Ericsson Technology Review

#8, September 2026

Technology and telecom trends 2026

Ericsson CTO Erik Ekudden presents four transformational shifts shaping the network into an intelligent fabric

Charting the future of innovation

ERICSSON TECHNOLOGY REVIEW I 2026

For 150 years, we've connected human intelligence. Now we're giving networks the intelligence to act

Author:
Erik Ekudden, CTO

For more than 40 years, network performance has defined mobile technology evolution, and it will continue to define ours. However, performance alone is no longer the defining measure of what a network can do. The intelligent fabric represents the next fundamental shift in network evolution. As infrastructure that is programmable, autonomous and AI-native by design, it lays the foundation for networks that don't just carry your data, but also sense, predict and act in real time.

Every mobile generation of connectivity has expanded what connected human intelligence can do. Ericsson pioneered the early exchange systems that first made it possible to route communication at scale, entering digital switching in the mid-1970s and mobile communication at the turn of the 1980s. With 5G now deployed across 180 countries with more than 3 billion subscriptions, and 6G on the horizon, we are entering a new era of connected intelligence.

The intelligent fabric represents a significant shift beyond the mobile network architecture, delivering a universal AI network framework that connects and seamlessly integrates distributed AI systems, cloud services, devices, sensors and edge computing across networks, cities and businesses.

It enables AI systems to interact continuously, learn collaboratively and operate at machine speeds while maintaining security, reliability and guaranteed performance through advanced 5G and 6G connectivity.

By combining strong wireless networks with cloud and edge computing, full network autonomy and sound data governance, the intelligent fabric allows networks to sense, predict and act across billions of simultaneous interactions in real time. This will support applications ranging from digital twins and real-time translation to autonomous driving and mission-critical first-responder services.

In doing so, the intelligent fabric will fundamentally redefine what networks can do, enabling AI to transform homes, cities, businesses and industries at a truly global scale.

ERICSSON TECHNOLOGY REVIEW I 2026

A trusted technology stack

The rapid advancement of digital infrastructure, with compute, memory/storage, and networking, is the foundation for any competitive company and country in the age of AI. This infrastructure needs to work as one system, forming a trusted technology stack.

In this stack, connectivity is evolving into an intelligent fabric, which together with cloud/datacenter and device domains, will provide the needed capabilities, both centrally in data centers and distributed nationwide across advanced connectivity and devices for the AI models layer. This, in turn, will be used by applications consumed by tomorrow's users.

Power generation and grid, semi-conductors, spectrum and raw materials, including rare earths, form a shared critical foundation for the technology stack. It will be based on open systems and open interfaces which connect the domains. Such an open and modular approach allows companies to both compete within their domain and cooperate. Across domains, enabling global scale for the full system.

The design principles for the tech stack are openness for innovation, trust for security and resilience, and flexibility and performance to cater for the world of AI.

Ericsson's vision for next-generation networks

To serve these new demands we are building networks based on fundamentally different architectural principles:

Open networks

We are architecting future networks for openness, with standardized interfaces that ensure innovation is never gated by access, and that any developer, enterprise or government can build, compete and grow on top of them.

Autonomous networks

We are architecting future networks for full agentic autonomy, where advanced services can be managed through intent rather than instruction, and the network self-optimizes in real time at a scale that no human operations team can sustain.

Energy-efficient networks

We are architecting future networks for energy efficiency. As AI scales globally, the cost of intelligence becomes a strategic variable, and the network that delivers the lowest energy cost per transaction holds a structural advantage that compounds over time.

Together, these properties enable guaranteed and differentiated services for consumers, enterprises and governments, accessed through open network application programming interfaces (APIs), managed through autonomous operations, and built on a platform that developers and innovators can use to create the AI-native applications of the next decade.

State of play 2026

The next era of connectivity is already taking shape. From best-effort consumer services to AI-native, mission-critical and sovereign applications, the demand profile of mobile networks is shifting structurally and irreversibly. The four trends examined in this article reflect what we see forming, and what we are designing to meet it.

We expect agentic AI to continue its strong trajectory into new categories

Most discussions about AI traffic focus on new generative demands. However, the real shift is behavioral. We are adding a new traffic category to networks optimized for human consumption – and AI-driven applications interact continuously, not intermittently. Agentic AI traffic is a fraction of overall mobile broadband volume today, but it is growing at a rate that makes its future network impact unpredictable. The shift from centralized AI models to distributed autonomous agents will place fundamentally different demands on uplink capacity and latency. We are designing for those demands now, not when they arrive.

Mobile is becoming the primary access layer for AI-native services, and the implications go beyond throughput

Smart wearables, personal agents and multimodal AI applications are converging on mobile as their primary connectivity layer, driven not by today's traffic volumes but by the high and continuous latency demands they will generate as adoption scales. Agents that must act in the world in real time require guaranteed, low-latency connectivity as a baseline condition, not a premium option.

A new class of network demand is forming and the window to design ahead of it is now

As autonomous AI agents scale across manufacturing, logistics, health care and public services, they will generate a distinct class of network demands that are continuous, latency-demanding and uplink-heavy, and they will require verified identity at machine speed. These demands do not yet dominate network traffic, but the architecture required to serve these takes years to build. The decisions made now will therefore determine whether networks are ready for these new demands when needed.

Post-quantum cryptography is standardized, available and ready to deploy

Large-scale quantum computing will eventually break today's public-key cryptography, and adversaries are already harvesting encrypted data in anticipation of that capability. NIST (The National Institute of Standards and Technology) published its first three quantum-resistant standards in 2024, and Ericsson is driving adoption across the 3GPP (the 3rd Generation Partnership Project), O-RAN (the Open RAN Alliance) and the GSMA (GSM Association), with full post-quantum cryptography (PQC) support confirmed for 6G from its first release. Quantum resilience must now be reframed as an infrastructural commitment, not a compliance exercise, with the need for communication service providers (CSPs) to ensure migration by 2030 at the latest.

2026 telecom industry trends shaping the intelligent fabric

In this CTO Technology Trends article, I examine the four technology trends that together define the intelligent fabric: what drives its adoption, what secures its foundation, what extends its perception, and what makes it computationally viable at groundbreaking scale.

ERICSSON TECHNOLOGY REVIEW I 2026

Trend 1: Physical AI

The network's primary user is changing. Alongside consumers and enterprises, autonomous agents are becoming significant actors on the network: generating traffic, consuming capabilities and making decisions at machine speed. This is giving rise to a new commercial reality and a new set of demands that existing network architecture must evolve to meet.

