded into the operational and business workflows that run the service provider, so AI can move closer to execution.

The fourth step is orchestrating end-to-end autonomous processes across domains. A more autonomous operating model depends on closed-loop systems that can detect, predict and act. That requires AI-driven capabilities to connect network, operations, customer experience and business functions so decisions and execution can

move across the enterprise rather than stopping at organisational or system boundaries.

This also changes the way people work. Instead of moving through static applications, comparing data across screens and executing each step manually, users will increasingly define the outcome they want to achieve and the policies that should guide it. AI agents can then help determine the right actions and coordinate them across systems and domains. Human expertise remains essential, but it shifts toward higher-value work: setting direction, governing automation, managing exceptions and improving how work gets done.

For most CSPs, the transition will be incremental. Hybrid operations will remain the reality for years. The providers that make progress fastest will be the ones that stop treating AI as a set of experiments and start operationalising it inside the systems and workflows that run the business.

Why this becomes a competitive issue

The move toward a more autonomous operating model is shifting from long-term ambition to near-term operational priority. Delaying the shift will only widen the gap between what current operating models can support and what the market will require.


Service providers can no longer compete on network quality or price alone. They also need to compete on experience, cost efficiency, speed and the ability to maximise the value of network and IT assets. That requires more than isolated automation. It requires the ability to turn fragmented data into intelligent workflows across the business. Oracle's role is to help make that connection practical: bringing network and telecoms data, applications, AI and cloud infrastructure closer to the workflows where decisions are made and work gets done.

The providers that move fastest will connect intelligence across network, operations, customer experience and business systems instead of treating them as separate worlds. Competitiveness will increasingly come down to operational adaptability.


The shift will not happen all at once. The direction is clear. The next phase of transformation cannot be about better automation alone. It is about building communications environments that can continuously understand conditions, coordinate decisions and improve outcomes in real-time.

That is the real test for the industry. AI will not deliver its full value if it is applied to operating models that remain fragmented, manual and slow to adapt. The service providers that lead will be the ones that move beyond AI as a collection of tools and make it part of how the business runs. That is the real promise of the autonomous telco: not a network that automates more tasks, but a service provider that can understand, decide and act in real-time.

www.oracle.com



Service providers can no longer compete on network quality or price alone. They also need to compete on experience, cost efficiency, speed and the ability to maximise the value of network and IT assets



TECHLED Agile Telco

OPTIMISED OPERATIONS FOR
MODERN CONNECTIVITY PROVIDERS

WWW.THEAGILETELCO.COM



Martin Rueckert
Tallence

From copilot to crew: How multi-agent AI is replacing the telco BSS middle layer

Most telcos have spent the last three years adding AI to their BSS stack the same way you add a car sticker: it looks like something changed, but underneath everything is the same. A chatbot in front of the IVR. A churn prediction model that outputs a score into a dashboard nobody checks. An anomaly detection layer that fires alerts into a ticket queue already full of alerts. Point solutions, layered on top of systems that were never designed to be operated by machines.

This approach has a ceiling. It is becoming visible now. The next architectural shift puts autonomous agent crews at the centre, crews that own BSS processes end-to-end, from insight to action, across system boundaries, writes Martin Rueckert, the chief AI officer at Tallence.



SPONSORED ARTICLE

In an agent crew, the process logic lives in a deterministic decision layer

The typical BSS environment at a mid-size European telco looks roughly like this: an order management system that dates back a decade, a billing platform that has been customised into something unrecognisable, a CRM that talks to both via a middleware layer that one person on the team truly understands, and somewhere between twelve and thirty integration touchpoints that were each added to solve a specific problem at a specific moment in time.

Why point AI hits a wall

When you drop a copilot or a narrow ML model into this environment, it can optimise within its own slice. The churn model fires. The recommendation engine suggests a retention offer. But getting from that insight to an actual action, such as updating the CRM, triggering a workflow in the order system or logging the outcome for the next model to learn from, requires a human to carry the signal across system boundaries. Or it requires yet another integration project.

This is the fundamental problem: AI at the task level does not solve the coordination problem at the process level. You still need something to orchestrate.

In traditional BSS architecture, that something was the middleware and the integration platform. When it worked, it was brittle. When it broke, it was expensive. And it was built for a world where processes are defined upfront, predictable and stable which describes neither telco operations today nor the speed at which the business needs to adapt.

What an agent crew does differently

The architecture that is starting to replace this model centres on autonomous agents that are given a goal and the tools to pursue it. The agent determines the sequence of steps at runtime, based on the current state of the task.

In a multi-agent framework, each agent has a defined role, a set of tools it can call and the ability to reason about what to do next. Agents can call APIs, read and write to systems of record, delegate subtasks to specialised sub-agents, and hand off to a human when they hit a decision, they are not qualified to make. They maintain state across steps and across system boundaries.

In an agent crew, the process logic lives in a deterministic decision layer, while LLM-based agents gather structure and context. Reasoning about context is what language models do well. Deciding the next action under audit is something else.

A crew of such agents, coordinated by an orchestration layer, can own a process the way a small team of specialists would. One agent handles the CRM interaction, another the provisioning system, a third monitors the outcome and triggers a follow-up if conditions change. The orchestrator decides sequencing, handles failures and routes work to the right specialist.

The open-source agent frameworks have matured significantly over the past 18 months. They are useful for prototyping. Production-grade carrier deployments require an additional layer of deterministic decision logic and standards-aligned integration, which is where the engineering work sits today.

Three BSS processes worth looking at first

Not every BSS process is equally suited for early agent adoption. The ones that benefit most share a common profile: they are high-frequency, they cross multiple system boundaries, and they currently require significant human coordination to complete.

Order fallout management fits this profile almost perfectly. When a service order fails in provisioning, the current response in most telcos is a ticket that lands in a fallout queue, where a specialist diagnoses the root cause,

If an agent crew can own the process layer on top of existing systems, the urgency of replacing the system underneath decreases

determines whether it requires a system fix, a configuration change, or a manual override, and then works through the resolution. This process takes hours on average. A significant portion of fallout cases are repetitive and involve the same provisioning errors and the same resolution steps. An agent crew trained on historical fallout data, with tool access to the provisioning system, the inventory layer and the order management system, can resolve a high proportion of these cases without human intervention. The ones outside its confidence boundary get escalated with a diagnostic summary already prepared.

Credit management and collection workflows are another clear candidate. The decision logic here is complex and covers which customers get a reminder, which get a suspension warning, which get an exception based on their tenure and payment history, but it is documentable and teachable. More importantly, it involves exactly the kind of cross-system coordination that today requires a human to navigate: checking the CRM for relationship context, the billing system for history, the dunning system for open actions, and then deciding what to do next. An agent crew can run this end-to-end across thousands of accounts simultaneously, with consistent application of the decision logic and a complete audit trail.

