Strategic Research and Innovation Agenda 2027-2035

European Technology Platform NetworldEurope

"SRIA 2026"

Draft for Public Consultation

2026

Foreword
Invitation for Comments

We are pleased to present NetworldEurope's Strategic Research and Innovation Agenda (SRIA) 2026, offering a forward-looking perspective for the period 2027–2035. This SRIA outlines the key priorities for next-generation networked systems, building on and extending the direction established in the previous SRIA editions of 2022 and 2024. It reflects a shared European ambition to lead the development of resilient, secure and energy-efficient connectivity, tightly integrated with computing paradigms such as AI-native infrastructures, secure distributed intelligence and end-to-end system assurances.

Europe has long been at the forefront of telecommunications research and innovation, shaping global standards, fostering open ecosystems, and advancing technologies that strengthen a trusted and inclusive digital society. As networks evolve into the backbone of an AI-driven, sustainable and sovereign digital infrastructure, Europe's role in setting direction and values is more important than ever. Networks have become the foundation on which our digital society grows.

This SRIA is the result of a broad and open collaboration across Europe's network and telecommunications research and innovation landscape. The process of building this document started with the Visions for Future Communications Summit in Lisbon, in 2025, where we had challenging presentations from multiple stakeholders. This was followed by work led by the Expert Group of Networld Europe, collecting a large number of contributions from volunteer experts. We sincerely thank all contributors from industry, academia, SMEs and public institutions, as well as European and national stakeholders, for their continued engagement and commitment. It is this collective effort that enables Europe to remain a driving force in global telecommunications.

At this stage, we invite the community to provide comments on this document, offering a final opportunity for its views to be reflected in this collective effort. The process for commenting is detailed in the SRIA landing webpage in NetworldEurope site.

We hope that this SRIA will help guide the alignment of research, innovation, and policy in the coming years, consolidating and strengthening Europe's position as a global leader in next-generation networks and systems. We are confident that European technologists and policymakers alike will be able to develop the informed policies and activities needed to underpin the human-centred European society of the future.

Rui Aguiar   Jyrki Huusko   Artur Hecker

Table of Contents

Executive Summary

The European Strategic Research and Innovation Agenda (SRIA) 2026, developed by NetworldEurope, presents networks as a strategic foundation for Europe's digital economy.

The importance of communication networks extends beyond connectivity itself. In our society, the networking challenges have moved from connectivity into the realization of a substrate supporting an integrated, intelligent system underscoring our values. Networks are evolving into a Connected Digital Fabric that integrates communication, computing, artificial intelligence, software, data, devices, microelectronics, photonics, trustworthiness, and emerging quantum technologies into a unified operational continuum. This fabric strengthens Europe's digital ecosystem and enables the convergence of AI, cloud-edge computing, microelectronics, and quantum technologies. Europe's ability to create value from these domains depends on leadership in the infrastructures that connect and secure them.

The coming decade will see networks become increasingly AI-native, autonomous, programmable, and software-defined. Distributed intelligence, digital twins, cloud-edge continuums, integrated terrestrial and non-terrestrial infrastructures, sensing and communication convergence, and quantum technologies will reshape system design and operation. Security, trustworthiness, resilience, sustainability, and governance must become intrinsic design properties.

Four strategic priorities emerge across all research areas: strengthening European technological sovereignty, ensuring trustworthy and resilient infrastructures by design, advancing sustainability and energy efficiency, and translating research excellence into industrial competitiveness and value creation. To secure strategic autonomy and avoid dependency in an increasingly commoditized value chain, Europe must sustain strong, focused research in the networking layer. A dedicated network and connectivity research domain is essential to ensure that Europe retains control over its digital infrastructure, enabling system-level integration, resilience, and sovereignty rather than becoming primarily a consumer of externally developed technologies.

The SRIA 2026 will provide both a coherent evolution path for next-generation networks and a framework for exploring disruptive opportunities beyond the state of the art. It is organised in four tightly interconnected chapters that together represent an holistic view for the communication systems. These chapters include Systems Views and Architecture Orientations, Softwarisation and Intelligence in Networks, Protocols and Communication Paradigms, and Physical Layer Innovations. Together, these chapters address the future communication and computing infrastructures evolution from system design and operational intelligence to protocols, software platforms, and enabling communication technologies. The final Transversal and Disruptive Technologies chapter explores longer-term, higher-risk research directions, including emerging paradigms and breakthrough technologies that may redefine communication systems beyond current roadmaps.

The SRIA 2026, developed based in the previous work inside SRIA 2024, was developed through a collaborative effort involving 142 contributors representing 74 organizations from industry, academia, and research institutes, reflecting a broad European consensus on future communication and computing research priorities.

Ultimately, Europe's future competitiveness, security, and digital sovereignty will depend on its ability to design, control, and continuously evolve the Connected Digital Fabric. The SRIA provides a common framework to support that ambition and maintain Europe's leadership in trustworthy, sustainable, resilient, and sovereign communication and computing systems beyond 2030, aligned with global standardisation, regulation, and experimentation.

1 Introduction

1.1 Purpose and scope

The NetworldEurope's Strategic Research and Innovation Agenda (SRIA) 2026 defines a comprehensive European agenda and highlights the strategic research directions for future communication networks and services as a foundational pillar of Europe's digital, economic, and societal transformation. It responds directly to the evolving European policy landscape, including the conclusions and recommendations of the Draghi report [C1-01] on European competitiveness, the Letta report [C1-02] on the future of the Single Market, and the Heitor's "Align, Act, Accelerate" report [C1-03] on research and innovation policy, but reflects also the current global situation and turmoil emphasizing the resilience, sustainability, trustworthiness, and reliability of holistic communication systems as part of the European critical infrastructures and business landscape.

The SRIA 2026 takes a **system-level perspective of networks**, treating them not merely as connectivity infrastructure but as integrated, programmable, and intelligent systems that merge communication, computing, data, and control. Networks are becoming foundational platforms where these domains converge to deliver critical services across every sector of the economy and society. The ecosystems built on today's networks underpin virtually all modern innovation. Networks also serve as the vehicle for integrating emerging technologies — quantum communication, photonic networking, advanced sensing — into scalable, operational systems, translating their economic and societal value into reality through end-to-end connectivity, orchestration, and control. This positions network research as the system-level discipline that unites communication, computing, intelligence, and emerging physical technologies into infrastructures that are secure, efficient, reliable, and resilient.

These properties emerge as fundamental requirements also from the end-user perspective, as both private and commercial users ultimately depend on the ability of the system to consistently meet the expected quality of service (QoS) and experience (QoE). Such requirements span the full end-to-end chain, where any deviation in performance, availability, or predictability will directly affect perceived service quality. For Europe, the existence of a sustainable, dynamic and innovative communications sector, is a fundamental requirement for preserving the quality of life of the everyday citizen.

This SRIA document positions network research within the European strategic context. It makes the case for a **distinct and sustained research domain focused on networks**, and explains how the role of networks as strategic infrastructure is evolving. A sovereign and strong communications ecosystem is essential for bringing the multidisciplinary competences required for sustaining and evolving an advanced, safe and human-centered digital society. The document also lays out the benefits of continued research even as the landscape becomes increasingly commoditized, and serves as a bridge to the SRIA's structure and research pillars. Above all, it outlines the **key research challenges, topics, and timelines** for the years ahead, which can be used as a guideline for European policies.

The SRIA provides a comprehensive framework of research directions and topics as seen by the network research community, helping stakeholders prioritize in an informed way. It does not prescribe fixed priorities but enables the identification of "must-win" areas in response to evolving technological, industrial, and policy needs.

1.2 European Competitiveness and Strategic Context

The high-level European policy and expert reports converge on a clear diagnosis; Europe faces at the end of 2024 a big structural competitiveness challenge in the global digital economy. This challenge is characterized by lower productivity growth, fragmented markets, slower innovation scaling, and higher costs of key enabling infrastructures compared to global competitors.

The Draghi report highlights that research and innovation are central drivers of productivity and long-term growth, and that Europe must strengthen its capacity to translate technological capabilities into industrial leadership. The report points to structural weaknesses related to market fragmentation, insufficient investment scale, and gaps in infrastructure deployment. It calls for accelerated investment in broadband, 5G and 6G research, regulatory harmonization, and the creation of a true Single Market for telecoms. At the same time, it emphasizes the need to strengthen European capabilities across the connectivity value chain, from infrastructure to services, and to secure leadership in future network technologies and standards. These recommendations underline the strategic importance of networks as a system-level infrastructure strengthening Europe's digital economy, innovation capacity, and technological sovereignty.

Following Draghi's report, the Commission launched the Competitiveness Compass in January 2025, translating key recommendations into action. Within a year, approximately 90 % of proposed flagship initiatives were preliminary launched and over half delivered, mobilizing more than €1 trillion across innovation, clean tech, and security, including dedicated efforts to strengthen digital infrastructure [C1-04]. As part of this momentum, the Digital Networks Act (DNA) proposal was adopted in January 2026, aiming to modernize the regulatory framework for connectivity, harmonize the single market for network infrastructure, and reduce fragmentation addressing a strategic gap, which was also highlighted by Draghi and Letta.

The DNA proposal and the Commission's digital simplification package mark important steps, but more **explicit emphasis on network infrastructure and continuity of future network system research** is needed. Notwithstanding its much-needed improvements, the communication ecosystem is one of the ICT areas where Europe still retains strategic autonomy, and such competitiveness needs to be maintained (or better: reinforced) to be able to compete with the dynamic ICT ecosystems in China and USA. As a compelling example, China has just identified next-generation networks as one of the six major infrastructures for their upcoming Five-Year Plan. Networks consolidate every digital ecosystem: their programmability, intelligence, and cross-domain integration are critical to unlocking the full benefits of the Compass and related initiatives, and European initiatives.

In parallel, the Letta report emphasizes that the Single Market must evolve to ensure the free flow not only of goods, services, capital and people, but also of data and digital services. This has been addressed during the past years through several concrete Acts, including Digital Markets Act, Digital Services Act and the recent Digital Networks Act proposal. Communication networks and connectivity are positioned as a core component of this vision, effectively constituting a "fifth freedom" underpinning the functioning of the European economy. In a similar voice, the Heitor report "Align, Act, Accelerate" calls for coordinated, large-scale investments and for reducing fragmentation across Member States.

By extending officially defined freedoms into the digital realm, networks provide the operational foundation needed for their effective exercise. They support seamless cross-border activity, strengthen market integration, and enhance resilience in times of disruption. **Networks are not only critical enablers of European economic growth, but also a direct source of value creation,** supported by Europe's strong telecommunications sector and increasingly by other digital-related sectors, including defence and security, given the dual-use nature of network technologies.

Across these policy documents, a set of common priorities emerges: strengthening **European technological sovereignty**, **reducing dependency** on non-European infrastructure providers, achieving **climate neutrality and energy-efficiency targets**, ensuring **resilience** against geopolitical, cyber, and environmental risks, and **accelerating** the deployment and adoption of advanced digital technologies.

In all these areas, networks play a central role. Networks and connectivity are not merely enablers, but are fundamental infrastructures on which economic activity, public services, and societal functions increasingly depend. The old **concept of "public service" is now taking a radically different meaning**, even more relevant than the old telephony concepts.

Europe's future competitiveness, sovereignty, and sustainability depend on our ability to design, control, and evolve our digital infrastructure. In this context, networks and connectivity are not passive service layers but must become increasingly programmable, intelligent, and distributed systems that strengthen all sectors of the economy and society.

The relevance of the networking and connectivity ecosystem cannot be overstated. A strong and focused European effort in network and connectivity research is therefore essential not only to support other technological domains, but to enable Europe to maintain control over its digital future and to translate technological capabilities into economic and societal value.

The economic weight of connectivity makes this dependency consequential. Mobile technologies and services alone contributed approximately €1.2 trillion to EU GDP in 2025, up from €1 trillion in 2024, and are projected to reach €1.6 trillion by 2030, equivalent to 8.2% of EU GDP, while sustaining some 2.4 million direct and indirect jobs [C1-05] [1]. Networks have thus become critical infrastructure in the same sense as energy or transport. Networks are a foundation on which economic activity, public services, and societal functions depend, and whose disruption, or external control, carries systemic consequences.

At the same time, Europe is under-investing in this foundation. GSMA Intelligence estimates that around €475 billion of mobile-network investment is required over the decade to 2035 for Europe to reach best-in-class connectivity, of which only some €270 billion is expected to materialise under current market conditions, leaving an investment gap of approximately €205 billion [C1-06]. The symptoms are already visible. Capital expenditure per connection in Europe stands at roughly €35, against some €70 among global connectivity leaders. "Standalone" 5G (the "full featured 5G") reaches only about 2 % of the European population, compared with around 80% in Greater China [C1-06]. This underinvestment is not confined to deployment, and it propagates upstream. If Europe does not sustain the research and innovation pipeline that feeds these networks, and pursues an aggressive deployment strategy or our own innovations, it will increasingly import the technologies on which this value rests, ceding both economic returns and strategic control.

The direct return on network research is difficult to isolate because most measurable value appears later, at the deployment and service layers. However, research and innovation are the decisive upstream conditions: they determine whether the architectures, standards, and system-level capabilities of future networks are developed in Europe or imported from elsewhere. Sustained network and connectivity research should therefore be understood as both an enabler of downstream value and a critical safeguard against dependency. Strategic autonomy in connectivity is not a retreat into isolation, but a path to global leadership in ideas, excellence, and implementation. Europe must design and control, on its own terms and in line with its values, the connectivity foundation that other regions are already investing heavily to secure.

1.3 The Evolving Role of Networks: Connectivity as the Integrating Layer

The transformation in economy and society fundamentally changes the role of network technologies. Networks have become the integration platform where communication, computing, data, and control converge. They are responsible not only for transporting data, but for enabling real-time coordination of distributed systems, supporting AI execution across edge and cloud, and enforcing security, trust, and policy across heterogeneous environments. As such, networks directly determine the performance, resilience, and sustainability including energy efficiency of Europe's digital ecosystem.

Over the past decades, communication networks have undergone a profound transformation, from the telecommunication infrastructures of the 1980s and 1990s to today's pervasive integration of communication, computing, and storage capabilities across all sectors of the economy and society. Networks now form the backbone of the Digital Society, enabling the operation of digital services and the functioning of virtually every industrial vertical.

---

[1] Using again the Chinese example to confront different expectations, the China Academy of Information and Communications Technology claims that 5G commercialization in China has directly boosted total economic output by about 720B€, and indirectly contributed to 1800B€, with projected digital contributions to reach 8.3% of GDP by 2030.

Developments in the underlying systems, technologies, and fundamental methods have seen vast investments that have reached trillions of Euros. In recent years, cloud and even more recently large-scale AI systems have dwarfed the former drivers of infrastructure renewal, namely the various generations (xG) mobile systems and the Internet itself. This bears the danger to lose focus in where and what investments are required to uphold the tremendous speed we are seeing in the overall industry.

With rapidly increasing investments in artificial intelligence (AI), large-scale computing clusters, and hyperscale cloud infrastructures, there is a need to shift the perspective on the role of network and connectivity infrastructures. The investments for national and international fixed and mobile networks, which connect billions of users and devices to these services, should not be viewed merely as standalone revenue-generating sectors, whose scale may appear small compared to the massive investments in AI, but rather as critical enablers of value creation across the entire digital ecosystem. How can a society fully explore AI capabilities without the interconnect layer that will bring societal realities into the digital realm, where they can then be intelligently analysed?

In this context, the success of investments in AI and cloud fundamentally depends on the availability, performance, and scalability of the underlying connectivity infrastructure. This perspective highlights the need to consider the Digital Society as an integrated system, where communication networks, computing platforms, and data infrastructures co-evolve. In this context, investments in network and connectivity technologies are not only driven by their direct impact, but also by their essential role in **enabling and scaling value creation across the broader ecosystem**. Thus, connectivity is no longer one technology domain among several, but an integrated fabric that brings together communication, compute, to enable intelligence and trustworthiness. The radio, optical, and satellite and non-terrestrial networks merge with the network cloud, and increasingly converge with cloud and edge computing, artificial intelligence and data into a single operational continuum. Without corresponding investments in networks and connectivity, the expected economic value and societal benefits of AI and cloud services cannot be fully realized.

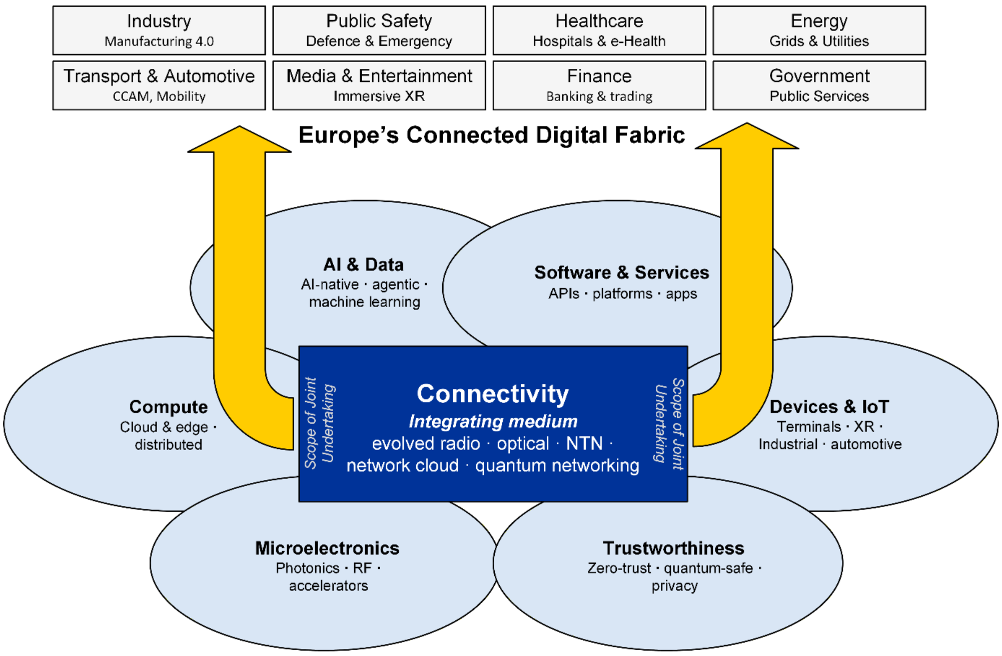

Figure 1-1: Connectivity as the integrating layer of the European digital stack.
Adapted from the 6G-IA position paper for the next MFF (2028–2034).

In this view, connectivity is the foundational layer of the European digital stack. Connectivity is the pillar on which the other digital domains depend and through which their value is realized. None of the digital societal layers can be considered genuinely sovereign if the underlying connectivity fabric is not. This system-level perspective is consistent with, and complementary to, the recognition of networks as a distinct research domain. It is precisely because connectivity integrates the other domains end to end that it cannot be subsumed within any one of them.

On a different, increasingly important, dimension, the role of networks is expanding beyond traditional civilian and commercial applications. The increasing ***dual-use*** **nature of network technologies** means that the same infrastructures and capabilities reinforce both civilian services and security- and defense-related applications. This further elevates the strategic importance of networks, as requirements for robustness, security, reliability, and operational control become more stringent and must be addressed at system level. Ensuring that communication networks remain operational under disruption conditions and can autonomously recover and reconfigure themselves when needed is critically important for maintaining reliability, resilience, and continuity of services. Furthermore, this same dual-use nature strengthens the need for European dominance in this same technologies, reducing potential unwanted externalities of what should be understood as critical for the European sovereignty.

Against this background, strengthening **European research and innovation in networks and connectivity is not optional, but it is a strategic necessity.** High-level European policy analyses consistently identify digital infrastructure as a critical bottleneck for productivity, innovation scaling, and strategic autonomy. At the same time, Europe faces increasing risks of dependency on non-European technologies in key parts of the network stack, including software platforms, control systems, and infrastructure components.

Network and connectivity research constitutes **a structurally distinct and indispensable domain**. While closely related to microelectronics, artificial intelligence, cloud, and energy systems, it cannot be effectively subsumed under these domains or addressed through large, generic partnerships. Network systems are inherently end-to-end, cross-layer, multi-domain, and multi-stakeholder. Their design requires the simultaneous optimization of physical layer technologies, communication protocols, distributed software systems, and real-time control mechanisms under strict constraints on latency, reliability, scalability, and energy consumption. These system-level challenges are not addressed within microelectronics, which focuses on components, nor within artificial intelligence, which focuses on models and data. They require dedicated research that integrates all layers of the network stack into a coherent whole, and increasingly act as the integration layer that binds societal computing needs in a effective tool for societal development.

1.4 Strategically positioning Networks within European R&D Ecosystem

The system-level perspective adopted in this SRIA reinforces the importance of networks and connectivity as a distinct and focused research domain within a broader European innovation landscape. The Digital Society relies on cross-domain solutions that address complex societal and technological challenges: interpenetration of areas such as artificial intelligence, microelectronics, and cloud is critical. However, ensuring that networking aspects receive sufficient focus is key to fully realizing the value of these cross-domain efforts. Furthermore, as networks evolve into programmable, intelligent, and distributed systems, they also become **a primary source of innovation**, where system-level performance, resilience, and efficiency are determined.

A dedicated network and connectivity research domain therefore strengthens, rather than competes with, other domains. It provides the necessary depth and continuity to address core system-level challenges, ensuring that advances in other domains can be effectively integrated, scaled, and deployed. In any large, multi- and interdisciplinary initiatives and research agendas, network and connectivity functions have always been considered as enabling infrastructure, under which all others are positioned. In this way, a focused approach to network research enables Europe to retain and expand its capability to design and operate advanced infrastructures end-to-end, while maximizing the impact of investments across the wider digital ecosystem.

This need is further reinforced by the evolving security landscape. Security can no longer be treated as an application-layer feature or an add-on to existing systems. It must be embedded into the fabric of networks themselves. Modern communication infrastructures must support zero-trust architectures, distributed identity, secure data flows, and resilience against large-scale cyber-physical threats. These properties depend fundamentally on network design, including protocols, architectures, and control mechanisms. As such, security research is inseparable from network and connectivity research and must be addressed as an intrinsic, system-level concern. Secure, trustworthy, and resilient communication systems require integrated, domain-specific research.

Challenging, the increasing commoditization of network components and technologies creates a paradox. While standardization and software-based implementations lower barriers to deployment, they also shift value towards system integration, orchestration, and control, and to the valuable services being deployed over these network infrastructures. Without sustained research, commoditization risks reducing Europe's role to that of a technology consumer. With sustained and focused research, it enables the creation of open, interoperable, and competitive ecosystems in which Europe can lead.

These reasons support the fact that future network and connectivity research must continue to be recognized as a strategic domain, complementary to, but not subsumed by, e.g. microelectronics, artificial intelligence, quantum, infrastructure building or energy. It is the only domain capable of delivering the cross-technology system-level innovation required to ensure that Europe's digital infrastructure is secure, resilient, sustainable, and aligned with European values.

1.4.1 Beyond Traditional Microelectronics

Devices and circuits, especially networked devices have critical and central role in connectivity and network technologies and research. While traditional microelectronics and microelectronics design provides the fundamental building blocks of digital and analog (RF, Optical etc.) systems, it does not address directly the system-level challenges inherent in networks including (e.g.): End-to-end performance optimization (latency, throughput, reliability, adaptability); Resource orchestration across distributed environments; Integration of heterogeneous technologies (such as System in Package, SiP) and domains; Energy efficiency at system scale; and Network-wide security and resilience.

Physical layer innovations, although closely linked to hardware, require co-design with higher-layer protocols and system architecture orientations. The performance of networks, especially those built for AI infrastructures, cannot be derived solely from component-level improvements, although the components and manufacturing are in central and critical role in connectivity and network technologies. In SRIA 2026, devices and circuits are therefore considered as a transversal theme across the entire network and connectivity research domain, rather than being limited to physical layer innovations. Devices, from the network infrastructure elements to end-user equipment and AI infrastructure (such as Data Centers) all enable and constrain system-level capabilities.

Devices must evolve to support new architectural choices, communication paradigms, trustworthiness, and intelligence embedded in networks. At the same time, network-level innovations, including softwarisation, AI-driven control, terrestrial and non-terrestrial integration, end-to-end orchestration, and quantum technologies, are creating new requirements for device capabilities, programmability, energy efficiency, and security. This bidirectional dependency means that device design, network architecture, and system operation must be addressed jointly.

As networks become increasingly distributed, software-defined, and intelligent, devices play a critical role in enabling end-to-end performance, trust, and efficiency. This affects not only system design, but also component technologies, manufacturing processes, and supply chains. Research on devices must therefore be closely integrated with system-level network research, ensuring that hardware, architectures, protocols, and intelligence co-evolve in support of Europe's strategic objectives in connectivity, competitiveness, and sovereignty.

Future network systems cannot be effectively designed if network and connectivity research is treated only as a subset of microelectronics, photonics, or component-level development. Although progress in these areas is essential, it is not sufficient to address the end-to-end, cross-layer, and system-level challenges of modern networks. A dedicated network and connectivity research domain is therefore needed to translate component-level innovation into functional, efficient, secure, and scalable networked systems.

1.4.2 Beyond Generic Artificial Intelligence

Artificial intelligence and softwarisation are key enablers of future networks, but they do not replace dedicated network and connectivity research. Their adoption fundamentally increases the complexity and criticality of networking systems.

Generic AI and software research often assume that data movement, distribution, and execution environments are available and scalable. In practice, these capabilities are constrained and governed by the underlying network. As AI becomes embedded in network infrastructures (“AI for networks” and “networks for AI”), new system-level requirements emerge: real-time distributed inference across heterogeneous nodes; AI integration into control loops; continuous observability; and compliance with European requirements on trust, security, and data governance. These are network-specific challenges that generic AI approaches alone cannot address.

Networks are becoming intelligent, software-driven systems. This creates a co-evolution: AI imposes new demands on networks, while networks require distributed, resource-aware, and trustworthy AI methods. Similarly, softwarisation enables programmability and automation, but depends on network-aware abstractions, orchestration, and control frameworks. This coupling means that networking, AI, and softwarisation cannot be effectively addressed only within broad, technology-centered R&D initiatives. Without dedicated focus, networking risks being treated as an implementation issue, leading to weaker system design, greater complexity, and strategic dependencies.

For these reasons, network and connectivity research must be recognized as a distinct strategic domain, ensuring that AI-enabled, software-driven infrastructures are functional, efficient, secure, sustainable, resilient, and aligned with European values and objectives.

1.4.3 System level integration

The central characteristic of network research is its system-level nature. Future communication infrastructures must be designed and operated as integrated systems, where performance, security, and efficiency result from the interaction of all layers and components. This includes cross-layer design from physical technologies to applications and services; interaction between communication, computing, and control; coordination across heterogeneous technologies, platforms, and domains; and operation under dynamic, uncertain, and large-scale conditions.

These challenges are interdependent and cannot be reduced to isolated component or algorithmic problems. Optimizing individual elements, such as chips, AI models, or software components, does not guarantee correct or efficient system behaviour. Without system-level integration, such optimizations may create inefficiencies, instability, or security vulnerabilities at scale. Network research therefore acts as the integration domain that ensures advances in microelectronics, AI, and software translate into functioning, efficient, and secure systems at European and global scale.

A dedicated network and connectivity research domain, grounded in a system viewpoint, is therefore essential. It provides the framework to address cross-layer trade-offs, end-to-end behaviour, and large-scale operational constraints together. Without this focus, system-level challenges risk remaining unresolved, leading to fragmented solutions, increased complexity, and reduced European control over critical infrastructures.

1.4.4 Strengthening Strategic Autonomy

From a strategic autonomy perspective, control over network and connectivity infrastructures is essential for European sovereignty. Europe’s strength also depends on a balanced research and innovation landscape, where microelectronics, AI, quantum, space, cloud, and energy systems are critical and mutually reinforcing.

As networks become programmable, intelligent, and distributed, control shifts from individual components to the design and operation of the whole system. Advances in adjacent domains only create real-world impact if they are integrated, orchestrated, and operated through network infrastructures.

Dependence on external technologies at any layer of the network stack creates systemic risks for security, resilience, and economic value capture. These risks cannot be addressed solely within other domains, since digital sovereignty ultimately depends on end-to-end control of connectivity and system behaviour.

A dedicated network and connectivity research domain therefore complements other European partnerships by providing the system-level integration needed to turn focused technological excellence into secure, scalable infrastructures. This is especially important given Europe’s (still) global leadership in telecommunications and network technologies, which represents an existing strategic industrial asset.

Maintaining networks and connectivity as a strategic research domain enables Europe to sustain this leadership and maximize investments in AI, microelectronics, and software by integrating them into sovereign, resilient, and high-performance systems.

1.5 Commodity Paradox of Networks and European Value Chain

Communication networks are a general-purpose infrastructure that strengthens the European digital economy. Their performance, cost, and availability directly affect industrial competitiveness, innovation, public services, and social inclusion. Economically, they enable digital markets, productivity gains, and facilitate the scaling of digital services and platforms.