Trend 2: Security in transformation

Trust is not a feature of the intelligent fabric; it is the condition on which everything else depends. The security architecture of next-generation networks must be built to protect today's data from threats that do not yet exist. AI is transforming how vulnerabilities are found and exploited. Quantum computing will eventually break the public-key cryptography that today's networks rely on. Both demand a proactive response and building quantum-resilient, AI-native security into the intelligent fabric is our industry's collective imperative.

Trend 3: The programmable world

The world's networks are becoming a programmable platform. Through open APIs, integrated sensing and cognitive digital twins, the network is evolving into a platform that perceives, exposes and acts on the physical world. For developers and enterprises, this means building directly on network capabilities at scale. For industries, it means a real-time model of the physical world that adapts as conditions change.

Trend 4: Heterogeneous compute

The intelligent fabric will run on a heterogeneous compute architecture, an optimized combination of purpose-built and general-purpose accelerators, each deployed where it delivers the greatest efficiency. Ericsson's custom silicon is optimized specifically for network processing tasks and AI inference, ensuring that the most demanding workloads run with the performance, energy efficiency and cost structure that make global-scale deployments viable.

ERICSSON TECHNOLOGY REVIEW I 2026

Trend 1: Physical AI

Physical AI is reshaping the network's primary user: From people consuming data to autonomous agents sensing, deciding and acting at machine speed

The new dynamic

Exponential agentic growth will generate demand that existing networks were not designed to serve.

Alongside business-to-consumer and business-to-business, a third commercial relationship has emerged: business-to-agent, where software systems buy services, negotiate capabilities and execute tasks autonomously on behalf of humans. The growth of agentic AI traffic has been exponential. In this new business-to-agent economy, competitive position is not secured by contract. It is earned continuously, as agents evaluate and switch based on performance, capability and cost at machine speed.

The next wave is already visible: AI-enabled wearables such as smart glasses place unrelenting demands on networks with the ability to continuously sense surroundings, transmit contextual data to cloud or edge models and act on behalf of users in real time, at any place.

Physical AI is accelerating this transition further. In 2026, we expect to see the first commercial humanoid deployments operating in controlled industrial settings, performing structured, repetitive tasks in logistics and manufacturing environments, with broader use in residential and urban outdoor environments remaining experimental. These systems offload critical algorithms, including real-time vision and control, to 5G edge cloud infrastructure, with onboard silicon augmented by frontier capabilities delivered over the network. What is genuinely new is not their physical capability but their perception. Computer vision, depth sensing and AI reasoning let these systems interpret their environment in real time, enabling autonomous behavior that fixed automation cannot match. Once an agent learns something new, every connected agent in the world has access to that knowledge. In doing so, the network will serve as the shared nervous system of a distributed, continuously self-improving system – and in that sense, an operating system for the physical world, abstracting the complexity of sensors, compute and AI into a single programmable layer.

