White paper

Beyond connectivity: A vision for the future of AI-powered networks

December 2025

Verizon: Jyoti Sharma, Chris Bramble, Hao Wang
Cambridge Consultants: Frank Long, Aashu Virmani, Amy Stephens

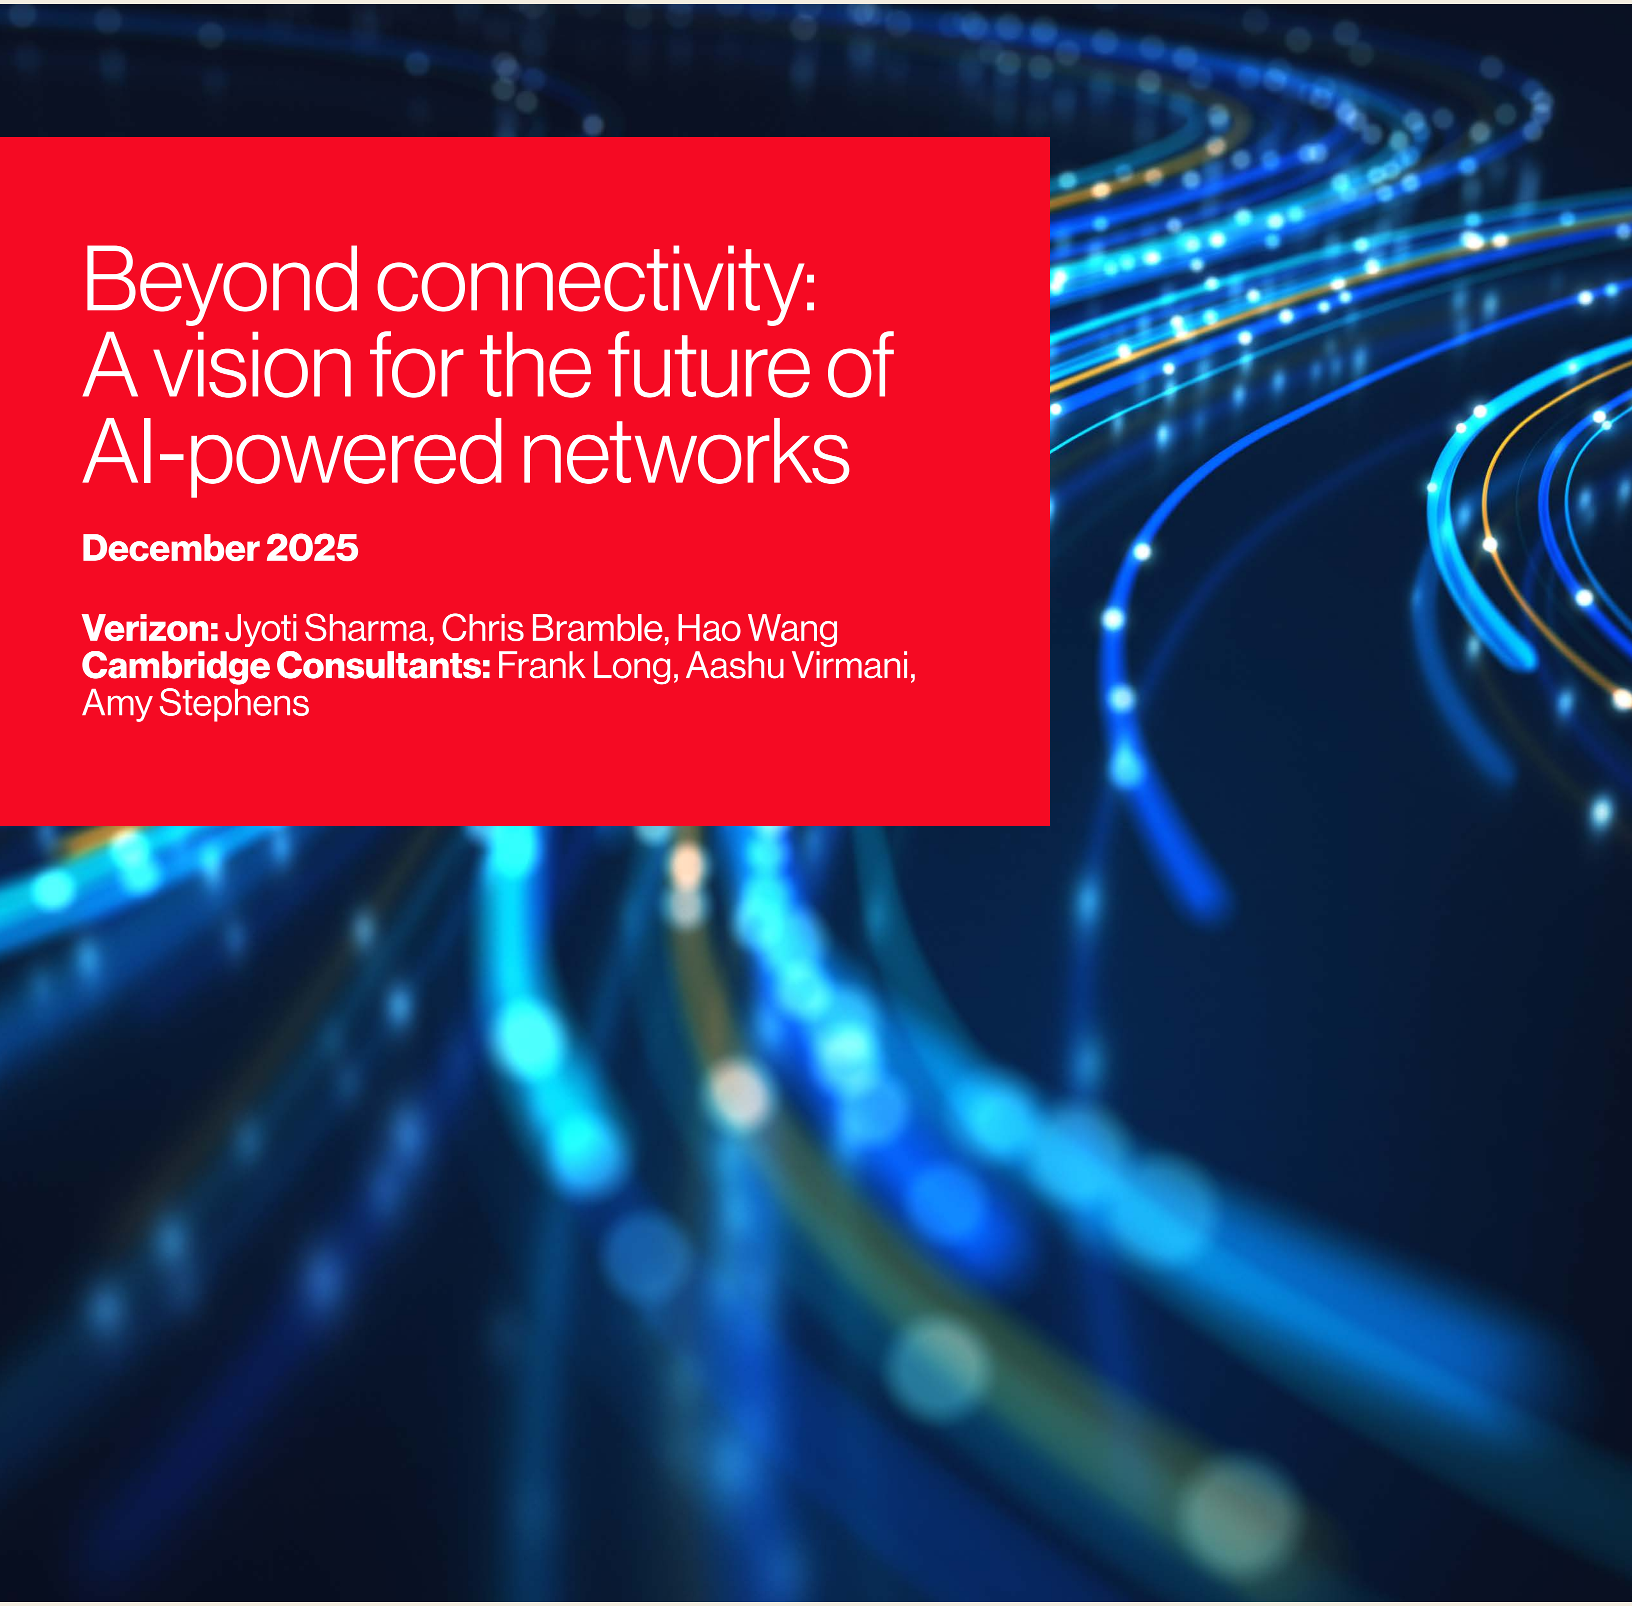

verizon business

Contents

1. Abstract **2**

1.1. The current state of telecom networks. 2

2. Imagine... . . . . . . **3**

3. Technologies that make this vision a reality **4**

3.1. Shifting from static to dynamic capacity management 4

3.2. Implementing fully autonomous operations 5

3.3. Creating a hierarchy of intelligence. 6

3.4. AI radio improvement 7

4. The opportunity of AI **8**

4.1. The telco AI value chain. 8

4.2. The internal network improvement opportunity 9

4.3. The B2B opportunity. 10

4.4. The B2C opportunity 11

5. The cost of inaction **12**

6. AI governance: Governance-by-Design rooted in risk-based principles. **13**

7. Early proof and call to action **14**

8. Meet the authors **16**

1. Abstract

Telecommunication networks are evolving into a dynamic, intelligent ecosystem, driven by artificial intelligence. This transformation enables fully autonomous operations and the dynamic allocation of resources, shifting the Mobile Network Operator's (MNO's) role from a simple connectivity provider to an orchestrator of hyper-personalized, SLA-bound intelligent services. This joint white paper from Verizon and Cambridge Consultants (CC), the deep tech powerhouse of Capgemini, outlines a vision where a hierarchy of AI, from centralized "brain" models to distributed agents, manages a marketplace of dynamic spectrum and multi-domain assets (terrestrial, aerial and satellite). By embracing this AI-powered future, operators can unlock new value streams, achieve significant operational efficiencies and redefine their market position. It is a strategic imperative to lead this change or risk commoditization of the network itself.