Europe also faces structural challenges in this domain, including higher connectivity costs than some global competitors, fragmented markets and regulation, and limited investment scale and deployment speed. Addressing these issues requires coordinated research and innovation to improve efficiency, reduce costs, and enable new capabilities across the network ecosystem.

As networks become more pervasive, their components and functions increasingly need to be standardized and widely available. This commoditization is driven by open standards and interfaces, software-defined architectures, virtualization, and global supply chains. While it lowers deployment barriers and accelerates adoption, it also reshapes how value is created in networking. At first sight, commoditization may suggest less need for research, as network technologies become widely accessible. In reality, the opposite is true. **As networks become more critical and embedded across society, performance demands intensify, energy efficiency becomes a key constraint, security threats grow, and system integration becomes significantly more complex.**

This direction shifts value from individual components to systems, system-level optimization, orchestration, and innovation. It creates a “commodity paradox”: as network elements become more standardized and accessible, the need for advanced research increases to integrate them into high-performance, secure, resilient, and sustainable systems.

Without sustained investment in network and connectivity research, commoditization becomes a risk, increasing dependence on external solutions (developed with lower production costs) and limiting European actors to lower-value segments of the global value chain. With targeted research and innovation, however, commoditization can support European leadership by enabling open ecosystems, strengthening system integration capabilities, and capturing value through services, platforms, and standards.

By reinforcing system-level expertise and leveraging its strong telecommunications industry, Europe can move beyond adopting externally developed technologies and capture higher-value activities in network design, operation, services, and standardization. The commodity paradox therefore reinforces the need for **a strong European effort in network and**

connectivity research, ensuring that widely available technologies become strategic assets for competitiveness, sovereignty, and long-term value creation.

1.6 Benefits of Network Research for Europe

Investing in network and connectivity research delivers mutually reinforcing benefits aligned with European objectives on competitiveness, sustainability, security, and strategic autonomy. As networks become programmable, intelligent, and distributed, they move beyond connectivity to become central drivers of economic and societal value, while enabling the introduction of technologies such as quantum communication, photonic networking, advanced sensing, and distributed computing.

From a competitiveness perspective, advanced networks support European growth by reducing operation and maintenance costs, improving efficiency, and enabling productivity gains across all economical sectors. They also create direct value through Europe's telecommunications industry, new digital services, and scalable ecosystems. High-performance networks are essential for industrial automation, XR/VR/AR services, data-driven innovation, quantum-enabled services, and next-generation hybrid computing.

In **sustainability**, network research is essential to meet Europe's climate and energy targets. Future infrastructures must improve energy efficiency across all layers, from physical transmission to system operation. Research is needed to enable novel energy-aware architectures, efficient protocols, and optimized management, reducing the carbon footprint of digital infrastructures.

Security and trust must be embedded by design, not added afterwards. Network research will address infrastructures with built-in protection, secure data flows, and trustworthy control systems. This is vital in critical infrastructures, providing privacy and data protection compliance, user trust, and future paradigms such as quantum-safe communication and advanced cryptography.

Resilience and reliability are also fundamental. Research will enable networks that withstand failures, cyber-attacks, and large-scale disruptions while maintaining predictable performance and service availability. This includes terrestrial and non-terrestrial networks integration, dynamic reconfiguration, crisis management, emergency services, and unified end-to-end architectures for heterogeneous technologies.

Network research strengthens Europe's innovation ecosystem by promoting open platforms, interoperable interfaces, and software-driven architectures. This lowers barriers for startups and SMEs, accelerates technology transfer, fosters competition, and supports new business models across the Single Market.

Finally, **network and connectivity research is a cornerstone of European sovereignty**. Control over critical infrastructures reduces dependency on external providers, supports European values and regulation, and allows Europe to influence global standards, strengthen its industrial base, and capture value in the global digital economy.

Taken together, these benefits highlight that investment in network and connectivity research is seen not only a technological priority, but as a strategic necessity. It enables Europe to translate innovation into economic value, to ensure secure and resilient infrastructures, and to maintain leadership in an increasingly competitive and interconnected world.

1.7 SRIA 2026 Overview and Structural Changes

The SRIA is structured around a set of interrelated research pillars that together define the network and connectivity research domain. Rather than representing independent areas of research, the pillars form complementary layers of a unified future communication system and infrastructure view, where progress in one domain directly influences and constrains developments in others. The chapters are defined as follows:

- **System View and Architecture Orientations**, addressing the holistic design, integration, and operation of networked systems across heterogeneous domains and infrastructures

- **Softwarisation and Intelligence in Networks**, enabling programmability, automation, and AI-driven control of network resources and services
- **Protocols and Communication Paradigms**, developing new models for efficient, scalable, and semantic interaction between distributed entities
- **Physical Layer Innovations**, advancing fundamental technologies for capacity, energy efficiency, sensing, and new communication media
- **Transversal and Disruptive Technologies**, exploring and introducing emerging technologies and enabling their integration into future network architectures

These chapters span the full stack of future networks, from physical infrastructure to system-level intelligence and orchestration. Together, they form a coherent research framework for the design and operation of next-generation communication systems, reflecting the complexity of end-to-end system design. This holistic perspective reflects the reality that networks and connectivity are becoming among the most critical infrastructures for Europe's future competitiveness, resilience, and sovereignty.

We emphasize cross-cutting system design parameters that must be addressed consistently across all research areas and throughout the system, rather than as isolated topics. These include sustainability and energy efficiency, security and trustworthiness, resilience and robustness, and performance and scalability..

These dimensions are not independent objectives, but constraints that shape future network design. Energy efficiency must be addressed across physical technologies, protocols, and AI-driven control. Trust and security must be embedded in devices, communication protocols, and orchestration mechanisms. The integration of terrestrial and non-terrestrial networks adds requirements for reliability, latency, and global coverage. Likewise, evolving devices, from antennas, optical systems, and programmable hardware to advanced IoT and cyber-physical systems, both enable and constrain system-level performance.

The SRIA structure highlights the strong interdependencies between research pillars (the SRIA chapters). Advances in physical layer technologies, such as new spectrum usage, photonics, and integrated sensing and communication, directly affect protocol design and system architectures. At the same time, softwarisation and AI-driven control require new communication paradigms, protocol-layer evolution, enhanced observability, and programmable hardware support. The system view provides the framework to coordinate these developments, ensuring that local innovations translate into overall system performance and efficiency.

Each SRIA chapter is inherently interconnected:

- The *System View and Architecture* chapter interfaces with all other chapters, defining the overall system design, requirements, and trade-offs. It integrates AI-driven system behaviour, decision-making, and orchestration logic, while being constrained and enabled by physical capabilities, protocol behaviour, and implementation feasibility.
- The *Softwarisation and Intelligence* chapter both benefits from and enables all other chapters by providing programmability, automation, and AI-driven closed-loop control. It relies on observable and controllable capabilities exposed by the physical layer Innovations, devices and protocols, while enabling adaptive, autonomous system operation across the full architecture.
- The *Protocols and New Communication Paradigms* chapter translates system-level objectives into operational behaviour, enabling coordination of communication, computing, and intelligence. It is increasingly shaped by AI-native operation, to allow support of semantic, goal-oriented, and agentic communication, while being constrained by physical-layer performance limits and system-level requirements.
- The *Physical Layer Innovations* chapter defines the fundamental performance envelope (capacity, latency, reliability, sensing) including advances in hardware, directly shaping feasible architectures, protocols, and AI-driven control loops. At the same time, it is increasingly driven by requirements for programmability, real-time observability, and AI-assisted optimisation.

Each of these technical chapters defines a set of research objectives and topics that reflect the future network and connectivity domain. These topics are intended to provide a **structured and forward-looking view of key research directions when considering the holistic system view for the future network system**, supporting the research community, industry, and policymakers in identifying opportunities for innovation and collaboration.

The SRIA does not aim to prescribe or fix a rigid set of priorities with strict timeline and schedule. The selection of "must-win" areas and investment focus is inherently dynamic and must be determined through continuous interaction between industry needs, market developments, policy objectives, and technological maturity. In this context, the role of the SRIA is to provide a common framework and shared understanding of the research landscape, enabling informed prioritization rather than dictating it. By presenting the full spectrum of research topics and their interdependencies, the SRIA supports the identification of high-impact areas where Europe can build leadership, while maintaining the flexibility required to adapt to evolving technological, economic, and geopolitical conditions.

Consequently, the purpose of the SRIA is to identify meaningful research and development topics for future work programs. Different stakeholders, including industry, Member States, and European institutions, will and should prioritize specific areas based on their strategic objectives, capabilities, and timing. Future European R&D program initiatives should consider these technical views, regardless of the specific funding and governance models pursued.

1.8 Concluding Remarks

Europe stands at a critical juncture in its digital transformation, and its ambitions must be framed within a rapidly evolving global and competitive landscape. Major regions are advancing coordinated long-term strategies for next-generation communication systems, including China's successive Five-Year Plans, such as the upcoming 15th Five-Year Plan (2026–2030), as well as large-scale public–private initiatives in the United States and other regions. These developments underline the strategic importance of networks as critical infrastructure and intensify the global competition for technological leadership.

Networks are **a recognized strategic asset and a key domain of constant innovation** while supporting operational reliability levels unseen in other domains. Networks are not a background utility, nor a immovable asset, nor an outdated technology. The ability to design, deploy, and control advanced network infrastructures will play critical role in determining its future competitiveness, sovereignty, and societal well-being in this context of increasing geopolitical competition. Recognizing and supporting network research as a distinct and essential field is therefore a prerequisite for achieving Europe's long-term objectives.

Nevertheless, considering the global nature of communication system, its development is closely tied to international standardisation efforts in organisations such as 3GPP, ITU, ETSI, and IEEE. European research must both align with and actively contribute to these global processes, ensuring interoperability, reinforcing European influence in global standards, and supporting the competitiveness of European industry while maintaining strategic autonomy in key technologies.

The SRIA 2026 provides the foundation for such efforts, aligning technological and research priorities with policy goals and setting the direction for coordinated, impactful research and innovation in Europe, which can and should be pursued by multiple coordinated initiatives, including a specific and distinct networking and services research initiative.

2 Systems View and Architecture Orientations

2.1 Introduction

Continuing our opening argument of the introduction, this chapter will detail the many technical challenges that systems overall face to become fit-for-purpose in the emergence of the tremendous investment into new AI superclusters, relevant chip and hyperscaler technologies [C2-01], the interworking of which relies on the enabling capabilities of our continuously transformed telecommunication systems to connect to the billions of end users and industry stakeholders alike to bring the expected value to society overall.

From a systems perspective, it is key to understand that such continued technological transformation is important as well as needed, going beyond the societal and economic argument of the introduction. Here, we identify as the next driver beyond the current establishment of AI superclusters that those superclusters will require more than just connections to customer for prompting token input to their quasi-centralised ultra-large language models (LLMs). More concretely, we assert that the very notion of the **interconnect backbone**, which is core focus in building super-scale AI systems, is being pushed out of those superclusters to enable truly collaborative, distributed, sovereign and even more scaled AI capabilities, incorporating edge compute capabilities in a truly seamless continuum of storage, compute, and communication capabilities – which fully aligns with the EC's 3CN vision of **Connected, Collaborative Computing Networks** (cf. EURO-3C Project).

We observe that the current Information and Communication Technology (ICT) landscape is undergoing a structural phase transition from deterministic, stakeholder-controlled systems to an autonomous and probabilistic agent-based paradigm. However, in and for Europe, this paradigm shift, driven solely by stakeholder expectations in the future commercial success of Artificial Intelligence (AI), must respect the societal necessity for digital sovereignty and the existing privacy laws [C2-01]. For this reason, such paradigm shift in system design and operation calls for a deep albeit prudent integration of AI/ML support into the existing infrastructure at various levels of the system, enabled by the programmability and elasticity, continuously integrated in the existing infrastructure as studied and implemented by previous research and already commercially deployed by our community.

As stated in the introduction, in order to get to the forefront of cutting-edge developments, it is furthermore needed to invent and elaborate on new, **truly European AI understanding, concepts and underlying systems**. European AI-native infrastructures must adhere to what Europe requires and should free themselves of the non-technical (e.g., purely commercial) constraints of the current architectural concepts and their resulting artefacts – the research therefore must posit the question as to what the right concepts and abstractions are that would unfold the power of AI and its enabling programmability technologies at all levels of the system and not just at the "ends" of it.

To handle the resulting complexity, we here propose a vision for future AI-driven ICT systems building on our previous work and structured in Autonomous AI Systems. While this vision seamlessly extends our previous SRIA vision [C2-02], it is meant to build the foundation for the future AI-reigned yet self-determined European ICT. Driven by this vision, we identify critical research gaps in the overall problem space.

2.2 Chapter Vision: Sovereign, Resilient, Distributed, Green ICT with AI

The ICT landscape is currently undergoing a structural phase transition. We are moving from a deterministic, data-centralized era toward an **Agentic, Distributed Intelligence paradigm**. This shift is driven by the dual pressures of AI-led commercial competition and the rising necessity for digital sovereignty. As AI and Machine Learning (ML) evolve from specialized tools into the foundational substrate of all system implementations, deeply embedded in the infrastructure, we must address the critical gaps in multi-stakeholder orchestration, verification, and generative infrastructure on a common, **globally connected, locally regulated** and, hence, multi-dimensionally shared resource pool.

2.2.1 The Fundamental Impact of AI/ML

AI/ML is not merely a transient technological revolution; it represents a fundamental advancement in how complexity is managed within ICT. Traditionally, systems were governed by explicit heuristics and rigid architectures custom-built from standard and non-standard modules by respective stakeholders. The integration of LLMs, diffusion models and neural reasoning signifies a move toward probabilistic, adaptive logic that mirrors biological cognitive efficiency. This is not just a technological revolution but a leap in how digital complexity is managed. By moving from explicit, hand-coded heuristics to probabilistic, neural reasoning, we have unlocked the ability to solve non-linear problems at scale.

From a commercial perspective, AI has already become the primary driver of value. It is no longer an optional optimization but a prerequisite for societal progress and global competition. The ability to process vast unstructured datasets and generate high-fidelity output (text, code, video) determines the economic and geopolitical standing of states. Consequently, research in this domain is not just timely—it is a critical requirement for maintaining the functional integrity of modern digital societies.

However, the real, long-term societal impact of AI/ML goes beyond the impressive advances in (virtual or digital) content generation that we witnessed in the past decade. AI/ML could be the game changer in the struggle to make our combined digital-physical world a better place for the society as a whole, meaning that it has the theoretical cognitive skills to strive for cleaner, fairer, safer interworking of all the spheres than any human organization could possibly achieve. In this vein, to guarantee inclusion of all actors with respect for their rights, while maintaining high trustworthiness of systems and sovereign operations of the latter, we foresee a much tighter integration of AI/ML models with the physical world, leading to the emergence of a **truly embodied AI**, which we believe is capable of correctly (i.e., safely, with situational awareness, and with compliance to rules and regulations) steering digital control and data traffic through intelligent infrastructures with the goal to steer millions of intelligent devices correctly through the physical world.

2.2.2 The Dissolution of Legacy ICT Paradigms

As AI/ML matures, existing paradigms for system implementation—such as microservices, cloud-native deployments, and standard application programming interface (API) protocols—face obsolescence in their visible impact. While they will still be used internally, in the near future, we expect systems to be self-assembling, self-optimizing and self-healing diminishing the relevance of the latter.

The scope of intent-based service invocations, the traditional "Request-Response" cycle is being extended by "Intent-Result" cycles from the view of applications. In this new world, the owner, the end-user or, in more modern system, an application, does not program a system to perform a task; rather, they respectively provide a goal to an intelligent substrate that synthesizes the necessary modules, protocols, and APIs in real-time. This requires a radical rethink of software engineering, shifting from code maintenance to the **management of "emergent behaviours" within neural architectures**. While we do not concentrate on the actual capabilities of the models and AI/ML methods, changes are required in the infrastructure to accommodate intelligent, adaptive systems.

More so, the dynamic nature, enabled by distributed agents acting at different levels, parts, planes, layers, etc. of the overall system, and the programmability that fundamentally enables those agents themselves and their unhindered systemic interactions, begs the question on how suitable existing abstractions for ICT systems of the future really will be, going ahead.

Reasoning over and capturing possibly better suiting concepts of building systems that inherently marry computation, communication and storage through a joint, yet distributed AI logic has a tremendous impact on the mechanisms that need to be devised, ultimately boiling down to protocols. Whether it is the alignment with policy (and the needed check of compliance), the desire to identify the objects and purposes of interaction, or to discover the capabilities of entities partaking in the exchange - all will be impacted by how the system is being designed, what are its building blocks, which separations and isolations are dictated by the separation of concerns and the presumed entity relationships (e.g., planes, subsystems, stakeholders), and what abstractions are to be followed.

The tussle that comes with the dissolution of legacy paradigms lies in how existing paradigms ought to be balanced against new paradigms that may be better suited to build a truly AI-native system albeit leading to possibly larger disruption – a perfect area for research to provide some answers.

2.2.3 Sovereignty and the Multi-Stakeholder Vision

Both the operational “autonomy” instrumentation and the emergent behaviours must respect the legal constraints and the interests of the respective stakeholders. In other words, the actions of both the agents (interworking with other agents and usage of tools and models) and the generated service elements must respect the interests of the principal as well as those of the data and infrastructure owners on whose systems they operate. In particular, environmental sustainability should be considered, contracts and laws must be abided to, and respective security policies must be enforced.

In this sense, a central tension exists between the compute-intensive nature of AI, making it a capability of the few, and the fundamental right to self-determination and data sovereignty of every (legal) entity, as stipulated in the European laws. The current concentration of intelligence within a few "hyper-scalers" poses an existential threat to:

1. *National Sovereignty:* Dependence on foreign AI models for critical infrastructure. Such critical infrastructure will not just execute AI workloads, but potentially be controlled from the latter. This is a paradigm change, as it represents an unprecedented potential for loss of control.
2. *Economic Self-Determination:* Corporate entities losing control over proprietary operational data and their own processes. Modern enterprises abide to Management By Systems principles, where the IT system is tightly integrated with various enterprise processes, from human resources, over accounting, sales, stock management to manufacturing. In the sense of the daily operations, a modern enterprise is its IT system.
3. *Democratic Integrity:* The risk of opaque, centralized algorithms influencing public discourse, legal frameworks and decision making. Beyond the dangers of uncontrolled social media and the quest for establishing ground truth in the era of generative AI, this might take a new form with the ongoing progress of digitization and further adoption of eGovernment principles.

To mitigate these risks, the ICT field must move toward **Autonomous AI Systems**. In this vision, data does not have to migrate to a central "brain." Instead, intelligence is localized. **This is especially necessary for Europe, built up from many sovereign Member States.** Hence, we foresee a fragmented but interoperable ecosystem, where data ownership and regulations (e.g., GDPR, EU Data Act, EU AI Act) are respected through strict data, model, training and inference boundary enforcement.

2.2.4 The Rise of Agentic AI Systems

To enable this multi-stakeholder world, we propose a Multi-Domain Agentic AI framework. In this model, the internal implementation and raw data of a stakeholder remain "hidden" within a secure enclave, which acts as an **Autonomous AI System**. Akin to the Internet routing and the autonomous system notion of the latter, the “boundary” of each such Autonomous AI System can be defined by the common **Policy** with respect to data treatment, exchange, storage, resource usage, etc. What is then exported to the world outside of the Autonomous AI System are **Services** (i.e., classical services, AI capabilities and skills), represented and offered through **AI Agents**. These AI Agents act as intelligent, purpose-bound proxies, offering (public) services such as:

- *Content Synthesis:* Real-time generation of specialized text, code, audio, and cinematic-quality video respectively from various multi-modal inputs.
- *System Analysis and Operation:* Continuous autonomous module, protocol and API redesign and optimization of operational systems with situational awareness.

- *Entity Proxying:* AI agents acting on behalf of individuals and commercial entities for role-limited yet potentially complex negotiations, scheduling, project management.
- *Modular Intelligence:* Offering sub-components (e.g., a specific "Translation Skill" or "Negotiation Logic") as a service without exposing the underlying training sets.

We believe this will become the driving framework for the future ICT systems, starting deployment at the commercial layers and slowly propagating over technical parameter negotiation and dynamic interworking layers to become a predominant user-system interaction paradigm. However, we also stress that it will hardly penetrate into each layer and each subsystem, as many tasks, including system-internal tasks such as security enforcement or policy routing, simply require strictly deterministic execution. Such can be implemented and used as classical services.

In our view, the overall ICT system becomes an ecosystem of smartly interworking Autonomous AI Systems, each of which acts as a (dynamic) recollection of any resources submitted to the same policy. We use instruments from previous SRIA documents and tools/capabilities of individual technology domains to maintain such a policy domain in a coherent form in spite of potential geographical spread, heterogeneity of resources, different constraints under which they operate, etc - cf, e.g., RFC8799 [C2-03]. The mature virtualization and isolation technologies already support efficient and secure resource sharing. The logically unified controllability layer allows a service (provider) to glue disparate resource islands to one system supporting smart flexible instantiation and adaptive, elastic and correct execution of any service on resources supporting such execution (Figure 2-1). The resources per se will stem from all system players (typically from mobile network operators, but also from cloud providers, non-public network providers, hyperscalers, etc), and might include generic hardware, virtualized equivalents and specialized devices such as various terminals, sensors and actuators, accelerators, DC servers and storage units, special radio, optical or transport network equipment, particular non-terrestrial network (NTN)-specific resources (e.g., multi-orbit constellation satellites with regenerative payloads, high-altitude platform systems (HAPS) and unmanned aerial vehicle (UAV) nodes, optical and radio inter-satellite links, but also ground gateways) – that all contribute to connectivity, computing, storage and sensing capabilities. In turn, service providers could be mobile system operators, particular service providers, vertical industries, enterprise networks, non-terrestrial network operators, IoT service providers or private subscribers.

The introduced consolidation layer in Figure 2-1 is meant to seamlessly integrate different resources, while keeping the apparent heterogeneity towards services low. Its main role is to map the heterogenous resources to a list of capabilities. In some cases, it can try to “complement” the capabilities missing at the resource layer, e.g., as some form of software capability. In other words, we architecturally embrace the fact that from a point of view of one given system, system resources used for service execution (including the AI agents) might, per se, be provided as services from some other system. In this vision, all system participants are potentially both resource providers and service consumers. In this situation, the properties of the service, including its security policy, resilience, etc. must be, in general, enforced regardless of (or even in spite of lacking) assurances at the resource layer. The layers are also not necessarily clearly cut and can be intertwined, when one mixes different services and systems.

AI agents will help to negotiate contractual modalities, technical parameters and even create missing wrappers or adapter functions to overcome remaining heterogeneity. Various task-dedicated AI agents will further package available functionalities in form of searchable skills and capabilities and will reason and resolve conflicts autonomously. Together with the statement that any resource can be provided as a service, which can per se constitute a resource, such AI Agents therefore will not be merely endpoints, but entities integrated at any depth of the overall system (e.g., different planes, layers, services, individual subsystems, even single nodes). Therefore, in that system view, AI Agents potentially can be integrated in every aspect; the decision on whether and where to integrate them depends on the suitability.

As we recognized in the previous SRIA editions [C2-02], the key challenges that the Smart Green Network Unified controllability and Consolidation layers must solve are: the aspects of control over multiple general-purpose, distributed, network control operating systems; the availability of powerful abstractions from resources to services; new naming schemes for virtualised resources; dynamic and automated discovery; structurally adaptive logical interconnection; multi-criteria routing in networks of different densities; intent-based open APIs and highly configurable policies to control the resource and service access as well as dynamics; isolation of application's execution environments and performances; efficient scheduling of requests to resources; a high degree of automation and support of self-* principles (self-driving networks); and distributed yet trustworthy ML instrumentations.

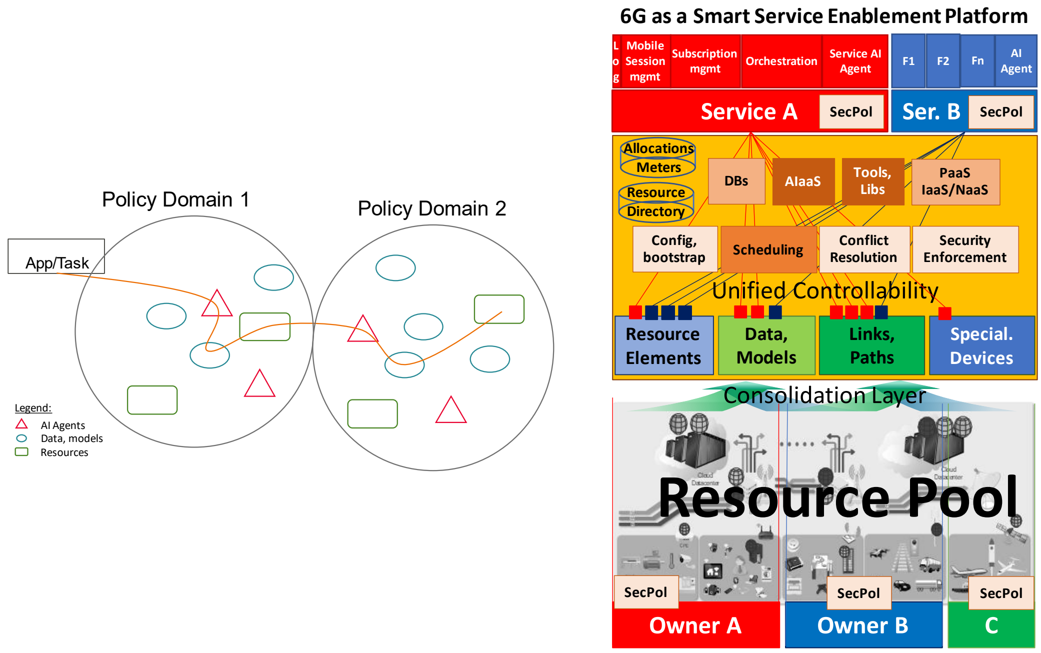

Figure 2-1 – 6G Ecosystem serving as Sovereign Platform for Smart Green Services

Especially in the basic infrastructure operations, time-proven traditional algorithm design approaches will still be required to provide provable and understandable behaviour. However, the Smart Green Networks concept will also integrate AI/ML algorithms, either in the currently typical way, or, ideally, as new, network-suitable, distributed AI/ML, to implement data-driven closed control loops that can enable cognitive and comprehensible system behaviour. The training and validation of such technologies require the availability of cross-technology and cross-sectorial datasets and of some form of agreed test. The networking research community needs to build those datasets, agreeing how they are generated, accepted and accessed.

Overall, citing from the previous SRIA 2024 [C2-02], it is imperative to:

- Allow dynamic pooling of resources from diverse participating (sub)systems, devices and objects;
- Enable seamless allocation and operation of complex IT objects (ranging from atomic modules to complex services, with the possibility of reusing allocated objects) over some selected as well as over panoplies of such objects, i.e. APIs, interfaces, control hooks, etc.;
- Support seamless fusion of logical and physical worlds, by providing, at each layer, data describing the situation of resources with respect to their involvement into the execution of allocated IT objects (e.g., consumption of energy of a session, cumulated

duration of the task execution, memory used by a distributed IT object, volume of transmitted data, number of involved resource elements);

- Integrate autonomics to enable both self-organized, resilient programmability and elastic, correct execution of such IT objects;
- Offer programmable analytics and cooperative machine learning to the service layer through open interfaces.

The points above are discussed in full length in the previous version of SRIA [C2-02] that this new SRIA seamlessly builds upon. In addition to the system architecture considerations in the previous versions of SRIA, partly still not answered by the research community, novel areas required for the realization of this vision are discussed in the next sections of this document:

- *The actual AI agent toolchain*, i.e., questions how such AI Agents can interact in a way to reasonably finish the actual task (section 2.3);
- *Trust questions and security and privacy considerations* specifically required to enable controlled and sovereign operations (section 2.4);
- *Resource efficiency and more general sustainability questions* required to enable Green ICT (section 2.5). This was already treated in the previous version, including the question of resource-aware AI. These points are still mostly relevant, and this version takes some of them directly from the previous version. Some points specifically targeting agentic AI have been added.
- *Device, terminal and other specialized device and subsystem* considerations (section 2.6).

It should also be noticed that it is important that the directions captured in this section be **accompanied by the appropriate economic and policy work** in future research to make way for the envisioned new services that go beyond what is currently supported by access networks, 5G, transport networks and data centres. Broader user involvement and incentives for more local and aware resource usage, notably at the (deep) edge, is highly recommended.

In the following, we outline specific research to better articulate our vision in resulting concepts, paradogms, and embodying architectures, before we discuss several technological "moats" that must be bridged to realize this vision.