Network-triggered customer actions sit at the intersection of network operations and BSS. When a network event degrades service quality for a customer segment, the current process in most telcos involves a handoff from NOC to commercial operations, where someone decides whether to send a proactive communication, apply a credit, or escalate to a field team. Each handoff takes time, and by the time the communication reaches the customer, the incident is often already resolved. An agent crew that monitors network event feeds and has the authority to trigger customer communications and apply service credits within defined parameters can compress this loop significantly.



The BSS investment question this changes

The implication for BSS investment strategy is easy to miss, and worth examining carefully.

Most large-scale BSS modernisation programmes are justified on the basis that the current system is too rigid to support new business requirements, and that a modern platform will be more flexible and easier to change. This is true. But the timeline and cost of a full BSS replacement programme for a mid-size telco are large enough that the programme is a bet, not a plan.

The agent crew architecture changes the calculus. If an agent crew can own the process layer on top of existing systems, the urgency of replacing the system underneath decreases. You can deploy the agent layer now, on top of the systems you have, and get most of the operational benefit today.

Modern platforms are genuinely better in many dimensions across data models, API design and cloud-native architecture. The decision to replace them can be made on technical merit and platform capability, at a pace the organisation can absorb. The “we cannot operate the business without this” argument for accelerating a replacement programme weakens when the agent crew is handling the process layer adequately.

For telcos currently in the planning phase of a BSS transformation, this is worth taking seriously. An agent-first approach may deliver 70% of the business value at 20% of the programme’s cost and timeline. The remaining 30%, the cases where you genuinely need a modern data model, or where the legacy system’s limitations constrain what the agent can do, become smaller, better-scoped programmes.

The operations teams that currently manage BSS processes are skilled at something specific

pressure to expand the scope quickly, because the results are good and the cases seem analogous. Expanding scope without re-evaluating the risk boundaries and adding appropriate guardrails is where things go wrong at scale.

The third is observability. A multi-agent process that spans six system calls and three agents is harder to audit than a workflow in an integration platform with structured logging at each step. Getting good visibility into what the agents actually did and why they did it requires investment in the observability layer that is easy to underestimate.

All three are manageable with deliberate engineering. Treating the deployment as a serious production programme, with proper scope, guardrails and observability from the start, is what separates successful rollouts from failed ones.

The transition is already happening

The telcos that have moved furthest in this direction built operational agility in a new layer, without waiting for a BSS replacement programme to get there first.

The architecture is sound. The tooling is available. The economics are compelling. What it requires is a team that understands both the BSS domain and the agent engineering discipline, which is still uncommon, though the talent pool is growing faster than most people expected.

For telcos running a brownfield BSS environment and looking for the most direct path to operational agility, the practical question is how to scope the first deployment so that it delivers credible results without creating new risks that undermine confidence in the approach.

That is a solvable problem, and it is a considerably more tractable one than a full BSS replacement programme.

predictable steps across systems and starts spending most of its time on the cases the agent flags as outside its confidence boundary, on refining the agent's behaviour based on what it observes, and on expanding the agent's scope to new process areas.

This is a different job profile. It requires people who can evaluate agent behaviour, write and refine agent instructions, and diagnose failures in a reasoning system. Telcos that are serious about this transition will need to invest in building that capability in their operations teams as much as in the technology itself.

Where things break

Multi-agent architectures also introduce failure modes that BSS teams are not accustomed to managing.

The most common failure mode is hallucination in the decision path. Mitigation is partially possible through structured tool outputs and human-in-the-loop gates, but the more durable answer is architectural: keep language models in the context layer and put decisions on deterministic controllers with explicit rule sets. That separation is what makes the system auditable under frameworks such as the EU AI Act.

The second failure mode is scope creep in agent autonomy. Teams that start with a conservatively scoped agent deployment, such as "only resolve these five fallout types", and face

What the ops team becomes

There is a personnel dimension to this that tends to get avoided in architecture discussions.

The operations teams that currently manage BSS processes are skilled at something specific: navigating system complexity, knowing which system to check first or understanding the undocumented business logic that lives in nobody's head except the people who have been there for ten years. When an agent crew takes over the repetitive process execution, this knowledge gets encoded into the agent's tools, its decision logic and its escalation criteria.

What changes is the ratio of execution work to oversight and exception-handling work. The team stops spending most of its time executing

Telecoms industry makes big moves to get to grips with agentic AI

Communications service providers are continuing their digital transformation journeys by trying to get to grips with agentic AI. Antony Savvas looks at the key areas they are focusing on, including AI tooling for cloud migration, telco-specific agentic care solutions and network operations, enabling service providers to deliver smarter, more intuitive customer experiences.

Much of the activity involves finding key technology providers and striking important alliances. Customer delivery and experience in the AI space is now front of mind and T Wholesale, the wholesale arm of Deutsche Telekom, has completed a major business support system (BSS) transformation in partnership with Enxoo to deliver it.

Enxoo is a Salesforce consulting and implementation partner for the communications industry. Built on Salesforce's Agentforce Communications, and powered by Enxoo for Wholesale, T Wholesale's new agentic AI-based platform unifies its global sales and delivery processes, enabling faster customer response, streamlined operations and accelerated service launches across its international footprint.

"This transformation is a significant step in our ongoing journey to simplify and digitise our global operations," says Silke Hoesch, the senior vice president of digital, data and billing at T Wholesale. "By unifying our BSS, we're giving teams the agility and transparency needed to deliver exceptional service to customers and partners worldwide."

The two-year programme reimaged T Wholesale's international lead-to-cash process, from quoting and ordering to delivery and billing, enhancing automation with features like attribute-based pricing, serviceability checks and instant quote generation.

"This project shows what's possible when carriers embrace true digital transformation," says Krzysztof

Zych, the chief executive of Enxoo. "By combining Agentforce Communications Cloud with Enxoo for Wholesale, T Wholesale now has the agility to innovate and scale at the speed the market demands."

The new platform establishes a future-ready foundation aligned with TM Forum open APIs and the MEF's Lifecycle Service Orchestration (LSO) Sonata standards, delivering measurable gains in speed, scalability and customer experience, and setting a new benchmark for digital agility in global wholesale operations, says T Wholesale.

LSO Sonata standards, developed by the Metro Ethernet Forum (MEF), define the automated, business-layer APIs and processes between different network service providers. They allow communication between buyers and sellers to streamline commercial and operational functions, like quoting, ordering and inventory management.

Sovereign AI

In another development, network infrastructure provider Cornerstone, in partnership with Stonesthro, a UK-based edge cloud provider, and Zadara, a global player in fully managed cloud services, is exploring new ways of bringing compute and data closer to where they are created and used.