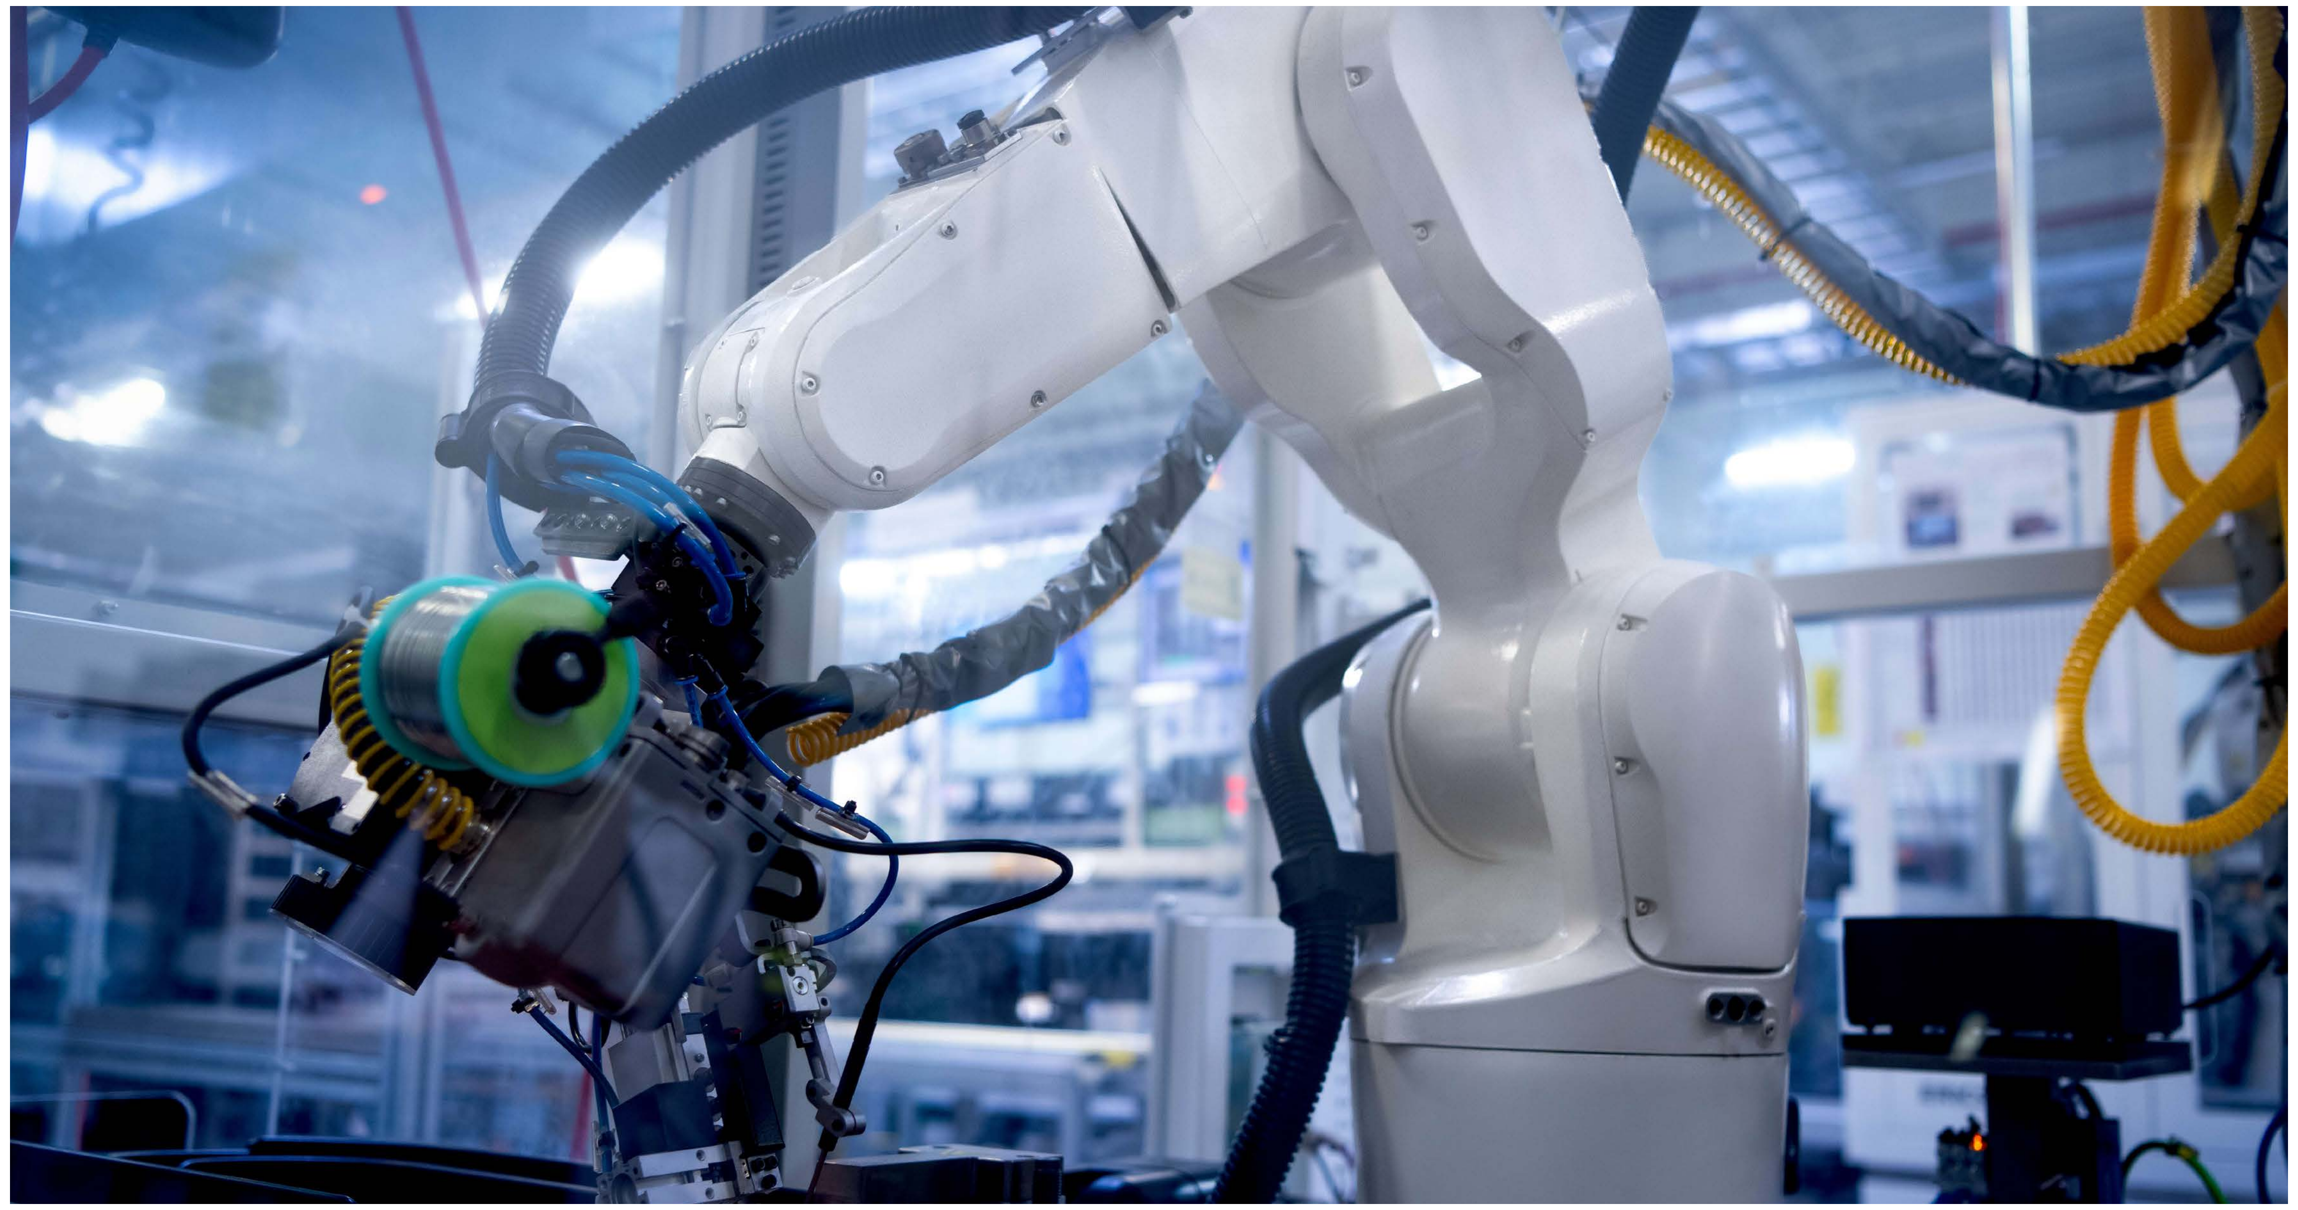

Architectural trends

The networks will evolve into an intelligent fabric that senses, decides and acts at machine speed.

Continuous sensing by agents and devices is reshaping traffic patterns, increasing the strategic importance of uplink capacity, low round-trip latency, edge-cloud integration and, critically, security and trust. Networks will evolve into platforms where agents are authenticated, authorized and granted differentiated treatment, with CSP capabilities exposed through standardized, machine-readable interfaces. As a founding member of the Linux Foundation's Agentic AI Foundation, Ericsson is working to enable autonomous agents to interoperate securely across vendors, industries and national boundaries, building the open and trusted foundation the agent economy requires.

Next steps for the industry

Communication service providers that are first to make their networks available to machines will gain a significant edge in the agent economy.

Agents will increasingly write their own integration code against CSP APIs, discovering and invoking network capabilities autonomously, with no human engineer in the loop.