2.2.5 Synopsis of Required Research

Although Section 2.2 aims for outlining the vision, driving not just the following technological “moats” but also the following chapters of this overall document, we do foresee necessary research to better articulate our vision through the required architectural considerations that ultimately guide the development of protocols (discussed in Chapter 4), utilising the softwarisation capabilities of the system (discussed in Chapter 3), utilising the physical resources available in the overall distributed system (discussed in Chapter 5).

Research Theme	**Architecture(s) for Sovereign, Distributed, Green ICT with AI**		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Evolving and Emerging Architectures Frameworks, concepts, abstractions, and architectures for truly AI-native systems to guide (i) development of suitable mechanisms, algorithms, and protocols and (ii) to reason over efficacy,	*Mid-term*	1. Overall set of concepts, abstractions, and architectures with roles and stakeholders 2. Set of tools and formalisms to allow for reasoning over architectures in terms of assembling, merging, and compliance checking systems The key outcomes will need evaluation against the compliance with stakeholder needs, particularly on security, privacy, sustainability, and	Paramount as per Section 2.2.2 The set of concepts and abstractions should allow for formulating suitable alternatives to existing abstractions, such as the OSI layer model, to better utilise the power of emerging AI and softwarised systems. The formalisms must enable architecture-level methods to compile systems at runtime (even in an generative manner), reason over their

deployment, and sustainability		sovereignty but also with implementability and socio/techno-economic pathways to adoption.	efficiency and policy-compliance, and guide the (likely AI-driven) generation of suitable enabling protocols.

2.3 AI Agent Toolchain

2.3.1 Agent Identification, Governance and Operations

If the world is to be populated by millions of AI agents acting autonomously on behalf of different types of stakeholders at different levels, we require a robust infrastructure for Agent Deployment/onboarding, Identification, Discovery, Selection, Skill Verification, with possibilities for failure tracking and problem resolution after the incidents. These subjects are partly mentioned here and partly summarized in the section on sovereignty and security below.

- *Agent governance and identification*: When deploying an AI agent for a particular task, what would be an identification scheme for agents? In particular, if we do not limit use cases to agents representing "big services" of today (main sites, which could use, e.g., DNS and URLs just like web services), but also consider agents working on various intents on behalf of end-users? What would be the relationship to user identity, if any? Regardless, how can users delegate their rights to an AI Agent, while limiting their own exposure and responsibility in case of problems? The question of responsibility is paramount and research should develop a methodology, which sets both expectations and boundaries for all (commercial) roles involved into the overall process, e.g., Agent developers, Agent integrators, Agent service providers and Agent users, such that, even before operations start and potential incidents happen, all involved parties can make reasonable and reliable risk estimations.
- *Semantic Decomposition:* How can a requester AI Agent, knowing a number of available serving AI agents with a set of skills, e.g., from Agent discovery or from former knowledge, generally and autonomously decompose any of its tasks to a subset of these agents and their advertised skills, such that the task is solved within functional and extra-functional (quality, robustness, budget, energy, security, etc.) constraints? Important aspects here include the ontological heterogeneity of skill descriptions, in particular for agents stemming from different systems of agents, and questions of runtime selection in case seemingly comparable skills are offered.
- *Data and Output Validity:* How can we ensure that the generative output is not hallucinated but grounded in verifiable fact or the stakeholder's specific data context? How can we be sure that the advertised skill is actually rooted in really available data and/or experiences? How can we estimate the expected quality of what will be generated, e.g., if two different agents claim the same skill but with potentially different prices/energy costs?

Given that AI Agents will often behave as non-deterministic software, it is crucial to be able to find the agent responsible for some potential glitch, wrongdoing or error, which leads to the question of auditability and failure tracing of the whole agentic AI toolchain. While asking for accountability or auditability, it is paramount in this scope to avoid privacy violations of users and providers, which would happen if one were to keep all traces of agents, in particular when they act with user identities.

2.3.2 Generative Infrastructure and the "Generative CDN"

In terms of AI processing capabilities both for training and for inference, mobile and edge devices remain the biggest bottleneck. While some aspects of required research are summarized in the section on devices below, even with another decade of Moore's Law, handheld or IoT devices will likely lack the thermal and battery capacity to run trillion-parameter models locally. Often, such devices will continue forming the interface from the subscriber or the physical world respectively to the digital world, yet will continue struggling to perform massive local generation without catastrophic battery drain. Overcoming this necessitates a hybrid model: Generate Locally, Ask Remotely (incl. Download Pre-cached).

To support this while guaranteeing inclusion, mobility and self-determination, the overall ICT architecture must transition to a hybrid model providing:

1. *Local Inference:* Small-scale models for privacy-sensitive, low-latency, energy-constrained tasks.
2. *Remote Synthesis:* Delegating heavy generation to "Trustworthy Elements" within the larger system/network.

For this reason, we foresee the emergence of Generative CDNs as evolution of current Content Delivery Networks (CDN). Even though GCDN derives its name from the well-known and successful CDN on the Internet, the suggested GCDNs are profoundly different from classical CDNs in many regards:

- Instead of caching static files, the future Generative Content Distribution Networks (GCDN) will "cache" model weights and "pre-cache" personalized generative content e.g. at the network edge (e.g., MEC) or at any suitable position in the overall system, where offloaded AI generative capabilities might be required.
- Also, instead of being positioned at some application layer only for terminals like current CDNs (which drive content distribution towards the end-users), the GCDN should allow general offloading of AI tasks from any system subpart to some other system subpart, presumably more capable in terms of AI processing/inference capabilities for that particular task – this aspect has a tremendous impact on the protocols used to communicate with the GCDN (for inferencing) and within it (for distributing the suitable knowledge across), since we cannot assume a rather narrow-waisted GCDN architecture (and its enabling protocols) that we can find in CDNs around HTTP(S) as its main enabling protocol.
- GCDNs could replace the current caching/storage/content distribution model of existing CDNs by generating the content as opposed to caching it. Besides:
    - while doing so, the GCDNs could automatically generate locally-compliant (per region, per requestor) and regionalized content (translations of voice, images, data formats, etc) and therefore actively help global content providers to respect locally vastly different legal situations (data clearing house functionality) and to attract local customers as considered more suitable by marketing experts. With GCDNs, this includes adjusting content accordingly to local culture, e.g., changing movie plots, punch lines, jokes, etc. seamlessly. Besides, with the evolution of chips and AI compute pipelines and architectures to be considered in this scope, and, on the other hand, with intelligent scheduling with respect to the availability of “green energy”, such generation could, over time, become more sustainable than the current “eternal storage” model.
    - GCDNs could allow utterly different types of content (e.g., to answer/process requests of AI Agents sitting anywhere in the system, incl. complex decision making such as security enforcement, routing, congestion, and any global decisions made on a base of a scattered, partial and local view).
    - GCDN should also introduce new forms of entertainment for end-users, not possible with classical CDN. For the latter, one could imagine gamified movies, where a user could take over the control of a protagonist in a movie and play his role towards a potentially different outcome; GCDNs could also support virtual visits / tourism, including sightseeing with friends and on-demand concerts of favourite artists in auditoria, venues, museums of user’s choice. Immersive live sports experiences are another thinkable attractive event, where a user could actively participate (play) in a live sports event against famous participants, leading to a different (simulated) outcome in this viewer’s personal experience. Such new use cases are important for the future of the ICT domain, in particular today, when the expected AR/XR/VR boom largely failed, while smartphones are at the end of their innovation cycle and, for the last 10 years, have been limited to gradual improvements only. Most importantly, it connects

this technological topic from the distributed systems/AI domain to the still very successful European gaming and content industry.

In this picture, first, mobile and transport networks (6G and beyond) must start treating "intelligence" as a primary service feature, similar to data rate, data volume or latency/jitter and, what is most important, make it available to both subscribers and business partners alike. This is especially true with the growing understanding that 6G service portfolio will include native sensing and positioning services, which might become paramount and a critical feature for new types of terminals such as autonomous cars and robots. The amount of network-generated and mobile network operator (MNO)-owned data related to this will increase, as will the requirements for treatment and storage. This in turn requires a redesign of terminals, potentially of radio access network (RAN), certainly of mobile core networks, but also of interfaces to transport/optical networks and clouds, as such treatment will be better done cooperatively, and bidirectional offloading should be allowed. While a deeper dive into business models and the future monetization of intelligence-as-a-service (e.g., through tokens) would be helpful, technologically, mobile networks should enable seamless bidirectional offloading of generative tasks while respecting the rights of all involved stakeholders yet maintaining a sufficient task execution quality at a low energy footprint for the terminal and business partner side. The NTN segment of the overall 6G system might play a distinctive and growing role in this picture. Orbital determinism can be used to turn the Generative CDN problem into a predictive 3C (Communication-Computing-Caching) topology-management problem: model weights and personalised generative content can be pre-staged along anticipated contact windows, compute tasks migrated between orbital and terrestrial tiers, and cached artefacts repositioned proactively across the GEO / MEO / LEO / HAPS / aerial / terrestrial stack. SWaP-aware on-orbit AI can complement heavier model execution at terrestrial edges or cloud tiers, forming a single distributed AI pipeline.

Second, mobile networks (6G and beyond, but also transport systems, subsystems and terminals) should be re-engineered to be able to start consuming and offering intelligence at each and every suitable and reasonable subsystem, plane or system part of themselves. In this area, one important topic would be to start thinking about use cases and, therefore, actionable ontologies of each suitable system part that could be (or should be) made available through agentic means. While one immediate use case would be to offer it at the respective service level for contract negotiation, in an operational reality, such ontologies could often reside on non-commercial technological interworking level.

2.3.3 Synopsis of Required Research

Research Theme	***AI AGENT TOOLCHAIN***		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
AI Agent Governance and Operations (AGO26) General concept and architecture for highly scalable agent deployment, onboarding, naming / identification, runtime control advertisement, discovery, selection, semantic task (de)composition,	*Mid-term*	1a: Overall concept and architecture with roles and stakeholders, who onboards agents where, how and in which environment agents interact. 1b: Concepts and methods for interaction of heterogeneous systems of agents. 2. Agent governance, agent naming and identification schemes. 3. Agent runtime control, agent commanding and operational reporting and logging. 4. Ontologies and common expressions of skills, capabilities, policies, etc. Information models, dictionaries, mechanisms and	Paramount as per vision in Section 2.2. The overall concept should support self-determined, sovereign execution of a group of agents by each individual and commercial entity, not just by major providers and/or big tech and states. Regardless of the owner, these agents need to be able to find and cooperate with other agents as per

reporting and logging.		algorithms for (potentially verifiable) agent and agent skill discovery and advertisement supporting semantic (de)composability. 5. Methods for semantic task decomposition into available skills and subsequent selection of agents with the required skills for the overall task execution within the constraints of the given task. The concept should support individual agents from different systems of agents and extreme scaling, the number of agents supported by the architecture should be over 100 billion. Herein, the robustness of the overall agentic task execution and the respect of the task constraints such as latency/duration, accuracy, costs, energy, policies should be enforced.	policy. Semantic decomposition and robust mapping of the intent at hand to the skills of other suitable agents is crucial in the setting, where agents operate in autonomy, under different legislations, run by and on behalf of different authorities.
Generative CDN (GCDN26)	*Long-term*	1a. GCDN Use cases for end-users: Novel types of content such as gamified movies, virtual tourism and visits, participatory televised sport events, etc. 1b: GCDN use cases for technological inter-domain system-task offloading such as scheduling, congestion, routing, security, energy saving decisions offloaded from some agents to others. 1c: Business models (e.g., token and energy economy, DRM, right licenses and copyright questions on generated content, multi-vendor systems) to be taken into account for both 1a and 1b. 2. GCDN Concept and architecture: roles, rights, responsibilities and mechanisms for GCDN, in particular as opposed to CDN, including dedicated energy-efficient AI generator architectures, trimmed specialized small models, non-LLM/non-FM driven generation of decisions, etc. 3. Hardware architectures, interfaces to power generation and novel devices / chips for accelerated low-power AI/ML content generation and presentation. 4. Concepts such as steganography watermarking, etc. to be selected or developed for a clear indication of generation / realism / truthfulness of the content.	Enabling new types of entertainment, user devices and interaction / presentation means. Enabling general AI task offloading is at heart of agentic AI and reinforces the requirement for self-determined digital inclusion. Offloading might be often more efficient than foreseeing local capabilities for content generation. Besides, long-term effect of generation vs. storage/caching can be positive in terms of sustainability, when one considers hardware evolution, green energy supply, etc.

		All the proposals to be evaluated with respect to: - Seamlessness of the delivery of generated content sessions across a set of devices, in particular with user mobility (in cars, trains, cruise-ships, planes): evaluations should cover device mobility, pause/resume from/on different devices, in play device-to-device and device-to-screen handoffs. - Quality of Experience (QoE) and quality of service (QoS) considerations for generated content such as accuracy and potential repetitiveness of content (free of artefacts and failures, replay possibility), latency of content availability (from request to screen), smoothness of interaction. Studies involving users to shed some light on the relationship of QoE vs. QoS for novel use cases would be a big plus; - QoS-preserving, or at least QoE-preserving, AI/ML concepts reducing energy/CO2e requirements for content generation by factors 10-20, validated on different types of content (text, video, images, audio, news, games, multimodal, new types of content); quality vs. energy in comparison to classical CDN for these different types of content. - For generated content, extension of battery life of the handheld devices in comparison to local generation; - Consideration of local laws/constraints;	

2.3.4 Recommendations for Actions/Instruments

Research Theme	*AGENTIC AI TOOLCHAIN*	
Action	**AGO26**	**GCDN26**
International Calls	**X** Naming and governance might require intl. alignment	
International Research		**X** Involvement of Asian and American actors would be highly beneficial for fundamental aspects
Open Data		
Large Trials		**X** Should be tested at scale

Cross-domain research	X Technology, legal and ethical implications	

2.4 Sovereignty: Trust, Security and Related Mechanisms

If we intend to establish sovereign, self-determined, independently controlled operations of services and infrastructures, security mechanisms must enforce the security policies of the relevant stakeholders.

However, in the distributed environment of autonomously operating entities, sketched above, the cooperation between such entities is key. Hence, isolation is not an option. Rather, such cooperation will require some form of multi-party runtime trust management and will result in active exchanges before, during and after the actual operation. While such exchanges come handy to poll available services/skills, select serving endpoints, adjust services, set technical parameters, perform verifications or set contracts and commitments, they should not expose too much of confidential information of any of the involved parties, even if used repeatedly.

While the situation might recall autonomous agent research/heterogeneous distributed systems, the operational reality is profoundly different. AI Agents are not merely software workloads, in the sense that they are expected to adapt to their operational environment (and learn), i.e., also as a result of legit, authorized interactions with other parties. The complexity of this adaptation, at heart of what modern AI technologies, e.g., based on LLM and foundation model (FM) paradigms, can achieve, is tremendous, as such models and the corresponding AI Agents can essentially “mimic” any realizable function from the available training data, including functions, for which in the classical ICT no deterministic algorithm was known (e.g., for the so-called “common sense” – like recognition of voices, faces, correct/incorrect social behaviors, content generation, etc). Because of this, it is very hard to correctly foresee the metamorphosis that a given Agent could undergo given a particular set of interactions. The probabilistic nature of these changes is what makes the situation different: a given AI Agent might behave differently in the same situations, as it might behave differently over time. Besides, it might sometimes provide wrong results (e.g., because of hallucinations), even though it is in principle operating correctly, etc.

2.4.1 Multi-Party Dynamic Trust Management

In the world of ever-changing entities, solely relying on verified identities will generally not suffice. For that reason, classical trust management based on pre-established trust (trust anchor) in some entities doomed as good enough for particular services is generally questionable in this model. Indeed, e.g., the very same AI-based service might provide satisfying output for some (class of) input, yet very questionable results for some other (class of) input, even though both input classes might be listed within the skills of the corresponding agent (or be understood as the same input class). Generally, the output quality might change with input, over time, under load, with price, after re-training, even based on history of exchanges, etc. This is not limited to AI Agents and AI model interactions: AI Agents can also actively reprogramme / change existing service elements, hence altering their behaviour as well.

In a nutshell, **putting deterministic trust on a non-deterministic machinery is not always reasonable.** To cope with that, **new dynamic trust management solutions are required**. These solutions should be able to establish trust starting from zero, but should be able to dynamically nourish (maintain) the trust relationship, e.g., based on interaction experiences and verified interaction results. For each established service identity, such trust management might be used to alter, e.g., the selection of that service for particular tasks according to the past experiences.

Such trust management does not primarily employ existing trust anchors to conclude on the trustworthiness of the service, but rather dynamically constructs and maintains operational confidence levels into the execution result expected of a given entity.

With growing maturity, one could also imagine that to result in more complex security management solutions, capable of adjusting operational security policies in the process of interactions, according to the result satisfaction levels. Since resources here can be pooled together to work cooperatively on services, in many deployments, resource-level security policies (as derived from resp. local laws and the resp. stakeholder wishes, etc), may be different from service-level security policies.

The situation is not limited to but can be nicely illustrated in an NTN context: when assets from different operators co-exist in overlapping orbital shells and, e.g., share inter-satellite-link capacity under federated arrangements, perimeter-based, network-membership trust models become inadequate: the set of trusted peers literally changes with orbital geometry. Besides, as the satellite changes its position over ground, it crosses international borders, hence so may the policies that need to be respected at the service provisioning level.

Research into solutions for trust and security management to follow the dynamic topology and respective usage constraints in these dynamic environments is therefore required. Generally, programmable, machine-readable encoding of jurisdictional, contractual and service level agreement (SLA) constraints, with cryptographically attested records of enforcement actions, will need to be considered to enable cross-border/cross-system agreement enforcement when transactions traverse overlapping and sometimes conflicting regulatory or policy regimes.

2.4.2 Security and Privacy of AI/ML

To enforce the respective security policies, the identification of all involved entities classically needs to be supported by mutual authentication and fine-grained, implementation-specific authorization (e.g., as delegated from the user to the agent or expressed by the service/resource owner). However, given the ephemeral and ever-changing nature of many of such entities (with orchestration: created and elastically maintained by software; with AI: developed or changed in behaviour on the fly), identity, and with it, security management becomes highly dynamic. While such “agentic exchanges” could be reasonably implemented today in a “closed club” trust setting (e.g., of a single enterprise) with the famous authentication, authorization and accounting (AAA) model, it would be hard to do so in a setup where autonomous, sovereign, ever-changing entities interact for unforeseen tasks. How would one verify identities of ephemeral entities created on the fly for a particular task and destroyed after the task, especially if privacy constraints exist? How to deal with new and expired identities? How would one express or check authorization for a yet unknown task? How to avoid that composition of authorized tasks results in an unauthorized behaviour? How to group identities stemming from different realms for various tasks? Finally, how to deal with changing behaviours of an entity with a confirmed identity?

For these reasons, in this model, it is not only hard to verify identities and check authorizations, it is also required to continuously verify the status of the claimed capabilities, including for otherwise authenticated and authorized entities.

In a nutshell, in this system, every block is in a constant feature upgrade mode. Providing operational assurance for any function in such a setting is very hard. Therefore, we will need to proceed with additional validations, but, given the complexity of exchanges and typical multilateral setups, we will also need a solid accountability framework. Both must prudently weigh off security properties against privacy needs, ideally never putting them in competition. Overall, the following problem areas arise:

- *Validation protocols:* How can a requester verify that a resource, tool, service, and in particular, an AI agent, actually possess the skill they claim to have while preventing cheating yet without violating agent’s justified model and data confidentiality? Research into applicability of, e.g., zero-knowledge proofs (ZKP), oblivious transfer techniques, (partially) homomorphic encryption or PIRS (private information retrieval and search) technologies as some potentially suitable form of automated skill testing is required. Before a transaction (payment) occurs, the agent must prove its capability through verifiable benchmarks while respecting the privacy of the request and without leaking its proprietary model details (e.g., raw data, weights).

- *Justified Accountability:* Related to the more general governance question, we need research enabling us to make sure that any involved legal entity behind a given agent can be made accountable for that agent's actions without violating its legal privacy expectations.
- *Metamorphic AI:* As AI agents change in time, space and operations, how do we approach AI agentic integrity vis-à-vis of service continuity? This question is paramount, among others, to the question of delegation of rights to the agents: As agents are constantly changing, they might morph into something that does not necessarily justify the granting of a particular authorization anymore. While it could be possible to consider approaches trying to limit this morphing (e.g., constant resetting), it is in the nature of AI agents to adapt and adjust. Therefore, new methods and concepts might be needed to secure the systemic interactions of such entities. This point should also be considered when suggesting strong mutual authentication frameworks for agents. Rather, it should be understood that an identity of an agent can be used for accountability, but, different from classic systems, it alone might be insufficient to derive expected behaviour of the latter (e.g., as a basis for granting rights).
- *"Repudiation" of corrupted/infected systems of agents:* since the agents constantly change, they might eventually get "corrupted". Such a corruption can be an effect of a natural ongoing learning (adaptation) activities of the agent, eventually transforming it into a "misbehaving entity" from the point of view of some users. It can also be adversarially inflicted, e.g., by targeted "infections" of the raw training data (which requires access to the raw data or the supply channels for the latter) but also through legit interactions with the agent. The latter effect of "infection through interaction" is akin to currently prominent attacks against ICT systems through the user interface (using specially crafted messages exploiting known implementation vulnerabilities, e.g., out-of-bound write, buffer overflow, format string, cross-site scripting attacks, etc). However, with the Agentic AI, such infection attacks go deeper and current defence mechanisms (sanity checks, service-level gateways/proxies, rapid update cycles) will not suffice in this scope: even an ideally implemented agent without any implementation vulnerability will be naturally prone to these attacks, as long as it takes the context of the interactions into considerations for, e.g., fine-tuning or long-time learning. While correctly estimating the effect of such learning, being able to trace it back (explainable AI) or undoing it is rather something for fundamental AI studies, systemically speaking it is important to be able to detect and remove such corrupted agents from the system, and, more importantly for users, to be able to repudiate the actions of the latter executed "after" such an infection and/or corruption.
- *Expression of an end-to-end service-/application-level security policy*: it is currently an open question how a stringent, end-to-end security policy can be expressed under the aspects of potential (de)composition of the task to the agents stemming from potentially different systems of agents with potentially conflicting security policies. This is partly related to the semantic (de)composition problem (cf. research topic AGO26), as the security policies of agents (e.g., from different systems of agents) can be written in completely different ontologies, yet it goes beyond this, because the relevant security policies might hinder interactions required to be able to understand the present ontologies correctly.

2.4.3 Synopsis of Required Research

Research Theme	*Sovereignty: Trust, Security and Related Mechanisms*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Sovereign AI Agent Operations (SAO26) Agent identification, authentication, authorization, skill	*Mid-term*	1. Agent naming, identification, authentication, authorization schemes in open environments. 2. Methods for privacy-respecting validation of search and of	Agents as non-deterministically operating software need to be held accountable; it must be

verification, with privacy-respecting possibilities for failure tracking and problem resolution after the incidents		advertisement of skills and service. 3. Privacy-preserving methods for auditability and accountability. Naming and identification of agents with respect to user and service provider identification and naming. Authentication, authorization and logging schemes for AI Agents and their actions. How can users and service providers delegate limited scope roles to agents? Enforcement for such limited scope delegation in an open environment, with slowly adapting/changing agents. How can this delegation be non-repudiable and auditable? Identification, authentication, authorization and validation have to respect local laws and privacy expectations of both users and service providers. Exploration of possible proofs/validation methods to verify that an advertising agent actually possesses the skills that it advertises, without the agent having to disclose its raw data, its confidential model parameters (privacy of provider), and without the agent learning what precise modalities of the skill the requestor is interested in (privacy of requestor). To be shown/proven/demonstrated in a typical suitable model (e.g., random oracle model). Efficient failure tracking and good accountability in an open environment, all that without full identity disclosure and avoiding unnecessary traces.	possible to check to which extent they really support the advertised capabilities but also to understand and pinpoint what happened, and whether the execution was indeed as requested by the user or rather a glitch within the agentic execution pipeline. Herein, the fine balance between accountability and privacy is key to maintain digital self-determination while proposing reasonable levels of security to promote adoption of this technology, without laissez-faire leading to unacceptable criminality and risk to democratic foundations.
Metamorphic AI Security (MAIS26)	*Long-term*	1. Concepts and methods for actionable definitions of integrity baselines for AI agents including e.g. methods for operationally measurable agent status estimations, service-specific accuracy or precision benchmarking suites, differential behaviour measurements, etc, where the main target is the capability to distinguish “corrupted” from “acceptable”	The natively mutable nature of AI agents does not allow us to continue reasoning of the latter like of “workflows”. Therefore, classical identity verifications and identity-based authorizations will not be alone sufficient for the enforcement of any

		agents from the point of view of some user/task. 2. Concepts, methods and security mechanisms capable of detecting, mitigating and preventing AI Agent metamorphosis beyond some acceptable boundary or deviation as per point 1, preserving agent adaptability, i.e., without reverting to solutions resulting in continuous agent reset. 3. Methods for undoing or at least repudiating actions of corrupted agents after a corruption has been detected.	practically viable security policy. If the security is still to be enforced (or, alternatively, if Agentic AI is supposed to handle security critical applications), then new mechanisms are required in this space.
Expression of E2E Security Policy in Agentic AI (SPAI26)	*Mid-term*	1. Concepts, methods and approaches for (de)composition of task/app- or service-specific security policy from(on) a chain of agents involved in / selected for the execution of that task/app or service. Selection / discrimination of agents with their shown/known/presumed security policy under the security policy definitions of the app; estimations of the impact on the overall resulting security policy (feasibility, coverage, redundancy, shadowing) when composing from individual security policies of involved agents. 2. Definition of formally verifiable elastic security policies capable of expressing boundary conditions of acceptable deviations from the stated policies. 3. Impact of composed security policies on the feasible AI services. Presentation of the overall system capable of (de)composing security policies from(on) selected agents from different systems of agents, including reconciliation of security ontologies. Efficient heuristics (polynomial time) for feasibility and coverage checks of security policies.	Users should understand the impact of engaging chains of agents, and in particular, they should understand the potential impact on the security policy feasible for their intended application and the opposite, users should understand the limits in feasibility of the application once the security policy of the latter has to be enforced without any elasticity.

2.4.4 Recommendations for Actions/Instruments

Research Theme	***Sovereignty: Trust, Security and Related Mechanisms***		
Action	**SAO26**	**MAIS26**	**SPAI26**
International Calls			

International Research	**X** *For private skill validation*		
Open Data		**X** *For benchmarks*	**X**
Large Trials	**X** *For latter, higher TRL*		
Cross-domain research	**X** *Applied Cryptography*	**X** *Applied Cryptography*	**X** *Security, services, AI*

2.5 Green ICT: Resource Efficiency and Sustainability

Sustainability has become a major design preoccupation of future ICT. Emerging compute platforms and telecommunications systems follow the principle of *sustainability by design*. Technologically however, it is mandatory to clarify, which new technologies or technological advances are required and lead ICT research into that direction.

2.5.1 Towards Green ICT

In recent years, the ecological conscience has generally increased in Europe. Backed by political and economic initiatives both by the Commission (e.g., Renewable Energy Directive, Green Deal) and the Member States (e.g., German *Energiewende*), the main trend is to *reduce the dependency on conventional energy sources* (nuclear, fossil) *to the advantage of renewable energy supply* (wind, photovoltaics, hydroelectricity). Given the decreased flexibility in the energy production of the latter, *this shift must be accompanied by smart energy demand management functions*, resulting in a strong push for Smart Grids in the energy sector. That is where ICT is generally regarded as an important enabler. However, swapping power sources does not address the energy consumption of the consuming infrastructure as such.

Previous SRIA [C2-02] has identified many required energy efficiency improvement methods yet also clarified that energy efficiency of distributed compute facilities is paramount yet insufficient. It identified requirements for novel methods and demanded that future research explore and develop approaches:

- **Continued investment into ICT energy efficiency methods** enabling elastic resource management beyond single DC in addition to the current trends, potentially using smart grid's demand management and the "recycling" of heat.
- **Research on new methods for CO2e reduction**, which require interfaces to the smartgrid but also transversal interfaces (we need to know, in runtime, at the service layer, which mix of current each involved serving infrastructure element has; we also need to know suitable alternatives). In addition to such interfaces, methods and algorithms should be able to re-route service requests to ecologically optimal resources, addressing all potential prerequisites (knowledge of energy mix and loads) and problems that might stem from such re-routing (affinity, context transfers). Ultimately, though, reliance on eco-current is subject to the external dependency on the energy sector. Technologically speaking, nothing can be done from within ICT to get more eco-current or eco-current at more locations, more often, etc.
- **Research on new explicit service user involvement methods** as a novel and promising direction for energy consumption and CO2e reduction in modern ICT. In particular, with conjunction with Agentic AI, it is to be considered throughout the whole service function chain (SFC) [C2-04] so as to enable three distinct features:
  1. *Service user awareness* of the ecological impact that the used service is responsible for;
  2. *Service user incentivization*, i.e., means to push users to rational decisions improving the overall ecological posture of the used services;
  3. *Service user actions*, i.e., means for the user to act upon the intended or used service, both prior to and during the service usage.