The partnership aims to provide solutions for national challenges, from improving data sovereignty and security to supporting low-latency AI applications and strengthening the

UK's digital backbone. By combining Cornerstone's national infrastructure, Stonesthro's sovereign edge capability and Zadara's cloud technology, the partnership will examine how edge computing can unlock new performance, sustainability and resilience benefits across industries. As demand for data-intensive and AI-driven services grows, organisations across sectors, including the public sector, healthcare, finance and critical national infrastructure, face mounting pressure to balance latency, sovereignty, scalability and sustainability.

Pat Coxen, the chief executive of Cornerstone, says: "This partnership brings together three complementary strengths: national infrastructure, digital innovation and artificial intelligence, to explore how the UK can build greater trust and resilience into its digital systems. It's about finding practical solutions to some of the UK's most pressing connectivity and compute challenges."

Ashley Sellar, the chief executive of Stonesthro, adds: "By combining the flexibility of edge with the assurance of UK sovereignty, we have an opportunity to redefine how organisations approach performance, compliance and sustainability in the cloud."

Yoram Novick, the chief executive of Zadara, explains: "By expanding access to Zadara-powered sovereign AI clouds, we are enabling UK organisations to build scalable and secure IT foundations that unlock the power of AI, while ensuring resilience and digital sovereignty."

Silke Hoesch
T Wholesale



Pat Coxen
Cornerstone



A more efficient edge

Multiverse Computing, a specialist in AI model compression, has released the LittleLamb open-source model family, to make operations easier at the edge. Designed for real-world AI deployment in a smaller footprint, the three models available are LittleLamb 0.3B, a general-purpose model, LittleLamb 0.3B Tool-Calling, a specialised variant optimised for tool use and agentic workflows, and LittleLamb 0.3B Mobile, a deployment-focused variant built for on-device and mobile applications.

Using Multiverse’s CompactifAI technology, each model reduces the base architecture by around 50%, enabling efficient deployment across edge, mobile, and offline environments. The compressed models improve system throughput, latency, output speed, and TTFT benchmarks. Time to first token (TTFT) is an LLM performance metric measuring the time it takes for a model to process a user prompt and generate its very first output token. It dictates a system’s initial responsiveness and determines how snappy a chatbot or AI application feels.

“The launch of LittleLamb continues our mission to make efficient AI available across every deployment environment without losing the flexibility and accuracy developers need,” says Enrique Lizaso Olmos, the chief executive of Multiverse Computing. “With CompactifAI, we’ve demonstrated that compression doesn’t require sacrificing intelligence or



Enrique Lizaso Olmos
Multiverse Computing

capability. This model family shows that compact models can do far more than lightweight chat and can run in environments where traditional models are simply too large or too dependent on cloud infrastructure."

Cloud alliances

Amdocs has struck a multi-year strategic collaboration with Amazon Web Services to advance AI-driven modernisation across the telecoms industry. Amdocs is combining its aOS agentic operating system, which orchestrates scalable, multi-agent enterprise workflows, with AWS' AI and cloud technologies. Together, the two companies say they will empower service providers to streamline core operations, enhance customer experiences and drive innovation at scale with measurable business outcomes.

The collaboration is said to be already delivering tangible value, with a large-scale modernisation programme currently underway at a "leading North American Tier-1 communications service provider".

"Every enterprise is sitting on decades of technical debt that stands between them and an AI-ready future, and the cost of inaction has never been higher," says Fabio Cerone, the managing director of the EMEA telco business unit at AWS. "By combining our agentic AI services, built on nearly two decades of migration experience, with Amdocs' deep telecoms expertise, we're giving operators a faster, lower-risk path to cloud transformation, unlocking new sources of innovation and value. What once demanded years of costly manual effort, now takes weeks, fundamentally changing what's possible."

Anthony Goonetilleke, the group president of technology and head of strategy at Amdocs, adds: "We're expanding our collaboration with AWS to bring transformative cloud, data and AI solutions to leading communications service providers around the world, generating measurable business impact. We are combining Amdocs' deep industry expertise and verticalisation in Amdocs aOS with AWS's advanced AI and cloud capabilities to help our customers supercharge the achievement of key business outcomes."

In another important alliance in the consulting space, Tech Mahindra is

collaborating with NVIDIA to introduce an AI-powered Telco Network Operations Reasoning Agent. The solution is designed to help CSPs accelerate their journey toward "Level 4+ (L4+)" autonomous networks, by transforming traditional network operations centres (NOCs) into intelligent, closed-loop operations.

TM Forum Level 4 (L4) autonomous networks represent a major shift from traditional human-defined automation to true AI-driven, intent-based independence. At this level, communications service providers (CSPs) use agentic AI and predictive analytics to self-manage, self-optimize, and self-heal complex network domains without requiring constant human oversight.

The alliance's solution is delivered through Tech Mahindra's Orion platform, combining AI engineering, domain expertise and human-in-the-loop design. In addition, it ensures enterprise-grade trust without using customer or personally identifiable information (PII), allowing CSPs to safely operationalise artificial intelligence, while maintaining data privacy and governance.

Amol Phadke, chief transformation officer at Tech Mahindra, says: "Network operations centres still rely on rule-based, open-loop workflows with significant manual intervention. Engineers continue to spend considerable time correlating alarms, logs and performance data across systems, impacting resolution times and operational efficiency. Our collaboration with NVIDIA addresses this challenge by embedding AI-driven reasoning into the core of network operations. By enabling intelligent, closed-loop execution, we are helping communications service providers move decisively toward simplified operations aligned with L4+ autonomous network ambitions."

The AI reasoning agent within the solution enables CSPs to move beyond conventional automation by embedding contextual intelligence directly into network operations. It supports autonomous alarm validation, root-cause analysis and resolution across operations support systems (OSS) and business support systems (BSS) environments, enabling faster incident handling, greater operational consistency and a better customer experience. A key differentiator of the solution is its extensibility, as once

CSPs deploy a foundational large telco model, they can build and scale domain-specific reasoning agents that integrate with their existing data lakes, tools and workflows.

Using NVIDIA AI Enterprise software, Tech Mahindra customised a reasoning model with synthetic and anonymised data using NVIDIA NeMo, and deployed it as a NVIDIA NIM inference microservice. The model is said to have provided a 2-3x improvement in accuracy, when compared with a non-fine-tuned model.

"Network operations demand rapid decision making across complex, real-time environments," says Chris Penrose, the vice president for global business development for telecoms at NVIDIA. "By combining NVIDIA's AI software stack with Tech Mahindra's deep telecoms expertise, this collaboration enables CSPs to deploy reasoning-based AI systems that can act, adapt and learn within live NOC environments. Together, we are laying the groundwork for a new era of autonomous, resilient and intelligent networks."

AI RANs

Testing firm Keysight Technologies and Samsung Electronics are collaborating to demonstrate an end-to-end AI radio access network (AI-RAN) testing and validation workflow with NVIDIA. Keysight's AI-RAN solution streamlines the validation of AI-driven RAN modules using the physical uplink shared channel (PUSCH) channel estimation use case, with the aim of bringing data generation, AI/machine learning training and performance benchmarking into a single automated workflow to accelerate innovation and reduce deployment risk for operators and vendors.