1.1. The current state of telecom networks

Before exploring the impact of AI, it's essential to understand the current state and challenges of the telecom network, generically. Networks were historically designed for peak usage, incorporating high redundancy, resulting in underutilization, often as low as 20-50%. AI is due to drive an increase in network traffic – a 2024 Nokia study predicts it will cause a fivefold increase in wide area network traffic.[1] Large accompanying increases in utilization, driven by AI, are forecasted for data centers, even as hyperscaler growth is being impacted by power availability.

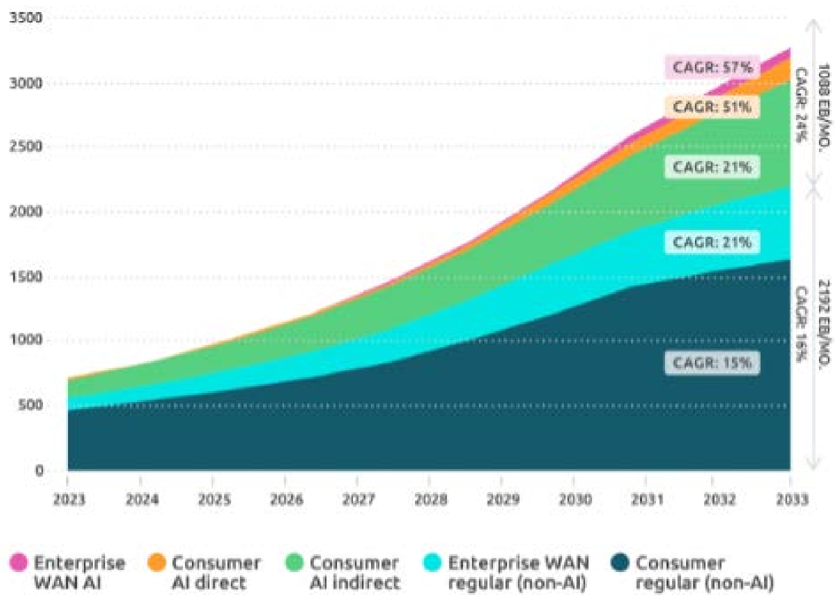

Figure 1: Global WAN AI traffic growth (exabytes/month), Nokia Bell Labs

The trend towards autonomous networks is clear. The majority of MNOs are targeting at least TM Forum Level 4, representing a shift from human defined to true autonomous decision making. To bridge the gap from Level 3 to Level 4 requires AI; at Level 4, networks need to operate without human oversight, they need to self-manage, self-optimize and handle complex tasks. The levels will be further defined later in section 3 (see Figure 3).

According to a 2025 TM Forum survey of 91 companies globally, 96% of network operators are currently at Level 3 or below.[2] Over 60% of North American respondents say they'll achieve Level 4 autonomy by 2030, which represents a rapid transformation of networks, with the top priorities being anomaly detection and resolution, service assurance and network change.

The physical nature of the network is also changing. More operators are integrating satellite services by default into their network,[3] drones can be used to temporarily provision new capacity in critical situations[4] and private 5G networks are being deployed across major enterprise customers.[5]

AI is transformative across all domains of telcos. While generative AI and large language models (LLMs) are the primary driving force behind AI's impact, AI is not monolithic. AI is a mix of different underlying technologies and techniques, depending on application and enabling technologies. AI applies across all major domains within telecoms networks, as illustrated (see Figure 2).[1]

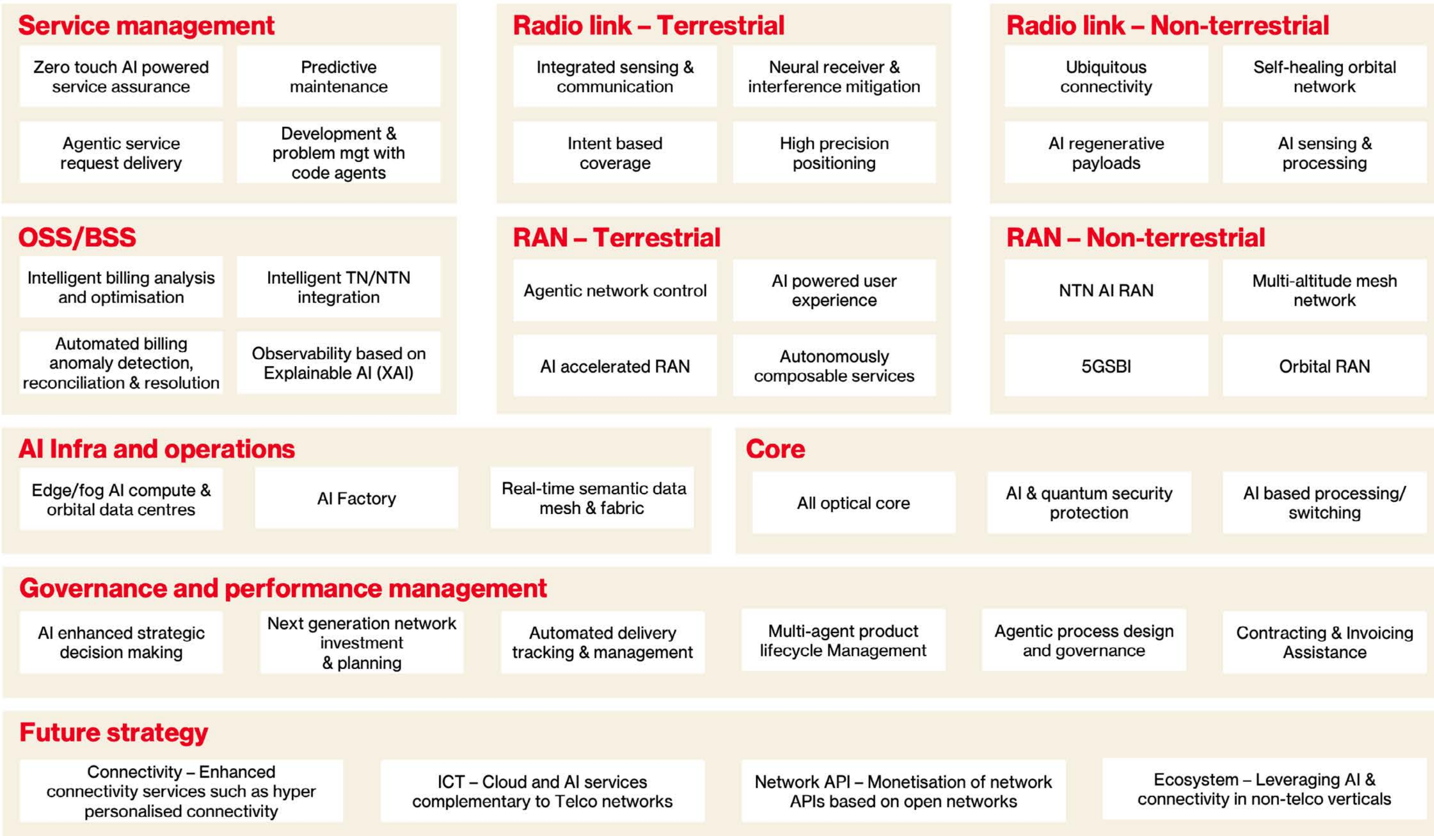

Figure 2: AI powered future telco vision, Cambridge Consultants

The nature of the network will change; it will become physically more dynamic, providing zero touch, self-service autonomous operations. AI technologies will drive improvements to the air interface.

AI is creating new capabilities and ways of thinking about the services networks provide. This will require new thinking about the nature and purpose of networks, shifting from providing content to enabling intelligence. Content distribution networks push the same content out to thousands of nodes. In contrast, AI will provide context-aware intelligence, adapting nodes to local context and needs, e.g. distinct AI systems for intelligent urban services in New York versus Los Angeles.

One of the most significant changes will be the obsolescence of human-powered Radio Frequency (RF) planning and cell-site placement. In the future network, AI will lead and handle most of these tasks, as the network itself becomes physically shape-shifting to meet demands.

This paper explores some of the key enabling technologies and ideas, and how they apply to an AI-powered vision of the future of telecoms networks.

2. Imagine...

The year is 2030. N*AI*ture, a pioneering AI-focused wildlife conservation company, is embarking on a critical mission. After a prolonged dry spell, sensors across California's iconic Sequoia National Park indicate an alarming increase in stress levels among the ancient giant sequoias such as the famous General Sherman Tree, signaling a heightened risk of disease.