Note that the user here is not necessarily the end-user, but the respective service user. As per general chapter vision (Section 2.2 above), any of such users could be represented by an AI Agent, negotiating and deciding on behalf of the actual stakeholder. The embodiment of such agents by providing them with the awareness of the respective resource situation allows for broader implementations of the so-called embodied AI, which could be a game changer on a way to more efficient, more inclusive, green ICT of the future.

Service User Awareness

Service User Awareness simply tries to account for the prevalent outsourcing nature of ICT: ICT does not work in a compartmentalized fashion, and it is very common to extend a local ICT system through capabilities brought by remote services. As more and more companies and industries migrate towards cloud computing, one could even say that the local ICT system dematerializes completely. This could be rapidly miscounted as reaching net zero locally, however the truth is that the CO2e emission corresponding to the same service still occurs; it just occurs somewhere else. It remains debatable whether the centralization and specialization in this particular case yields an overall CO2e gain; on the one hand, the cloud provider certainly has more competence and all interest (TCO) to optimize his own resource footprint to best match the cumulated demand; on the other hand, the remote nature of such service usage makes it hard to account for all CO2e emissions in the SFC, as, for a fair comparison, one needs to consider the use of access and transport networks as well. All this is difficult today, because all ICT domains so far only support domain-wide energy consumption metering. **Service-level energy consumption and CO2e emission metering is hardly possible anywhere.**

While for the society at large it does not matter, where the CO2e is accounted for, it does matter for optimizations and reductions, and, hence, it might matter for legal entities, regions or states: indeed, if the emissions occur with the respective service provider, and the service user is left unaware of the emissions due to his service usage, then optimizing this service usage, the configuration of the service, limiting unnecessary service flows, incentivizing such limitations and other similar measures would be altogether useless, as they all would represent an economic externality for the service user: such optimizations would reduce service provider's OPEX by using less of that provider's carbon credit (e.g., in the EU-Emission Trading System/EU-ETS regulatory carbon market) without any visible gain for the service user.

To illustrate this, consider as example a user backing up local data to a remote cloud. This user could easily use local data compression (prior to applying any security measures) known to be highly effective to reduce the data volume transmitted through the network. Alas, such compression would increase user's own CO2e emissions, as doing so requires more effort on the user side (for compression and decompression, respectively). However, if the CO2e emission of the network is attributed to the network provider, especially under the flat rate pricing, compressing data would actually harm user's own CO2e footprint. Note how this example shows that a global optimum cannot be reached due to externalities.

That is why we strongly suggest to **raise service user awareness at all stages of the overall SFC**. Hence, each involved service provider should be able to gather energy consumption and, ideally, corresponding CO2e emissions at the service layer, so as to attribute actually measured quantities to each service instance of interest. This represents a serious departure from the current practice and constitutes a novel type of requirements on the participating service providers. Once this **ecodata** (energy consumption and CO2e emissions) is gathered, it needs to be exposed to the user. Depending on whether the end-user has knowledge (e.g., through an explicit contract) of a given service function in the SFC or not, ecodata might need to be aggregated or forwarded by various SFC stages, so that it can reach interested stakeholders, for instance end users.

Service User Incentivization

Equipped with the ecodata of used services, the user can make informed decisions, whether local or remote service realization represents a better ecological choice, where and when a service should be used, etc. Future business models could take this into account, so as to provide concrete motivation for rational users to act in an ecology-friendlier way. While the incentivization in the economic sense is not a technological method, ICT needs to support this. A typical example would be **exposure of greener alternative services**, ecologically better **service eco-variants** (reducing quality, changing regions, postponing usage in time, etc.) or ecologically less impactful **service eco-configurations** to the user, as the user cannot know what levers exist in the service realization of the provider, while provider cannot always use these levers under the existing SLA. In general, only the user can decide to go for a service below the promised SLA; while the incentivization has to find ways, why a service variant would be acceptable or even interesting for the user, the technology has to cater for the existence of such variants and the means for timely exposure of the latter to the users along with their expected, comparative ecological costs, when it matters (prior, during or post session, etc).

<u>Service user actions</u>

The support for the corresponding actions from the service user is the last piece of the overall service user involvement puzzle. **ICT technology must provide technological means for the users to rapidly (seamlessly) switch to a greener service alternative or variant.**

2.5.2 Resource- and Energy-Aware AI

After several years of pursuit of performance of AI, the pursuit of **Energy-Aware AI** becomes increasingly urgent as classical data-center optimization levers – such as virtual-machine consolidation, improved resource utilization strategies, and advances in power-usage effectiveness (PUE) – have largely plateaued after a decade of refinement. The point is that modern AI workloads already operate in such highly optimized DC environments, meaning that further efficiency gains cannot rely on naïve orchestration alone (e.g., re-allocations). Instead, the field must shift its attention toward the **foundational layers of AI systems**, where the largest untapped potential for energy reduction remains. This includes the design of more efficient **model architectures**, the development and strategic integration of accelerators at the right stages of the computational pipeline, the development of smaller task-specific models, and advanced inference-time optimizations (potentially, with dedicated hardware accelerators). Collectively, these efforts define a new research direction: reducing the **energy cost per token**—or an equivalent metric—**without loss of task performance or accuracy**, a constraint that elevates the challenge from engineering to fundamental science.

Related to this, research can consider particular hardware accelerators, potentially also quantum machines, Quantum AI, systemic optimization of the pipeline, including of the Agentic pipeline. This defines a "**systemic** optimization research topic", strongly bound to how AI operates and which resources are used for it, when and how.

2.5.3 Energy Considerations for Agentic AI Systems

While early AI systems were dominated by graphics processing unit (GPU)-centric computation, particularly during training, the landscape is shifting with the rise of **Agentic AI**. Training workloads remain relatively straightforward to measure: GPUs run at high utilization for long, uninterrupted periods, enabling precise estimation of energy consumption. In contrast, agentic systems introduce a qualitatively different computational profile. Instead of merely generating outputs, agents **take actions**, orchestrating models, tools, memory stores, APIs, and networks to achieve goals. This results in a far more heterogeneous and distributed workload, with significantly increased reliance on CPUs, memory, and system-level operations. The energy footprint becomes fragmented across multiple services, machines, and even administrative domains, rendering traditional measurement approaches difficult and orchestration overhead too high compared to the possible gains. Consequently, a second major research frontier emerges: how to measure and optimize the ecological impact of Agentic AI. Unlike monolithic training runs, agentic workloads are dynamic, context-switching, and geographically scattered. They involve interactions across cloud services, edge devices, and third-party systems, each contributing to the overall energy cost in ways that are difficult to isolate. Developing robust methodologies for tracking, attributing, and minimizing this distributed energy consumption is essential for ensuring that the next generation of AI systems remains both powerful and environmentally sustainable.

2.5.4 Inclusion vs. Efficiency vs. Privacy

Within a given Autonomous AI System, by their nature, various devices will make different observations and will require different insights (e.g., different terminals – different situations, different types of resources – different preoccupations, etc). Inclusion of such observations is key to model generality and accuracy, in particular in the training phase, or otherwise the model could be biased by some "default" conditions, as the considerations of "specially located" devices could be ignored, making the model globally less accurate. Besides, such a model would hardly be capable of producing reasonable actionable outputs for such special devices.

Unfortunately, with privacy considerations, observations cannot be always processed remotely; rather, privacy considerations often dictate that some operations must be done locally on the device. Because of a strongly skewed distribution of, e.g., computation capability, this can lead to incomplete inclusion, a bias towards strong devices (excluding so-called stragglers) and end up entailing a strong limitation of accuracy.

Therefore, to include observations from resource-constrained devices, either we need to accept a generally higher exposure of such terminals (and their principals), or we will need to accept an overall slower system, as we will have to wait for these devices to finish.

2.5.5 Synopsis of Required Research

Research Theme	**Green ICT: Resource Efficiency and Sustainability**		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/ Value**
Resource-awareness of the service layer, articulation of needs and provisions from the system to the user/applications (EXP26)	Mid-term	Methods, concepts and APIs to establish resource-situational awareness at the respective service execution level. Service level needs to become aware of security, reliability, mobility, load, energy, neighborhood, and other contexts of the given used resource. Protocols, algorithms, architectures and solutions for application-to-system, user-to-service, agent-to-resource or generally digital-to-physical world interfaces, i.e. exposing available resources, their capabilities, requested/authorized status pertaining to the executed application/service and, in particular, fetch ecodata attributed to the requestor, to the applications/users and getting requirements from applications/users explicitly or implicitly. Final solutions should demonstrate: a. Support for legacy applications as much as possible, including support of applications that use end-to-end traffic encryption; b. Collection and attribution of relevant resource status data (e.g., ecodata, lower layer triggers, physical world parameters, security indicators, etc.) on a service-level, and, if necessary, forwarding that from uplink services, or aggregating from the remote	- Improving sustainability and universality - Improving transparency, trustworthiness, governance

		services used to provide the service instance in question. c. Reuse (extending, integrating, mimicking, maintaining compatibility to) existing methods, where applicable, such as network exposure, ETSI MEC, etc. d. Runtime negotiations with compatible, novel applications.	
Resource-Aware AI: Foundational Aspects (AIF26)	*Long-term*	1. New model architectures leading to substantial energy gains. 2. Dedicated, task-specific models. 3. Definition and agreement of AI-specific energy metrics (incl. CO_2e if possible). Demonstration of at least factor 3 gains in energy without losses in accuracy, precision or generality. In other words, the fundamental progress should strive to maintain the quality while substantially lowering energy costs.	Progressing the foundations of the AI workloads in the sense of the energy requirements is expected to have a multiplicative effect, with a potential for huge energy and hence carbon savings.
Energy Posture of Agentic AI (EAI26)	*Mid-term*	Proposals, concepts, methodologies, approaches, studies on mechanisms, protocols, architectures for assessing ecological impact (energy, carbon) of agentic workloads in operations in closed (one company) but, most importantly, in typical multi-domain, multi-agent environments (“agentic Internet”). Agreement on some baseline agentic AI use cases. Capability to precisely assess agentic AI’s energy consumption per use case in realistic setups, in particular for the use phase (no life cycle assessment (LCA), just the usage), to be able to judge service level costs (e.g., per task, per prompt, per token, etc) to be demonstrated at TRL4 in the first and TRL6 in the second phase. Establishing a public catalogue of agreed energy expenditure baselines for service level Agentic AI energy information as benchmarks for the future.	Understanding the ecological impact of the usage of the agentic AI services is key both for the estimation of the real societal costs (beyond the pure energy costs as the provider sees them) and for the future technological improvements.
Resource-Aware AI: *Systemic Aspects (AIS26)*	*Long-term*	1. New system designs for AI computing leading to substantial performance vs. energy gains, using smart allocations, runtime scheduling, dedicated accelerators, new compute architectures including quantum computing for suitable tasks/tools/etc., quantum AI, security-specific hardware, etc. Demonstration of at least factor 5 operational gains either in performance for the same resource (e.g. energy) budget or in resource footprint for the	Introduction of specific accelerators and more efficient, dedicated architectures for rapid AI learning and inference are key for the further progress of these technologies in line with the European values.

		exact same performance (same quality, same speed, same goodput) for realistic AI workloads and systems. In other words, the systemic improvements should strive to improve the performance of the system while substantially improving energy efficiency.	

2.5.6 Recommendations for Actions/Instruments

Research Theme	**Green ICT: Resource Efficiency and Sustainability**			
Action	**EXP26**	**AIF26**	**AIS26**	**EAI26**
International Calls		**X** Energy metrics		
International Research		**X** Fundamental AI research	**X** Accelerator design	
Open Data		**X** Benchmarking	**X** Benchmarking	**X** Measurements, evals
Large Trials	**X** Support for legacy apps			**X** Tests with real use cases
Cross-domain research	**X** Incentives and regulation aspects	**X** Mathematics	**X** Physics, materials, quantum	**X** Smartgrids/energy domain

2.6 Resource, Terminal and Device Capability Considerations

2.6.1 Generic Capabilities

As per Figure 2-1 above, all various and heterogeneous (groups of) resources must be seamlessly integrated over unifying logical abstractions (consolidation) in the overall unified controllability layer, which then implements various control and usage possibilities in a higher layer abstraction (e.g., IaaS abstractions, but also particular PaaS primitives as per Figure) over these resources.

Herein, it is important that:

- The consolidation correctly handles the scalability aspects. Depending on the type of resources and the intended services, the visibility granularity should be adjusted from the individual resources to the groups of the latter.
- The "seamless integration" has to respect the security policy of the respective resource: the latter can define access and authorization limitations, which should be expressed to the higher layers, such that the allocated service can reasonably select the best suitable resources.
- While the service layer will use the resources (the resources will be used for the allocation of service-specific elements through the controllability layer), the integration must consider the opposite direction: it is paramount that all resources (recursively from any atomic resources to logically bound and abstract virtual resources) expose their particular dynamic state situation to the service layer over suitable APIs. While the precise API is out of scope and can be implementation specific, the expected information which should be made available to the service layer includes (but is not limited to): available resource capacity; current load; energy situation, if possible per

capability and with carbon information; security context, including security policy and other constraints, along with resource security-relevant indicators, where available; local context, including potentially visible neighbours, mobility-relevant triggers, etc.

This latter point is paramount to fuse the physical and digital to one (cf. cyberphysical, critical systems), and it is the only way to further develop current AI capabilities beyond the purely virtual capabilities (such as content or code generation) to a **truly embodied AI**.

All specific resources and subsystems in the following chapters can use their own, resource and subsystem-specific architectures, striving for the performance without forgetting the capability to provide the required resource status information as discussed above.

2.6.2 Terminal Considerations

Agentic AI on Mobile and Embodied Terminals

As per architecture presented in Section 2.1, terminals and other devices are considered as resources but also as execution environment within a broader system in a seamless manner. A terminal may therefore participate simultaneously in several Autonomous AI Systems, with AI agents from these systems being executed locally (i.e. on the device) but also cooperating with other agents from the infrastructure.

While current research primarily focuses on deterministic inter-agent communication across distributed edge-to-cloud infrastructures, deploying AI agents directly on mobile devices, robots, vehicles, and IoT terminals introduces a new set of challenges. These extend beyond classical resource management and include agent heterogeneity, cross-layer protocol adaptation, physical safety, and network/application security.

Unlike agents operating within server infrastructures, agents deployed on physical terminals are embedded in the real world, denoted as a truly embodied AI paradigm (see Section 2.1.1). The variability of wireless connectivity and on-device computational resources can directly impact agent correctness and system integrity. These are part of and in relation to Resource-Aware AI (AIF26) and Situational Awareness (EXP26) as defined in section 2.5.5. Therefore, future agentic systems must operate under fluctuating radio conditions, intermittent connectivity, mobility-induced disruptions, handovers, varying network loads, and dynamic availability of edge and cloud resources, including mono and multi-domains scenarios. Particularly in mission-critical domains such as autonomous transportation, industrial automation, remote healthcare, and smart infrastructure management, these constraints can adversely impact inter-agent coordination, latency, planning reliability, and service continuity. For instance, an autonomous vehicle, which is a heavy object weighing more than two tons, can lead to fatal accidents due to misaligned agent behaviour, incorrect planning instructions, or slow control responses caused by network/resource issues or poor coordination among agents. Similarly, a group of robots cooperating in a shared environment, can lead to danger if their controlling agents receive incomplete, faulty, conflicting, or intentionally misleading instructions. In these contexts, the ramifications of agent failures go beyond service disruption or data loss, they can result in injury or damage to human life and physical surroundings. Moreover, the variability of agent toolchains and cross-layer Quality of Service (QoS) pose additional challenges: communications between terminals and from terminal to edge in AI wireless networks can worsen the difficulties associated with the AI Agent Toolchain mentioned in Section 2.3. These communications require resilient, adaptable solutions over rapidly changing execution environment (including connectivity, resources, security, ...) while maintaining all anticipated aspects of the toolchain, particularly the ability to connect different types of agents.

The network can play an active role in preventing critical incidents by delivering timely risk data to terminal agents. However, this dependency makes the network itself a potential point of failure: disconnection, interception, or adversarial injection of messages through the wireless channel can propagate failures directly into safety-critical control loops. A compromised or degraded communication channel is thus not merely a connectivity problem but a physical safety hazard. This motivates architectures in which terminals remain autonomously safe by design. Terminal agents must incorporate built-in recovery behaviors and local trust mechanisms that allow them to maintain safe operating envelopes even when the network becomes unavailable, unreliable, or adversarially manipulated.

The resulting research agenda is multi-dimensional. It encompasses: lightweight AI agents capable of dynamically adapting behaviour to resource availability while preserving formal safety guarantees; intelligent task partitioning, workload redistribution, and agent migration between terminals, edge nodes, and cloud infrastructure; Reinforcement Learning-based autonomous decision-making for service placement, resource allocation, and reconfiguration based on real-time observations and predictive network analytics; formal verification techniques and runtime safety monitors enforcing bounded safe-operating envelopes; policy-driven cross-layer protocol switching tightly integrated with agent toolchains to manage heterogeneous QoS requirements over dynamic radio links; and cryptographic attestation, anomaly detection, and local safety arbitration to protect against network-borne threats. Together, these directions aim to produce agentic systems that are simultaneously adaptive, heterogeneity-aware, network-resilient, and physically safe — capable of operating robustly across the full spectrum of resource, connectivity, and adversarial conditions encountered by mobile and embodied terminals.

The research agenda should focus on the following:

- Lightweight AI agents capable of adapting to resource constraints while maintaining formal safety guarantees
- Intelligent task partitioning, workload redistribution, and agent migration across terminals, edge nodes, and cloud infrastructure
- Reinforcement learning approaches for dynamic service placement, resource allocation, and system reconfiguration based on real-time and predictive signals
- Formal verification methods and runtime safety monitors ensuring bounded safe operating conditions
- Policy-driven cross-layer communication protocols integrated with agent toolchains for heterogeneous QoS management over dynamic networks
- Cryptographic attestation, anomaly detection, and local safety arbitration mechanisms to mitigate network-based threats

These are summarized through two main actions (see 2.6.4).

2.6.3 NTN Considerations

NTN plays a key role in the vision of sovereign, distributed, agentic ICT. While the preceding sections establish the conceptual framework for multi-stakeholder federation, dynamic trust, and energy-aware intelligence, NTN research topics address technical and architectural imperatives making this vision globally inclusive and resilient.

The transition toward agentic AI systems assumes a seamless integration across heterogeneous resources spanning terminals, edge nodes, and cloud infrastructures. TNs alone cannot achieve the global connectivity required by this paradigm, nor they can support the latency-asymmetric, dynamically reconfiguring topology that multi-orbit constellations inherently present. NTN fills this gap by extending the autonomous agent ecosystem to regions where the ground infrastructure is sparse, expensive, or non-existent.

The multi-stakeholder sovereignty framework depends critically on the NTN ability to support federated asset pooling across national boundaries. NTN constellation operators, TN providers, and enterprises must be able to dynamically negotiate policy-driven resource sharing not as temporary bilateral arrangements, but as an integral system behaviour. The NTN research topics below address this: on-board distributed control planes enabling real-time consensus across orbital tiers, cryptographic policy domains enforcing data sovereignty despite intermittent connectivity, and unified service orchestration frameworks presenting a coherent abstraction to agents and users despite the radical heterogeneity of the resources.

Furthermore, NTN amplifies the green ICT imperative. Orbital determinism transforms the generative content distribution problem into a unique opportunity: pre-staging model weights and cached artifacts along anticipated contact windows, migrating compute tasks between orbital and terrestrial tiers according to energy availability and SWaP constraints, and optimizing the 3C topology proactively. The research topics that follow establish the foundations for this vision: from unified TN-NTN service models that harmonize billing and QoS across orbital and terrestrial layers, to AI-native forwarding that grounds routing decisions in semantic relevance and energy efficiency rather than IP-layer abstractions alone.

In essence, NTN is not a specialized subsystem to be integrated a posteriori: it is a fundamental component for realizing the distributed, policy-driven, energy-aware agentic infrastructure envisioned in this Chapter.

Joint multi-layer/-stakeholder architecture design, optimisation, and security.

A fundamental gap in the current NTN is the absence of a framework that jointly optimises coverage, capacity, latency, energy, service quality, mobility, and security across the full multi-layer topology under dynamic conditions. Today, 3GPP procedures treat TN and NTN segments largely independently through incremental adaptations rather than holistic co-design, a limitation that is especially marked for heterogeneous RATs, variable backhaul, and multi-operator deployments. R&I activities shall therefore develop co-design frameworks that treat TN and NTN as a single optimisation domain, supported by INL-aware routing, multi-orbit system-level simulation, enhanced AMF and AI-assisted predictive handover meeting 6G latency and reliability targets, and Zero Trust Architecture with micro-segmentation tailored to orbital SWaP constraints and dynamically changing trust relationships. Notably, different domains may also indicate different stakeholders or owners, and thus harmonisation among multiple SLAs. R&I activities should develop **programmable adaptation mechanisms** capable of dynamically adjusting the system behaviour in accordance with regulatory, contractual, and service-level constraints.

AI-Native Orchestration for NTN.

Despite the recognition of multiple handover types in multi-layer NTN, protocol solutions meeting 6G latency and reliability targets for all of them do not exist. Satellite-to-terrestrial and ISL handovers are particularly challenging in dense multi-layer deployments. Make-before-break handover requires simultaneous multi-beam operation and predictive link availability estimation that current standards do not support. AI-assisted beam pointing and predictive handover algorithms show promise but have not been validated at the constellation scale and traffic diversity of operational 6G deployments. The enhanced AMF required for multi-layer architecture mobility has not been designed or standardised. R&I studies shall focus on **developing NTN-specific AI applications**, e.g., using the O-RAN framework (xApps and rApps), to be validated under multi-orbit constellation dynamics and NTN channel conditions. The placement of AI inference functions (on-board satellites/HAPS, at ground gateways, or in cloud infrastructures) should be jointly optimised with the latency requirements of the corresponding control loop. Distributed and federated AI training mechanisms adapted to the intermittent connectivity and heterogeneous compute capabilities of NTN nodes should be developed as foundations for truly autonomous constellation operation.

Distributed and Federated Control Planes.

Ground-centric control architectures can neither respond rapidly enough to the dynamic topology changes of multi-layer NTN, nor efficiently manage the highly distributed and intermittently connected nature of satellite constellations. Distributing software defined networking (SDN) controllers on-board (for real-time link selection, QoS negotiation, and resource federation across organisational and orbital boundaries) is a recognised technology direction, but it lacks standardised APIs, inter-controller communication protocols, and consensus mechanisms robust to the NTN latency and packet loss environment. Space-as-a-Service (SaaS) architectures that enable multi-operator asset federation are still at conceptual stage without implementation frameworks or agreed interoperability specifications. R&I activities should thus develop such **on-board SDN controller architectures with standardised APIs** across multi-operator, multi-orbit environments. This includes consensus mechanisms for distributed controller coordination that are robust to NTN latency and intermittent connectivity, inter-controller communication protocols that minimise ISL overhead, and Plug-and-Play network architecture specifications enabling SaaS deployments. The interaction between ground-based O-RAN non-RT RIC functions and onboard near-RT or distributed controllers should be explicitly addressed.

UE Mobility, Context Management, and Protocol Extensions.

3GPP mobility management procedures for NTN were designed under the assumption that gNB functions are at fixed ground locations, with satellite acting as a relay. In distributed architectures, both the CU and DU may be on-board moving satellites, requiring new approaches to UE context management that maintain CU-DU associations across orbital passes, ISL reconfigurations, and handovers. The interaction between end-to-end delay, underlying transport protocol behaviour, and the functional split configuration has not been characterised for multi-hop ISL topologies. R&I shall focus on developing extensions to handover and UE context transfer procedures for environments where both CU and DU functions are located on-board moving satellites. This includes CU-DU association management protocols that maintain service continuity across orbital passes and ISL reconfigurations, end-to-end delay characterisation for multi-hop ISL transport under functional split configurations, and an assessment of whether legacy 3GPP solutions are sufficient or whether fundamental protocol redesign is required. AI-assisted predictive context pre-positioning, migrating the UE context to the next NTN node before handover, should be investigated as a latency reduction mechanism.

2.6.4 Synopsis of Required Research

Research Theme	***Special Resources: Terminals***		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Lightweight, Safety-Guaranteed Agentic AI on Mobile & Embodied Terminals (TERAIS)	Mid-term	Resource-adaptive lightweight AI agents for terminals; intelligent agent migration and task partitioning across the terminal–edge–cloud continuum; RL-based dynamic service placement and resource allocation; formal verification methods and runtime safety monitors enforcing bounded safe-operating envelopes under fluctuating connectivity and resources. The other way around, specific accelerator and security hardware for terminals to enable better, more efficient Proposals should demonstrate: 1. Reference designs and safety-by-design methodologies of both hardware and software enabling certifiable deployment of AI agents in mission-critical terminals (autonomous vehicles, robots, healthcare devices); 2. Reduced risk of agent-induced physical harm, not limiting demonstration to	Bringing embodied AI onto the future smart objects will help elevate such “remotely controlled objects” into the paradigm of “autonomous intelligent subjects”, which will be more and more able to act for the good of the human society, in which they will execute their tasks. Using trustworthy infrastructures as trust anchors and additional verification means for safe operation of robots, automated vehicles, etc. will result in better, safer and cheaper operation than trying to stuck every capability on the device. This will translate into better value for society.

		(constructed) particular use cases; 3. Foundational building blocks feeding 6G/6G+ standardization on terminal-resident AI agent management to be brought in the standardization pipeline.	
Network-Resilient & Secure Agentic AI for Mission-Critical Terminal Operation (TERCAI)	Mid-term	Policy-driven cross-layer protocols enabling heterogeneous agent toolchains to interoperate over dynamic or degraded networks and computational resource; cryptographic attestation, anomaly detection, and local safety-arbitration mechanisms countering disconnection, interception, and adversarial message injection; "autonomously-safe-by-design" terminal architectures maintaining safe operation independent of network availability. Proposals should demonstrate: 1. Standardizable security and resilience frameworks for 6G+ (3GPP, ITU-T, ETSI) decoupling physical safety from network dependency; 2. Trust mechanisms lowering certification barriers for autonomous transport, industrial automation, and remote healthcare; 3. Interoperable cross-vendor protocol baseline for multi-domain Autonomous AI System ecosystems. 4. Correct operation facing adversaries, and correct operation facing presence of “weak” terminals.	Cooperative embodied AI will allow for even more efficient, potentially safer and more rapid response and reaction from autonomous intelligent subjects such as robots, drones or autonomic cars. Creating comprehensible, reliable and verifiable designs and insisting on comprehensible, resilient and verifiable operations is key to better trustworthiness of such constructs.
Research Theme	***Special Resources: Non-Terrestrial Networks***		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Joint multi-layer/-stakeholder NTN/TN architecture design, optimisation, and security (NTNJAO).	Mid-term	1. Unified TN-NTN co-design framework with multi-orbit INL-aware routing. 2. Enhanced core network for 6G latency/reliability targets. 3. Zero Trust Architecture tailored to orbital SWaP constraints. 4. TN-NTN policy-aware dynamic adaptation mechanisms.	True global connectivity is guaranteed and digital divide is prevented. The multi-stakeholder co-design preserves sovereignty of all operators and nations.

		[KPIs]: Routing convergence time <200 ms Routing signalling overhead <1% CN E2E latency <500 ms Integrate 2+ services in the architecture	
NTN AI-Native Orchestration (NTNAIO).	Mid-term	1. Predictive management validated at constellation scale and enhanced AMF for multi-layer handover. 2. Distributed federated learning mechanisms adapted to intermittent connectivity. 3. Validation on multi-orbit operational 6G deployments. **[KPIs]:** Handover success rate >99% Handover latency <100 ms (TN-NTN support) or <1 min (no TN-NTN support) Control loop latency (failure reaction) <500 ms	Human intervention reduction through autonomous operations. It enables reliable, predictive services critical for emergency response and connectivity to remote regions.
NTN Distributed and Federated Control Planes (NTNDCP).	Mid-term	1. On-board SDN controller architecture specifications and standardised APIs across multi-operator, multi-orbit environments. 2. Consensus mechanisms robust to NTN latency/intermittent connectivity. 3. Inter-controller communication protocols minimising ISL overhead. 4. Plug-and-Play network architecture for SaaS deployments. **[KPIs]:** SDN routing convergence <200 ms ISL signalling overhead <1% of ISL capacity Multi-operator SLA compliance >99% Distributed CP availability >99.9% Policy adaptation reaction time <500 ms	It distributes control authority across multiple operators and regions; it prevents single-point-of-failure. New business models and market entry for smaller players are enabled.
NTN/TN UE Mobility, Context Management, and Protocol Extensions (NTNMCP) Ensures equitable global connectivity; prevents	Mid-term (5y)	1. Extended 3GPP handover procedures for distributed on-board CU/DU architectures. 2. CU-DU association management across orbital passes and ISL reconfigurations. 3. End-to-end delay characterisation for multi-	Human intervention is reduced through autonomous operations. Reliable, predictive services critical for emergency

digital divide. Multi-stakeholder co-design preserves sovereignty of all operators and nations.		hop ISL under functional splits. 4. AI-assisted context pre-positioning mechanisms. **[KPIs]:** Handover requirements as per “AI-Native Orchestration for NTN” CU-DU association stability >99% Context pre-positioning hit rate >85%	response and remote regions are enabled.