As AI moves deeper into the RAN, engineers must validate algorithms across diverse real-world network conditions. Data collection, model training and performance benchmarking often occur in separate environments, making results difficult to compare, reproduce and trust prior to deployment, particularly for physical-layer functions, such as channel estimation, that directly impact throughput, reliability and user experience.

To address this, the Keysight AI RAN Simulation Toolset orchestrates an integrated, end-to-end workflow that automates realistic dataset



Anthony Goontilleke
Amdocs



and scenario generation, AI/ML model training and repeatable benchmarking, enabling engineers to compare approaches consistently and gain clear performance insights before field testing. The effort integrates Keysight's AI RAN Simulation Toolset with the NVIDIA Aerial Testbed, an end-to-end, over-the-air AI-RAN research testbed running on platforms including NVIDIA GH200 and NVIDIA DGX Spark, the NVIDIA Aerial Omniverse Digital Twin (a network digital twin), and commercial radio hardware such as Analog Devices' Titan O-RU platform. AI/ML models were developed by Samsung, NVIDIA and Keysight, and are trained and benchmarked within a unified test workflow to deliver actionable performance insights.

Soma Velayutham, the vice president of AI and telecoms at NVIDIA, says: "As AI moves deeper into the radio access network, validating complex algorithms across diverse, real-world conditions becomes even more critical. By using NVIDIA AI Aerial platforms into a unified workflow, Keysight is simplifying the end-to-end data pipeline, essential for

training, validating and deploying AI-native 5G and 6G networks."

Charlie Zhang, the executive vice president of Samsung Research America, adds: "Our partnership with Keysight and NVIDIA is revolutionising network deployment by delivering integrated AI-RAN validation, bridging the gap between theoretical models and real-world implementation. This collaboration empowers the industry to confidently adopt AI-driven RAN solutions, ensuring robust and commercially viable foundations for the future of 6G networks."

Balaji Raghothaman, the chief technologist for 6G at Keysight, explains: "AI in the RAN only delivers value when it can be validated with confidence. Working with NVIDIA and Samsung, we're demonstrating a streamlined, automated workflow that unifies data generation, AI/ML training and benchmarking, helping operators and vendors accelerate deployment, reduce risk and make more informed decisions as they introduce AI-driven RAN capabilities."



Chris Penrose
NVIDIA

What's next for telcos at Autonomous Network L4?

As global digital and intelligent transformation enters a critical phase, the communications industry stands at a key inflection point, shifting from the connectivity era to the AI-native era. With the large-scale commercialisation of 5G-Advanced, accelerated R&D of 6G, and the maturity of new technologies such as AI large models and digital twins, autonomous networks (AN) are rapidly evolving from automation-assisted operations and maintenance (O&M) towards autonomous decision-making. This is more than a tool upgrade, empowering global operators to shift from cost reduction to revenue generation, writes Peng Zheng, the general manager of Service and Data Intelligence at ZTE.

According to TM Forum, more than 60% of the world's leading operators have deployed TM Forum Autonomous Networks Project Level 3 (L3) autonomous capabilities. Pioneers including China Mobile, Deutsche Telekom, Vodafone and China Telecom have achieved L4 closed-loop in multiple high-value scenarios, signifying that autonomous networks have moved from concepts to large-scale deployment. By

the end of 2026, both the proportion of operators meeting L4 deployment criteria and the global penetration rate of L4 scenarios will climb further. However, the excessively long static payback period (SPP) for L4 has become a major bottleneck hindering large-scale rollout. Reducing SPP has thus become one of the key factors in overcoming it.

SPONSORED ARTICLE



Peng Zheng
ZTE

Industry practice points to three breakthrough paths: first, standardising specifications to reduce development and integration costs and shorten deployment time; second, splitting complex workflows through refined process management; and third, tackling core technologies to enhance model accuracy. It is important to note that the key to advancing these paths lies in the fusion of three technical paradigms. As the cognitive brain, large models overcome the limitations of traditional rule engines by enabling natural language understanding, expert experience accumulation and long-horizon reasoning. As the execution unit, agents are goal-oriented with closed-loop perception, decision, execution and learning, completing complex tasks such as cross-domain fault self-healing. As the decision sandbox, digital twins provide a safe environment for validating AI decisions through simulation.

Based on the aforementioned technical paradigms, TM Forum standards and the practices of global leading operators, this article identifies three core development trends for autonomous networks over the next 3-5 years.

Trend 1: System architecture evolves from partial intelligence to full-stack AI

Currently, most vendors' autonomous solutions focus on AI at the network management system (NMS) layer, resulting

in delayed perception, slow response and cross-domain fragmentation. The essence of full-network L4 lies in network element-native intelligence, where every network element (NE) becomes an intelligent agent node with local decision-making, forming a distributed intelligent collaboration system.

Trend 2: Operations leap from O&M efficiency to revenue growth

Against the backdrop of diminishing traffic dividends and intensifying homogeneous competition, the focus of network operations is shifting towards experience monetisation, with L4 autonomous network capabilities serving as a key enabler. Three major monetisation paths are gradually taking shape. First, monetisation through differentiated experience packages. For example, customised services, such as low-latency guarantees for cloud gaming, are offered to different user groups, with charging based on experience quality. Second, monetisation via vertical industry empowerment. In sectors such as the low-altitude economy and the industrial internet, 5G-A integrated communication and sensing networks can provide 'communication + sensing + navigation' integrated services for drones. Third, monetisation through data elements. After desensitisation, network data can empower scenarios such as automaker site selection and financial risk control, unlocking data value under compliance.

At the single-domain level, taking performance optimisation as an example, agents are embedded into the operator's O&M system processes

Trend 3: Ecosystem transforms from single-vendor closed-loop to open collaboration


L4 autonomous networks span multiple domains including radio access, core network, transport, IP, services and business. No single vendor can provide full-stack capabilities, so open collaboration has become essential. Global standardisation is accelerating. TM Forum is promoting multi-agent collaboration and the A2A-T protocol to address coordination challenges. The China Communications Standards Association (CCSA) is advancing standards for large models and agents, building a 'general + specialised' framework. 3GPP is enhancing AI security specifications, providing a foundation for AI-native security. In the future, building an open agent factory, shared test suites and a unified governance framework will lead the ecosystem of autonomous networks.

ZTE's approach to advancing high-level autonomous networks

Facing the technical challenges and business demands of Autonomous Networks L4, ZTE has built a strategic framework around three pillars: co-scenarios, co-technologies and co-ecosystems, which work in concert to turn capabilities into value.

Co-scenarios: Promoting high-value effectiveness breakthroughs

Working with global operators, ZTE advances high-value scenario practices through a progressive path – from single-scenario automation to single-domain network autonomy, and ultimately to cross-domain end-to-end autonomy. This drives L4 autonomy from pilot validation to large-scale deployment.