A small group of researchers are setting off from their San Francisco headquarters, traversing the Central Valley and heading deep into the park. As they leave San Francisco during morning rush hour in their autonomous pickup, heavy traffic results in high local network traffic from the autonomous vehicles now populating the streets of San Francisco. In response, a small fleet of drones acting as mini cell towers are provisioned to provide extra network capacity.[4]

There's an accident on Interstate 5 that threatens to disrupt their journey. The emergency services rapidly get situational awareness by using insight provided by cameras located on the highway, interacting with local compute to create a report on the location, severity and recommended responses less than a minute after the accident happens. The autonomous vehicle uses intelligent sensing from onboard network sensors, in addition to its standard sensing equipment to safely navigate the accident scene.

As the team from N*AI*ture continues its journey through the Central Valley, their network experience agent detects a blackspot of coverage along their route. The team's network experience agents determine that at their velocity, direction and standard patterns of travel, the team will hit the blackspot shortly. Their network experience agents submit a request to the network control agents to mitigate that upcoming network blackspot. The request is analyzed by a regional network control agent, checking the network digital twin for any issues, and handed off to a local network control agent, which reactivates Band 5 at one of the local cell towers and adjusts all its neighbors for the next few minutes. When the team passes by, the network reverts to its previous configuration.

As they set up camp in the mountains, the team sends up a few drones to survey the local area; their vehicle acts as a private network, providing an access point to the drones, with a direct satellite backhaul connection. The sunroof on their vehicle is a large, phased array antenna, using digital beamforming to provide a high bandwidth connection to an AI factory where the sensor data and drone footage is analyzed in detail. N*AI*ture's planning agents provide them with an optimized hiking route for them to visit the most critically affected areas in person to collect samples and gain crucial insights about the health of the forest.

Making this futuristic scenario a reality requires a shift in how we think about networks, from providing simple connections to providing layered connectivity. Internally, AI becomes the fundamental engine driving predictive network planning, highly autonomous operations and the dynamic provisioning of hyper-personalized services. By embracing open ecosystems and advanced APIs, the telco integrates a diverse array of access technologies – from various specialized providers, satellite and drone operators to vehicle manufacturers – creating new business opportunities and ways of working. The focus moves decisively from selling data packages to delivering guaranteed, intelligent and adaptive connectivity experiences, making the network an intuitive and indispensable partner in each user's activities.

3. Technologies that make this vision a reality

The realization of the AI-led, automated and hyper-personalized network hinges on the maturation and strategic integration of several cutting-edge technologies accelerated by AI. AI will drive fundamental change to the network, changing it in four main ways:

- Shifting from static to dynamic capacity management
- Implementing fully autonomous operations
- Creating a hierarchy of intelligence
- Significant changes to the air interface through AI radio improvement

3.1. Shifting from static to dynamic capacity management

Networks are moving away from fixed capacity deployments to dynamic, on-demand resource allocation. AI-driven systems will continuously monitor traffic patterns, predict demand fluctuations and reconfigure network resources in real time.

3.1.1. AI-Driven on-demand resource allocation

The first layer of dynamic capacity management involves using AI to intelligently orchestrate a diverse set of network resources. AI-driven systems, functioning as the network's cognitive core, continuously monitor traffic patterns, predict demand fluctuations and reconfigure resources to meet needs as they arise. Instead of relying solely on fixed cell towers, this new model leverages a multi-altitude, multi-domain ecosystem of assets:

- Non-terrestrial networks (NTNs): High altitude platforms (HAPS) and Low Earth Orbit (LEO) / Geostationary Earth Orbit (GEO) satellites provide broad, resilient coverage. AI will be essential to manage handover and traffic between terrestrial and non-terrestrial networks.[6]
- Aerial and ground assets: Drones acting as "mini cell towers" and even autonomous vehicles serving as mobile hotspots ("vehicle-as-tower") can be dynamically provisioned by AI agents to provide temporary, hyper-localized capacity.[4] Managing the RF planning, provisioning and orchestration of on-demand drones or vehicles acting as towers requires AI to manage the complexity doing so introduces to local networks.

3.1.2. Dynamic spectrum arbitrage

The most profound shift lies in reimagining how spectrum – the lifeblood of wireless communication – is managed. The current model, involving large, upfront acquisitions of statically allocated spectrum, creates artificial scarcity and monetizes a finite resource over decades.

A more novel, AI-driven marketplace approach treats spectrum as a dynamic commodity.[7] This vision, "dynamic spectrum arbitrage," mirrors the deregulation seen in the utility industry, where power generation was decoupled from distribution. Key elements include:

- **Open access and pay-per-use:** Regulatory bodies could make spectrum available not through static auctions, but on a pay-per-use model, by treating the space-time-geography cube as the underlying axes and allowing the spectrum as a marketplace commodity. As an example, it would be possible to make available an additional 100 Mhz of spectrum over a certain geography for a set period of time, to provision extra local capacity for a large sports event.

- **AI-enabled marketplace:** This open access marketplace, enabled by AI-led planning and rich network APIs, would allow any telco network to dynamically expand and contract its capacity. An MNO could transparently buy capacity from this marketplace to handle peak demand for a local event and, just as easily, sell its own excess capacity during periods of low use to an enterprise or another terrestrial or non-terrestrial operator.

A few years ago this idea would still have been valid but would have met with great difficulty in implementation – largely because of the static nature of network design and planning and the need to place assets to cover a geography and develop the additional back-office support to operate the network. However, with AI led network planning, a rich set of APIs exposed by peer networks and the underlying slicing capabilities in the standards, we are not far from the vision when any telco network can buy and boost its capacity by tapping into an open-access, transparent marketplace and acquire additional capacity – not unlike one goes to a hyperscaler to scale CPU capacity on-demand.

3.2. Implementing fully autonomous operations

The ultimate vision is for a self-optimizing, self-healing network driven by agentic and physical AI. This entails a progression from current AI-enhanced operations (TM Forum Level 2/Level 3), which involve basic automated network configuration management and genAI for IT ticketing support, to fully autonomous lifecycle management and sophisticated closed-loop automation driving intent or experience based operations.[8]

Autonomous levels	Level 0	Level 1	Level 2	Level 3	Level 4	Level 5
	Manual operation and maintenance	Assisted operation and maintenance	Partial autonomous networks	Conditional autonomous networks	Highly autonomous networks	Fully autonomous networks
Execution	P	P/S	S	S	S	S
Awareness	P	P/S	P/S	S	S	S
Analysis	P	P	P/S	P/S	S	S
Decision	P	P	P	P/S	S	S
Intent	P	P	P	P	P/S	S
Applicability	N/A	Selected scenarios				All scenarios

Figure 3: Levels of Network Automation TMForum

Figure 3 illustrates the types of advancements in use cases possible with increased levels of autonomy in the network, powered in a large part with AI adoption within the operations.

3.2.1. From anomaly detection to proactive self-healing

Achieving this level of autonomy requires a network that can perceive its environment, analyze complex situations, make intelligent decisions and execute actions without human intervention. Networks must move beyond reacting to failures to predicting and preventing them. This is being achieved by:

- **Predictive anomaly detection:** By leveraging AI to analyze thousands of network metrics, sophisticated "heat maps" of network behavior can be created. This allows identification of previously invisible leading indicators of outages, shifting maintenance from a reactive, costly exercise to a proactive, data-driven discipline.

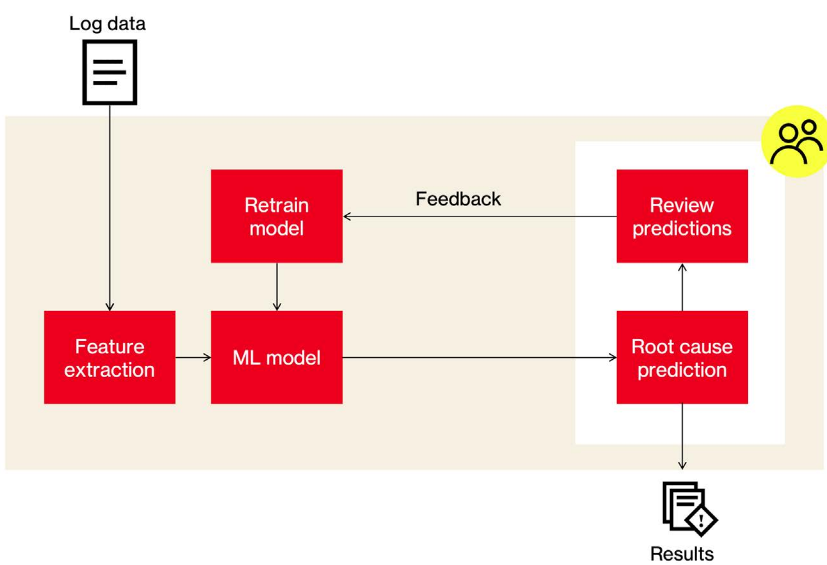

Figure 4: Predictive anomaly detection implementation, Cambridge Consultants

- **Predictive resource management:** AI-powered cell load predictors can forecast traffic demand with very high degrees of accuracy, enabling the network to autonomously power down capacity layers during off-peak times. This translates to significant energy savings and OPEX reduction without any degradation in performance. This predictive capability, working with network management agents, could provide coverage to blackspots before users encounter them or provision a cell drone to meet peak demand traffic.
- **User QoE using network experience agents:** The challenge with true user-experience is that networks are focused on parameters such as throughput and latency, while the user experience has to do with the real perception of quality at the UE. Network experience AI agents can infer the user experience using standard RAN parameters, to understand application type and user context. These agents can work with network control agents to proactively trigger autonomous network actions – like traffic prioritization or handovers – to prevent poor experiences before they happen, increasing satisfaction and retention.
- **Network control agents:** AI RAN will drive an AI-native approach to RAN management, using large numbers of agents to provide intent based RAN management.[9] Cooperative multi-agent reinforcement learning (MARL) is a decentralized agentic AI technique that deploys many lightweight, collaborative agents across the network. Instead of a single master AI, these agents learn from experience to collectively optimize the system. They can take on various roles, manage real-time tasks like RAN power saving and configuration, network routing and network protection.[10] Early models suggest at least a 20% reduction in power savings with an improvement in customer experience, using as little compute as a Raspberry Pi.
- **Self-healing in large phased array antennas:** As an example of AI's applicability in self-healing networks, AI can be used to compensate for failed elements in satellite and large MIMO antennas in a terrestrial network. This is achieved by using AI techniques to automatically adjust the weights on the other elements of the antenna to compensate for the failure. Large, multi-element antennas can be successfully trained to self-compensate against the first few failures without service degradation, enabling self-healing phased array antennas.[11]

3.3. Creating a hierarchy of intelligence

The landscape of AI is undergoing a profound transformation, with the emergent intelligence of LLMs and agentic AI driving rapidly improving reasoning and planning capabilities. These capabilities are driving a shift to a more distributed, multi-tiered architecture, to address the diverse and often conflicting demands of modern AI applications, including stringent requirements for low latency, optimized power consumption, robust data privacy and national data sovereignty.