2.6.5 Recommendations for Actions/Instruments

Research Theme	***Special Resources: Terminals***	
Action	**TERAIS**	**TERCAI**
International Calls		
International Research	**X**	
Open Data		**X**
Large Trials		**X**
Cross-domain research	**X**	**X**

Research Theme	***Special Resources: Non-Terrestrial Networks***			
Action	**NTNJAO**	**NTNAIO**	**NTNDCP**	**NTNMCP**
International Calls	All research aspects are partially addressed by SNS calls, but more calls are needed to: i) increase the TRL; ii) bring together relevant, strong, and innovative EU playersin their fields; iii) support and drive standardisation efforts in relevant areas (3GPP, ORAN, GSMA, IETF, …).			
International Research	In line with the International Calls, International Research efforts are fundamental and beneficial for the design and assessment of the technologies and system-level solutions described in Section 2.6.6. Moreover, considering the multi-stakeholder nature of NTN, MNO/SNO collaborations shall be pursued.			
Open Data	Orbital TLE, traffic data, channel measurements			
Large Trials	Need for constellation scale, multi-operator data (TLE, channel measurements, resource usage, traffic) for design and testbed validation purposes.			
Cross-domain research	Need to bring together technology providers and domains involved in emerging use-cases to develop business models and to solve ethical key questions (e.g. use of data and privacy) right from the start. Given the nature of the research aspects covering system-level solutions, cross-domain research with Air Interfaces, Security, Protocols, and network orchestrations are key factors to be considered.			

2.7 Chapter Summary

Taking into account the particular European constraints as discussed in the document introduction, this chapter proposes a novel systemic vision and identified gaps and required research on the way from where we stand towards a digitally self-determined, resource-aware, distributed and embodied agentic AI. The transition from a data-centric to an agent-centric ICT paradigm is inevitable. By moving toward a world of sovereign AI agents and by specifically acknowledging privacy, security and sustainability needs of the involved stakeholders and the society as a whole, we can preserve democracy and individual rights while reaping the massive productivity gains of the AI revolution. Our goal is an overall ICT ecosystem where intelligence is ubiquitous, yet power is distributed.

We have identified many individual research topics with a systemic impact on the way towards that goal.

3 Softwarisation and Intelligence in Networks

3.1 Introduction

Future communication and computing infrastructures will evolve toward highly distributed, programmable, intelligent, and increasingly autonomous systems capable of supporting a wide range of digital services across heterogeneous and multi-domain environments.

In this vision, infrastructure elements including devices, communication systems, computing platforms, data and service functions, and specialized hardware are progressively abstracted into resources characterized by capabilities, constraints, policies, and operational context. These resources will need to be dynamically composed across device, edge, cloud, and non-terrestrial environments.

Managing and orchestrating such complex ecosystems requires new architectural paradigms that enable flexibility, interoperability, automation, and continuous adaptation. At the same time, these paradigms must support federation across multiple stakeholders, trustworthy coordination across administrative and sovereign boundaries, and the exposure of capabilities to services, developers, and vertical domains.

Two key technological foundations reinforce this transformation: **softwarisation** and **intelligence**.

For the telecom sector, this transition should not be understood as a clean-sheet replacement of existing systems, but as an evolutionary path from current 5G and 5G-Advanced deployments toward more distributed, AI-native, and federated 6G-era architectures. This implies that future softwarisation and intelligence solutions must support brownfield integration, coexistence with legacy operational and management environments, and staged adoption across radio, transport, core, edge, cloud, and non-terrestrial domains.

At the same time, these solutions must satisfy carrier-grade expectations on service continuity, resilience, assurance, observability, auditability, and controlled fallback, so that innovation remains deployable and trustworthy for mobile network operators, telecom vendors, and the broader next-generation internet ecosystem.

Together, softwarisation and intelligence enable infrastructures to evolve from static and siloed systems into federated, adaptive, and value-generating platforms. This evolution must also be accompanied by strong security and sustainability considerations to ensure that future infrastructures remain trustworthy, resilient, resource-efficient, and aligned with broader European strategic priorities.

3.1.1 Softwarisation

Concept

Softwarisation enables the transformation of communication and computing infrastructures into programmable, composable, and evolvable systems in which heterogeneous resources are abstracted and controlled through software. In this paradigm, infrastructure elements including devices, network functions, compute platforms, storage, accelerators, and specialized hardware are represented through common data models that describe their capabilities, constraints, policies, and operational context. By decoupling services and functions from the underlying infrastructure, softwarisation allows dynamic discovery, allocation, and composition of resources across distributed environments spanning devices, edge, cloud, and non-terrestrial domains.

As this evolution progresses, softwarisation is moving beyond virtualization of isolated functions toward the software-defined composition of services across multi-provider and multi-domain environments. Future telecom platforms will need to support federated service layers across operators, cloud providers, edge platforms, and sector-specific infrastructures, while also embedding policy, trust, locality, compliance, and data-handling constraints directly into orchestration and execution workflows. In this sense, programmability alone is no longer sufficient: future softwarized infrastructures must also be governable, interoperable, and able to support service exposure and coordination across sovereign and administrative boundaries. This direction is closely aligned with the European ambition to build an interoperable, openly accessible, multi-provider cloud-to-edge continuum with federated, energy-efficient, and trustworthy data processing capabilities and services.

Softwarisation is also becoming a key enabler of AI-native infrastructures. Future telecom software platforms will need to orchestrate not only network functions and service chains, but also AI workloads, model lifecycles, feature flows, inference pipelines, accelerators, observability functions, and runtime assurance mechanisms across cloud-edge-telco environments. At the same time, they will increasingly expose communication, computing, sensing, assurance, and data capabilities through programmable APIs and developer-facing platforms, allowing third-party applications, vertical industries, and autonomous software agents to consume telecom resources as structured services. This expands telecom software from an internal operations toolset into the foundation of open service ecosystems and future marketplace-oriented digital infrastructures.

Key Research Challenges

A first challenge is the creation of interoperable abstractions and software models that can represent highly heterogeneous resources and capabilities across network, compute, storage, memory, accelerators, and service layers without losing operationally relevant semantics. This includes not only portability and interoperability, but also the ability to support service federation, cross-provider coordination, and capability exposure at scale. A second challenge concerns the orchestration of large-scale distributed infrastructures spanning device, edge, cloud, and NTN environments, where lifecycle management, control, and assurance must operate across multiple domains, stakeholders, and trust boundaries.

A third challenge is to make softwarised infrastructure policy-aware and governable by design. Future execution environments must respect data locality, sovereignty, privacy, contractual constraints, and sector-specific compliance requirements while still enabling flexible composition and efficient resource use. A fourth challenge is the need for AI-native software platforms capable of jointly orchestrating telecom functions and AI workloads under shared latency, energy, and trust constraints. Finally, future telecom software must become self-evolving and verifiable, supporting CI/CD/CT, automated certification, runtime assurance, rollback-capable upgrades, and safe in-service evolution of software, policies, and AI components. These challenges are central if Europe is to build trustworthy, sustainable, interoperable, and economically relevant telecom-cloud-edge platforms rather than remain dependent on fragmented or externally controlled software ecosystems.

3.1.2 Intelligence

Concept

Intelligence introduces autonomous and adaptive capabilities into communication and computing infrastructures, enabling resources to perceive system conditions, interpret operational objectives, and optimize their behaviour accordingly. Rather than relying solely on centralized control, intelligence is increasingly distributed across services, devices, and infrastructure components, enabling more scalable and resilient operation. Agent-based architectures are emerging as a key paradigm for realizing this vision. In such systems, services, devices, and infrastructure resources are represented by software agents capable of reasoning, collaborating, and coordinating actions through agent-to-agent interactions. These agents support advanced operational models such as intent-driven management, closed-loop automation, adaptive orchestration, and dynamic service composition across distributed environments.

As intelligent infrastructures evolve, the focus shifts from isolated AI functions toward richer forms of reasoning and coordination. Future telecom intelligence will increasingly rely on semantic interoperability across domains, so that agents can understand not only exchanged messages, but also the meaning of goals, constraints, policies, and capabilities expressed by actors in different technological and administrative environments. In parallel, telecom-tuned foundation models, knowledge planes, and multi-modal reasoning layers are emerging as important enablers of advanced intent handling, troubleshooting, anomaly explanation, predictive planning, and adaptive orchestration. This means that intelligence in future telecom systems will depend not only on learning from data, but also on combining semantic representations, world models, operational knowledge, and explainable decision support.

Beyond internal optimization, future intelligence must also support distributed and federated operation across operators, public infrastructures, industrial domains, edge platforms, and cloud environments while respecting privacy, sovereignty, and regulatory constraints. At the same time, intelligent capabilities will increasingly need to be exposed to services, developers, and stakeholders, so that gains in autonomy, resilience, and dependability translate into operational, service, and business value across the broader digital ecosystem. This direction is consistent with Europe's broader push toward stronger AI capability, more usable AI infrastructure, and trustworthy cloud-edge ecosystems.

Key Research Challenges

A first challenge is enabling scalable coordination among distributed agents operating across heterogeneous domains while maintaining coherent system-wide behaviour. Mechanisms for orchestration, collaboration, negotiation, and semantic alignment must ensure stability and efficiency in large-scale multi-agent environments. A second challenge concerns the reliability and trustworthiness of autonomous decision-making. Intelligent systems must operate under dynamic and uncertain conditions while ensuring transparency, safety, accountability, and the possibility of runtime validation and fallback when needed.

A third challenge is interoperability at the level of meaning, not only syntax. Agents representing services, devices, and infrastructure components must be able to exchange information, interpret intents, negotiate goals, and reason across common representations while respecting security, privacy, and policy constraints. A fourth challenge concerns the development of robust predictive models, world models, and digital representations capable of capturing the complex dynamics of distributed infrastructures and supporting decision-before-action paradigms at scale. Finally, future intelligence must support federated, privacy-preserving, and sovereign operation across domains, while also enabling trustworthy testing, verification, identity, discovery, and coordination mechanisms for embodied and digital agents. Achieving this combination will be essential for realizing intelligent future infrastructures that are not only autonomous, but also governable, dependable, and aligned with European strategic priorities.

Together, these paradigms enable infrastructures to move beyond static and domain-specific management models toward integrated, programmable, and self-optimizing systems. By combining software-driven infrastructure with distributed intelligence, future communication and computing infrastructures will be able not only to autonomously manage large-scale heterogeneous resources and improve resilience and dependability, but also to create value beyond internal infrastructure optimization through adaptive service delivery, intelligent capability exposure, trustworthy multi-agent coordination, and federated operation across the edge-to-cloud continuum. This chapter explores the key concepts, technological directions, and research challenges associated with these two complementary pillars shaping the evolution of communication and computing infrastructures towards 2030 and beyond.

3.2 Chapter Vision: where are we, hot trends, where do we want to be

3.2.1 Softwarisation

Where we are:
Today, communication and computing infrastructures are progressively adopting software-defined, virtualized, and cloud-native paradigms. Virtual network functions, software-defined networking, containerized deployments, and orchestration frameworks already enable modular and programmable components to be deployed across centralized and edge environments. Abstractions for devices, subsystems, and network functions are emerging, and open platforms increasingly support automation and lifecycle management across parts of the continuum. However, most deployments remain limited to specific domains, provider environments, or infrastructure layers, and federation across operators, cloud-edge ecosystems, and vertical stakeholders is still immature.

Current systems are also still largely optimized for internal infrastructure management rather than for governable multi-provider service composition, AI-native execution, or broad capability exposure. Policy-aware orchestration, sovereign execution, service-layer federation, and telecom developer ecosystems remain fragmented, while common abstractions for joint orchestration of telecom functions, AI workloads, and distributed computing resources are still emerging. In practice, softwarisation has advanced significantly, but it has not yet reached the level required for a truly interoperable European cloud-edge-telco continuum.

Hot trends:
Key trends include the extension of software-defined abstractions across the full device-edge-cloud continuum, the federation of multi-provider service environments, and the emergence of policy-aware execution frameworks that integrate sovereignty, locality, trust, and compliance into orchestration decisions. Another strong trend is the evolution toward AI-native infrastructure, where software platforms jointly orchestrate network functions, AI workloads, telemetry, accelerators, and model lifecycles across distributed environments. In parallel, runtime digital twins, semantic resource models, post-P4 programmability, and self-evolving software platforms are becoming important building blocks for more autonomous and trustworthy infrastructure management.

At the same time, telecom platforms are moving toward structured exposure of communication, computing, sensing, and assurance capabilities through programmable APIs, SDK-oriented integration, and developer-facing service layers. This connects softwarisation more directly to future ecosystem models, including capability marketplaces and cross-domain service composition. These trends align closely with broader European priorities around interoperable cloud-edge infrastructures, secure and climate-neutral edge deployment, stronger digital sovereignty, and the build-up of competitive cloud and AI capacities in Europe.

Where we want to be:
By 2030, softwarisation should transform telecom and distributed computing infrastructures into fully programmable, composable, governable, and interoperable environments spanning devices, edge, cloud, and NTN domains. The target is not only flexible automation of infrastructure, but the creation of a federated software fabric in which services are decoupled from physical infrastructure, resources are dynamically discovered and allocated according to capabilities and constraints, and execution remains aligned with policy, sovereignty, trust, and sustainability requirements. In this vision, telecom infrastructure becomes a core part of a broader European cloud-edge-service ecosystem rather than a closed operational silo.

By the same horizon, softwarisation should also enable AI-native and developer-facing telecom platforms. Operators and infrastructure providers should be able to orchestrate AI workloads and telecom functions jointly, expose capabilities through open yet governable service interfaces, and support cross-provider, cross-sector, and cross-border digital services at scale. The desired end state is a trustworthy and evolvable European cloud-edge-telco continuum that supports industrial deployment, innovation, resilience, and strategic autonomy, while remaining open, interoperable, and economically attractive for future service ecosystems.

3.2.2 Intelligence

Where we are:

Currently, intelligence in networks is mostly applied in centralized or semi-centralized contexts, such as network analytics, automated configuration, and limited intent-based management. Agent-based approaches and closed-loop automation are still in early adoption, often confined to specific network segments or service domains. Predictive models and digital representations, including network digital twins, are emerging but are not yet widely integrated across distributed infrastructures. Moreover, most current deployments focus primarily on internal operational optimization, with limited integration between infrastructure intelligence, service adaptation, cross-domain reasoning, and multi-stakeholder value creation.

Current intelligent systems also remain relatively weak in semantic interoperability, federated reasoning, and trustworthy multi-agent coordination across administrative and technological boundaries. Telecom-tuned foundation models, knowledge planes, AI testing and verification, agent identity and discovery frameworks, and more advanced NTN intelligence strategies are still emerging rather than established. As a result, intelligence has advanced beyond narrow automation, but not yet to the point of supporting a fully interoperable, federated, and trustworthy European cloud-edge-telco intelligence layer.

Hot trends:

Emerging trends include agentic architectures, agent-to-agent communication, intent-driven management, distributed closed-loop orchestration, and the use of semantic representations to enable cross-domain goal negotiation. Digital twins, predictive modelling, and world models are enabling proactive decision-making and decision-before-action paradigms, while telecom foundation models and knowledge-centric reasoning layers are becoming relevant for advanced intent handling, troubleshooting, anomaly explanation, and adaptive orchestration. Service, device, and infrastructure agents are increasingly integrated to support multi-agent reasoning, negotiation, and self-optimization across heterogeneous environments.

At the same time, intelligence is moving beyond internal optimization toward distributed enablers of adaptive services, business-aware automation, federated and privacy-preserving operation, trustworthy testing and verification, and stronger identity and coordination mechanisms for embodied and digital agents. Another important trend is the growing need for intelligence-native operation in resource-constrained and highly dynamic NTN environments, where decentralized coordination, hybrid computation, and frugal resource management are becoming essential. These trends also align with the broader European ambition to strengthen AI capability, data availability, trust, and strategic autonomy across cloud-edge infrastructures.

Where we want to be:

By 2030, intelligence should be pervasive, embedded across services, devices, and infrastructure layers, enabling autonomous and adaptive operation of distributed networks. Agent-based systems should coordinate seamlessly at scale, make decisions in real time while maintaining transparency, trust, and safety, and support cross-domain semantic interoperability across operators, cloud-edge platforms, vertical systems, and public infrastructures. Predictive and self-optimizing infrastructures should anticipate operational conditions, allocate resources efficiently, and respond dynamically to changing requirements through world models, digital twins, and trustworthy multi-agent reasoning.

By the same horizon, intelligence should also become federated, privacy-preserving, and sovereign by design, allowing operators and stakeholders to learn, reason, and optimize collectively without requiring full centralization of data, models, or control. Intelligent infrastructures should support verifiable AI-native operation, trusted participation of embodied and digital agents, and scalable coordination across terrestrial, edge, cloud, and NTN environments. The desired end state is a resilient, trustworthy, and strategically relevant European intelligence layer for telecom and distributed computing infrastructures, capable of supporting industrial deployment, adaptive service ecosystems, and multi-stakeholder value creation at scale. This evolution should be achieved through deployable and incremental migration paths that preserve interoperability with existing telecom systems and support phased adoption by operators and vendors.

3.3 Softwarisation

3.3.1 Softwarisation of Federated Multi-Provider Service Layers

Future telecom infrastructures will not evolve as isolated software stacks controlled by single operators or vendors, but as federated service environments spanning telecom operators, cloud providers, edge platforms, and sector-specific infrastructures. In this context, softwarisation must evolve beyond the virtualization of individual functions toward the software-defined composition of services across administrative, technological, and sovereign boundaries. This requires common abstractions for service capabilities, interoperable lifecycle management, and programmable coordination mechanisms that allow multiple stakeholders to contribute resources and functions to a shared service layer while preserving autonomy over their own assets. Such an evolution is particularly important for European telecoms, where long-term value will increasingly come not only from operating infrastructure, but from participating in a federated service ecosystem capable of supporting cross-operator, cross-country, and cross-sector digital services. Key research challenges include service interoperability at scale, policy-aware coordination across providers, and service assurance in multi-stakeholder environments.

3.3.2 Softwarisation for Sovereign, Policy-Aware, and Data-Space-Ready Execution

A defining requirement for future European telecom software platforms is that they must not only be programmable, but also governable. Softwarisation should therefore incorporate policy, trust, locality, compliance, and data-handling constraints directly into orchestration and execution workflows across the device-edge-cloud continuum. Placement, migration, scaling, and service composition decisions must be aware of regulatory requirements, sector-specific rules, data sovereignty needs, and cross-domain trust relationships. In this perspective, software platforms evolve from neutral execution environments into policy-aware control fabrics that can enforce where data is processed, how services are composed, and under which conditions resources may be shared. This is essential for enabling large-scale public-sector, industrial, and cross-border deployments in Europe, where trust and compliance are not optional add-ons but core architectural requirements. The main research challenges concern verifiable policy enforcement, integration of orchestration with governance frameworks, and trustworthy execution across heterogeneous and federated environments.

3.3.3 Softwarisation for AI-Native Infrastructure and Network-for-AI Orchestration

The next phase of telecom softwarisation should treat AI workloads, AI pipelines, AI-serving functions, and accelerator resources as first-class infrastructure elements rather than external applications running on top of the network. Future telecom platforms will need to orchestrate not only network functions and service chains, but also training jobs, inference pipelines, feature flows, model updates, observability pipelines, and accelerator-aware execution across distributed cloud-edge-telco environments. This requires software-defined control over compute, storage, networking, memory, accelerators, and telemetry, together with lifecycle mechanisms that support joint orchestration of telecom functions and AI workloads under shared latency, energy, trust, and service constraints. In this direction, telecom infrastructure becomes both AI for networks and network for AI, enabling operators to support internal autonomy while also acting as programmable infrastructure providers for distributed AI applications. Key research challenges include joint orchestration of telco and AI workloads, model placement and lifecycle management, efficient use of heterogeneous accelerators, and balancing autonomy with explainability, trust, and operational control.

3.3.4 Softwarisation of Telecom Capability Exposure and Ecosystems

A major future direction for telecom software is the transition from internal orchestration to structured external capability exposure. Networks will increasingly need to expose communication, computing, data, sensing, and assurance capabilities through programmable service interfaces that can be consumed by third-party developers, vertical industries, and autonomous software agents. Softwarisation must therefore support capability abstraction, intent-driven service exposure, programmable APIs, SDK-oriented integration, service discovery, SLA-aware invocation, and lifecycle management of exposed telecom functions. Over time, this will extend toward marketplace-oriented ecosystems where capabilities are not only technically exposed but can also be discovered, combined, negotiated, and traded across multiple providers. This transforms telecom software from an operations-centric toolset into the foundation of an open, developer-facing, and economically active service ecosystem. The main challenges are capability abstraction without oversimplification, secure and policy-aware exposure, interoperable discovery and composition, and alignment between technical APIs, operational APIs, and service monetisation models.

3.3.5 Softwarisation at the Optical Transport Level

Optical networks are becoming an increasingly programmable foundation for future communication and computing infrastructures, especially as AI training and inference, spatial computing, and highly distributed service fabrics drive demand for ultra-high-capacity and ultra-low-latency transport. In this setting, transport layers spanning access, aggregation, metro, and core cannot remain statically configured. They must support continuous reconfiguration, resilience, and optimization through software-defined control planes that monitor, predict, and adapt the optical data plane in real time. This includes the use of high-resolution streaming telemetry, physical-layer observability, and programmable photonic components such as reconfigurable optical switches, multi-band systems, and spatial-domain multiplexing. Softwarisation is therefore essential to integrate optical resources into a unified programmable environment and to support emerging “AI for networks and networks for AI” scenarios. Key research challenges include optical-specific abstractions and data models, real-time photonic control and automation, predictive failure handling, and efficient support for dynamic AI-era transport demands.

3.3.6 Softwarisation of Robotic Services into Communications Architecture

The evolution of connected robotics is pushing cellular architectures beyond their traditional role as transport substrates and toward a more active role as programmable service environments. Future robotic systems will depend not only on connectivity, but also on access to edge computing, positioning, sensing, coordination, and adaptive service control across highly dynamic environments such as factories, hospitals, logistics hubs, cities, and remote or safety-critical scenarios. Softwarisation becomes essential because these capabilities cannot be delivered efficiently through rigid, static architectures. Instead, they require software-driven environments where functions can be instantiated, exposed, and composed according to the specific needs of the robotic task. This includes the structured exposure of communication, computing, intelligence, and governance capabilities, potentially through layered architectures such as robotic service, intelligent service, network service, and data governance planes. Key challenges include capability exposure, intent-driven coordination, policy-aware control, service assurance, and the integration of trust, privacy, safety, and compliance into the operational fabric of connected robotics.

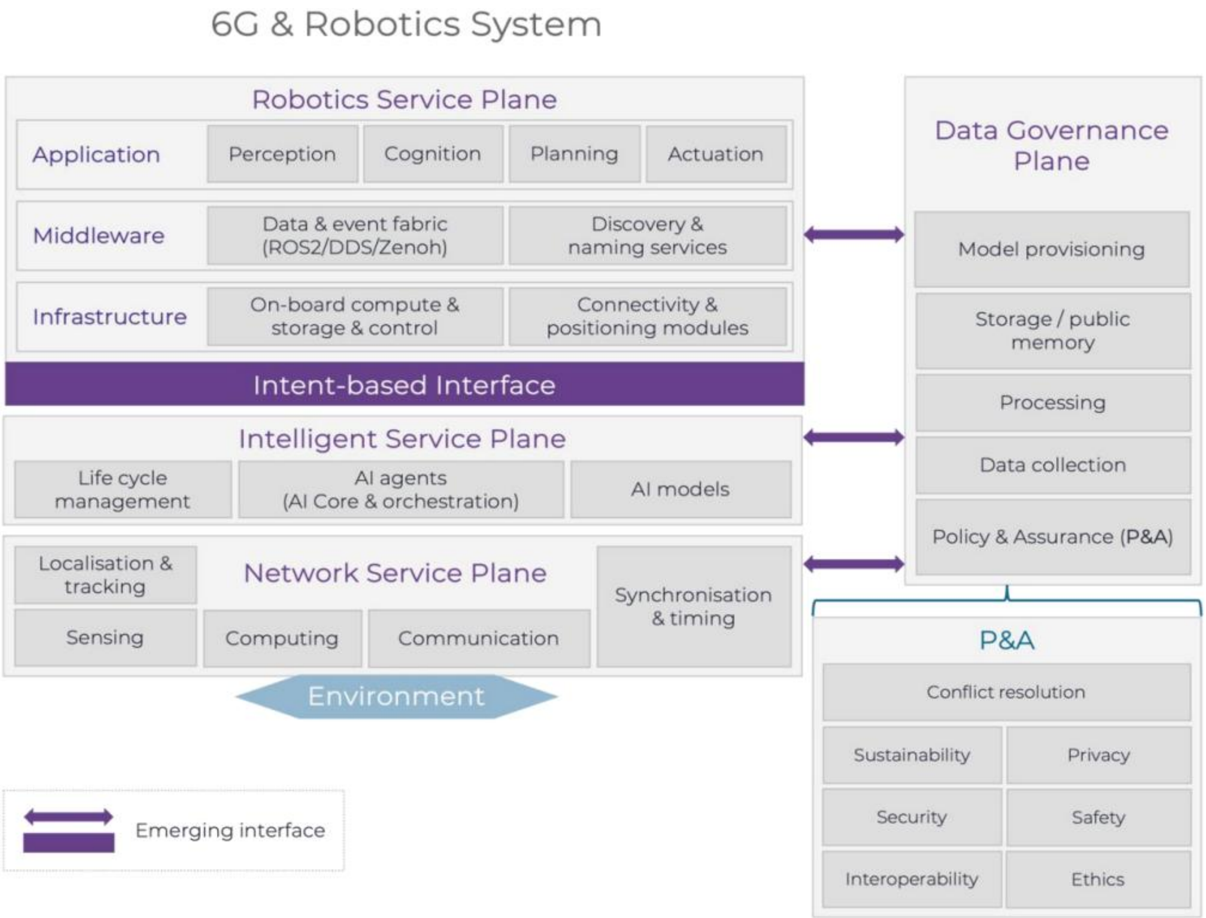

Figure 3-1: Proposed architecture to enable connected robotics [C3-01]

3.3.7 Softwarisation of the Edge-Cloud Continuum and Extreme Edge

The traditional boundaries between device, edge, and cloud domains are dissolving as networks evolve toward fully softwarized and distributed compute fabrics. Compute, storage, networking, and intelligence are no longer tied to fixed physical locations, but abstracted into a continuum of resources spanning centralized cloud infrastructures, distributed edge platforms, radio nodes, and extreme-edge environments embedded in devices and local systems. This transformation is enabled by virtualization, containerization, disaggregation, and programmable hardware, which allow heterogeneous resources to be dynamically composed into logical execution environments tailored to service intent and context. At the extreme edge, intelligence increasingly resides where data is generated and acted upon, enabling low-latency, bandwidth-efficient, and privacy-preserving services. The central challenge is to enable real-time orchestration across this continuum, where workloads are decomposed, placed, migrated, and recomposed according to context, performance objectives, and user intent. This requires joint compute-network optimisation, strong observability and state consistency, AI-native orchestration, resilience under failures and intermittency, and coordinated federation across multiple operators, providers, and administrative domains.

3.3.8 Open Platforms for Infrastructure Management

Over the next five to ten years, the role of software platforms for connectivity infrastructure will shift from enabling programmability and orchestration to supporting autonomous, trustworthy, and self-evolving infrastructure management, targeting higher levels of operational autonomy. Such platforms should span devices, networks, edge, cloud, and accelerators, and provide capability-aware control across heterogeneous environments. Research should address post-P4 programmability, semantic resource models, open APIs for control, telemetry and data exposure, cross-layer observability, runtime digital twins, and AI-native control planes capable of safe, explainable, and policy-compliant decision-making. At the same time, these platforms must support deterministic performance across virtualized heterogeneous environments and self-evolving, verifiable software life cycles through CI/CD/CT pipelines, automated certification, runtime assurance, rollback-capable upgrades, and safe in-service evolution of software, policies, and AI components. In the long term, the goal is a self-optimising, interoperable, and sustainable management platform that can continuously evolve while preserving resilience, trust, and operational control at planetary scale.