At Ericsson, we are building the open, standardized interfaces that make this possible. The CSPs whose networks are ready for machine-to-machine transactions first will set the commercial terms on which the emerging agent economy will run.

ERICSSON TECHNOLOGY REVIEW | 2026

Trend 2: Security in transformation

Building quantum-resilient, AI-native security into the intelligent fabric is our industry's collective imperative

The new dynamic

AI is reshaping the security landscape across multiple dimensions: accelerating both threat detection and exploitation, introducing new challenges around agent identity and access governance, and demanding a fundamentally broader security response. At the same time, large-scale quantum computing is expected to break today's widely deployed public-key cryptography. Both require a proactive industry response.

AI is simultaneously the most powerful tool available for network defense and the source of a fundamentally new class of threat.

AI-powered systems can now automate vulnerability discovery, detect previously unknown zero-day threats and generate remediation at machine speed. As the window between threat discovery and exploitation continues to narrow, these capabilities are becoming a baseline requirement for critical infrastructure. The challenge on the threat side is equally significant. As agentic AI proliferates across enterprise environments, it introduces attack surfaces that traditional security models were not built to address. Agents authorized to access systems and execute decisions autonomously require a rethink of authorization, intent verification and action auditing across the full operational chain.

The cryptographic challenge for our industry is structural and long-range. Telecom networks are long-lived, globally interconnected and carry the communications of critical services. Adversaries are already harvesting encrypted data today in anticipation of the quantum capability that will eventually make today's public-key algorithms obsolete.

Architectural trends

We are developing next-generation network architecture on a foundation of AI-native security and quantum-resistant cryptography, across every deployment model and product life-cycle phase.

Security is built into all Ericsson network products from the design stage, with AI-based automation and assurance integrated to enable real-time response to both known and emerging threats. To address the cryptographic transition, Ericsson has been an active contributor to the global PQC standardization process and is driving adoption across the 3GPP, O-RAN and the GSMA. Full PQC support is confirmed for 6G from its first release.

Next steps for the industry

The migration timelines are defined, the standards are in place and the tools exist. The decisions made in this period will determine the security posture of the networks on which the intelligent fabric runs.

Agentic AI is already proliferating across enterprise environments, and the threat landscape it introduces cannot be addressed by traditional security models.
The proactivity of CSPs, vendors, enterprises and governments in addressing this very real threat will be critical: reevaluating authorization, intent verification and action auditing before agentic deployments scale, not in response to the vulnerabilities they expose.

On PQC, high-priority systems should complete migration to quantum-resistant algorithms by 2031, with full transition by 2035. The geopolitical context demands earlier action on the most critical deployments. I expect PQC adoption to begin in the evolution of 5G and to be a baseline requirement in first-generation 6G systems. The networks carrying the world's most sensitive communications must be quantum-resilient before quantum computing capability matures.

ERICSSON TECHNOLOGY REVIEW I 2026

Trend 3: The programmable world

The intelligent fabric converges programmability, sensing and real-time simulation into a single network layer, making the physical world programmable at scale

The new dynamic

Network capabilities at scale, such as differentiated connectivity and integrated sensing, demand autonomous operations programmed on intents and supported by AI, thus transforming the network from a pipe into a platform. This will be further emphasized by agentic API exposure capabilities, enabling a new way to consume network services.

Network programmability will be critical to efficiently support open agentic network APIs, enabling applications to directly express their demands to the network. Early commercial traction in fraud prevention, identity verification and quality-on-demand has established the logic for what follows. The next step is intent-based networking, where CSPs declare desired outcomes and the network fulfills them autonomously, shifting the commercial model from selling connectivity to monetizing guaranteed outcomes.

Alongside this, integrated sensing and communication (ISAC) extends the network into environmental awareness, building on shared radio sites, spectrum and hardware, while introducing new sensing-specific signals, control mechanisms and processing capabilities to detect objects, map conditions and generate sensing data as a service across traffic monitoring, infrastructure protection and urban management.

Cognitive digital twins anchor this programmable layer to the physical world, evolving from passive simulations into self-adapting models that respond to real-world change in real time. On the factory floor of Industry 5.0, collaborative robots (cobots) and human workers collaborate, guided by twins that adapt continuously to production changes, with machines handling hazardous and repetitive tasks, while humans focus on higher-value work. Spatial computing extends this further, softening the boundary between physical and digital worlds across health care, logistics and manufacturing.

Architectural trends

We are building the programmable network layer by layer: from API exposure and intent-based management today, through cognitive digital twins and AI-powered autonomous operations, to full ISAC-sensing capabilities in 6G.

Autonomous networks are evolving from traditional rule-based automation toward systems capable of learning, adapting and making business-critical decisions in real time. Through global ecosystem partnerships, we are creating standardized, machine-readable network APIs through which developers, enterprises and AI agents can access and build on operator capabilities at scale.

As the primary consumers of those interfaces shift from human developers to autonomous agents, we are designing the authentication and validation frameworks that will determine who shapes participation in the programmable network and who captures the value it generates. ISAC is progressing through 3GPP standardization, and we are actively contributing to the architecture that will deliver its full sensing potential at 6G network scale.

Next steps for the industry

By 2030, the autonomous, intent-driven intelligent network will be the invisible yet indispensable fabric of a fully programmable, human-centered world. The commercial decisions being made today will determine who builds it and who benefits.

Near-term, the priority is API monetization at scale, building the developer and enterprise ecosystems that make differentiated connectivity commercially viable. The medium-term opportunity is sensing-as-a-service and cognitive digital-twin integration. Realizing this requires addressing the challenges that remain: cross-vendor interoperability, zero-trust security across programmable layers, and the skills development that widespread deployment demands. The CSPs and ecosystem partners that address these challenges systematically, rather than deferring them, will be those that define the programmable world rather than inherit it.