Highly capable reasoning or "brain" models are power-intensive, the most intelligent of which require specialized, water-cooled GPUs in centrally located data centers. AI factories provide the dedicated hardware and software to create, maintain and serve these massive models.

Lower powered "body" models provide localized, low latency intelligence and immediate, on-device adaptations, operating with lower or minimal power consumption for inference-only tasks, such as real-time monitoring, local agentic action or data processing, analysis and reporting based on local contextual data.[9] They can be run on lower power GPU or even CPU compute, with MARL being an example of CPU-based AI agents.

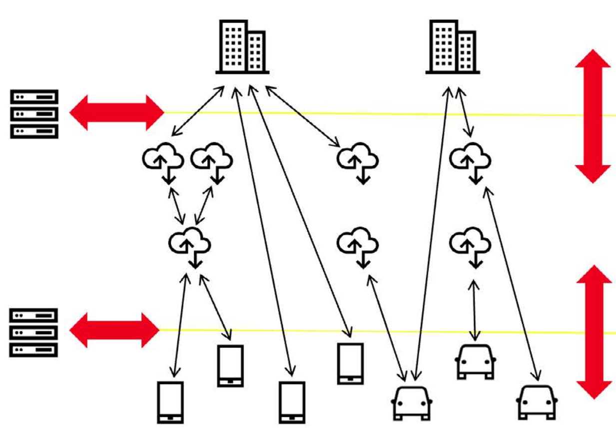

Brain models

Very large Language and Multi-Modal Models
- Used for powerful reasoning and cognition
- Models that require lots of GPU compute power for training and inference, e.g. Llama 3 405B
- Focused on large scale inference and training

Brain/body models

Mid-sized LMM and LLMs
- Used for mid-level powered cognition, offloading device intelligence, delivering services at the edge.
- Models that require mid-level of GPU compute power for inference, e.g. Llama 3 Instruct 3b
- Primarily inference, can be used with intelligent workload orchestration for effective fine tuning & training

Body models

Smaller language
- Used to deliver services and local intelligence to users
- Models that require lots of GPUs for training and minimal GPUs to run locally on the device
- Inference only

Figure 5: Hierarchy of intelligence, Cambridge Consultants

3.3.1. Large telecom models (LTMs)

These are specialized LLMs and multi-modal models (MMMs) tailored for telecommunications data and operational insights. Functioning as "brain models" LTMs leverage vast network data for powerful reasoning and cognition, automate complex network design, generate insights from operational logs and assist in strategic planning by processing unstructured data from diverse sources.

3.3.2. Collaborative, multi agent reinforcement learning

MARL is an established AI technique that builds upon multiple individual agents collaborating for a higher reward. A reinforcement loop punishes suboptimal actions and rewards good outcomes - allowing the model to improve naturally over time.

Techniques like collaborative MARL are particularly useful in a large network composed of various radio sites, where adjusting the parameters at one site affects its neighbors. While each of the agents can act autonomously, cooperative MARL ensures the changes in environment resulting from one agent's action can be acted upon by other agents.

This AI technique doesn't require a massive amount of capital investment in distributed GPU farms all throughout the network, and agents can be low-powered enough to run on Raspberry Pis. Verizon has a mature Mobile Edge Computing (MEC) deployment, VCP, 5G UPF's etc., MARL could be employed as a tool for network power optimization, leveraging the existing capital that has gone into building out this infrastructure.

3.4. AI radio improvement

AI supports the development and deployment of a few key enabling technologies for making improvements to the air interface.

3.4.1. AI for beam management

Beamforming is a core technique used in wireless networks to improve signal strength and reduce interference and power consumption by directing the transmitted signal towards the correct user. It is especially important at higher frequencies to compensate for increased path loss and low penetration through obstacles. To enable the use of beamforming, a beam management process is required which sweeps through beam directions, measuring and reporting the quality of the reference signals to determine the optimum beam pattern to use.

The problem with this process is twofold: the delay before the connection is configured adds latency, and the bandwidth required to transmit the reference signals is overhead that reduces spectrum efficiency. To mitigate this, 3GPP has been investigating the use of AI to streamline beam management in 5G-Advanced, and the standardization process is already underway.[12]

3.4.2. Neural receiver

A neural receiver is a revolutionary AI-based technology that replaces traditional signal processing algorithms. Instead of relying on rigid mathematical models, it uses a trained neural network to decode radio signals, intuitively filtering out real-world noise, interference and distortions.

This enhances performance in challenging environments, significantly boosting spectral efficiency, increasing network capacity and improving link robustness. It extends cell coverage and is a critical enabling technology for making higher frequency bands viable. By improving the signal at the most fundamental level, it provides a powerful, future-proof competitive advantage for the entire network – initial results show a 40% improvement for channel estimation in layer 1[13] and operate at lower SNR (by up to 4dB) than conventional receivers.

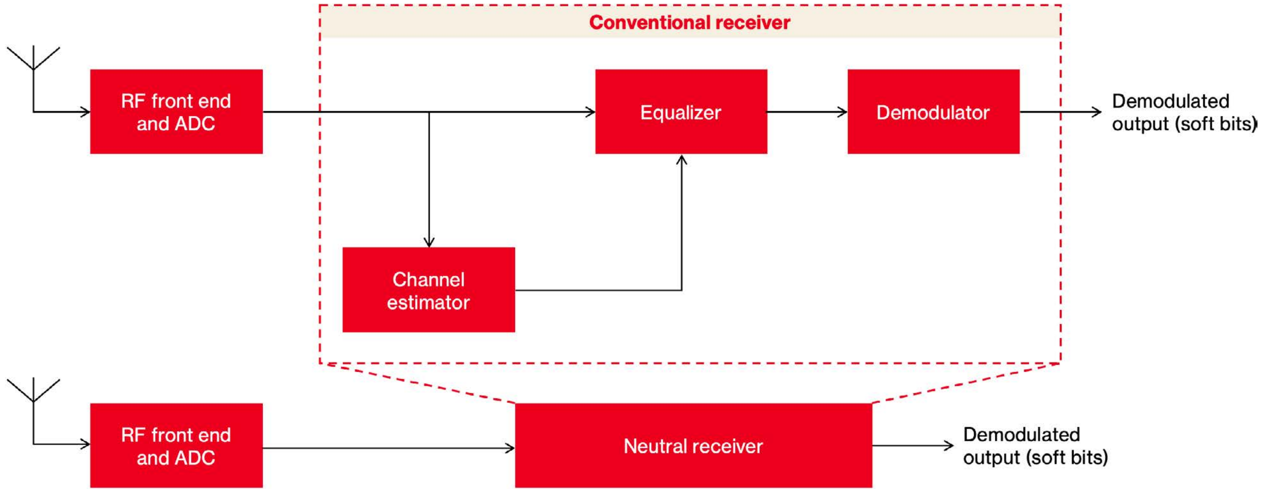

Figure 6: Neural receiver implementation, Cambridge Consultants

3.4.3. Reconfigurable intelligent surfaces (RIS)

RISs are passive or active surfaces embedded with metamaterials that can intelligently reflect, refract or absorb electromagnetic waves. RIS can dynamically steer signals, fill coverage holes (especially at higher frequencies like mmWave) and enhance signal quality, acting as "smart mirrors" for radio waves. They're particularly important in high density urban environments, providing line-of-sight high bandwidth connectivity at high frequencies in complex physical environments. They are crucial for ubiquitous, high-frequency deployments in dense urban environments and filling coverage gaps.

3.4.4. Integrated Sensing and Control (ISAC)

A technology that combines communication and sensing functionalities within the same hardware and spectrum resources, ISAC allows the network to not only transmit data but also to "sense" the environment (e.g., precise user positioning, object detection). In the N*AI*ture scenario, emergency services situational awareness, autonomous vehicles using network sensors and drone control are all direct applications of ISAC. The network is actively perceiving its environment, enabling rapid, context-aware responses. This capability enables new services like advanced Vehicle-to-Anything (V2X) communications and transforms the network into a broad sensor platform.

4. The opportunity of AI

The profound impact of AI on telecommunications is no longer a question of if, but when. Customer service and service management are already being transformed by generative AI. AI will transform all aspects of the network over the longer term as detailed in previous sections. This transformation opens opportunities for improving the network, and the services provided to both businesses and consumers.