3.3.9 Softwarisation in NTN Architectures

The evolution of 6G relies increasingly on software-defined control across multi-domain environments, creating programmable architectures that integrate terrestrial, non-terrestrial, and even quantum-secure infrastructures. Extending softwarisation into NTN environments requires solutions compatible with highly dynamic topologies and resource-constrained nodes, since traditional virtualization approaches are often too heavy for space platforms with limited power and memory. Future NTN softwarisation must therefore support lightweight execution environments, continuous telemetry-driven adaptation, and multi-vendor interoperability across orbital systems. As constellations scale, satellites will need to operate as autonomous clusters over selected geographical areas, enforcing spectrum and policy constraints with limited reliance on ground segments. This also drives the adoption of agentic AI for NTN orchestration, while creating new challenges around explainability, stateful coordination, conflict resolution, and mission-critical assurance. In parallel, cost-efficient multi-stakeholder federation will require standard APIs, secure policy-constrained information exchange, and interoperable controller hierarchies to enable resilient, scalable, and globally distributed “Satellite-as-a-Service” ecosystems.

3.3.10 Synopsis of Required Research

Research Theme	***Softwarisation***		
Research Topic	**Timeline**	**Key outcomes**	**Contributions / Value**
S1: Softwarisation of Federated Multi-Provider Service Layers	Mid-term	Common abstractions for service capabilities, interoperable lifecycle management, programmable coordination across telecom, cloud, edge, and sector platforms, policy-aware orchestration across sovereign and administrative boundaries.	Enables cross-provider service composition, strengthens European interoperability, and supports federated digital services across operators, countries, and sectors while preserving stakeholder autonomy.
S2: Softwarisation for Sovereign, Policy-Aware, and Data-Space-Ready Execution	Mid-term	Policy-aware placement, migration, scaling, and service composition; trust- and locality-aware execution workflows; verifiable governance of data handling and resource sharing.	Makes softwarized infrastructures governable and compliant by design, supporting public-sector, industrial, and cross-border deployments with strong sovereignty, trust, and data governance requirements.

S3: Softwarisation for AI-Native Infrastructure and Network-for-AI Orchestration	Mid/Long-term	Joint orchestration of network functions, AI workloads, feature flows, model updates, observability pipelines, accelerators, and inference/training placement across distributed cloud-edge-telco environments.	Positions telecom infrastructure as both AI for networks and network for AI, enabling internal autonomy, distributed AI services, and tighter convergence of connectivity and computing.
S4: Softwarisation of Telecom Capability Exposure, Developer Platforms, and Marketplace Ecosystems	Mid-term	Capability abstraction, intent-driven service exposure, programmable APIs, SDK-oriented integration, service discovery, SLA-aware invocation, lifecycle management of exposed telecom functions, marketplace-oriented service interfaces.	Turns telecom software from an operations-centric toolset into an open, developer-facing, and economically active service ecosystem supporting new B2B and B2B2C models.
S5: Softwarisation at the Optical Transport Level	Mid-term	Resource, device, and system data models including physical-impairment information and photonic device abstractions; optical network control architectures enabling digital twins and AI/ML-assisted automation; dynamic and reconfigurable optical networks for inter- and intra-data-centre transport, including support for AI training and inference fabrics.	Enables dynamic, resilient, and sustainable optical infrastructures with fine-grained programmability, predictive control, and support for future ultra-high-capacity AI-era transport.
S6: Softwarisation of Robotic Services into Communications Architecture	Mid-term	Open reference architecture for robotics-over-cellular systems; structured exposure of communication, sensing, positioning, edge computing, and intelligence capabilities; governable interaction across robotic, intelligent, network, and data governance planes.	Enables trustworthy, scalable, and interoperable robotic services over cellular networks while preserving safety, privacy, compliance, and human oversight.

S7: Softwarisation of the Edge-Cloud Conti-nuum and Extreme Edge	Mid/Long-term	Unified compute fabric spanning device, edge, and cloud with dynamic workload placement; context-aware orchestration across network, compute, and energy dimensions; autonomous workload placement driven by intent; fully decentralised collaborative edge intelligence; common abstractions for heterogeneous resources and accelerators.	Improves QoE, network efficiency, privacy, sustainability, and cost efficiency; reduces operational complexity; supports scalable infrastructures, digital sovereignty, and robustness in crisis or intermittent-connectivity scenarios.
S8: Software Foun-dations for Distributed Conti-nuum Intelligence	Mid-term	Common abstractions and data models for communication, computing, storage, and service resources across the edge-to-cloud continuum; programmable control and lifecycle mechanisms supporting distributed intelligent functions and agents; interoperable software frameworks for deploying, updating, and coordinating intelligence across heterogeneous environments; support for capability exposure and dynamic composition of infrastructure and service functions.	Provides the software foundation required to deploy and coordinate intelligence across heterogeneous infrastructures; improves programmability, interoperability, portability, adaptive service support, and more efficient use of distributed resources across the continuum.

S9: Open Platforms for Infrastructure Management	Long-term	Post-P4 universal programmability across network, compute, storage, memory, and accelerators; AI-native control planes with safety, explainability, and policy compliance; semantic infrastructure APIs; cross-layer observability and causality; runtime digital twins; trustworthy autonomy through formal verification, zero-trust orchestration, and confidential computing; deterministic softwarized systems; self-evolving and verifiable software life cycles through CI/CD/CT, automated certification, runtime assurance, rollback-capable upgrades, and safe in-service evolution of software, policies, and AI components; federated orchestration across operators and sovereign domains; sustainable infrastructure intelligence.	Creates autonomous, trustworthy, interoperable, and sustainable platforms for connectivity infrastructure, enabling continuous evolution, reduced operational complexity, stronger resilience, and higher levels of autonomy.
S10: SDN over Multi-Domain Networks (TN-NTN)	Short-term	Unified multi-layer programmability through standardized control interfaces; telemetry-driven dynamic orchestration via distributed SDN controller hierarchies; cross-domain coexistence mechanisms for diverse traffic and service types.	Breaks down technological silos, enables secure federation and resource sharing, improves mission-critical resilience during outages, and allows on-demand resource adaptation across TN-NTN environments.
S11: Softwarisation Solutions Compatible with Resource-Constrained Satellites	Short-term	Ultra-lightweight virtualization environments; alternative softwarisation technologies for space nodes; topological optimization of orchestration roles allowing adaptive control-plane role changes in orbit.	Maximizes on-board efficiency, improves autonomous orbital resilience, and supports agile, secure, multi-tenant “Satellite-as-a-Service” models.
S12: Autonomous Cluster Deployments in NTN over Selected Geographical Areas	Mid-term	Satellite-to-satellite clustering, dynamic geofencing frameworks, orbital control-plane handovers, autonomous regional operation with limited ground dependence.	Supports automated regulatory compliance, resilient and localized service continuity, and reduced handover latency through geographically scoped in-orbit orchestration.

S13: Multi-Stakeholder Federation of Multiple NTN Clusters toward Cost-Efficient Solutions	Long-term	Automated transfer interfaces, policy-constrained information exchange, and Federated Space Systems architectures integrating isolated operator-specific clusters into shared infrastructures.	Reduces capital expenditure through resource pooling, enables secure multi-operator services, and supports scalable marketplace-oriented “Satellite-as-a-Service” ecosystems.

3.3.11 Recommendations for Actions/Instruments

Research Theme	***Softwarisation***												
Action	**S1**	**S2**	**S3**	**S4**	**S5**	**S6**	**S7**	**S8**	**S9**	**S10**	**S11**	**S12**	**S13**
International Calls	x	x	x	x	x	x	x	x	x	x	x	x	x
International Research	x	x	x	x	x	x	x	x	x	x	x	x	x
Open Data			x	x	x		x	x	x				
Large Trials	x		x	x	x	x	x	x		x		x	x
Cross-domain research	x	x	x	x	x	x	x	x	x	x	x	x	x

3.4 Intelligence

3.4.1 Semantic Interoperability and Goal Negotiation for Agentic Systems

As telecom infrastructures evolve toward federated multi-provider environments, intelligence must move beyond isolated decision-making and basic message exchange toward semantic interoperability across domains, including operators, cloud providers, edge platforms, and vertical systems. Future agents will need to understand not only messages, but also the meaning of goals, constraints, capabilities, policies, and service-level commitments expressed by agents operating in different administrative and technological domains. This requires shared ontologies, machine-readable capability models, negotiation protocols, and reasoning mechanisms that support autonomous goal alignment, conflict resolution, and cross-domain service composition. Such capabilities will be essential if distributed intelligence is to underpin a European cloud-edge-telco continuum in which services are dynamically composed across multiple stakeholders while remaining interoperable, governable, and trustworthy.

3.4.2 Telecom Foundation Models, Knowledge Planes, and Multi-Modal Reasoning

Future intelligent telecom infrastructures will require more than task-specific AI components. They will increasingly depend on telecom-tuned foundation models and knowledge-centric reasoning layers able to integrate telemetry, topology, service state, policies, digital twins, operational history, and external context into a coherent decision fabric. This direction can support advanced functions such as intent decomposition, troubleshooting, anomaly explanation, predictive planning, and adaptive orchestration across heterogeneous infrastructures. Key research challenges include grounding foundation models in telecom-specific semantics, combining symbolic and data-driven reasoning, supporting multi-modal inputs, enabling efficient deployment across cloud-edge environments, and ensuring trustworthy and explainable outputs. This topic also aligns with the broader European push toward stronger AI infrastructure, better data availability, and more usable AI development environments.

3.4.3 Federated, Privacy-Preserving, and Sovereign Intelligence Across Domains

A major future direction is intelligence that can learn, reason, and coordinate across domains without requiring all data, models, or decision logic to be centralized. In a European context, this means enabling federated and privacy-preserving intelligence across operators, public infrastructures, industrial domains, edge platforms, and cloud environments while respecting sovereignty, locality, regulatory, and contractual constraints. Future research should address federated and distributed learning, secure collaborative inference, policy-constrained model exchange, provenance and trust for shared intelligence components, and coordination mechanisms that preserve privacy and autonomy while still enabling collective optimization. This is particularly relevant for Europe's goal of building federated, trustworthy, and energy-efficient cloud-edge ecosystems and for the expected development of a stronger internal market for data and AI-enabled services.

3.4.4 Imagination as an Agentic Tool

Imagination is emerging as a core agentic capability by embedding action-conditioned world models directly within the control loop. Unlike conventional digital twins, which primarily serve as synchronized mirrors for monitoring and replay, world models act as internal generative engines that capture causal dynamics and allow agents to explore counterfactual futures before taking action. This enables predictive deliberation, safer decision-making, and more efficient adaptation under uncertainty. The main challenge is to develop dynamical models that support fast, uncertainty-aware simulation within near-real-time latency budgets while remaining computationally efficient and sufficiently faithful to real system behaviour. If achieved, imagination-like capabilities could transform intelligent infrastructures from reactive optimizers into proactive systems capable of evaluating alternative courses of action before operational changes are made.

3.4.5 Intelligence at the Transport Network Level

Autonomous transport networks will become increasingly necessary as optical and packet infrastructures grow too large, too dynamic, and too mission-critical for human-driven management alone. With reconfigurable photonic components, dynamic slices, edge AI clusters, and orbital optical systems, future transport intelligence must support real-time performance optimization, fault correlation, and cross-layer adaptation under demanding physical-layer constraints. This requires control-plane intelligence capable of combining physical-layer telemetry, predictive analytics, digital twins, and multi-agent coordination to continuously tune paths, modulation formats, power levels, and service behaviour. The key challenge is to build transport intelligence that remains stable and trustworthy despite large numbers of micro-scale changes in the underlying optical and packet layers. Such intelligence is essential for future transport systems that must support both AI-enabled operations and growing "network for AI" traffic demands.

3.4.6 World Models for Ubiquitous Digital Twins

Digital Twins are evolving from static or offline replicas used mainly for design-time validation into continuously synchronized, operationally embedded models of physical systems. As they become enriched with learned or hybrid world models, they can move beyond monitoring to support predictive and prescriptive simulation, allowing operators and autonomous agents to evaluate actions before deployment in the real world. This enables "decision-before-action" paradigms in which failures, reconfigurations, policy changes, or optimization strategies can be tested in a virtual environment before execution. The core challenges are to support high-fidelity modelling of complex multi-domain systems, near-real-time synchronization across distributed infrastructures, scalable simulation, explainability, uncertainty quantification, and trustworthy decision support. Ubiquitous Digital Twins will be foundational for autonomous infrastructures spanning telecom, cloud-edge, physical systems, and broader socio-technical environments.

3.4.7 Distributed Continuum Intelligence for Autonomous Infrastructures

Future communication and computing infrastructures will increasingly depend on intelligence that is not confined to a single network segment or control layer, but distributed across devices, communication systems, edge platforms, cloud resources, and service environments. In this perspective, intelligence becomes a continuum-wide capability that enables the coordinated operation of heterogeneous resources across the edge-to-cloud continuum, supporting autonomous operations, dependability and value-generating service delivery. Key research challenges include architectures and mechanisms that allow service, device, and infrastructure agents to coordinate decisions across communication, computing, storage, and service resources in real time, while aligning local decisions with broader operational, service, and policy objectives. At the same time, such intelligence must remain dependable, resilient, and transparent under failures, uncertainty, and partial visibility. Ultimately, distributed continuum intelligence should support not only technical optimization, but also adaptive service delivery, intelligent capability exposure, and multi-stakeholder value creation.

3.4.8 Telecom AI Testing and Verification

The integration of AI into telecom infrastructure is accelerating faster than assurance and validation frameworks can follow. This creates a structural risk, since AI-enabled and AI-native functions are increasingly being trusted with operational decisions in complex, real-time environments. Future research must therefore establish methodologies, test environments, and tools for validating the safety, trustworthiness, compliance, and operational reliability of AI-in-telecom systems, including agentic AI architectures. The challenge is especially acute because such systems are probabilistic, adaptive, and often opaque, making them difficult to assess with traditional telecom validation methods. This topic is strategically important not only for operational risk mitigation, but also for sovereign capability development in Europe, where the ability to independently test, certify, and benchmark AI-native infrastructures will increasingly shape competitiveness, trust, and regulatory readiness.

3.4.9 Digital Identity for Embodied and Digital Agents

Future intelligent telecom systems will require a unified, scalable, and trustworthy digital identity framework for new identity consumers beyond traditional subscribers, including embodied and digital AI agents. This involves mechanisms for agent identification, authentication, authorization, and credential management that go beyond device-centric models to support autonomous and learning entities. A central challenge is enabling verifiable credentials for dynamically acquired skills and capabilities while ensuring safety, trust, and regulatory compliance. Interoperability across telecom, over-the-top (OTT), cloud, and enterprise ecosystems will be essential, requiring common identity models, cross-domain trust frameworks, and support for both SIM-based and non-SIM-based agents. In this context, network operators may evolve toward a stronger role as trusted identity and assurance providers for future agentic ecosystems.

3.4.10 AI Agent Registration, Discovery, and Communication

Future distributed intelligence will depend on scalable and interoperable mechanisms for AI agent registration, discovery, and communication across heterogeneous infrastructures. Agents will need to advertise capabilities, discover appropriate collaborators, and form dynamic relationships across administrative and technological domains. This requires more than a simple directory service: future frameworks must support intent-aware discovery, context-sensitive cooperation, group formation, and secure communication across different protocols, modalities, and ownership domains. A major challenge is to ensure that agent communication remains efficient, interoperable, and trustworthy at scale while also supporting flexible composition of temporary and persistent agent ecosystems. Such capabilities are foundational for multi-agent telecom environments in which coordination and collaboration increasingly replace isolated decision loops.

3.4.11 Quality of Experience for Agentic Traffic

The emergence of agentic traffic requires a substantial rethink of Quality of Experience concepts. Traditional QoE models are centred on human perception, but future networks will increasingly need to optimize for machine-driven interactions in which success depends on task completion, workflow latency, distributed inference reliability, and agent collaboration efficiency. This requires new QoE metrics and management mechanisms able to capture intent, context, and task-level value rather than only user-perceived service quality. The challenge is to integrate these new QoE concepts with AI-driven network control and application-layer intelligence so that resources can be dynamically allocated according to agent objectives and operational priorities. This will be essential for future infrastructures supporting large numbers of autonomous service interactions.

3.4.12 Dynamic and Generative Networking

Dynamic and generative networking introduces a shift from static provisioning toward the on-demand generation of localized, purpose-specific, and task-oriented connectivity environments. In this view, networks are no longer pre-established and only adjusted incrementally; instead, they can be instantiated, adapted, and dissolved in response to agent goals and changing conditions. This is a more radical form of network intelligence than traditional slicing, since it must support large numbers of short-lived, fine-grained, and outcome-specific network instances. Key challenges include intent-driven generation, autonomous lifecycle management, isolation, privacy, and seamless integration across heterogeneous infrastructures. If realized, this direction could enable highly agile network behaviours better aligned with future agentic services and dynamic digital ecosystems.

3.4.13 Agent Empowerment (Computing-, Sensing-, Data-as-a-Service)

Future networks are expected to evolve into distributed service platforms that empower devices and applications, including AI agents, with on-demand access to computing, sensing, and data resources. This requires intelligence capable of abstracting and exposing network and edge capabilities as standardized, discoverable services, together with secure and efficient invocation mechanisms aligned with intent and context. The challenge is not simply to expose resources, but to enable dynamic service composition across device, edge, cloud, and multi-provider environments while ensuring data trustworthiness, privacy preservation, and efficient resource use. This direction is important because it shifts intelligent infrastructures from optimizing internally to actively enabling new adaptive services, operational models, and ecosystem interactions.

3.4.14 AI Reliability, Trust, and Trusted Agentic 6G Architectures

Ensuring reliability and trust in AI-native network architectures is one of the most important research challenges in future intelligent telecom systems. AI-driven decision-making must be combined with strong verification, assurance, and monitoring mechanisms capable of continuously validating model outputs against objective criteria and operational policies. This is especially important in agentic 6G architectures where different AI components may interact, adapt, and influence network behaviour at runtime. A key challenge is the coexistence of trusted deterministic mechanisms and less predictable AI-driven components, with the ability to fall back to rule-based behaviour when needed. Future work must therefore address runtime assurance, continuous performance validation, scalable low-noise verification, and user- or policy-centric trust establishment for AI-enabled infrastructures.

3.4.15 Intelligence in NTN Architectures

The evolution toward 6G is pushing NTN systems toward intelligence-native operation in which multiple distributed agents cooperate under dynamic topologies, partial observability, intermittent connectivity, and strict resource constraints. Future NTN intelligence must therefore support decentralized multi-agent operation, robust coordination and orchestration protocols, and hybrid computation strategies that determine where to compute, what to share, and how to balance local and offloaded processing across device-edge-satellite-cloud environments. In addition, NTN intelligence must remain frugal, since communication, sensing, and computation compete for limited spectrum, energy, and processing resources. This makes intent-aware resource allocation, AI-traffic management, topology-adaptive reassignment, and coexistence mechanisms essential. The overarching challenge is to deliver scalable, resilient, privacy-preserving, and energy-efficient intelligence across NTN ecosystems while maintaining service continuity, mission-critical reliability, and cross-domain interoperability.

3.4.16 Synopsis of Required Research

Research Theme	***Intelligence***		
Research Topic	**Timeline**	**Key outcomes**	**Contributions / Value**
I1: Semantic Interoperability and Goal Negotiation for Cross-Domain Agentic Systems	Mid-term	Shared ontologies, machine-readable capability models, negotiation protocols, and reasoning mechanisms for autonomous goal alignment, conflict resolution, and cross-domain service composition.	Enables interoperable multi-agent coordination across telecom, cloud, edge, and vertical domains, supporting governable and trustworthy service federation.
I2: Telecom Foundation Models, Knowledge Planes, and Multi-Modal Reasoning	Mid/Long-term	Telecom-tuned foundation models, knowledge-centric reasoning layers, multi-modal inputs, and explainable decision-support mechanisms integrating telemetry, topology, policies, twins, and operational context.	Strengthens advanced intent handling, troubleshooting, anomaly explanation, predictive planning, and adaptive orchestration across heterogeneous infrastructures.
I3: Federated, Privacy-Preserving, and Sovereign Intelligence Across Domains	Mid-term	Federated and distributed learning, secure collaborative inference, policy-constrained model exchange, provenance, trust, and coordination mechanisms that preserve privacy and autonomy.	Supports European sovereignty, privacy-preserving optimization, and trusted cross-domain intelligence across operators, public infrastructures, industry, edge, and cloud.
I4: Imagination as an Agentic Tool	Mid-term	Action-conditioned world models, uncertainty-aware near-real-time simulation, and counterfactual reasoning support for agents inside operational control loops.	Enables predictive deliberation and safer, more proactive autonomous decision-making beyond reactive optimization.

I5: Intelligence at the Transport Network Level	Mid-term	Physics-driven ML models, digital twins, and multi-agent control for real-time topology reconfiguration, photonic tuning, layer-crossing optimization, predictive maintenance, and dynamic spectrum / wavelength / waveband / core / fibre allocation.	Supports highly autonomous transport networks, real-time self-calibration, fault correlation, and efficient Tb/s-scale transport for future AI-era infrastructures.
I6: World Models for Ubiquitous Digital Twins	Mid/Long-term	Near-real-time synchronized digital twins, what-if simulation APIs, predictive and prescriptive simulation environments, uncertainty-aware decision support, and federated twin architectures.	Reduces operational uncertainty, supports safer closed-loop automation, and extends digital twins toward multi-domain autonomous infrastructures.
I7: Distributed Continuum Intelligence for Autonomous, Dependable, and Value-Generating Infrastructures	Mid-term	Distributed intelligence frameworks spanning devices, network functions, edge, cloud, and service platforms; cooperative orchestration; dependability-aware control loops; value-aware capability exposure and coordination.	Makes intelligence a continuum-wide capability, improving autonomy, resilience, dependability, and multi-stakeholder value creation across heterogeneous infrastructures.
I8: Telecom AI Testing and Verification	Short/Mid-term	Formal verification methods for AI-enabled and AI-native networks, regulatory sandbox approaches, common VV&T methodologies, and sovereign testing and certification tools.	Builds trust in AI-native telecom systems, mitigates operational risk, supports sovereign validation capability, and strengthens European competitiveness.
I9: Digital Identity for Embodied and Digital Agents	Mid-term	Unified identity frameworks, verifiable credentials for dynamic skills and capabilities, cross-domain trust models, SIM-based and non-SIM-based support, and task-aware authorization.	Enables trusted participation of embodied and digital agents across telecom, OTT, cloud, and enterprise ecosystems.

I10: AI Agent Registration, Discovery, and Communication	Mid-term	Capability registration, intent-aware discovery, context-aware group formation, cross-network and cross-protocol communication, and scalable secure inter-agent interaction.	Enables interoperable collaboration among operator-managed and external agents across heterogeneous environments.
I11: Quality of Experience for Agentic Traffic	Mid-term	QoE metrics for machine-driven interactions, including task completion success, latency sensitivity, reliability of distributed inference, and efficiency of agent collaboration; intent-aware QoE management.	Aligns resource allocation with agent goals and future machine-driven service requirements rather than only human-centric service perception.
I12: Dynamic and Generative Networking	Mid/Long-term	On-demand domain networks, task-driven and localized connectivity environments, autonomous lifecycle management, privacy and isolation mechanisms, and dynamic instantiation / dissolution of networks.	Moves networking toward highly dynamic, purpose-specific, agent-oriented operation beyond conventional slicing.
I13: Agent Empowerment (Computing-, Sensing-, Data-as-a-Service)	Mid-term	Standardized abstraction and exposure of computing, sensing, and data capabilities; secure service invocation; dynamic orchestration and autonomous service composition across device-edge-cloud resources.	Turns intelligent infrastructure into an active enabler of adaptive services, ecosystem interactions, and distributed capability consumption.
I14: AI Reliability, Trust, and Trusted Agentic 6G Architectures	Mid-term	Runtime assurance, continuous validation, trusted monitoring frameworks, fallback mechanisms between deterministic and AI-driven behaviour, scalable low-noise verification, and policy-centric trust establishment.	Establishes trust in AI-native and agentic 6G architectures while enabling robust and controllable autonomous operation.

I15: Multi-AI-agentic NTNs for Decentralized and Autonomous Systems	Short-term	Distributed NTN entities operating without centralized bottlenecks, cooperative multi-agent coordination, adaptive response to evolving network conditions, fair and intent-aware resource allocation, and more energy-efficient operation. **[KPIs]:** Control-loop latency <500 ms to react to constellation-wide failure; Post-disaster NTN reconfiguration <10 min	Enables zero-touch decentralized NTN operation with improved scalability, responsiveness, and sustainability under highly dynamic topologies
I16: AI-agentic Management and Protocols for Coordination and Orchestration in NTN	Mid-term	Standardized cross-layer control interfaces, intent-aware decentralized frameworks, dynamic telemetry-driven orchestration, and coexistence / AI-traffic protocols across terrestrial and NTN domains. **[KPIs]:** Service activation <5 min; signalling reduction ≥50%; Routing convergence <200 ms; feeder-link switch-over <100 ms	Eliminates centralized bottlenecks, improves resilience and reliability, and enables seamless multi-domain integration and intelligent orchestration of NTN resources
I17: Hybrid Computation Strategy in NTN	Short-term	Dynamic task-offloading frameworks, goal-oriented data representation, cross-layer decision engines, and cooperative inference models across device-edge-satellite-cloud environments. **[KPIs]:** Offload-driven latency mitigation & device battery-life extension; share only necessary features	Extends battery life, mitigates NTN latency, and preserves privacy by intelligently distributing computation and sharing only necessary features
I18: Frugal and Dynamic Resource Management to Support Future NTN Intelligence	Mid-term	Intent-aware resource optimization, asymmetric AI-traffic protocols, coexistence mechanisms for shared infrastructure, and topology-adaptive resource reassignment under time-varying visibility and fast-moving nodes. **[KPIs]:** Signalling-overhead reduction ≥50% via intent-aware optimization; Energy-efficiency	Supports energy-efficient operation, removes signalling bottlenecks, harmonizes competing subsystems, and maintains mission-critical stability and responsiveness in intelligent NTN ecosystems

		improvement +20–30% from frugal resource management	

3.4.17 Recommendations for Actions

Research Theme	***Intelligence***																	
Action	**I 1**	**I 2**	**I 3**	**I 4**	**I 5**	**I 6**	**I 7**	**I 8**	**I 9**	**I 10**	**I 11**	**I 12**	**I 13**	**I 14**	**I 15**	**I 16**	**I 17**	**I 18**
International Calls	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
International Research	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
Open Data	x	x		x	x	x	x	x			x		x	x	x	x	x	x
Large Trials	x		x		x	x	x	x	x	x	x	x	x	x	x	x	x	x
Cross-domain research	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

3.5 Security Considerations

The security considerations of this chapter focus specifically on the security implications introduced by the softwarisation and intelligence directions discussed in Sections 3.3 and 3.4. They do not aim to cover security for the full SRIA in a generic sense. Instead, they identify the assets, mechanisms, services, and management requirements that become particularly relevant when infrastructures evolve toward federated multi-provider service layers, policy-aware execution, AI-native orchestration, programmable capability exposure, open infrastructure platforms, digital twins and world models, distributed continuum intelligence, trusted agentic architectures, and intelligent NTN operation. In this context, security must support not only protection of software-defined and AI-enabled infrastructures, but also trustworthy coordination across multiple stakeholders, domains, and trust boundaries.

3.5.1 Related/treated security assets & potential resulting security needs

Research Theme / Research Topic	Security asset	Explanation
Federated softwarized and AI-native control fabrics	Critical system / component; Critical data	Federated service layers, policy-aware orchestration, AI-native control loops, and open management platforms introduce critical assets such as service graphs, placement decisions, orchestration policies, telemetry streams, model outputs, and runtime assurance states. These must be protected against manipulation, unauthorized reconfiguration, leakage, poisoning, and integrity loss across multi-provider environments.
Agent identities, credentials, discovery registries, and coordination channels	Critical data; Critical system / component	Agentic architectures introduce identities, credentials, reputation/context information, capability descriptions, discovery mechanisms, and communication channels that become core trust anchors for distributed intelligence. These assets require strong authentication, authorization, attestation, anti-impersonation protection, and secure group coordination across domains.
Digital twins, world models, and knowledge planes	Critical data; Critical system / component	Continuously synchronized twins, predictive models, and knowledge-centric reasoning layers influence operational decisions and therefore become high-value security targets. Their data pipelines, learned models, synchronization states, and decision traces must be protected against tampering, drift abuse, adversarial manipulation, and unauthorized access.
Exposed telecom capabilities, APIs, and service interaction layers	Critical system / component; Critical data	Capability exposure, service APIs, and marketplace-oriented interfaces create externally reachable control and business surfaces. These assets require secure mediation, fine-grained authorization, policy enforcement, service integrity, SLA-aware protection, and traceable usage across operators, developers, verticals, and autonomous agents.
Intelligent NTN control, hybrid computation, and frugal resource management	Critical system / component; Critical data	NTN intelligence introduces decentralized decision logic, offloading choices, topology-adaptive control, AI-traffic handling, and resource coordination under strict constraints. These assets must remain resilient to spoofing, disruption, false telemetry, malicious offloading decisions, and trust breakdowns in dynamic and intermittent environments.