ERICSSON TECHNOLOGY REVIEW I 2026

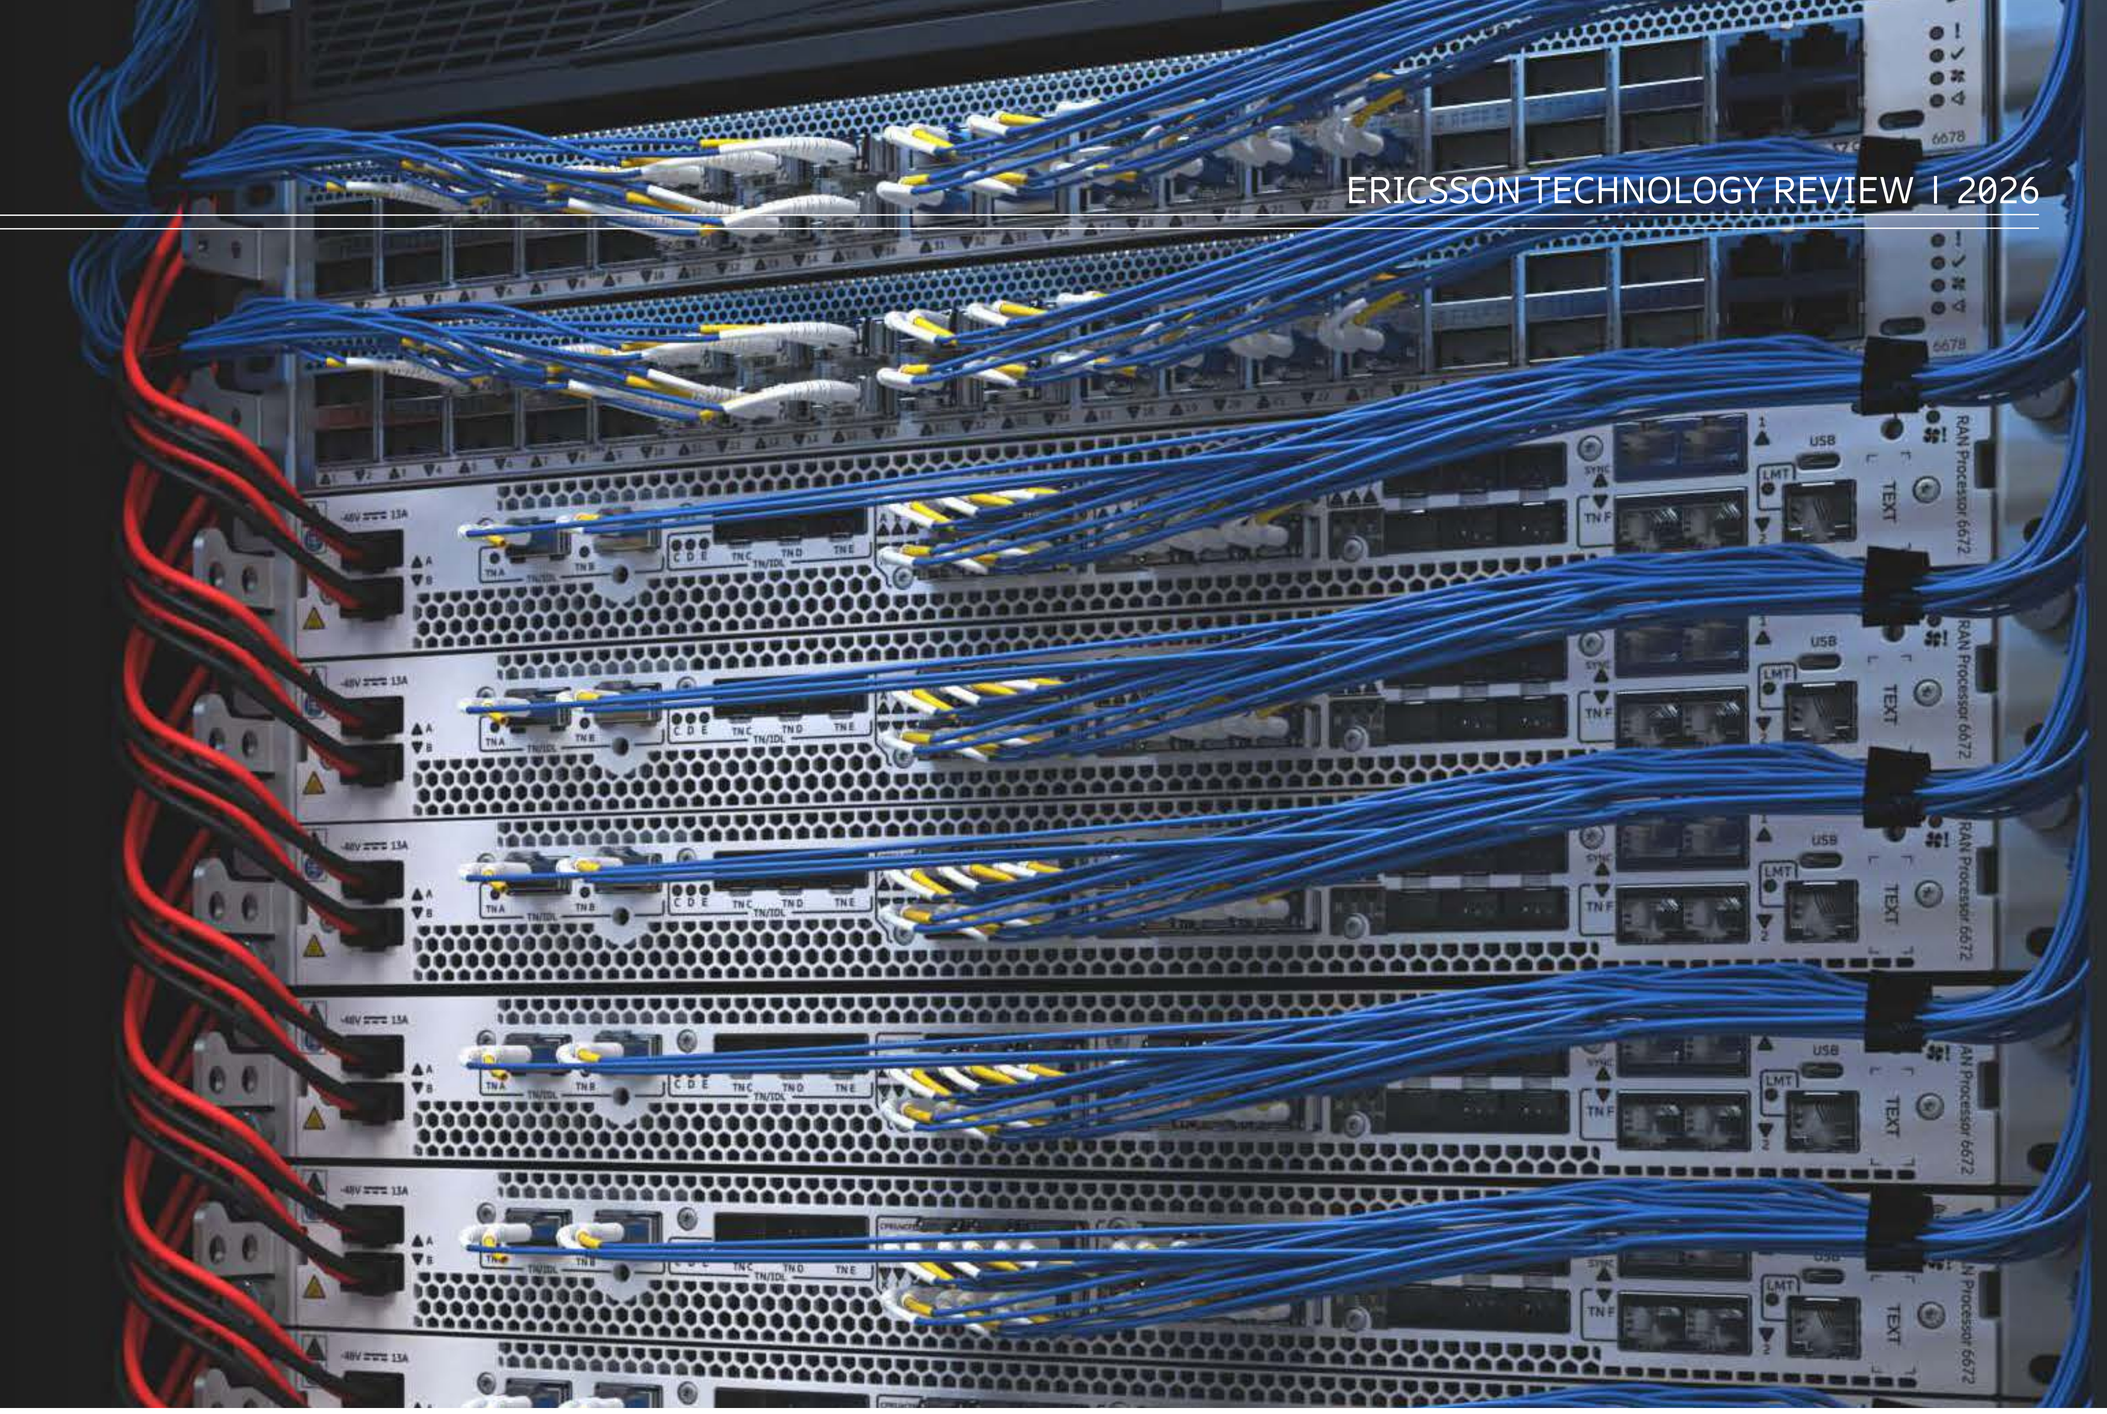

Trend 4: Heterogeneous compute

Heterogeneous compute, built on purpose-designed silicon, is the foundation on which a viable, scalable intelligent fabric will be built

The new dynamic

The shift to AI-native networks demands a compute strategy purpose-built to deliver the performance, real-time responsiveness and energy efficiency that CSPs require at global scale.

Advanced antenna systems in 5G mid-band marked a step change in compute demand that general-purpose architectures were never designed to meet; multiple antenna branches, wide-band beamforming and real-time signal processing in parallel demand a fundamentally different compute architecture. The evolution toward 6G and deeper AI integration across every layer of the RAN (radio access network) will increase that demand further. As AI inference workloads scale across the network, CSPs face a direct consequence: networks built on general-purpose hardware will face a widening gap between what they can deliver and what AI-native applications require, with rising energy costs that compound that disadvantage over time. The compute strategy that CSPs define today will determine their performance ceiling, energy economics and competitive positioning for the decade ahead.

Architectural trends

Purpose-built heterogeneous compute architectures, combining specialized silicon, tensor processing and telecom-specific accelerators, are the design foundation of AI-native RAN.

The compute architecture required for AI-native RAN combines specialized digital signal processors, tensor processing cores and telecom-specific accelerators, implemented through chiplet technologies and advanced packaging on leading-edge semiconductor nodes.

This heterogeneous approach enables AI workloads to execute directly within the RAN stack, at the microsecond-level latency that real-time radio functions demand, while minimizing energy consumption across the network footprint. Ericsson Silicon and the Ericsson Many-Core Architecture represent Ericsson's implementation of this approach, with neural network accelerators now integrated directly into Massive MIMO radios to enable on-site AI inference at scale. The Tensor Core and DSP (Digital Signal Processing) compute elements are designed specifically for increasing AI algorithm use across the RAN stack, providing the computational foundation for the AI-native 6G transition. Looking further ahead, neuromorphic and quantum computing architectures represent long-horizon research investments, with near-term focus on heterogeneous chiplet strategies that scale with evolving network requirements.