At the single-scenario level, for example, in the network optimisation scenario, the solution utilises structured large models in the network optimisation domain and multi-modal reasoning large models for quality degradation modeling and inference, automatically identifying poor-quality cells and performing root cause analysis, which effectively improves the automated disposal rate of poor-quality cells.

At the single-domain level, taking performance optimisation as an example, agents are embedded into the operator's O&M system processes. Based on intent recognition, root cause analysis and digital twin simulation, work orders are self-optimised, shifting performance optimisation from manual experience to data-driven intelligence.

At the cross-domain level, in fault monitoring scenarios, the end-to-end fault handling agents follow the standard closed-loop process of perception-analysis-decision-execution. By integrating innovative technologies such as graph search, they enable secure interaction between cross-domain and single-domain agents in the wireless and transport domains, achieving precise localisation of complex faults and significantly shortening demarcation time.

Co-technologies: Breaking through core autonomous capabilities

Built on the 1+N model matrix and the Co-Sight Pro Agent Factory, ZTE has developed two types of industrial-grade communication intelligent agents: Copilots and Experts, achieving agile development and safe, reliable operation. On this basis, ZTE has innovatively introduced self-evolution technology, enabling agents to autonomously identify capability gaps and update the tools, business knowledge and decision chains, continuously improving the analysis and processing efficiency in complex complaint scenarios.

ZTE and China Mobile have jointly established the Co-Innovation+ Autonomous Networks Open Lab, focusing on challenges such as high-precision AI decision-making and multi-agent collaboration. By deeply integrating knowledge graphs with strong-reasoning large language models, the lab constructs the highly precise, highly reliable network graph model, significantly enhancing autonomous analysis and decision-making capabilities in high-value scenarios.

In addition, the lab team actively contributes to the development of the A2A-T protocol, enhancing both the efficiency and security of cross-vendor, cross-domain agent collaboration. These innovative achievements have been verified in pilots across several regions such as Beijing and Guangdong. The results demonstrate cross-domain, multi-vendor fault root cause localisation capabilities, with root cause localisation accuracy exceeding 90%, which provides a replicable experience for large-scale L4 commercialisation.

Co-ecosystems: Building an open and collaborative industry

In 2026, ZTE systematically advances AN L4 from technology validation to industry-wide prosperity through open collaboration. At the standard level, the corporation deeply contributes to industry standards organisations such as TM Forum, 3GPP and CCSA, with a focus on frontier directions including large models and agents, and leads multiple standard proposals and patents. In global outreach, the corporation explores benchmark site construction, helping operators worldwide deepen autonomous network deployment and validating the cross-regional replicability of the solutions. For industry leadership, the corporation joins hands with industry partners to continuously make their voices heard together at global industry summits, releasing innovative technical solutions, autonomous network white papers and autonomous network rating certifications, thereby building industry consensus and accelerating ecosystem maturity.

L4 is not the destination, but the starting point for autonomous networks to evolve towards higher-level intelligence. At the inflection point of the AI era, ZTE is committed to working with global operators, vendors, standards organisations and vertical industry partners to deepen full-stack AI innovation, build an open and win-win ecosystem for autonomous networks and jointly shape a smart and sustainable future.

www.zte.com



Why DTW Ignite 2026 matters for telco leaders

Telco leaders will convene in Copenhagen in Denmark to accelerate AI-native transformation as TM Forum, the global alliance of telco and tech companies, brings its annual DTW Ignite Global Summit to the Bella Center on 23-25 June 2026.

The telecoms industry's path to AI-native operations isn't theoretical anymore. Operators are deploying Level 4 autonomy in high value use cases in production environments. CIOs are replacing legacy stacks with composable architectures. Chief AI and data officers are scaling AI from pilots to deliver measurable business impact. The question now isn't whether this shift is happening; it's how prepared you are to accelerate from vision to action, whether you are learning from the operators who are already doing it.

DTW Ignite 2026 brings 5,000 leaders from more than 100 countries together for a fundamentally different kind of industry event. This year's event focuses on one goal: getting you from where you are today to the future - faster - with the frameworks, case studies and expert guidance that compress months of trial-and-error into three high-impact days.



The AI-native telco is here so your next move matters

We're standing at an inflection point, writes Nik Willetts, the CEO of TM Forum. The telecoms industry faces a fundamental choice: evolve into AI-native operators capable of competing in an era of sovereign AI and hyper-personalised services – or get left behind by organisations that will. It's digital Darwinism.

Revenue growth is stubbornly flat for most operators while complexity spirals upward. AI pilots multiply, yet profit and loss (P&L) impact remains elusive. Meanwhile, customers, whether consumers or enterprises, expect the kind of intelligent, frictionless experiences that only truly autonomous operations can deliver at scale.

The gap between those who execute and those who deliberate is widening fast and it's not a technology gap anymore, it comes down to people, execution and trust. How do you build agents that you trust to run a regulated business, make decisions at scale and stand behind the outcomes?

From vision to action

The conversations I'm having with our global membership – including communications service providers (CSPs), hyperscalers, vendors and systems integrators – show that decision makers are done with theoretical discussions. They need to see what's working at scale, who's solved specific problems, and how to adapt those solutions for their own

organisation and customer market.

One CTIO recently said, "I don't have budget for more experiments. I need proven frameworks, real case studies and access to the teams who've actually deployed this, at scale." Another said, "Show me the path from autonomous network pilot to dark NOC. I need to see the business case, the skills gap we'll face and the governance model that makes it sustainable."

These are the kinds of challenges that will determine survival and growth. And collaboration is the only way to move fast and safely enough to stay in the race.

When theory stops and delivery starts

That's why, when 5,000 industry leaders from 100 countries unite at DTW Ignite in Copenhagen, it's not just another conference with AI-generated slides, demos and polite networking. The entire event, like the Forum behind it, is designed to move the industry from conversation to action by bringing proven solutions to solve real challenges, with hands-on support from genuine experts.

For almost 40 years, TM Forum has harnessed the power of collaborative problem-solving. We bring diversity of thought under one roof to solve common challenges because we've learned something fundamental: operators who




Nik Willetts
TM Forum

collaborate move faster than those who don't. When China Mobile and Telefónica share what's working in their networks, everyone can de-risk their own implementation. When Verizon, BT Group and Deutsche Telekom share their approaches to composable IT architectures, the entire industry compresses months of trial and error into actionable patterns.

The winners aren't the ones with the biggest budgets. They're the ones who move fastest by learning from each other, adopting proven frameworks and refusing to reinvent what's already been solved.

DTW Ignite is where innovation happens. This June in Copenhagen, join industry leaders who've chosen execution over deliberation. See what's working at scale, meet the teams who've deployed it, and build your roadmap to AI-native with the people walking the talk.