3.5.2 Proposed security mechanisms

Research Theme / Research Topic – Proposed security mechanism	**Security mechanism type / toolbox**	**Explanation**
Federated identity, attestation, and authorization for agents and platforms	Cryptography; System engineering; Ethical, legal, regulatory	Research should deliver cross-domain identity, attestation, credential validation, and fine-grained authorization mechanisms for software functions, AI agents, digital twins, and exposed capabilities. These mechanisms should support federated environments while preserving autonomy of individual stakeholders.
Policy-aware and zero-trust orchestration enforcement	System engineering; Cryptography; Ethical, legal, regulatory	Policy-aware execution and multi-provider service composition require zero-trust orchestration mechanisms that continuously verify identities, enforce locality and policy constraints, and validate whether orchestration decisions remain compliant and trustworthy at runtime.
Provenance, integrity, and auditability for models, telemetry, intents, and decisions	Cryptography; System engineering	Research should support end-to-end integrity protection, provenance tracking, secure logging, and auditability of telemetry flows, model artifacts, intent translations, policy decisions, twin synchronization, and autonomous actions.
Runtime assurance and secure VV&T for AI-native and agentic systems	System engineering; Ethical, legal, regulatory	AI-enabled infrastructures require runtime assurance, secure test and verification pipelines, regulatory sandbox support, and independent validation mechanisms for probabilistic and adaptive decision systems.
Resource-aware and adaptive security for edge and NTN environments	Cryptography; System engineering; Ethical, legal, regulatory	Security mechanisms must adapt to constrained compute, energy, and connectivity conditions, particularly at the extreme edge and in NTN. This includes lightweight protections, adaptive crypto usage, secure offloading support, and security reconfiguration according to resource availability and threat context.

3.5.3 Proposed security services

Research Theme / Research Topic – Proposed security service	**Security service description**	**Explanation**
Trusted coordination and policy-compliant assurance service	Authentication, authorization, attestation, policy validation, secure coordination, and auditable decision logging for agents, functions, and platforms	This service should support operators, service providers, infrastructure owners, vertical actors, and developers interacting across federated softwarized and intelligent environments. It should be interoperable, scalable, low-latency, and verifiable.
Federated agent identity and discovery assurance service	Secure agent onboarding, credential management, capability validation, trust establishment, and protected discovery/communication	This service is specifically needed for embodied and digital agents operating across telecom, cloud, OTT, and enterprise domains.
Model, telemetry, and twin integrity service	Integrity protection, provenance tracking, synchronization assurance, and tamper detection for telemetry, world models, digital twins, and knowledge layers	This service protects the reliability of decision-support and decision-before-action mechanisms that increasingly shape operations.
Secure capability exposure and API mediation service	Authorization, policy enforcement, SLA-aware access control, traffic inspection, and auditable interaction for exposed capabilities and APIs	This service protects telecom capability exposure, developer-facing interfaces, and marketplace-oriented service ecosystems.
Resilient NTN intelligence assurance service	Secure decentralized coordination, offloading assurance, topology-aware trust handling, and mission-critical fallback	This service is particularly relevant for dynamic NTN environments with intermittent connectivity and constrained resources.

3.5.4 Security management mechanisms and related requirements

Research Theme / Research Aspect	Security management type	Explanation
Federated agent and platform trust management	Key management; credential lifecycle; revocation; trust-anchor updates	Chapter 3 requires lifecycle management of identities, credentials, trust anchors, and permissions for software functions, AI agents, orchestration systems, and exposed services across multiple domains.
Secure lifecycle management of software, policies, models, and twins	Security updates; policy/model updates; rollback; mechanism reconfiguration	AI-native and self-evolving environments require secure update, validation, rollback, and reconfiguration processes for software, policies, model artifacts, telemetry pipelines, and digital twin components.
Incident response and runtime containment for distributed intelligence	Mechanism reconfiguration; incident response; trust re-establishment	Distributed intelligent control loops require rapid isolation, degraded-mode operation, fallback to deterministic behaviour, and coordinated recovery after failures, compromise, or anomalous model behaviour.
Resource-aware security adaptation	Key management; policy adaptation; mechanism reconfiguration	Edge and NTN environments require continuous adjustment of protection mechanisms according to energy, compute, latency, and connectivity constraints without collapsing the security posture.

3.6 Sustainability Considerations

The sustainability considerations of this chapter focus specifically on the sustainability implications introduced by the softwarisation and intelligence directions discussed in Sections 3.3 and 3.4. They do not aim to provide a generic treatment of sustainability across the full SRIA. Instead, they address whether the software-driven and AI-enabled infrastructures proposed in this chapter deliver a net positive effect once orchestration overhead, telemetry pipelines, model training and inference, digital twin synchronization, continuous validation, and security mechanisms are taken into account. In this context, sustainability must be evaluated not only in terms of reduced energy use during operation, but also in terms of resource efficiency, avoided data movement, reduced overprovisioning, longer infrastructure usefulness, and the ability to place capabilities where they create the most value with the lowest acceptable resource cost.

Research Theme / Research Aspect	**Actual impact**	**Relation to sustainability**
AI-native infrastructure and network-for-AI orchestration	Joint orchestration of telecom functions, AI workloads, and accelerators can reduce overprovisioning, improve workload placement, and better match compute and communication resources to actual demand. However, it also introduces significant model execution, observability, and orchestration overhead. Sustainability therefore depends on energy-aware placement, accelerator efficiency, selective telemetry, and lifecycle-aware orchestration.	Sustainability dependency
Softwarisation of the edge-cloud continuum and distributed continuum intelligence	Distributed placement of functions, intelligence, and services across device-edge-cloud infrastructures can reduce unnecessary data movement, support locality-aware processing, and improve the efficiency and usefulness of heterogeneous resources. At the same time, excessive coordination overhead, duplicated processing, and poorly designed multi-agent control loops can offset these gains.	Helps with sustainability
World models, ubiquitous digital twins, and AI testing/verification	Digital twins, predictive models, and simulation-first validation can reduce outages, failed deployments, unnecessary experimentation on live systems, and excessive safety margins. This can improve service continuity and reduce waste. However, these benefits depend on controlling the computational cost of synchronized twins, simulations, and verification pipelines.	Helps with sustainability
Agent empowerment and capability exposure	Exposing computing, sensing, and data capabilities as services can improve reuse of shared infrastructure and reduce duplicate deployments by enabling composable service ecosystems. Sustainability gains depend on efficient discovery, service composition, and avoidance of redundant invocation or excessive mediation overhead.	Helps with sustainability
Intelligence in NTN architectures and frugal resource management	Resource-aware NTN intelligence can improve energy efficiency, offloading choices, and spectrum/computation coordination in constrained environments. The chapter already points to frugal and dynamic resource management as a core need. Sustainability depends on intent-aware allocation, lightweight models, and careful balancing between local processing and offloaded intelligence.	Dedicated to sustainability
Resource-aware and adaptive security	Security mechanisms that adapt to available compute, energy, and connectivity conditions can reduce unnecessary overhead and help prevent security measures themselves from becoming a sustainability burden. This is especially relevant for	Helps with sustainability

	AI-enabled security, edge environments, and NTN contexts.	

3.7 Chapter Summary

The architectural vision presented in Chapter 2 can only be realized through communication and computing infrastructures that are fully programmable, intelligent, and capable of operating autonomously across heterogeneous, federated, and multi-stakeholder environments. This chapter has identified softwarisation and intelligence as the two complementary technological pillars enabling this transformation, evolving today's cloud-native and virtualized systems towards AI-native, agentic, and policy-aware infrastructures spanning the device-edge-cloud continuum and non-terrestrial domains. The research directions presented address the foundations required to build trustworthy, sovereign, sustainable, and resilient infrastructures that expose capabilities, orchestrate distributed resources, and support autonomous decision-making at scale. These advances provide the software and intelligence substrate upon which the communication protocols, interaction mechanisms, and coordination paradigms discussed in Chapter 4 can evolve to enable efficient, interoperable, and accountable operation of future communication and computing systems.

4 Protocols & Communication Paradigms

4.1 Introduction

Protocols are the enabling techniques that realise the behaviour of interaction between communicating elements in a networked environment, utilising the resources provided through this environment. Protocols embody underlying **communication paradigms**, thus the set of concepts upon which the protocols are built and the standard along which they are realised in an interoperable fashion to build the networked systems today and in the future.

The SRIA in its 2024 edition positions protocols as the "fundamental enablers for 6G systems" and this chapter will outline the continued enabling natures of protocols beyond.

Chapter 2 of this document positions those enabling protocols into the **system** considerations and architectures that will realise future communication & compute systems, while Chapter 3 postulates the importance that **software** and **intelligence** will play in building those future systems. In this context, protocols remain the foundational enablers of networked systems, but their role is fundamentally changing. Following the vision of Chapter 2, we can recognise that protocols in the post 5G/6G area no longer merely *connect endpoints,* they will increasingly *coordinate distributed compute, data, and intelligence across the IoT-edge–cloud continuum*, utilising a various enabling **communication paradigms** that are enriched with intelligence and softwarisation capabilities, such as semantic networking, goal-oriented networking and others. More so, **AI** in various forms, either driven by large-scale LLMs as well as agentic or generative versions, will be deeply embedded into the functioning of future systems as well as the resources utilising the system, as positioned in Chapter 2.

The remainder of this chapter will continue in this vein by outlining a vision for a **continued development of protocols** that will enable various communication paradigms, which in turn embody the vision of Chapter 2. We will position key trends and forces that drive their future development, before identifying key research themes and topics that need to be addressed for methods and solutions to be deployed in future communication and compute systems.

4.2 Chapter Vision: Revolutionising Protocol Development

Protocols have constantly evolved ever since digital communication has been utilised to build what we call today the ***Digital Society***. Whether it is the development and evolution of the Internet, the progress in mobile communication from 1G to currently deployed 5G networks and the anticipated 6G development, but also through the use of digital communication in the many industries from transportation over manufacturing, aerospace, and education to healthcare and defense, to name just a few.

While this evolution of protocols in itself requires constant investment into research and development of the key methods that underlie those protocols, we see a giant leap in the overall **process for protocol development** itself; a leap that will impact the very methods for protocols, their usage, but also their efficacy and the speed of deployment in ever more scenarios and aspects of our society.

This leap is driven by the latest telco-cloud evolution that increasingly converges collaborative networking, computing, and data into a **cognitive, software-defined continuum**. Key for this continuum is to preserve data sovereignty across an edge-to-edge Internet fabric, thus the enabling protocols must ensure that data generated, processed, and consumed within Europe can remain under European governance, even as services dynamically span heterogeneous infrastructures ranging from far-edge devices to hyperscaled cloud environments.

Relevant drivers in Europe in this context are Industrial IoT environments, manufacturing, energy/datacenters, smart cities, and logistics. These environments are already generating dense, heterogeneous, time-critical data flows that impose hard requirements on transport and network layer protocols, e.g., bounded latency, deterministic scheduling, zero packet loss, and seamless mobility across heterogeneous wireless and wired access technologies.

In IIoT environments specifically, this continuum is not an abstract architectural aspiration but an operational reality: field devices, embedded controllers, multi-AP wireless networks (e.g., Wi-Fi 6/7, TSN), and Edge-Cloud/hybrid Cloud-hosted AI inference engines must interoperate seamlessly within a single coherent data-compute-network fabric. The protocols governing this fabric, at layers spanning from the PHY and MAC layers through transport and into the application plane, must be co-designed as a system, not as independent stacks.

As a consequence, a stronger focus is needed on awareness for locality and jurisdictional **compliance**, posing constraints on data and computational **placement** as well as mobility, thus impacting protocols functions like routing and transport, while also necessitating verifiable **trust relationships** across administrative and technological boundaries.

At the same time, telco-cloud architectures are transitioning toward **cognitive, data-driven operation**, where real-time telemetry, Digital Twins, and AI models jointly steer communication behaviour across the continuum. This creates both an opportunity and a responsibility: protocols must expose native control surfaces for learning systems while preserving determinism, accountability, and compliance with European values.

Bringing those trends into a lifecycle of development and innovation, Figure 4-4-1 outlines the consequence of this leap in protocol development. The black elements illustrate the traditional cycle for protocol research and development, where underlying **Communication Paradigms** enable the desired facets of the **Digital Society**. Stakeholder engagements identify **Stakeholder Needs**, which are then translated into the **Requirements Space** to constrain the **Protocols**, which realise the communication paradigms.

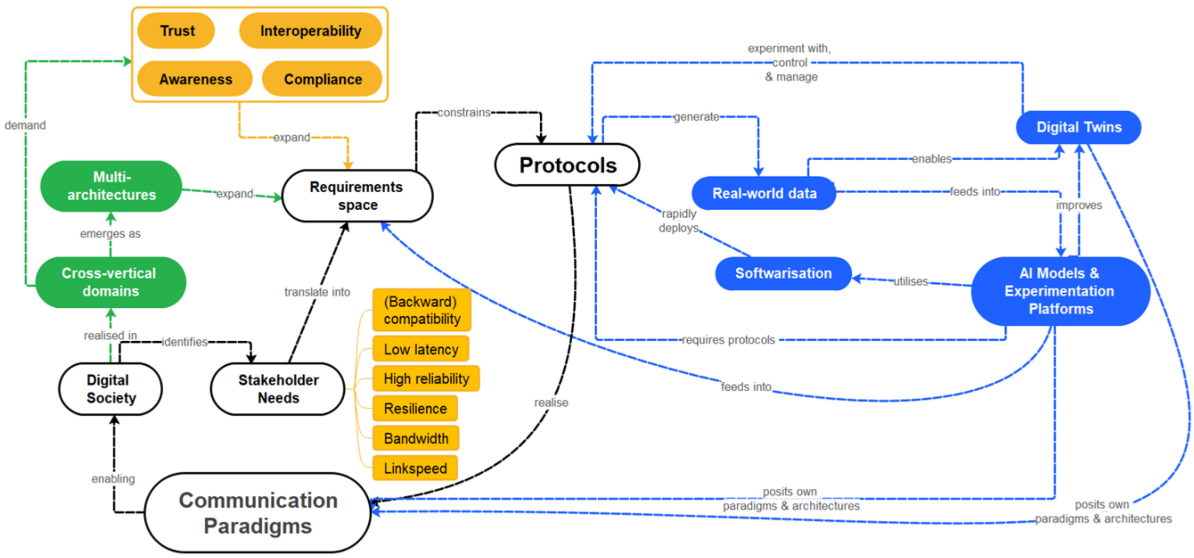

Figure 4-4-1: The Protocol Revolution: Softwarisation, Data, and AI turbocharging protocol development

We identify two key trends that extend this R&D cycle. Firstly, shown in green, the ever-growing penetration of digital technologies in our Digital Society is realised in a growing number of (cross-)**Vertical Domains**, going beyond the set of Internet technologies we are using every day. The authors in RFC8799 [C4-1] refer to those as *Limited Domains*, often relying on domain-specific extensions to the global interconnection fabric provided by the Internet. With it, architectures stemming from those vertical domains emerge into which enabling protocols are enabled. This **Multi-Architecture** nature of the digital communication system spanning across those domains needs understanding as to its common but also domain-specific principles, while the architecture specifics directly impact, often expand but also limit the requirements space, thus impacting the protocol R&D cycle on the one hand. This widening scope of communication systems expands the set of requirements in relation to compliance, the trust relationships between communicating entities, and the awareness for a wider set of digital environments, particularly for carbon- and energy-awareness, all while expanding on the need for interoperability but also backward-compatibility of this growing fabric of compute and communication, shown as orange components in Figure 4-4-1.

On the other hand, those protocols that do cater to the increasing multi-architecture nature of our changing Digital Society are also impacted by the growing role (real-world) **Data** plays to enable **Digital Twins** of the communication environment to experiments with as well as control and manage the protocol system, as illustrated by the blue parts in Figure 4-4-1. But this data also plays into the emerging role of large-scale **experimentation platforms and AI models**, e.g., enabling the training of LLM and other AI methods, in turn utilising the growing Softwarisation of communication systems (as discussed in Chapter 3 of the SRIA) to rapidly deploy revised protocols or entirely new ones, while also improving the Digital Twins that control and manage them. Conversely, the integration of AI and Digital Twins not just enables the emerging Internet of Agents, but it also ***necessitates own communication paradigms and architectures***, particularly when moving from today's largely application layer into lower layer functions of the overall communication system.

Many works have appeared in recent years that show the transformative power of this protocol revolution. Future R&D on protocols and communication paradigms must fuel and accelerate this revolution to increase the speed of development but also to ensure hardening the protocols against security, privacy, and resilience concerns. For this, we formulate the following requirements:

1. Emerging **cross-vertical domains and their requirements** MUST be captured through involvement of the suitable stakeholders in the Digital Society.
2. Suitable **architectures** MUST be formulated that facilitate those vertical domains.
3. Principles and invariants need identification to guide the requirements analysis and development of protocols.
4. The role of **real-world data** MUST be explored to not only ensure new methods but also better understand operation, security, resilience, and sustainability impact of existing ones. For this, data from industrial use MUST be unlocked for use in protocol research and innovation, facilitated through incentive-driven governance frameworks.
5. It MUST be investigated how **Digital Twins** can become a standard tool for understanding, controlling, and managing operation of individual and joint components of the larger communication system, moving from system-level management into the protocol-level control.
6. The likely **varying timescales** at which the phases in Figure 4-1 work, e.g., the training of models in relation to the functions where the knowledge is applied, SHOULD be properly understood to optimise operations of the protocol innovation cycle.
7. The combination of AI (models) and softwarisation advance needs a **foundational integration** into the protocol development cycle, including the handling of transition of functions for "in-flight" changes to operational systems.
8. Interaction across limited domains, ensured through **interoperability** mechanisms, MUST preserve the overall interlinked, often global nature of communication.
9. Deployments must consider greenfield and brownfield options, the latter recognising and considering existing investments and deployments by key stakeholders.

The key expected outcomes from addressing the above requirements, fueled by the research themes, topics, and actions discussed in the next sections of this chapter, will be

1. Improved **sustainability** of protocols with in-flight replacement of methods and entire protocols, backed up by insights from real-world operational data, with critical deployments providing the most demanding and data-rich validation environments, where energy consumption, computational cost of protocol processing, and network determinism are simultaneously observable and measurable across full protocol stacks.
2. **Secure** and **resilient** systems that are optimised for the specific critical communication need that the vertical domain represents in which they will operate.
3. A faster **speed of development** for protocols that is aligned with the abilities to retrieve suitable data and model operations, driven by softwarisation methods.
4. A growing set of **communication paradigms** that will drive the adoption of digital communication solutions for the Digital Society, where the insights on suitability and

efficacy of those paradigms can be adjusted increasingly in near-real time due to the linkage with the data and AI-driven extended development cycle in Figure 4-4-1.

5. A validated framework for **cross-layer protocol co-design** in a data-compute-network continuum, grounded in IIoT operational experience, providing the community with datasets, simulation tools, and Digital Twin environments that accelerate the innovation cycle from requirements to deployment across industrial vertical domains.

The following sections formulate an emerging albeit not exhaustive list of research themes that have been identified, driven by our vision in Figure 4-4-1 on how protocol evolution will change. With that, the reader will find considerations that are driven by the **extended environments** in which future systems will operate (Section 4.3) and the role that **data-driven operations**, e.g., through Digital Twins (Section 4.4), will play together with the **emergence of AI** as both a capability and a goal for communication (Section **Error! Reference source not found.**944.6), to allow for **optimising protocols** across control and data plane operations (Section 4.6**Error! Reference source not found.**), all while ensuring **secure**, **correct** (Section 4.7) as well as **sustainable**, energy-efficient operations (Section 4.899).

4.3 Evolving Protocols Across Multi-Domain Environments

The insights of decades-long research into Internet transport protocols but also the recent advances in other network domains, such as data centres or Non-Terrestrial Networks (NTN), on those protocols opens up the question how domain-specific protocols as well as interoperable cross-domain protocols could be most efficiently designed. For instance, can latency-optimised protocols for the Internet be utilised to bring new solutions for specific vertical network environments? What would be the key drivers? What could be an approach to design? How can the innovation lifecycle in Figure 4-1 impact such cross-domain applicability?

Conversely, the impact of multi-domain environments on protocol design needs deeper understanding, possibly revisiting design approaches that were deemed unsuitable in the Internet environment. For instance, the role of proxies in transport protocols is seen as crucial for environments such as aerospace, while Internet solutions largely avoid those designs. Greater understanding on requirements and resulting impact on design considerations for protocols is required. Similarly, in Industrial IoT environments, co-locating computing and networking functions at edge nodes — where protocol processing, data aggregation, and AI inference occur within the same fabric — challenges the classical separation between network and application layers, raising the question of how transport protocol design must adapt when the network node is no longer a passive forwarder but an active data-compute participant.

4.3.1 Even Lower Latency

The last decade and a half saw a shift of focus in network performance from throughput to latency, in support of the variety of latency-critical applications that are in use today. The need for latency reduction is not declining – if anything, emerging applications such as haptics and remote tactile feedback combined with immersive video feeds will increase it further.

For stakeholders in the digital society, the latency is perceived at the user or application level – the time to decode a video frame, the time for a cloud document to appear on screen after being, the responsiveness of clicking a button or the time to move a robotic arm.

State of the Art: Several latency-reducing technologies at the network level have been put forward – some examples are the FQ-CoDel and CAKE queuing algorithms from the "Bufferbloat" community, deployed in Linux and OpenWRT; L4S, deployed by Comcast within the network and Apple at the end systems; API changes that enable better control over the data instead of requiring it to sit in the network stack's buffer and cause delay within the sending host; "0-RTT" connection setup approaches in QUIC and TCP which eliminate the need of the initial connection handshake.

Despite the big wins through these technologies, their latency gain has a hard limit because their focus of latency reduction is on each individual packet along its end-to-end path. In contrast, often, latency perceived by users is related to the time it takes to finish a data transfer – which is influenced by the time a protocol needs to saturate the capacity that is available to it. In protocols such as QUIC and TCP, this time manifests as protocol handshakes.

What must be researched: A future solution needs to involve in-network elements in a different way, to i) intelligently determine the available rate across bottlenecks, possibly in a way that relates to the underlying link layer technology; ii) control data appropriately to feed the right information back to the sending application, iii) use localized caches of data as needed.

For (i), note that current approaches in the literature as well as deployment, including the upcoming IETF SCONE standard, only offer an upper limit – how to correctly give an *upward* signal to quickly make use of newly available bandwidth is a completely unsolved problem.

For (ii) and (iii), consider that, in principle, minimizing latency means to bring data from a sender as close to a receiver as possible, as quickly as possible. In a network with dynamic conditions, e.g. across wireless link with fluctuating capacity, it is beneficial to have data available close to that link when capacity suddenly becomes available, rather than requesting it from the origin. At least for data that an application may not want to update, such data should therefore be cached along the path – but not in a FIFO queue that delays other traffic. An application may want to choose to transmit newer data that can override data that may potentially still be in flight – but at the time of sending the original data, this is not known.

4.3.2 Hybrid Network Orchestration

The 6G landscape is not defined by a single communication medium but rather by the seamless integration of, e.g., a Space-Air-Ground Integrated Network continuum.

State of the Art: Integrated, e.g., TN-NTN, mechanisms of orchestration have been investigated in recent years. The dynamic and natural mobility of NTN nodes has been addressed by extending current MANO architectures with predictive mechanisms to anticipate service provision in target areas (e.g. ETHER). These novel capabilities offer the possibility to manage the lifetime of network functions satisfying user demands. Additionally, deploying SDN architectures in NTN has been investigated, addressing the controller placement problem. However, these solutions does not include the necessary intelligence to adapt their behaviour according to the network status, nor ensuring the necessary autonomy to efficiently operate.

What must be researched: Research must focus on the intelligent orchestration of hybrid networks, where autonomous controllers dynamically steer traffic across diverse layers and frequencies to optimize both performance and energy consumption. For instance, an AI-native routing agent should learn to dynamically segregate traffic: routing critical, sensitive data flows through high-throughput optical ISLs, while offloading delay-tolerant broadcast data (e.g., massive IoT or NB-IoT) to standardized, lower-power RF links. Protocols shall be developed to ensure that the control mechanisms in the orchestration procedure is achieved in the overall hybrid architecture.

4.3.3 Decentralized Trust and Entitlement-Based Multitenant Network Governance

This topic explores the transition from centralized network brokerage and bilateral contracts to a ***multi-domain, zero-trust framework*** based on machine-verifiable digital entitlements. It focuses on embedding explicit time- and space-bounded resource rights, regulatory constraints, and service-level agreements (SLAs) directly into cryptographic tokens, managed via blockchain smart contracts. This decentralization is critical for securely governing multi-operator TN-NTN environments, enabling continuous zero-trust authentication through Access Control Contracts (ACCs), and decoupling user access from traditional session-based cellular architectures.

State of the Art: Decentralisation is a key focus of existing systems to enable a *permissionless* participation but also to allow for *scale*. Aspects like *latency* (to converge the distributed consensus process needed for the entitlement verification) and *costs* have been mainly secondary in the current works. For realising distributed trust systems, distributed ledgers and their protocol mechanisms often rely on an iterative diffusion mechanism, executed by autonomous peer participants of the trust network. Blockchains and distributed ledger technologies in general, rely on such unicast-based solutions, leading to significant issues in converging the consent to a stable state in a reasonable time but also creating significant costs at the level of packets sent around the network [C4-12]. Multicast as an underlying communication paradigm is mostly used in intra-domain networks, therefore not providing a solid and scalable basis for multi-tenant and multi-domain. As [C4-13] has shown, decentralisation, as intended by the system designers and realised in its protocol mechanism (e.g., through the aforementioned diffusion), is also heavily impacted by the deployment of resulting systems in hyperscaler environments, raising availability, resilience, and security issues, while questioning the very use of autonomous diffusion mechanisms in the first place.

What must be researched: Cost and latency must be elevated as key requirements to decentralised trust mechanisms to counter the observed waste in current mechanisms [C4-12] and thus ensuring sustainability of solutions but also to enable new use cases at significant scale where verifications may be needed in much shorter time than current systems allow for. Multicast here is a key enabler and its use beyond single network segment seems a promising avenue of research [C4-14], not just reducing costs for control and data plane traffic significantly but also enabling much shorter consensus times and therefore enabling solutions for decentralised trust verification at the scale of future 6G and beyond systems.

4.3.4 Synopsis of Required Research

Research Theme	*Evolving Protocols Across Multi-Domain Environments*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Even Lower Latency	*Mid-term (finished in 5y)*	i) IETF-standardized transport layer mechanisms that use in-network signals to increase the sending rate faster; ii) IETF-standardized mechanisms to transfer upper-layer objects faster across links with a fluctuating capacity via intelligent use of in-network buffers. Surpasses current per-packet limits (e.g., L4S, QUIC 0-RTT) by using intelligent in-network signalling to utilize bandwidth faster **[KPIs]:** at least 30% reduction of latency for short application objects.	Enables realistic tactile haptics and immersive cloud services by achieving deterministic latency through in-network capacity saturation. Surpasses current per-packet limits (e.g., L4S, QUIC 0-RTT) by using intelligent in-network signalling to utilize bandwidth faster. Provides inclusive "remote awareness" applications to assist citizens with reduced physical abilities.
Hybrid Network Orchestration	*Mid-term*	i) ETSI-standardized orchestration mechanism for natively integrated TN-NTN, ii) orchestration intelligence for quick adaptation, iii) IETF-standardized protocol stack to support control signalling among cross-domains, iv) adaptive and intelligent routing protocol according to global dynamics. **[KPIs]:** i) at least 50% reduction of reconfiguration time against network events; ii) 30% increase of service	Reduction of complexity through native integration of NTN infrastructure in orchestration mechanisms, ensuring domain-agnostic operations. Improved efficiency, and thus sustainability, through intelligent routing protocols that understand global

		provision; iii) 30% reduction of network consumption.	network dynamics and traffic fluctuations.
Decentralized Trust and Entitlement-Based Multi-tenant Network Governance	*Mid-term*	i) protocols for distributed consensus at ultra-large scale, utilising multicast at network and below level, with an inter-domain service model **[KPIs]:** i) 90% reduction of convergence latency (ii) 80% reduction of data plane costs and 90% reduction in control plane costs (iii) 10 times increase in scale, measured by participating members of the trust network.	Increased sustainability of decentralised ledger networks through reduced costs and waste Improved resilience Improved economic value through wider range of application use cases

4.3.5 Recommendations for Actions/Instruments

Research Theme	*Evolving Protocols Across Multi-Domain Environments*			
Action	**Even Lower Latency**	**Modernizing Space Protocols for Softwarised Payloads**	**Hybrid Network Orchestration**	**Decentralized Trust and Entitlement-Based Multi-tenant Network Governance**
International Calls		**X** *Enable the required international collaboration for the desired solutions*	**X** *Enable the required international collaboration for the desired solutions*	
International Research	**X** *Target global IETF standardization*		**X** *ETSI standardization*	
Open Data	**X** *Backing efforts with open data will benefit IETF standardization*			
Large Trials		**X** *Demonstration of regenerative payloads as enablers*		
Cross-domain research				

4.4 Data-Driven Protocol Operations through Digital Twins

While digital twins are often positioned at the management and control plane, Figure 4-1 also foresees a direct impact of Digital Twins on data plane behaviour of new protocols, complementing the tasks of network modelling and emulation done in existing approaches. Problems such as generating and obtaining reactive traffic, twinning realistic existing and future anticipated network conditions as well as the overall emergence of a “live platform” that allows for data ingestion, twinning, and ultimately training of suitable AI models. Digital Twins must additionally capture the interdependency between network state, computational load, and data placement decisions twinning not just the communication system but the joint behaviour of the continuum as a whole. Experiences from IIoT deployments demonstrate that a Digital Twin of the network layer alone is insufficient: accurate twinning requires co-modelling of edge compute resource availability, in-network processing latency, and data flow provenance, all of which directly influence protocol behaviour.