Next steps for the industry

The organizations that define their compute strategy for AI-native networks now will be significantly better positioned to absorb the demands of 6G than those that address it reactively.

Purpose-built heterogeneous platforms enable the growing demands of AI inference to be absorbed within the existing network footprint, avoiding the costly reinvestment that reactive approaches require. As AI workloads scale, architectures that deliver the lowest cost of intelligence will generate an operational cost differential that widens over time. The transition to 6G will be the definitive test of these decisions, and CSPs with purpose-built compute foundations will be best placed to meet it.

Conclusion

My vision for the next decade is a world where the network does more than transport the data. It senses, decides and acts at the speed of intelligence itself, supporting humans, enterprises and autonomous systems alike.

The four trends examined in this article are not parallel forecasts. They are interrelated dimensions of a single transformation, each a precondition for the next. Agentic AI defines what the intelligent fabric must serve. Security and post-quantum cryptography define the conditions under which it earns trust. The programmable world defines what it can perceive and expose. And heterogeneous compute defines what makes it computationally and economically viable at global scale.

In my view, the intelligent fabric is to the AI economy what the best-effort mobile broadband service was to the internet economy in the past: the connective layer that makes everything else possible, and the differentiator that determines who leads and who follows. It is not built by extending what exists. It is built by designing for what is forming, with openness, programmability and trust as foundational properties from the outset.

Ericsson will not build this alone, and we have never intended to. The intelligent fabric will be shaped through deep collaboration across CSPs, enterprises, developers, governments and academia, all aligned on a common architecture and a shared direction. The decisions made in the next two to three years will determine how this transformation unfolds, and who is positioned to define it.

We are committed to leading that effort, and wherever you sit in the ecosystem, we invite you to shape it with us.