The AI-native era won't wait, and neither can you. See you in Copenhagen!



New for 2026, Mission Garages are embedded at the heart of each summit

What to expect at DTW Ignite 2026

TM Forum's work centres on three interconnected Missions that define the key communities within the AI-native telco: Composable IT and Ecosystems, Autonomous Networks, and Trustworthy AI and Data. At DTW Ignite 2026, each Mission gets its own dedicated summit, providing a multi-day journey from strategic vision to hands-on execution.

Our summits each tackle the high-stakes questions shaping the industry right now: How do we address the complexity that is killing speed and agility? And how do we rewire for accelerated execution? How do we move from network automation to true end-to-end autonomy? And what does agentic AI mean for control, governance and trust?

These questions converge within this year's theme, enabling the telco industry to reach 'The Future. Faster.'

Opening keynotes feature a strong C-suite cast, including Colin Bannon,

CTO at BT Group, Andrea Fogueiras, global CTIO at Telefonica, Scott Petty, CTO at Vodafone, Pietro Labriola, CEO of TIM, Chen Hong, CIO at Singtel, and Ankur Kapoor, chief network officer at T-Mobile. These aren't aspirational talks. They're candid assessments of what's working in production, what assumptions proved wrong and what they'd change if starting today.

Throughout the event, you'll hear from operators living this transformation: Antonietta Mastroianni from Etisalat on CIO-level IT reinvention, Charles Molapisi from MTN on scaling autonomy across markets, Keri Gilder from Colt on sovereignty and the path to profitable growth, and dozens more sharing unfiltered insights from the front lines.

Mission Garages: the innovation engine

New for 2026, Mission Garages are embedded at the heart of each summit. These aren't expo booths or theatres for slideware. They are working hubs where

theory stops and delivery begins.

Each Mission Garage combines multiple elements in one collaborative space:

- **Mission overviews and deep dives:** Fast, focused briefings that explain what each Mission solves, why it matters now and what to do next
- **High-value scenario sessions:** Short sessions showing real business problems, the approach taken and measurable impact
- **Innovation Hub demos:** Live demonstrations of AI-enabled Open Digital Architecture (ODA) in action, with practical patterns you can adapt
- **Catalyst tours:** Curated walk-throughs of the most relevant proof-of-concept projects, with direct access to the teams delivering them
- **Techlab immersive networking:** Hands-on problem-solving with peers and experts enabling quick connection, shared challenges and real collaboration

When you're at DTW Ignite, you're not watching slide decks about future possibilities or having casual conversations over coffee

A record Catalyst showcase

Each year, our Catalyst showcase – culminating in the closing award ceremony – is a true highlight and a celebration of innovation through collaboration. This year, we're bringing around 60 Catalyst projects to Copenhagen; the largest collection of telco innovation proofs of concept ever assembled at a single industry event. More than half of the projects focus on tackling the real-life challenges of deploying and harnessing the business value of autonomous networks, with Catalysts addressing everything from agentic AI for network optimisation to robotic canine companions for healthcare support, to achieving the 'dark NOC' operating model. Others tackle composable IT transformation, AI governance at scale and data monetisation strategies.

When you're at DTW Ignite, you're not watching slide decks about future possibilities or having casual conversations over coffee. You're seeing production-scale implementations, talking to the architects who built them, and understanding the specific technical and organisational challenges they solved.

Future Skills Arena: Closing the talent gap

The Future Skills Arena tackles this head-on with active learning aligned

to the three core Missions, including hands-on masterclasses and blueprint clinics that translate ODA, autonomous networks and AI governance into day-to-day execution. Role-based talent tracks for developers, architects, product leaders and executives tie to clear Mission outcomes. The aim isn't passive knowledge transfer, but practical problem-solving so you leave with usable skills, defined next steps and confidence to lead AI-native transformation.

Wildcard themes: stress-testing your strategy

Beyond the three core Missions, DTW Ignite 2026 introduces wildcard theme sessions. These are fast, provocative conversations on emerging frontiers that sit outside the main agenda but directly shape AI-native operating models.

These sessions challenge assumptions and push into topics like quantum networking, physical AI and robotics. They're designed for leaders who want to stress-test their three-year strategy against what's on the horizon – not a decade from now, but in the next 18–24 months.

Why this matters now

Operators face a stark reality: revenue growth remains flat while complexity compounds. AI pilots proliferate but profit and loss (P&L) impact stays elusive. Customers expect intelligent,

frictionless experiences that only autonomous, AI-enabled operations can deliver at scale.

The distance between operators who execute and those who experiment is widening by the quarter. DTW Ignite 2026 is a chance to learn directly from the teams solving the problems you're facing, see implementations operating at scale and build a roadmap that moves you from ambition to measurable outcomes. When you're watching a Mission Garage demo showing 40% faster fault resolution through autonomous operations, you're de-risking your own deployment. When you're in a Future Skills masterclass with the architects who've built production-ready composable IT, you're compressing months of exploration into actionable next steps. When you're in CAIDO Circle sessions with data leaders who've moved AI to P&L impact, you're getting the blueprint that separates winners from wishful thinkers.

Copenhagen, 23-25 June 2026. Three days. Three Missions. One goal: accelerating your AI-native strategy to reach The Future. Faster.

Register now using the Agile Telco promotional code **AGILE-MEDIA15** for a **15% discount on tickets**.

Five things not to miss at DTW Ignite 2026

Copenhagen is where the AI-native future gets built. Here's what you need to know:

- **Mission Garages** – Spend time in the hands-on hubs where Catalysts, live demos and real-world delivery collide. Theory stops here.
- **C-level keynotes** – Hear from industry leaders on tackling complexity, agentic AI, why leadership beats algorithms and more.
- **Future Skills Arena** – Masterclasses and blueprint clinics that turn AI strategy into knowledge you can take back to the office.
- **60 Catalyst projects** – Real-life challenges solved by our Catalyst collaborations. Come for the demos, stay for the awards.
- **Wildcard sessions** – Quantum networking, physical AI and sovereignty. Stress-test your strategy against what's next in the world of tech.



ACCESS LEADERS OF CHINA AND APAC'S TECH MARKETS

MWC26 Shanghai is Asia's largest and most influential connectivity event. It's the place to access China and APAC's dynamic digital economies and meet leaders across business, policy and vertical industries.

Uniquely, it brings all three major Chinese mobile operators together, in one place. Be here to drive the future of transformative technologies.

Join us in Shanghai

Get involved:





From humanoid robots to high-performance connectivity – experience the next phase of mobile innovation at MWC26 Shanghai

As the mobile industry enters a new phase of development defined less by standalone breakthroughs and more by the convergence of technologies, MWC26 Shanghai offers a timely lens on the trends shaping what comes next. Across artificial intelligence (AI), advanced connectivity and emerging network architectures, the direction of travel is clear: intelligence is becoming embedded, infrastructure is becoming integrated and innovation is accelerating at their intersection. This year's event further underscores that shift, combining immersive demonstrations, real-world use cases and a diverse global speaker line-up that reflects how connected technologies are moving rapidly from concept into deployment across industries, writes Sihan Bo Chen, the head of Greater China at GSMA.