With opportunities to participate at multiple levels in the AI value stream, a critical question to answer is: "Where should one invest?"

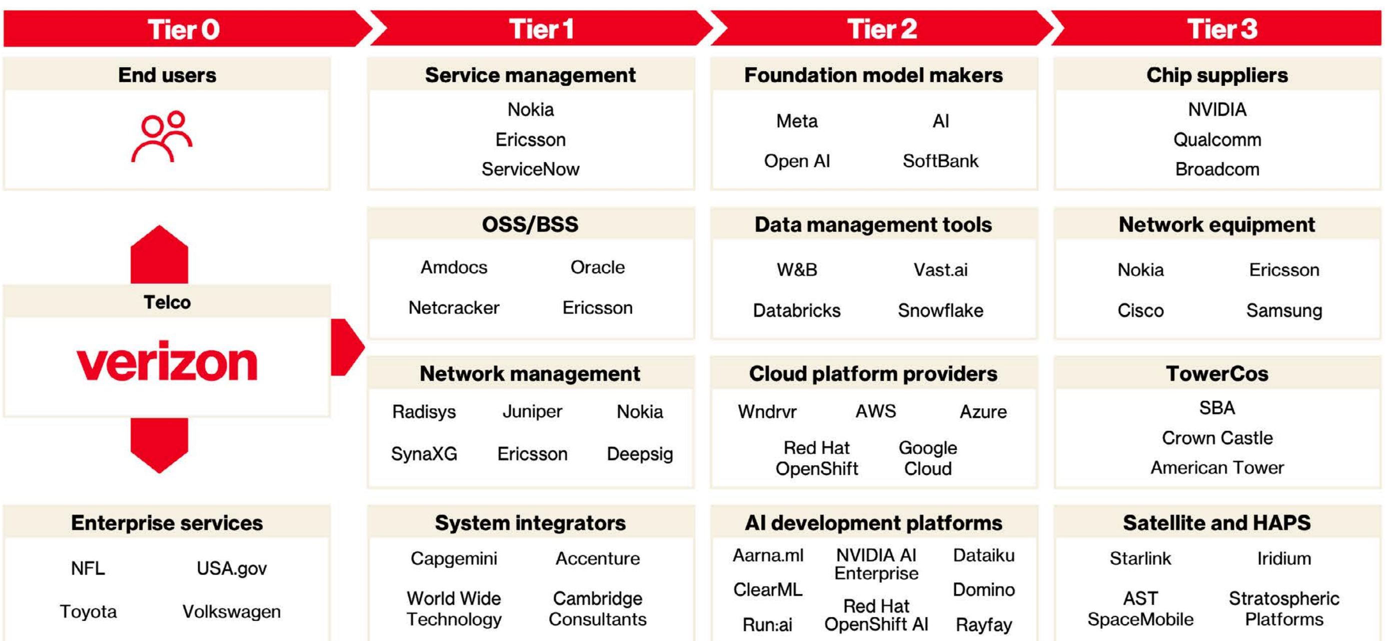

Figure 7: AI-powered telco value chain – not intended to be exhaustive

4.1. The telco AI value chain

4.1.1 Tier 0

End user and enterprise AI applications that provide direct value to customers by providing new opportunities and solving challenges.

An example of this from the N*AI*ture scenario is providing sensing as a service to vehicle manufacturers, consuming their sensor data and providing results that leverage local MEC for offloaded processing of complex data.

Verizon is already building services and applications at this layer and is likely to be the key source of differentiated advantage due to existing strengths in business and consumer services and solutions.

4.1.2. Tier 1

Key services and systems that built telco network services, enhanced by AI.

AI services are essential to achieve Level 4/5 autonomous networks, in the N*AI*ture scenario, the agentic approach to networks is not possible without AI.

To achieve the objectives in highly autonomous networks, MNOs will both have to work with partners, and develop Tier 1 AI solutions, to address their specific opportunities and challenges. MNOs have many different suppliers that integrate to create its services and systems, AI will be required to effectively orchestrate and automate the integration between vendor tools in a way that is unique to the specific landscape of each network. This integration is unlikely to be available off the shelf as no two networks have the same combination of partners and vendors.

4.1.3. Tier 2

AI platform components that can be used to create new telco AI solutions.

To create AI solutions, it's necessary to have the compute, data and software tools to build and run solutions. It ranges from using vast amounts of compute and data to create a new foundation model, such as a large telecom model, or consuming platform-as-a-service components to create a retrieval augmented generation pipeline to extract insights from existing enterprise data.

There's a complex ecosystem of AI development tools that have their own strengths and weaknesses, depending on whether they are needed for enterprise, R&D or GPUaaS, and if they're run in hyperscaler cloud, on private cloud, or part of an AI factory.

4.1.4. Tier 3

Physical hardware that enables AI solutions.

Hardware manufacturers, equipment providers and infrastructure asset owners are embracing AI, from the specialized silicon to high-bandwidth fiber and mobile networks using AI to optimize performance, design and transmission techniques.

MECs are an example of where MNOs are invested in this tier and must make key investment decisions as to whether to equip edge locations with AI enabled hardware.

In the N*AI*ture example, MECs are used as part of enterprise agents, ISAC for automotive and network control agents. Other use cases for AI-focused MEC sites include local AI for hyper-contextualized and adaptive smart cities, personalized predictive vehicle health and agentic cognitive twins for supply chains. An example of the latter is using a cognitive twin to optimize electric vehicle routes and appointments, ensuring that charging is not blocking productivity.

4.2. The internal network improvement opportunity

AI will fundamentally reshape internal telco operations, driving network improvements, operational enhancement, augmented customer experience and new capabilities to customers.

The shift from algorithmic management to agentic AI, through to fully autonomous intent-based management and service provision, will reduce the operational burden while accelerating time to market for new capacity, capabilities and offerings.

	Implementable now	**Near term** 2-5 yrs	**Future vision** 5+ yrs
Network improvement Performance, cost and energy savings	• Initial predictive maintenance • GenAI for Network Analysis, Anomaly Detection, Root Cause Analysis, Energy Analysis • Network APIs (Internal Initial)	• Comprehensive Agentic AI RAN • GenAI for Network Design • Integrated Sensing & Comms • Transparent Antennas • Advanced digital beamforming	• Adaptive, Autonomous Network Reconfigurable Intelligent Surfaces, Pervasive Neural Receivers • Multi-altitude meshing • Vehicle-as-tower
Network operations Enhancement of operational processes with AI	• TM Forum L3 automation: Network Monitoring, GenAI for IT Ticketing Support • Basic Automated Network Config Mgt using Large Telecom Models	• TM Forum L4: Zero-Touch Service Assurance, Automated SLA/SLO/SLI management • Agentic + Physical AI for field tech support, remote inspection	• TM Forum L5: Fully autonomous lifecycle, Agentic & Physical AI for complex tasks • Automated Repair & Maintenance
Customer experience Improved customer network experience	• AI Agents monitoring of user network experience using standardised data • GenAI-powered sales/support agents and generative search	• Experience agents and RAN agents cooperating to provide optimised network experience • Personalised network agents supporting network assistance	• AI-Native hyper personalized network based on user intent. • Seamless Quality of Experience across a mesh of terrestrial and non-terrestrial networks
New capabilities Additional network abilities and services	• AI Factory & AI RAN based AI/ GPU as a service • AI-powered management and optimization for early HAPS/NTN services	• Advanced AI-managed private networks and monetizable external Network APIs • Initial ISAC data and high-precision positioning services	• Evolve into an AI ecosystem enabler for various verticals • Fully autonomous composable B2B solutions

Figure 8: Network high level opportunity roadmap

4.3. The B2B opportunity

The ability for businesses to create their own services offerings in an automated manner, creates new services based on network APIs; these APIs, in turn leverage the hierarchy of intelligence provided by an AI powered telco and enables previously unforeseen opportunities.

Transforming the network from a provider of network connectivity to a provider of intelligence enables MNOs to build on their inherent physical advantages to become the intelligence provider of choice for enterprises, especially when combined with private 5G offerings.

	Implementable now	**Near term** 2-5 yrs	**Future vision** 5+ yrs
Customer experience Improved customer network experience	• GenAI for B2B Support: RFP Answer Generators, B2B Chatbots, Document Processing for Contracts, Analytics & Data Query Generation for B2B data • Network APIs (Internal for B2B service provisioning)	• Automated SLOs & SLAs with AI Visibility • Agentic Service Request Delivery (standard B2B services) • GenAI for B2B Strategic Insights, Contract Management • Customer Service Agent Augmentation for complex B2B support	• Fully Intent-Driven B2B Services: Define outcomes, AI network delivers • Co-managed AI-driven Network Ops Portals • Zero-Touch Service Assurance for complex B2B solutions • Digital Human/Support team for high-value B2B interactions
New capabilities Additional network abilities and services	• LLM as a Service (B2B) • Basic Private Network Optimization (AI-assisted) • AI-Accelerated Edge Platforms (Fog Compute/AI Factory) Sovereign Cloud & LLM (GenAI)	• Advanced Private Networks (AI-Managed with xApps, AI-RAN) • AI Factory & Fog Compute enabling AI on RAN • Monetizable Network APIs (External for B2B) • Integration of HAPs & B2B NTN solutions • Autonomous Nav of UAVs/ Vehicles • Highly Distributed Federated Learning for B2B analytics	• Telco as “”AI Ecosystem Enabler”” for Verticals (ISAC data, NTN logistics, AI Slicing) • Autonomously Composable B2B Solutions via Agentic & Physical AI & advanced Network APIs • Globally Orchestrated Autonomous Supply Chains • Dynamic Spectrum Allocations for specialized B2B industrial networks