About the author

As Group CTO, Erik Ekudden is responsible for Group Strategy and Technology. His extensive experience of working with technology leadership globally influences the company's strategic decisions and its investments in 5G, 6G, edge computing, artificial intelligence, augmented/virtual reality and the Internet of Things. Ekudden's leadership builds on his decades-long career in technology strategies and industry activities. He joined Ericsson in 1993 and has held various management positions in the company, including Head of Technology Strategy, Chief Technology Officer Americas in Santa Clara (US) and Head of Standardization and Industry. He is also a member of the Royal Swedish Academy of Engineering Sciences and the publisher of Ericsson Technology Review.

ERICSSON TECHNOLOGY REVIEW | 2026

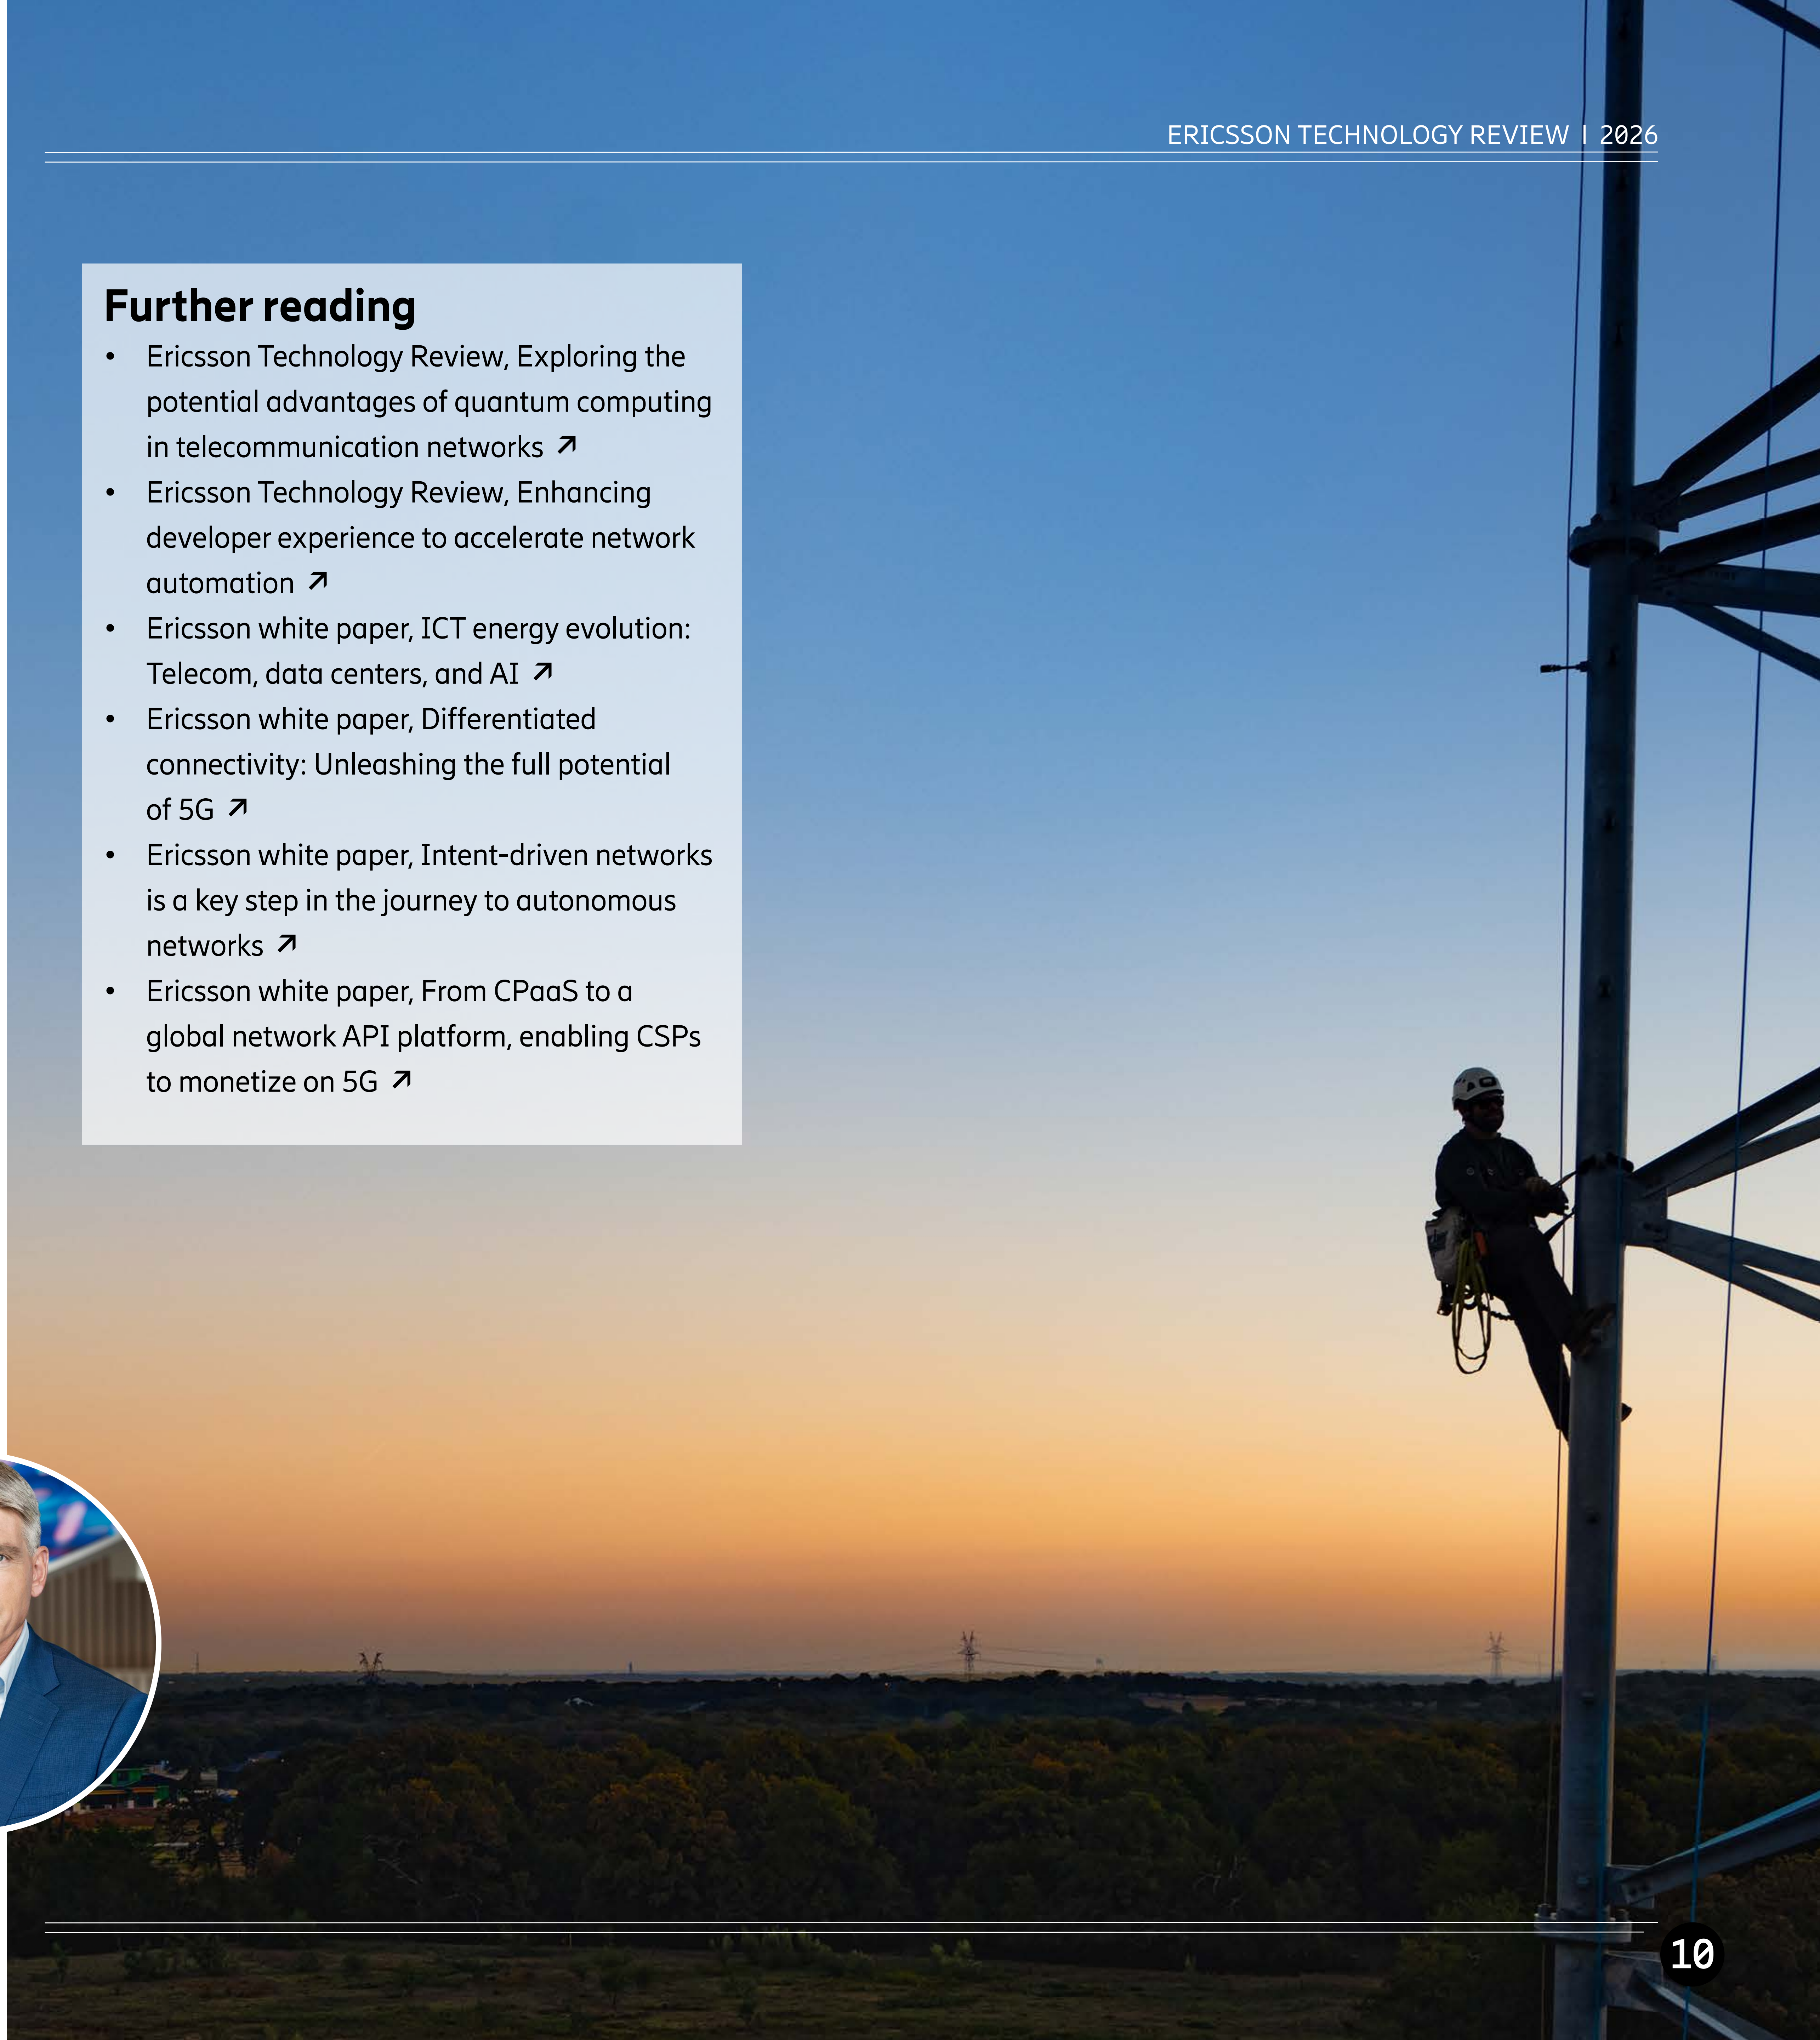

Further reading

- Ericsson Technology Review, Exploring the potential advantages of quantum computing in telecommunication networks ↗
- Ericsson Technology Review, Enhancing developer experience to accelerate network automation ↗
- Ericsson white paper, ICT energy evolution: Telecom, data centers, and AI ↗
- Ericsson white paper, Differentiated connectivity: Unleashing the full potential of 5G ↗
- Ericsson white paper, Intent-driven networks is a key step in the journey to autonomous networks ↗
- Ericsson white paper, From CPaaS to a global network API platform, enabling CSPs to monetize on 5G ↗