Trend 1: **Embodied AI moves into the real world**

One of the most visible shifts is the rise of mobile-native and embodied AI. What was once largely confined to applications and cloud-based services is now moving into devices that can sense, decide and act in complex, real-world environments. This transition – from digital intelligence to physical capability – is exemplified by the debut of the Humanoid Robot Football Penalties Challenge at MWC26 Shanghai. While visually engaging, the showcase points to a deeper transformation underway: AI systems are becoming increasingly dependent on real-time interaction with the physical world.

This shift requires more than just smarter algorithms. It depends on low-latency connectivity, edge computing infrastructure and continuous data exchange to support autonomous decision-making. In effect, robots and intelligent machines are becoming endpoints within a broader, connected ecosystem, where sensing, computation and communication must operate seamlessly together.

Market momentum reinforces this trajectory. The humanoid robotics sector is projected to grow from US\$2.92 billion in 2025 to US\$15.26 billion by 2030, representing a rapid compound annual growth rate (CAGR) of 39.2%, reports MarketsandMarkets. By 2035, the market is expected to reach

approximately US\$38 billion. These figures reflect increasing adoption across industries such as manufacturing, healthcare and logistics, where robotics can augment human capabilities and improve efficiency. As these systems scale, embodied AI will become a fundamental layer within the digital economy.

Trend 2: **Satellite and terrestrial networks converge**

Closely linked to the rise of intelligent devices is the growing need for ubiquitous, resilient connectivity. This is driving the convergence of terrestrial mobile networks with satellite and non-terrestrial infrastructure, reshaping how connectivity is defined and delivered. Rather than relying solely on ground-based systems, operators and technology providers are exploring hybrid architectures that combine terrestrial and space-based capabilities to deliver seamless coverage.

The increasing focus on non-terrestrial networks (NTN) and direct-to-device satellite connectivity reflects this shift. These technologies aim to extend coverage into previously unreachable areas, enhance network resilience and support new applications ranging from remote industrial operations to global IoT deployments. As a result, satellite-mobile integration is evolving from a niche capability into a strategic component of next-generation network design.

Trend 3: **6G signals a move toward intelligent networks**

5G-Advanced continues to expand globally and the industry's priority remains completing the 5G journey – accelerating the transition to standalone and 5G-Advanced networks to fully realise the technology's economic and societal benefits. However, preparation has started for the next evolution of connectivity. Early visions for 6G suggest networks that are not only faster, but inherently more intelligent and adaptive. Rather than functioning solely as data transport layers, future networks are expected to integrate communication, sensing and computing into a unified platform.

This shift represents a significant evolution in what networks can do. Networks will increasingly interpret data in real-time, optimise performance dynamically and support AI-driven decision-making across distributed systems. In this context, 6G is less about incremental improvements in speed or capacity and more about enabling entirely new capabilities. It points towards a future where connectivity itself becomes an active participant in digital operations, rather than simply a passive enabler.

Trend 4: **Digital transformation becomes systemic**

Alongside advances in AI and network infrastructure, digital transformation is becoming more systemic and deeply embedded across industries. Sectors such as manufacturing, healthcare, automotive and finance are no longer simply adopting digital tools – they are being fundamentally redesigned around connected intelligence.

From AI-enabled automation on the factory floor to data-driven healthcare delivery and intelligent mobility systems, connectivity is becoming central to how organisations operate. At the same time, cities are emerging as critical testbeds for this transformation. By integrating connectivity, data platforms and interoperable systems, urban

environments are evolving into smart ecosystems that deliver more efficient and sustainable services.

What distinguishes this phase of transformation is its scale and integration. The focus is no longer on isolated use cases, but on re-architecting entire systems to operate in a more connected, responsive and intelligent way.

Trend 5: **Connectivity meets performance in real-world environments**

A notable addition to MWC26 Shanghai is a stronger emphasis on how advanced connectivity performs under real-world conditions. The introduction of a live Formula E showcase highlights the role of 5G in high-performance, data-intensive environments. Featuring a Gen3 Evo demo car equipped with a 5G modem, antenna and live onboard video streaming, the demonstration illustrates how networks operate in scenarios where speed, reliability and latency are critical.

This goes beyond a technical demonstration. It reflects a broader industry shift towards performance-driven connectivity, where networks are designed not only to provide coverage but to deliver consistent, real-time responsiveness in demanding environments. Similar requirements are emerging across autonomous systems, industrial automation and smart mobility, where split-second decisions depend on reliable data transmission.

By bringing these capabilities into focus, MWC26 Shanghai highlights how connectivity is evolving to support mission-critical applications with tangible, real-world impact.

The foundation is scalable connectivity enablers

Underlying these trends is a set of foundational technologies that enable scale and flexibility. eSIM and IoT, while less visible than headline innovations, play a critical role in supporting billions of connected devices across industries. They enable seamless provisioning, remote management and interoperability, which are essential for scaling AI-driven and connected systems.

As these technologies continue to mature, they will further support the expansion of intelligent devices, distributed networks and integrated platforms, acting as the backbone of the connected ecosystem.

From convergence to integration at scale

Taken together, these developments point to an industry moving decisively towards integration at scale. AI is being embedded across networks and devices, infrastructure is extending beyond traditional boundaries, and innovation is increasingly driven by ecosystem collaboration.

In this context, MWC26 Shanghai is less about individual technological breakthroughs and more about how converging capabilities are shaping the future of mobile. It reflects a shift toward a world where connectivity is seamless, intelligence is pervasive, and digital and physical systems operate as one.



How Africa is rewriting its connectivity story

As more Africans come online than ever before, the continent faces a defining choice; will connectivity remain a story about coverage or become a strategy for competitiveness?

For years, the story of African connectivity was one about coverage gaps – but this is changing, quickly. According to the GSMA's Mobile Economy Africa 2025 report, only around 13% of the population in Sub-Saharan Africa now remain outside mobile broadband coverage – a transformation built over a decade of sustained network investment across the continent.

The pressing question of today – and the one that will define the next decade – is what people do with the coverage they have. Today, according to the GSMA's State of Mobile Internet Connectivity 2025 report, only around 25% of people in Africa actively use mobile internet, even when networks reach them. The gap between coverage and usage is now the central challenge, and closing it is less of an infrastructure problem than a policy one.

Connectivity is becoming the foundation for economic transformation across the continent. It is enabling digital commerce, expanding financial inclusion, supporting access to healthcare and education, and powering entrepreneurship. More recently, it is laying the groundwork for intelligent technologies and AI-enabled services that will shape the future of African economies.