Figure 9: B2B high level opportunity roadmap

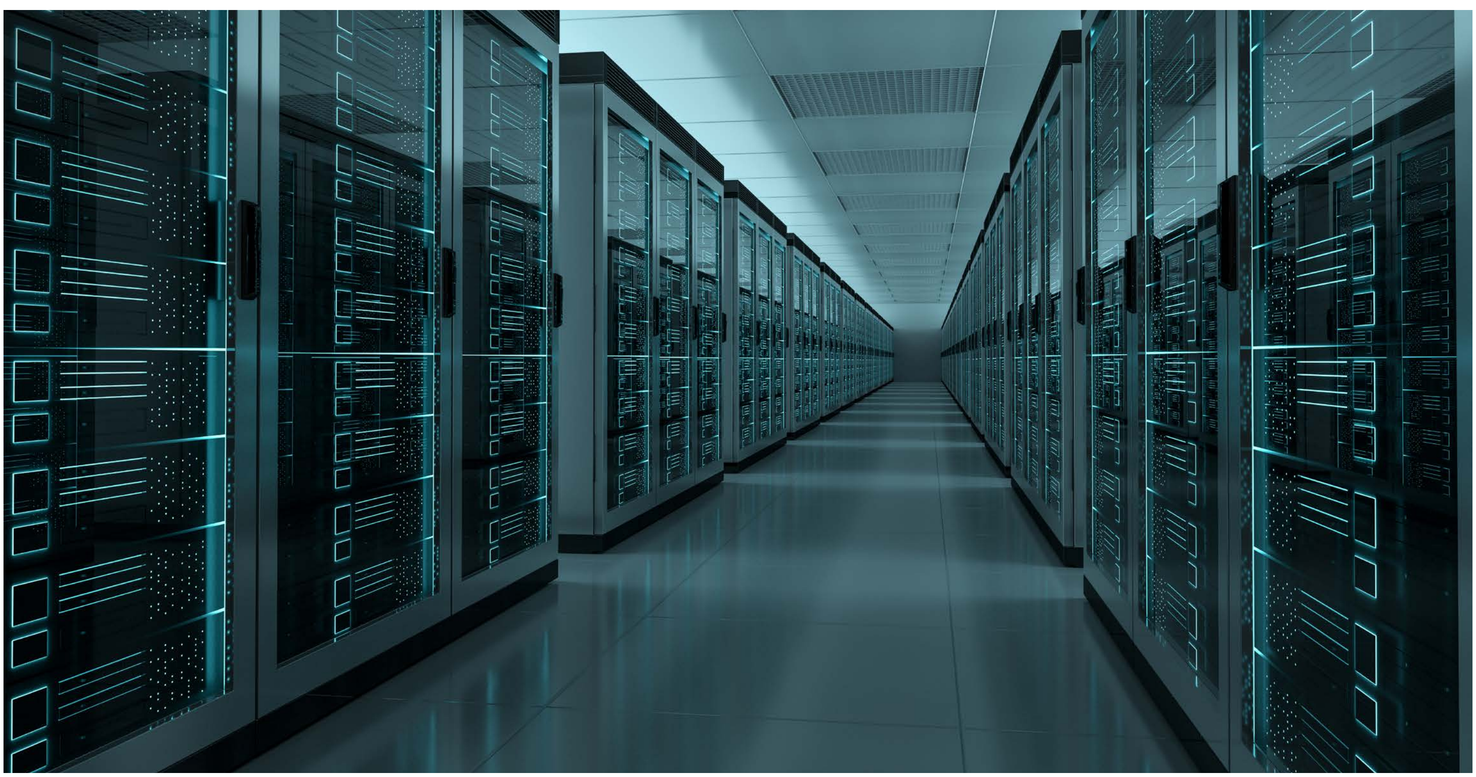

4.4. The B2C opportunity

AI powered networks can deliver on the promise of intent-based networks, with the network self-adjusting to provide hyper-personalized connectivity, shifting dynamically to meet the precise needs of customers. The augmentation of on-device AI with low latency edge intelligence and high-powered AI factory intelligence makes possible experiences that are unimaginable today.

We can imagine a world where on-device AI agents monitor and locally process biometric health data, edge intelligence maintains a local, personalized intelligent view of health, reviewing individual health data on an hourly or daily basis, supported by long running digital twins of individuals maintained by health clinics within an AI factory.

It's clear that there's a wide range of possible use cases for end consumers arising from a future AI-powered telco, especially by enabling low latency, high power intelligence at the edge.

	Implementable now	**Near term** 2-5 yrs	**Future vision** 5+ yrs
Customer experience Improved customer network experience	• Agentic & RAG powered Customer Sales/Support Chatbot/Voicebot, Generative Search & Support FAQ, Customer Call Insights, Personalized Content, Customer Call Live Translation • Network APIs (Indirect CX via 3rd party apps)	• Enhanced GenAI CX: Customer Service Agent Augmentation, AI Phone/Assistants (GenAI) • Proactive AI-driven issue notification • Highly Distributed Federated Learning for privacy-preserving personalized experiences via federated fine tuned SLMs • Brand Avatar (Early)	• Hyper-Personalized "Personal Connectivity Concierge": Agentic & Physical AI understanding intent, curating experiences • Digital Human as primary interface • Proactive, "invisible" problem resolution • Selling "connectivity w/SLA" over selling "connections/SIMs"
New capabilities Additional network abilities and services	• Basic Value Added Services (B2C) using AI • Automatic (real-time) P2P language translation • Roaming onto Satcomms networks (user-aware)	• Advanced AI Phone/Assistants (GenAI/Web 3.0) • Personalized Network Slices (Pilots) • Integration of HAPs & Multiple Paths for improved ubiquitous connectivity experience • Early ISAC-enabled consumer services (e.g., enhanced local info, safety alerts) • Advanced V2X (BVLOS) • Automatic (user-unaware) roaming onto Satcomms networks	• Full Hyper-Personalized Connectivity via Agentic & Physical AI • Immersive AR/VR/XR via edge-rendered, AI-optimized network • Autonomous Nav of UAVs/ Vehicles (e.g., personal drone deliveries, autonomous taxis) • Dynamic Spectrum Allocations enabling new types of flexible B2C services • Highly Distributed Federated Learning for deeply personalized new services

Figure 10: B2C high level opportunity roadmap

5. The cost of inaction

The question is not whether AI will have an impact, but one of timing. Should we be leveraging AI to improve (a) the network itself, (b) services towards enterprises and (c) services toward consumers? Or do we wait to see what comes outside of the network, affecting their experiences from services to OTT network operations? In essence giving up the ability to monetize and direct on these impacts.

Telecoms globally are in a red-ocean market, which suffers from a constant amount of churn (MNOs grabbing customers from each other) and high customer acquisition costs.

The global CAGR for telecoms is around 3.6%, as stated in a recent report by STL partners.[5] There are pockets of growth, however in services to enterprises, which is a small market ($1.2bn in 2025) but is expected to grow to $21bn in 2032 at a 49% CAGR.

The MNOs, and the network equipment providers (NEPs) are together expected to capture 85% of this market, though interestingly, the NEPs serving the enterprises is the larger chunk (55%). The implication is that the MNO isn't the top choice in servicing the private connectivity needs for the enterprise.

As such, it stands to reason that there needs to be an urgency to change behavior. The supply chains at large Telcos were all but locked in by large vendors and any new innovation required the participation of both the MNO and the NEP.

AI represents one such shift in the entire ecosystem where the legacy NEPs are also at an equal footing as new entrants to the market.

AI is going to yield OTT and significant enterprise value-add market opportunities by providing intelligent network services. Currently MNOs have a positional advantage, however market disruption is coming, do we want to be the disruptors or followers?

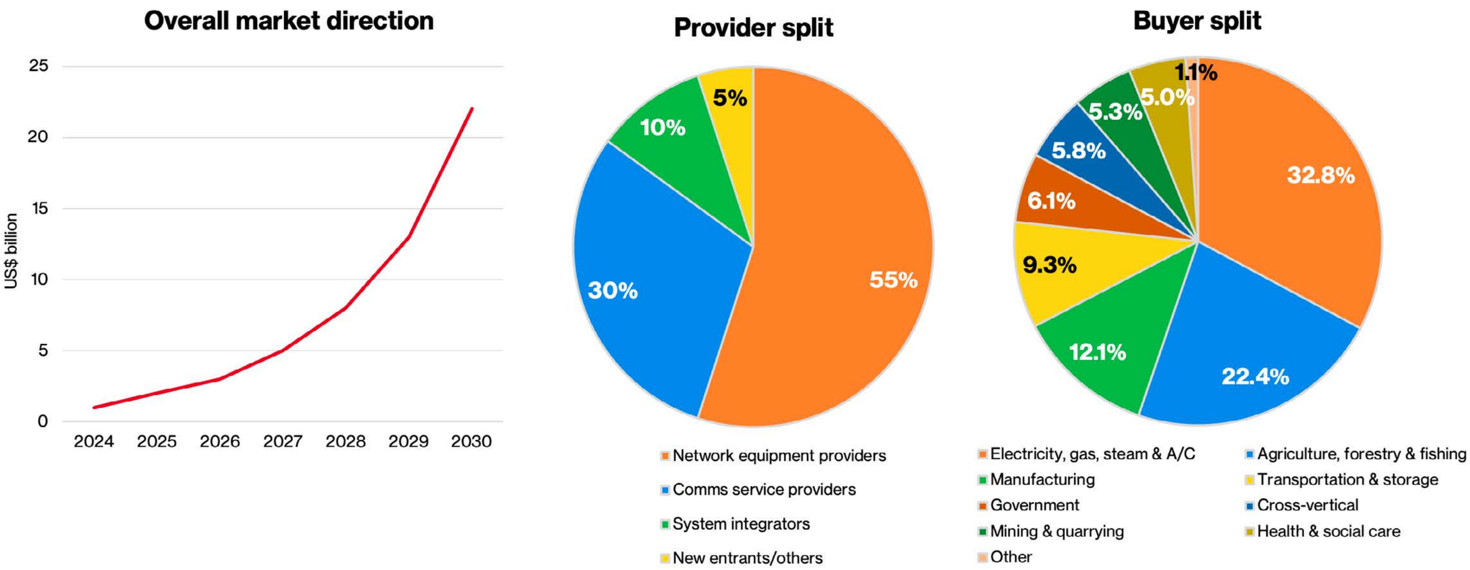

Figure 11: Enterprise private connectivity growth trend, STL Partners

6. AI governance: Governance-by-Design rooted in risk-based principles

The integration of AI across telecommunications networks will present transformative opportunities but concurrently introduce complex risks that require structured oversight. AI is still an emerging field, characterized by rapid technological advancement, and so its deployment should be guided by meaningful governance principles (which are also evolving and should be frequently reassessed).

Why governance matters

AI governance encompasses the process and policy considerations necessary to manage and monitor the development, deployment and use of AI systems. This includes establishing clear lines of accountability, defining organizational policies for AI development and deployment and enabling cross-functional processes and collaboration among engineering, legal, security and ethics teams. Importantly, the organization should employ an appropriate risk framework to address the potential impacts of using AI, informed by the following:

Governance-by-design: AI systems, particularly those that manage core network functions, should be engineered to be safe (preventing unintended harm to people or infrastructure) and secure (resisting malicious interference, data poisoning, or model theft). AI in telecommunications impacts everything from network reliability to user data privacy. Governance builds trust by establishing guardrails against security threats, and system failures, and unintended consequences.

Mitigate systemic risk: Automated network optimization, for instance, could lead to widespread outages if an AI model encounters unforeseen data or operational conditions. Governance ensures mechanisms for human oversight and auditability.

Uphold responsible use of AI: Use of AI systems should be subject to appropriate oversight to reduce potential bias and avoid discriminatory outcomes, to evaluate the system's ability to generate accurate and reliable outputs, and to promote transparency and explain how AI outcomes are reached.

A risk-based approach

AI is based on probabilistic technologies and techniques, which creates an inherent risk in operations if not mitigated appropriately. A cornerstone of responsible AI implementation is a risk-based framework. This principle dictates that the level of governance, rigor and control applied to an AI system should be proportional to the potential impact of its failure or misuse. The organizational mandate for security and safety is supported by leading global and U.S. frameworks, such as the NIST AI Risk Management Framework, the EU AI Act and the OECD AI Principles.

The confluence of these standards highlights a shared global perspective that, due to the nature of AI as a rapidly evolving and complex technology, proactive, systematic risk management and a Governance-by-Design philosophy are the most effective ways to encourage innovation while safeguarding critical interests.

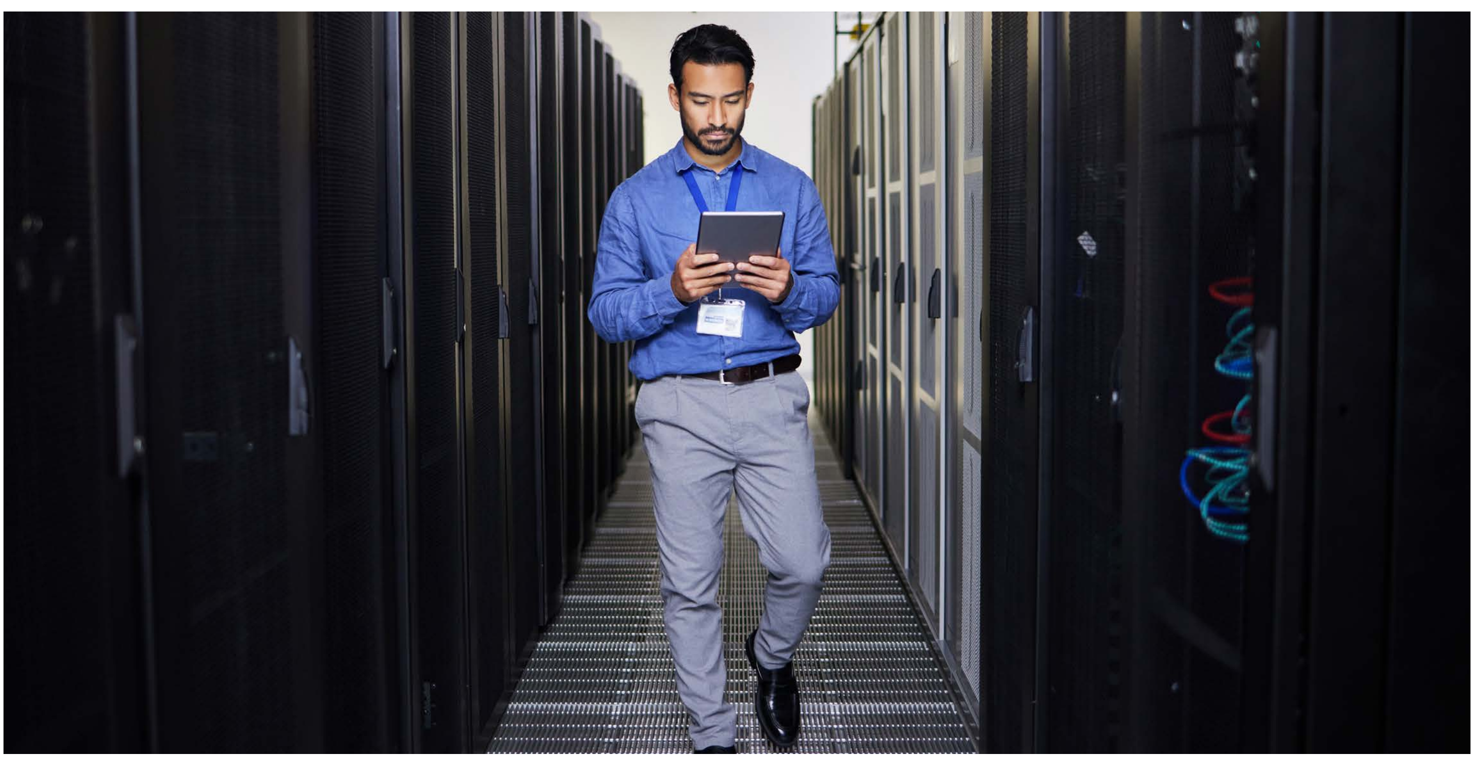

7. Early proof and call to action

The technologies mentioned in this paper are rooted in real use cases where AI has been applied at various points in the network for deriving value - network autonomy, lower power consumption, better end-user experience or enabling a new capability. All of them have been implemented in the real world, some of them in live production systems, some of them at the experimental stage.

Figure 12 showcases seven different early proof points in the network, each with its own application of AI, and the resulting benefit:

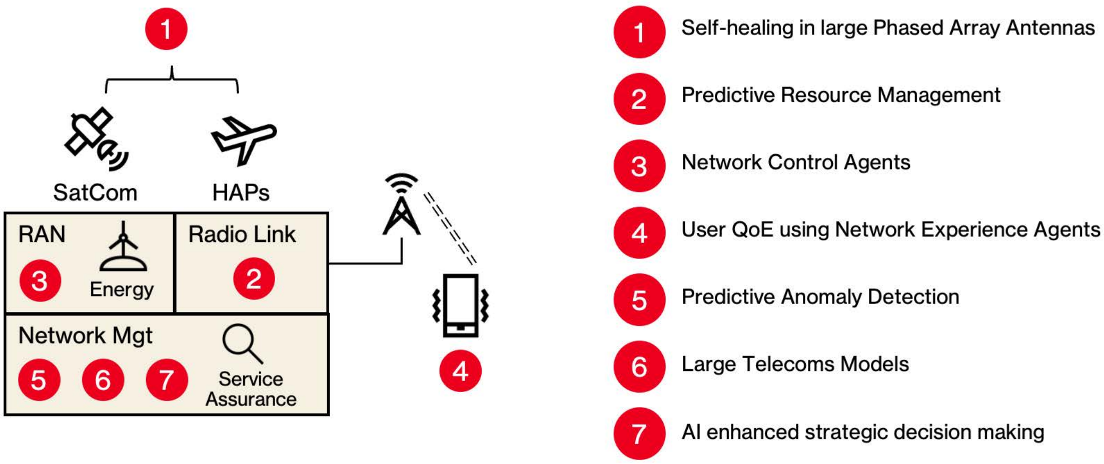

Figure 12: Early AI proof points, Cambridge Consultants

7.1.1. Self-healing in large phased array antennas

Large phased array antenna technology is vital for both satellite and future terrestrial systems. Cambridge Consultants has successfully used AI to autonomously recalibrate and compensate for the failure of individual elements in a satellite antenna. Individual elements of a phase array antenna were simulated to fail, and the weights on the other elements of the antenna were adjusted to compensate for the failure. We discovered that large, multi-element antennas can be successfully trained to self-compensate against the first few failures without service degradation.