Connectivity as economic infrastructure

This shift changes how we think about connectivity itself. Infrastructure conversations in Africa typically focus on roads,

ports, rail and energy – the physical foundations of economic activity. Connectivity is an increasingly critical part of the infrastructure mix which drives economic development. When connectivity is treated as economic infrastructure, it unlocks new pathways for economic growth across Africa. The expansion of 4G and 5G networks, combined with enterprise digitalisation and the rise of AI-enabled services, is already accelerating this transformation. As intelligent technologies become more integrated across industries, resilient connectivity becomes foundational to economic performance itself.

The numbers support this. The GSMA's Mobile Economy Africa 2025 report estimates mobile technologies generated US\$220 billion in economic value in 2024, equivalent to 7.7% of Africa's GDP. This figure is projected to rise to US\$270 billion by 2030, driven by the expansion of 4G and 5G networks, enterprise digitalisation and the early scale of AI-enabled services across industries. This is being realised in practice too. Mobile and digital infrastructure is enabling small businesses to reach new markets and access digital financial services. It's powering entrepreneurs

building digital businesses at scale, and healthcare and education to reach places they previously couldn't.

The value of connectivity must therefore be measured not only in population coverage but in its ability to unlock economic participation and opportunity at scale. Ultimately, it determines who can transact, who can innovate and who can compete in the digital economy.

Africa's mobile-led innovation model

Africa has already demonstrated the transformative potential of connectivity when digital infrastructure is built around local realities and market needs. Across the continent, some of the most impactful digital innovations have emerged from using mobile connectivity to solve uniquely African challenges at scale.

Standout examples of Africa's mobile-led innovation have been born from the need to leapfrog significant infrastructure gaps. The rise of mobile money in the face of the limitations of traditional banking systems is a prime example.

What began as a tool for basic peer-to-peer transfers has evolved into one of the world's most successful financial inclusion models. Innovators in Africa used mobile connectivity to deliver essential services directly to millions of underserved people. This mobile-first landscape turned infrastructure constraints into opportunities for digital innovation.

According to the GSMA's Mobile Economy Africa 2025 report, Sub-Saharan Africa now accounts for more than 1.1 billion registered mobile money accounts and approximately US\$1.1 trillion in annual transaction value. This makes the region the global leader in mobile money adoption, with services extending far beyond payments into savings, credit, insurance, payroll and merchant services.

Across African markets, mobile connectivity has transformed basic devices into tools for economic participation. Simple

In many ways, Africa's mobile-first ecosystem showed that digital connectivity is a powerful driver of economic opportunity

feature phones became powerful tools for economic empowerment, enabling small businesses to accept digital payments via basic SMS and unstructured supplementary service data (USSD) technology. Individuals gained access to financial services without traditional bank accounts, and digital platforms expanded access to commerce at scale.

In many ways, Africa's mobile-first ecosystem showed that digital connectivity is a powerful driver of economic opportunity.

The same opportunity now exists with artificial intelligence and next-generation digital services. As AI applications expand into areas such as language processing, healthcare diagnostics, enterprise automation and digital public services, the scale of the opportunity for financial and economic growth is increasingly clear. However, connectivity infrastructure and how that connectivity is used – and by who – will determine which economies participate in the AI era and which watch it happen elsewhere.

The remaining barriers to inclusive connectivity

Despite the progress made in expanding connectivity, some barriers to inclusive digital participation remain across the continent. The challenge, however, is no longer simply whether network coverage exists, but whether people can meaningfully participate in increasingly digital economies.

The primary barrier in many markets is no longer network availability, but affordability of devices and services, limited digital literacy, energy constraints, spectrum inefficiencies and regulatory fragmentation. For many Africans, the cost of a smartphone is the major obstacle to accessing digital services. At the same time, regulatory fragmentation, spectrum policy challenges and lack of investments across markets slow the pace of digital infrastructure expansion.

The next phase of Africa's connectivity story will depend on how effectively the continent closes the usage gap. With policies that reduce sector-specific taxation, improve spectrum allocation efficiency, create predictable investment environments and measures to address energy constraints, support affordable devices and strengthen digital skills, coverage across the continent can translate into adoption at scale.

Achieving this will require coordinated action across governments, operators, investors and technology partners to build resilient, inclusive digital infrastructure capable of supporting long-term economic growth.

The conversation is moving beyond expanding access toward defining how connectivity can power Africa's next era of economic transformation. Africa now has an opportunity not only to extend networks, but to redefine what connectivity-enabled growth looks like in practice. The choices made today will determine how fully the continent participates in the future global digital economy.

www.gsma.com/about-us/regions/africa

AI and IoT: a marriage made in heaven?

The impact of Artificial Intelligence is being felt throughout the technology world, and beyond. There has been a lot of discussion about the impact that it will have on the Internet of Things, as illustrated by the surfacing of new terms such as 'Physical AI' and 'AIoT'.

In a webinar on 17th June, Transforma Insights analysts will share their perspectives, including the following aspects of the intersection of AI and IoT:

- **AI a driver for IoT adoption.** Given AI's ability to crunch larger volumes of data more effectively, what use cases are enabled or expanded by the use of AI, and what new business models might be enabled?
- **AI on IoT devices.** Where does it make sense for AI instances to be deployed directly onto IoT devices, and what are the implications of managing AI on those constrained AIoT devices, and what platforms are required to do so?
- **Cloud-to-edge orchestration.** AI orchestration will need to span a range of locations from the edge device through to the cloud. What will be the optimal approaches?
- **Impact on the IoT hardware ecosystem.** The need to support AI on edge devices drives a new market for hardware equipment with value-added capabilities.
- **Impact on supporting infrastructures.** Support for AI in the field will place new demands on supporting platforms and infrastructures, including evolved IoT platforms and also cellular core networks.
- **Using AI to support IoT operations.** Use of AI for customer support, task automation, and coding is common. We will explore how IoT suppliers use AI for security, compliance, connectivity, churn prediction and more.
- **The logistical challenges of harnessing AI.** The arrival of the new functionality presents a number of operational challenges, for instance in adapting business processes or managing a mixed estate of devices.
- **A focus on video analytics.** We will use the example of AI-enabled video analytics to illustrate the various evolving components of IoT in an AI context, and the associated new market dynamics.



Learn more about the webinar and other Transforma Insights events at transformainsights.com/events

Autonomous networks built for the AI era

As AI becomes the dominant workload on networks, static infrastructure can't keep up. You need a programmable, AI-native platform designed to anticipate changing traffic patterns and operate at machine speed.

Optimized at every layer for deterministic performance, end to end. Built to turn AI workloads into revenue.

We deliver the complete foundation for autonomous network operations, across all domains and technologies. And we are trusted by the world's most advanced telecommunications providers to apply AI to networks, ready for the journey to L4+ autonomy.

**The AI era won't just run on networks.
It will be defined by them.
Build yours with Nokia.**

NOKIA



Discover more