7.1.2. Predictive anomaly detections

The first step towards autonomy is prediction. Cambridge Consultants has been able to deploy both supervised and unsupervised AI models on live networks to detect root causes of failures. Over a thousand network metrics were used to create an unsupervised "heat-map" of the metrics that correspond to a network failure, and this resulted in the network operations team learning about previously unnoticed metrics that led to unplanned outages.

7.1.3. Predictive resource management

The future network must efficiently manage resources while optimizing for perceived quality. In Cambridge Consultants' project, a cell-load predictor was developed that uses prior 12 hours of data within geographical squares, to predict the next nine 10-minute windows within the same areas. The results are then used to power down capacity layers of a cell outside of peak loading. In actual results, the system came to within 1% of a system with perfect (post-facto) knowledge.

7.1.4. User QoE network agents

The challenge with true user-experience is that networks are focused on parameters such as throughput and latency, while the user experience has to do with the real perception of quality at the UE. In this project, Cambridge Consultants created an AI agent that infers the user experience using standard RAN parameters. The agent runs deep within the telecom network, making it possible to do user-centric network optimization without direct access to user equipment.

7.1.5. Network control agents

The vision of a network managed by collaborative agents is essential to a fully autonomous network. MARL has been employed in simulations to create a RAN power usage model that yields a 20%+ reduction in power savings with an accompanying improvement in customer experience. Crucially, these agents are lightweight enough to run on existing infrastructure, proving that advanced autonomy does not require a complete overhaul of network compute. MARL is an agentic technology that is used by organizations such as Waymo for their self-driving car technology.[14] MARL based AI agents developed for the UK's Defense and Science Technology Laboratory to protect autonomous vehicles in hostile environments can be run on a Raspberry Pi.[10]

7.1.6. Large telecom models (LTMs)

The “brain” of the hierarchical intelligence model is already being deployed by pioneers in the industry. Softbank, for instance, is using its proprietary LTM to automate the initial configuration of new cell sites with 94% accuracy and to propose optimized configurations for special events. This validates the feasibility of using centralized, high-level AI to manage complex, distributed network operations.[15]

7.1.7. AI enhanced strategic decision-making

The strategic value of AI is also being realized. A large European operator is leveraging AI to create a smart Capex investment system, enabling it to analyze investment approaches, technology readiness and potential ROI. This is considered one of their highest-return AI initiatives, confirming that AI’s impact extends from operational efficiency directly to financial performance.

7.2. Getting involved

If you’d like to understand or dive deeper about the ideas within this white paper, please reach out to the authors, whose contact details are in the following section.

8. Meet the authors

Cambridge Consultants

Frank Long
Associate Director,
Intelligent Services Business Unit
frank.long@cambridgeconsultants.com

Frank has a background in strategic consultancy, mobile handset electronic design, cloud and AI. He has deep telcoms experience, having led architecture & design teams for some of the UK's largest networks, led the creation of one of Europe's first AI factories and provided strategic technical leadership for large global satcoms, and automotive businesses.

Frank holds an ME in Electronic Engineering from University College Cork, Ireland and an MSc in International Business Economics from Oxford Brookes University, UK.

Amy Stephens
Senior Consultant, AI,
Intelligent Services Business Unit
amy.stephens@cambridgeconsultants.com

Amy is a multidisciplinary technical lead with over 20 years of experience in technology development, spanning embedded software, signal processing, communications and artificial intelligence. She works on projects ranging from early-stage research through to real-world deployments, often serving as a system architect.

Her focus is applying AI and machine learning to advanced telecoms challenges, including neural receivers and AI-driven beamforming.

Amy loves applying a collaborative problem-solving approach to taking complex requirements and translating them into innovative, deep-tech solutions. She holds a Master of Engineering in Electrical and Information Sciences from the University of Cambridge, UK.

Aashu Virmani
Senior Vice President of Business Development, Telecoms Business Unit
aashu.virmani@cambridgeconsultants.com

Aashu is passionate about simplifying the complex and articulating the business value of an idea. He has been a Business Development, M&A and Product Strategy Executive for over 20 years. Formally trained in data science, he has invested over two decades in the world of communications.

He's responsible for finding areas of deep tech innovation within the communications space and allowing CC's customers to adopt such innovation to create a sustainable and competitive advantage.

Aashu has a Ph.D. in data-mining/ML from Rutgers University and a Bachelor's in computer science from Indian Institute of Technology.

Verizon

Jyoti Sharma
Associate Director of Network Planning, Technology, Architecture and Planning

jyoti.sharma1@verizon.com

Jyoti is responsible for defining the architecture and strategy for Verizon's Network, Enterprise and Consumer products and services. She has over 25 years of industry experience in technology and telecommunication roles, including wireless systems engineering, system performance and new technology introduction. Jyoti leads a team defining technology strategy for emerging technologies and represents Verizon on the 5G Automotive Association (5GAA) Board.

Jyoti is a senior member of IEEE and currently serves as the Chair of Women In Engineering at the IEEE North Jersey Section. She has also worked as an Adjunct Professor at Fairleigh Dickinson University and DeVry University. Jyoti earned a Ph.D. in Electrical Engineering from the Indian Institute of Technology, Delhi, India, a Master's in Telecommunications from the Asian Institute of Technology, Bangkok, Thailand and a Bachelor's in Electronics & Communications Engineering from the Delhi Institute of Technology, Delhi, India.

Hao Wang
Distinguished Engineer of Network Planning, Emerging Technology,

hao.wang@verizon.com

Hao is responsible for the strategy and planning of Verizon's private network and its vertical applications. He possesses extensive industry experience in telecommunications, encompassing Wireless Networks, Data Science, CV2X, Cloud Computing and Artificial Intelligence.

He has earned a Master of Engineering degree in Electrical and Computer Engineering from Carnegie Mellon University

Chris Bramble
Senior Consultant, AI,
Distinguished Engineer of Network Planning, Emerging Technology,

chris.bramble@verizonwireless.com

Chris is enthusiastic about the latest and greatest technology and enabling it for everyone's utilization. This helps with his responsibilities for assessing the impact of future technologies on today's network and aligning strategy across multiple teams.

He has over 10 years of experience in technology development, with work in device virtualization & testing, network management and network application development & deployment. He has worked on the Emerging Technology team in Private Wireless Networks applications, Non-Terrestrial Networking, Vehicle-to-Everything Communications and other up-and-coming technologies.

His current concentration is on the impact of AI both in and on the network, and how Verizon can support these workloads moving forward.

1. Nokia Bell Labs, “Global Network Traffic Report,” [Online]. Available: onestore.nokia.com/asset/213660
2. TMForum, “A Regional Guide to Autonomous Network Progress,” [Online]. Available: inform.tmforum.org/research-and-analysis/reports/a-regional-guide-to-autonomous-networks-progress
3. GSMA, “Growing momentum behind partnerships between telcos and satellite companies,” [Online]. Available: gsma.com/solutions-and-impact/technologies/networks/5g/growing-momentum-behind-partnerships-between-telcos-and-satellite-companies/
4. Parrot, “Another first for ANAFI Ai the 4G/5G connected robotic UAV,” [Online]. Available: parrot.com/en/newsroom/une-autre-premiere-pour-anafi-ai-le-drone-robot-4g5g
5. STL Partners, “Private networks market forecast 2024 – 2030,” [Online]. Available: stlpartners.com/research/private-networks-market-forecast-2024-2030/
6. Cambridge Consultants, “5G’s future is hybrid – the non-terresterial opportunity,” [Online]. Available: cambridgeconsultants.com/wp-content/uploads/2023/11/5Gs-future-is-hybrid-the-non-terrestrial-opportunity-whitepaper.pdf
7. Crampton et al, “An open-access market for global communications,” [Online]. Available: cramton.umd.edu/papers2020-2024/cramton-et-al-open-access-global-communications-market.pdf
8. TMForum, “Autonomous Networks Framework v2.0.0 (IG1218F),” [Online]. Available: tmforum.org/resources/introductory-guide/autonomous-networks-framework-v2-0-0-ig1218f/
9. NVIDIA, “The Telco AI Renaissance Is Here,” [Online]. Available: https://www.nvidia.com/en-us/on-demand/session/gtc25-s72984/
10. Cambridge Consultants, “Pioneering AI cyberattack response for the UK military,” [Online]. Available: cambridgeconsultants.com/project/pioneering-ai-cyberattack-response-for-the-uk-military/
11. Cambridge Consultants, “Seize the future: 5 Deep tech areas driving long term growth,” [Online]. Available: cambridgeconsultants.com/five-deep-tech-areas-driving-telecoms-growth/
12. Ericsson, “The next wave of 5G – 3GPP Release 19,” [Online]. Available: ericsson.com/en/blog/2023/12/3gpp-release-19
13. Deepsig, “DeepSig Announces Commercial Release of Industry’s First Neural Receiver Software for Open RAN 5G Networks,” [Online]. Available: deepsig.ai/deepsig-announces-industrys-first-neural-receiver-software-for-open-ran-5g-networks/
14. Waymo, “Improving Agent Behaviors with RL Fine-tuning for Autonomous Driving,” [Online]. Available: waymo.com/research/improving-agent-behaviors-with-rl-fine-tuning-for-autonomous-driving/
15. Softbank, “SoftBank Corp’s Large Telecom Model: Shaping the Future of NetworkOperations with Generative AI,” [Online]. Available: softbank.jp/en/corp/technology/research/story-event/079/?adid=sbn