State of the Art: As protocol developments are brought closer to reality, more complex scenarios are commonly used in simulations, and large-scale testbeds allow for a richer and more realistic set of test scenarios, frequently amended with additional tests to demonstrate that a proposed mechanism can indeed work in a production network. In a sense, such a final demonstration is often just "icing on the cake", not part of the development. Meanwhile, tech giants, commonly located in the US, use a very different approach which is enabled by their massive customer base: ***A/B testing***. Once a protocol mechanism is seen to work well in controlled test environments and there are good reasons to believe that customers will not be hurt by it, it is rolled out to customers in a treatment group (B), with relevant metrics being compared against customers in a control group (A) that does *not* experience the new mechanism. With millions of customers, meaningful statistics can be obtained even when the set of environment conditions is very diverse, and the result of such a test is much more convincing than experiments in the meagre conditions that are available to academics and SMEs in Europe. This gives these tech giants a distinct competitive edge, further widening the gap between their own technology and any European technology development.

The key benefit of large-scale A/B testing is that it encompasses the full complexity and heterogeneity of a production network. This includes, for example:

- A realistic mix of network traffic, which is *reactive* (if the network is congested, mechanisms in an OS or application may react, and if the Quality of Experience suffers, users will react).
- A combination of heterogeneous conditions, including various on-path devices and heterogeneous links. E.g., a video communication between two persons may traverse a Wi-Fi network on one end and a cellular connection on the other, with a wired network employing a traffic shaper in between.
- Realistic patterns of user behaviour, including the duration and type of network usage and user mobility.

These factors are all active at the same time, causing a conglomerate of complex interactions that is never matched by simulations or testbed experiments.

The goal of this research theme is to produce a Digital Twin of the multi-domain environment that realises the Digital Society – an environment that realistically replicates conditions such as the ones mentioned above (realistic *reactive* background traffic, fully heterogeneous networks with all their influencing effects, realistic usage patterns), and possibly exceeding them. This will allow experimentation akin to large-scale A/B testing, but there will be *further, greater benefits*: e.g., different from A/B testing, there is no need to be cautious in a Digital Twin (no customers are hurt if an innovation fails). Environment conditions can be made to be deterministic, such that results can be replicated and better understood. Investigating effects can be done at a finer granularity, e.g. by going back in time to understand a certain effect at a specific time instance.

4.4.1 Realistic network traffic generation

Internet traffic modelling has a long history, but such models are never reactive. Reactive traffic generators that can produce realistic traffic, at scale, are needed. For example, a household's traffic mix can consist of streaming video on demand, real-time video communication, web surfing and interactive game playing at the same time; such communication is between members of the household and various peer nodes in the network. Technically, these are multiple applications with different traffic patterns in different directions, some of it being reactive to congestion, with a mix of congestion control mechanisms, and control loops that operate at different time scales. From the bottleneck's point of view, it is an aggregate of incoming and outgoing traffic that will react to congestion in some way.

What must be researched: A way to model this aggregate behaviour is needed – possibly by modelling the individual applications and systematically aggregating the joint behaviour. ML-based approaches (e.g., Reinforcement Learning) are obvious candidates to replicate the reaction of a traffic source to various environment conditions. If a general, systematic way to aggregate such models can be found, this would allow to scale up, e.g. from one application to a household, from one household to many, from a group of households to a city area. Corporate networks could be modelled in a similar fashion.

4.4.2 Realistic network conditions

What must be researched: Research topic 1, network traffic generation, needs to be based on a multitude of *combined* network conditions. A part of a network path may be a Wi-Fi or 6G link, possibly using line-of sight technologies such Visible Light Communication or sub-THz, typically requiring multipath and hand-over schemes. Another part of the same path may use a different such technology or communicate with a LEO satellite. In between, traffic shapers or Performance Enhancing Proxies (PEPs) might be used, and routing may be dynamic. The entire setup is influenced by human mobility and typical human behaviour regarding connectivity. These conditions need to be modelled in a way that makes them replicable for model training so that an ML-based traffic generator can be prepared to operate in the Digital Twin under realistic conditions.

4.4.3 A platform with live data

Traffic modelling approaches, reactive or not, risk becoming stale. Similarly, modelling network conditions must be based on real, up-to-date network behaviour, or risk becoming irrelevant as the network environment evolves.

What must be researched: A Digital Twin of the Internet therefore needs a platform within which models can be trained with up-to-date information which must be collected from various sources, live, and fed into the platform in a uniform fashion.

This must capture fine-grain timing but does not overload production systems, neither in terms of computation nor traffic overhead, while staying up to date. Data must be anonymized without losing this granularity.

4.4.4 Synopsis of Required Research

Research Theme	*Data-Driven Protocol Operations through Digital Twins*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Realistic network traffic generation	*Mid-term (finished in 5y)*	A systematic framework and toolset for generating *reactive* network traffic. Uses ML/Reinforcement Learning to replicate traffic source reactions to congestion; supports multi-scale aggregation (e.g., from a single household to a city area). Moves beyond traditional static models to traffic that responds dynamically to network conditions and user Quality of Experience. **[KPIs]:** Increase traffic model fidelity (measured by correlation with real-	Closes the competitive gap with non-EU tech giants by enabling European SMEs to perform *A/B testing equivalents* without needing a massive customer base. Supports European data sovereignty by providing a sovereign environment for hardening protocols against security and privacy concerns.

		world telemetry); scalability with 5 times higher number of concurrent reactive flows supported.	
Realistic network conditions	*Mid-term (finished in 5y)*	A comprehensive library of replicable multi-technology path models. Captures the interaction between heterogeneous links, dynamic routing, and in-network elements; incorporates human mobility and connectivity patterns. Transitions R&D from isolated simulations to a deterministic Digital Twin that replicates the full complexity and heterogeneity of a production network. **[KPIs]:** At least double telemetry correlation accuracy (how closely the twin matches real-world link behaviour); Increase technology coverage by at least 50% (number of 6G-relevant access types modeled).	Directly fulfills the requirement to explore real-world data and establish Digital Twins as standard protocol R&D tools. Deliver a validated library of path models within 5 years that replicates 5G+/6G handover scenarios with 90% fidelity to real-world production telemetry.
A platform with live data	*Mid-term (finished in 5y)*	A unified platform for live data ingestion, twinning, and AI training. Captures fine-grain timing without overloading production systems; ensures data anonymity while maintaining granular detail; provides up-to-date network behaviour. Beyond SotA: Eliminates the risk of stale or irrelevant models by replacing static datasets with a platform that evolves alongside the production environment. **[KPIs]:** At least double data freshness (i.e. reduce latency from production to twin); improve ingestion throughput by 50%; Double model drift detection rate.	Fulfills the mandate to explore real-world data to understand the "operation, security, resilience, and sustainability impact" of protocols. Supports European data sovereignty by ensuring data processed within Europe remains under European governance and complies with jurisdictional rules. Establish a functional platform within 5 years capable of synchronizing live telemetry from at least three different network domains (e.g., IIoT, 6G, Satellite) to support real-time protocol adaptation.

4.4.5 Recommendations for Actions/Instruments

Research Theme	*Data-Driven Protocol Operations through Digital Twins*		
Action	**Realistic network traffic generation**	**Realistic network conditions**	**A platform with live data**
International Calls		**X** *Ensure involvement of international actors*	**X** *Ensure involvement of international actors*
International Research	**X** *Much of the competence resides outside Europe*	**X** *Much of the competence resides outside Europe*	
Open Data	**X** *Mandatory for training Reinforcement Learning-based traffic generators*	**X** *Datasets for training heterogeneous link-layer models*	**X** *Unlocking industrial telemetry for research*
Large Trials	**X** *Capture the conglomerate of complex interactions from real production networks*	**X** *Validate complex production-scale link interactions*	**X** *Ingesting live data from production systems*
Cross-domain research		**X** *Integration of 6G, satellite, and human mobility*	

4.5 AI as a Capability and Goal for Communication

With Figure 4-1 outlining a growing role of AI methods, including agentic AI, the question arises how to build a system BASED on such methods, while providing the key methods for those agentic methods to perform the needed communication in the first place? Particularly, while the notion of Internet of Agents is mainly placing agentic capability at the application level, thus requiring a fully functioning communication system underneath, the use of agentic AI at lower levels, such as for congestion control, routing control, and other functions, must see agentic communication embedded more natively at the lower layers of the system. If positioning the agentic capabilities as a key control capability, this will lead to re-thinking the design of a future agentic control plane for the overall communication system.

4.5.1 Decentralized Agentic Networking and AI-Native Routing

The 6G framework established by early 3GPP Release 20 studies and consortiums like the AI-RAN Alliance dictate a profound shift toward an “AI-native” architecture, wherein artificial intelligence is embedded directly into network control loops rather than applied as an external overlay. Future networks must operate as distributed, autonomous multi-agent systems.

State of the Art: Early work on AI-Native Routing focused on using AI to improve traditional routing in Software-Defined Networking (SDN), where a centralized controller applies Machine Learning (ML) models for congestion prediction, traffic engineering, or path selection. However, the current AI-native shift goes further, with routing formalized as a sequential decision problem and the network being treated as a learning environment, where agents continuously adapt routes based on feedback (latency, loss, congestion). This implies shifting from the centralized model to distributed approaches, with each router acting as an autonomous agent, and federated learning allow coordinating the global behaviour. This vision is naturally finding his way in 6G research. Unlike 5G, where AI is often applied at the management or orchestration layer, 6G envisions AI-native networking where routing, resource allocation, and mobility management will be continuously learned and adapted in real time.

What must be researched: Research must advance decentralized routing algorithms that allow satellites to independently compute optimal paths, negotiate routing tables, and manage state synchronization without continuous coordination from terrestrial gateways. A critical area of focus is the implementation of Deep Reinforcement Learning for Energy-efficient Routing (DRL-ER). Research has demonstrated that DRL-based routing agents can independently learn policies that balance strict QoS constraints against the severe energy limitations of solar-powered LEO satellites. Furthermore, research must explore Two-Timescale Semantic-Agentic designs. This advanced paradigm combines semantic data compression with cooperative multi-agent control. In this model, “fast-loop” AI agents onboard the satellite proactively shape traffic and prioritize high-utility packets to avoid acute network bottlenecks, while “slow-loop” semantic updates provide stable resource budgets that prevent oscillatory network behaviour. To train these sophisticated models without overwhelming capacity-constrained feeder links with raw telemetry data, research must leverage Federated Learning (FL). FL allows AI agents to train models locally at the edge (onboard the satellite) and only transmit updated mathematical weights back to the terrestrial SMO, significantly enhancing data privacy and minimizing backhaul congestion.

4.5.2 Agent Discovery

The Internet of Agents (IoA), is developing into a key infrastructure for next-generation intelligent systems, supporting seamless interaction, automatic agent discovery, adaptive task orchestration, and collective reasoning among large populations of virtual and embodied agents. As the number of such agents grows, mechanisms are required to allow agents and systems to discover relevant agents dynamically.

State of the Art: While AI are really flourishing, current approaches for their discovery remain quite basic, mainly based on static registries, which differ significantly in their architecture, metadata formats, discovery models, and trust assumptions. This diversity makes interoperability difficult and prevents the emergence of a global agent discovery infrastructure. Furthermore, Agents are increasingly expected to operate across organizational and network boundaries. Yet current systems are quite fragmented and do not consider agents and registries beyond the local domain.

Agent discovery needs to go beyond identification schemes and naming. In future eco-systems, agents will interact in a semantic way, and discovery queries will be based on complex human readable expressions. Simple keyword-based discovery is often insufficient for AI agents that provide complex capabilities. A common framework or semantic ontology is paramount to enable semantic discovery that integrates with agent registries and discovery protocols. For instance, a simple keyword search for agents labeled “currency” differs significantly from recognizing that a request for “foreign exchange rates” should also retrieve agents associated with “forex,” “FX,” or “international money transfer.”

What must be researched: An effective agent discovery system must fulfill several key requirements. It should be capable of locating agents across different administrative domains while accurately identifying their capabilities. In addition, it needs to evaluate relevant policies and constraints, ensuring that any discovered agents meet operational and governance requirements. Security and trustworthiness must also be integral to the discovery process. Finally, the system should support scalable indexing and search to efficiently handle a growing number of agents.

4.5.3 Cross-Domain Agent Authentication and Authorization

As AI agents become more autonomous and exert greater influence over both digital and physical environments, ensuring trust, safety, and ethical compliance within the IoA becomes increasingly important. As a result, the IoA demands reliable regulatory mechanisms to verify agent identities, govern their behaviour, and protect against malicious activities.

State of the Art: Current AI agent protocols and platforms typically offer transport-level security and local permission controls, but they lack a consistent cross-domain foundation for verifiable agent identities, accountable delegation, interoperable credential management, and reliable audit trails. As a result, interoperability is often achieved through custom integration that vary in security quality and have ambiguous life-cycle semantics. Moreover, in scenarios where agents delegate tasks to other (sub-)agents, if each sub-agent must individually authenticate and request permissions from the authorization server, the server is required to verify identity and permissions of each one separately and issue corresponding access tokens. This approach is not scalable, as it results in frequent interactions, increased computational load on the authorization server, and reduced efficiency in the token issuance process.

What must be researched: Emerging scenarios and use-cases require novel frameworks for authenticated, authorized, and auditable delegation of authority to AI agents, where human users can securely delegate and restrict the permissions and scope of agents while maintaining clear chains of accountability. Such new frameworks have to take into account the ephemeral nature of AI agents. Existing authorization models are ill-suited for this dynamic environment. They typically bind permissions to stable, long-lived identities. Agents, by contrast, are often ephemeral, created in large numbers for specific, short-lived tasks. A robust framework must support hierarchical delegation of authority, ensuring a sub-agent's permissions are a strict subset of its parent's, and providing a clear chain of accountability. When an undesirable action occurs, tracing the root cause across a chain of agent interactions remains still a daunting challenge.

4.5.4 Semantic Communication for AI-Native Networks

Conventional communication protocols are designed around the reliable transfer of bits, packets, and sessions, assuming that the communication objective is the accurate reconstruction of transmitted symbols at the receiving endpoint. Many emerging applications across the IoT-edge-cloud continuum, including collaborative AI inference, distributed sensing, industrial control, robotic cooperation, digital twins, and agentic coordination, do not fundamentally require exact bit-level reproduction of all transmitted data. Instead, they require that the receiver obtains the relevant meaning or task-effective representation needed to act correctly and efficiently. This shifts communication design from a ***syntactic paradigm to a semantic one***.

Semantic communication, enabled by foundation models and domain-specialised AI models, introduces a new protocol paradigm in which communication endpoints, intermediate nodes, or network functions can exploit shared models, context, and task awareness to compress, prioritise, reconstruct, or infer information according to its utility for the receiving task. In such a paradigm, communication may move from transmitting full data objects toward transmitting compact semantic descriptors, such as prompts, metadata, or latent representations, from which the receiving side reconstructs or generates the required content locally. This fundamentally redefines the role of protocols: from transporting content to transporting the instructions and context needed to recreate it, where mechanisms can be integrated into existing protocol stacks with minimal changes, enabling backward compatibility through fallback to conventional transmission when semantic capabilities are unavailable or insufficient.

In Industrial IoT, cyber-physical systems, and other critical environments, semantic communication could enable substantial reductions in communication load while preserving control fidelity and operational awareness. Rather than transmitting complete raw sensor streams, protocols may exchange latent representations, feature summaries, model updates, event abstractions, or uncertainty-bounded semantic descriptors. Likewise, agentic systems may communicate intentions, goals, constraints, and task states instead of only payload objects, thereby improving efficiency and coordination across distributed control loops. In broader digital services, local generation capabilities may further enable context-aware and personalised content delivery, adapting outputs to user context or device capabilities.

State of the Art: Initial research has demonstrated the potential of task-oriented encoding, joint source-channel coding with neural models, and AI-assisted compression. Recent approaches show that transmitting semantic representations instead of full content can achieve orders-of-magnitude reduction in data volume for certain modalities, while maintaining task performance. Early system prototypes indicate that such approaches can significantly reduce network load and storage requirements and enable new interaction models where content is generated at the edge or client side rather than transmitted directly . However, current work remains fragmented across application-specific solutions, with no standardised protocol abstractions, capability negotiation mechanisms, or interoperability frameworks.

What must be researched: Semantic communication must be elevated to a protocol-native capability across the network-data-compute continuum. This includes defining semantic utility and task relevance as protocol-visible metrics, and designing mechanisms for negotiating semantic capabilities, context, and reconstruction fidelity between communicating entities.

A key research area is the management of ***shared models and context alignment***, including model version negotiation, compatibility verification, and drift handling across heterogeneous domains. Closely related is the need for semantic interoperability, enabling semantic representations produced in one domain or model ecosystem to be interpreted and validated in another. Protocols must also support ***adaptive operation*** modes, allowing seamless switching between semantic and conventional communication depending on capability, reliability, or policy constraints. This includes explicit fallback mechanisms to ensure robustness in critical systems. Another important direction is ***in-network semantic processing***, where network nodes perform filtering, aggregation, or inference to reduce end-to-end communication and improve responsiveness. This requires new governance, trust, and verification models, especially in multi-domain and sovereignty-constrained environments. Finally, the system-level ***trade-offs between communication, computation, and energy*** must be rigorously understood. While semantic communication can drastically reduce transmitted data volumes and storage needs, local or in-network generation introduces computational overhead and latency. Early evidence suggests that generation costs can dominate transmission costs in some scenarios, highlighting the need for optimised model placement, hardware acceleration, and protocol-level decision-making to achieve net gains in efficiency and sustainability. At the same time, semantic communication introduces ***new risks***, including semantic ambiguity, bias, manipulation, and uncontrolled personalisation effects. Research must therefore ensure trustworthiness, traceability, and alignment with societal and regulatory requirements, preventing misuse such as echo chambers, misinformation amplification, or adversarial manipulation of generated content.

4.5.5 Inference Defined Networking

Inference-Defined Networking (IDN) is an emerging paradigm in which the primary aim of the network is oriented towards efficiently supporting Machine Learning (ML) inference. Rather than treating AI applications as traffic flows to be routed according to traditional metrics such as IP destination addresses or aggregate load, an inference-defined network dynamically reconfigures its resources in response to the distributed, latency-sensitive, and compute-intensive demands of live inference pipelines.

IDN is fundamentally enabled by the paradigm of 3C Networks — networks conceived as a unified fabric combining Communication, Computing, and Caching. Communication provides the substrate for passing intermediate inference artefacts among distributed nodes with ultra-low latency. Computing, co-located with network devices, enables split-inference and model-layer offloading, so resource-constrained devices need only execute locally a portion of a neural network. Caching allows large model weights, quantised model replicas, and KV-cache snapshots to be pre-positioned at topologically appropriate devices, eliminating repeated retrieval from remote data centres and decoupling inference latency from backhaul variability. IDN is particularly relevant for Non-Terrestrial Network (NTN) segments where space-ground integrated networks introduce asymmetric link capacities, time-varying topologies, intermittent connectivity, and propagation delays. Distributing inference across terrestrial edge devices, LEO satellite payloads with on-board processing capabilities, and ground stations requires inference-aware routing that accounts for orbital mechanics, contact schedules, and heterogeneous compute availability across the space-ground continuum.

State of the Art: Early work on inference-oriented networking largely emerged from edge Computing communities, which recognised that latency requirements for real-time ML tasks could not be met by centralised cloud inference alone. Pioneering contributions on split computing formalised the idea of partitioning Deep Neural Network (DNN) inference across device-edge-cloud hierarchies, optimising partition points based on available bandwidth and compute resources. Complementary work on early-exit architectures allows inference to terminate at an intermediate network device as soon as sufficient confidence and trust is achieved, reducing unnecessary data transmission to deeper infrastructure layers.

More recent research has extended towards multi-model and multi-tenant inference pipelines. Systems such as Cocktail, Clipper, and Triton inference servers require model multiplexing, and adaptive model selection to maximise inference throughput and service-level objective compliance. In the networking plane, proposals for inference-aware traffic engineering have been explored, where routers or SDN controllers maintain awareness of model deployment locations and route inference-related flows towards nodes with appropriate model caches. The integration of network telemetry with inference workload monitoring has also received attention, enabling feedback-driven adaptation of the network and the inference pipeline. In the space domain, on-board processing payloads on LEO satellites — used for tasks such as Earth observation inference, anomaly detection, or data compression prior to downlink — represent early practical deployments of network-embedded inference, though these remain largely isolated from ground-based orchestration systems.

What must be researched: Despite significant progress, IDN raises a set of open research challenges that span network architecture, control theory, security, and protocol design. Coordinating inference pipelines across heterogeneous, multi-domain infrastructure — spanning user devices, edge nodes, satellite payloads, and cloud backends — demands orchestration frameworks that can jointly optimise model placement, resource allocation, and traffic steering in near real-time. Existing orchestration systems designed for containerised microservices (e.g., Kubernetes-based MEC orchestration) lack native constructs for expressing inference-specific constraints such as model version consistency, tensor format compatibility, or KV-cache locality. Extending intent-based networking paradigms to encode inference service-level objectives as network intents is an important area of investigation.

As inference pipelines span administrative domains and autonomous network segments, establishing trust in provenance, integrity, and outputs of distributed inference nodes becomes critical. Decentralised trust mechanisms, potentially leveraging distributed ledger technologies or cryptographic attestation frameworks, are needed to verify that inference is being performed by authentic, unmodified models, particularly in the context of decentralised agentic networking, where AI agents embedded within network devices autonomously negotiate resources, route inference tasks, and adapt policies without centralised coordination. Ensuring safety, auditability, and alignment of such agents with network-wide objectives constitutes a fundamental open problem, in alignment with Section 4.5.3.

Traditional routing protocols are designed to forward packets based on syntactic addressing — IP addresses, port numbers — and are agnostic to the semantic content of the data being transported. IDN motivates a transition towards AI-native routing, as already discussed in Section **Error! Reference source not found.**.1 and semantic communication, discussed in Section 4.5.4.

A core architectural challenge for IDN is the mismatch between inference workloads — which are inherently service-oriented, addressed by model identity, version, and capability — and the location-centric addressing of the current Internet. Service-based addressing schemes, akin to those explored in Information-Centric Networking (ICN) and more recently in proposals for 6G service-layer routing, offer a natural fit: a device can request inference from a named model without knowing its physical deployment location, and the network resolves this to an appropriate replica. Integrating such schemes with existing inter-domain routing infrastructure, and extending them to the space-ground domain where contact opportunities are intermittent, remains a substantial open challenge.

Distributed inference is computationally intensive, and its energy footprint across network infrastructure is non-trivial. In terrestrial networks, the co-location of GPU-equipped edge servers with radio access nodes raises questions about dynamic power management and inference load shifting to minimise energy consumption without violating service-level objectives. In the space domain, satellite on-board processing is severely power-constrained, requiring energy-aware scheduling of inference tasks relative to solar power availability and battery state. Research is needed on energy-proportional inference orchestration that jointly optimises inference quality, latency, and power consumption across the full terrestrial-space stack, including mechanisms to defer or simplify inference during low-energy periods whilst maintaining acceptable service continuity.

Collectively, these challenges indicate that IDN is not merely an application-layer concern but demands co-design across routing architectures, control planes, trust frameworks, and physical resource management — particularly as the Internet evolves towards a heterogeneous fabric integrating terrestrial and non-terrestrial domains.

4.5.6 Synopsis of Required Research

Research Theme	*AI as a Capability and Goal for Communication*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Decentralized Agentic Networking and AI-Native Routing	*Mid-term (finished in 5y)*	Agentic nodes to ensure autonomous TN-NTN networks which decides operations according to dynamic environment. Efficient and fast routing protocols to manage fluctuating global traffic and unpredicted network events. **[KPIs]:** i) >50% reduction of traffic reconfiguration time; ii) increase of network robustness by at least halving downtimes.	Network autonomy will ensure its availability for end-user services. Intelligent routing decision will ensure energy-efficient infrastructure, and fast reaction against unpredicted events (e.g. natural disasters).
Agent Discovery	*Mid-term*	Global Agent Discovery Solution (not necessarily based on one single unified protocol). The system must be compliant to local regulations, governance while meeting strong requirements of security and trustworthiness. **[KPIs]:** achieve 90% usage of discovery solutions in agentic systems	Agent discovery will have strong economic impact as Internet of Agents is the natural evolution from micro-services to autonomous service of the future digital society.
Cross-Domain Agent Authentication and Authorization	*Mid-Term*	Remote attestation technology able to guarantee the identity of the agents, what are they authorized to	Agent authentication and authorizations is paramount to EU digital sovereignty. This includes

		perform/access, their level of trustworthiness, as well their accountability. **[KPIs]:** achieve >90% usage of attestation solutions in agentic systems	the ability to trust agents and guarantee that data is autonomously processed locally according to EU laws and regulations.
Semantic Communication for AI-Native Networks	*Mid-term*	Algorithms to encode semantic purposes of communication in for use in signalling protocols to initiate more efficient data exchange. Solutions need to be evaluated in terms of goodput improvements in defined test cases. **[KPIs]:** identify at least 5 realistic test cases with >95% reduction in throughput	Improve on sustainability through more efficient communication and support for AI-native knowledge-driven operations. Increased security by basing communication on defined semantic metadata.
Inference Defined Networking	*Long-term (finished in 7y+)*	- KV-cache and model-weight caching at the edge - Service-based addressing for inference - Intent-based orchestration for inference SLOs - Decentralised trust and usage across multi-domain inference pipelines - Decentralised Agentic Networking - Full terrestrial-space IDN integration **[KPIs]:** (i) 95% reduction in inferencing latency at network level (ii) 80% reduction in system-wide inferencing costs compared to DC.only model (iii) at least 10x more system throughput compared to DC only model	- Improved sustainability - Secure and resilient systems - Faster development - Digital sovereignty - Fairness and equity - Ethical AI deployment

4.5.7 Recommendations for Actions/Instruments

Research Theme	*AI as a Capability and Goal for Communication*				
Action	**Decentralized Agentic Networking and AI-Native Routing**	**Agent Discovery**	**Cross-Domain Agent Authentication and Authorization**	**Semantic Communication for AI-Native Networks**	**Inference Defined Networking**
International Calls		**X** *Collaborative project aiming at standardization.*			

International Research	**X** *Aiming at global standardization.*		**X** *Aiming at global standardization.*		**X** *Aiming at global standardization.*
Open Data					**X** *Faciliate evaluation based on open data*
Large Trials			**X** *Experimentation on large cross-domain testbeds.*		**X** *Evaluate scale, alongside early pilots*
Cross-domain research					

4.6 Optimised Capability Discovery

A capability that, not just, agentic AI relies on is discovery (of other agents) for the transfer of data between them. Anticipating agentic capabilities to emerge in many functions and capabilities of the communication system thus poses the question on how to optimise discovery and data transfer away from separate look-up steps, followed by a channel-oriented transfer (as done in today's Internet through explicit domain name system (DNS) lookups with transport layer security (TLS)-based connections to follow). Joining those mechanisms may point to approaches such as information-centric networking but its impact on the existing routing infrastructure has often been cited as its biggest drawback albeit not negating the need for joint mechanisms that may utilise existing transport network technology.

4.6.1 Backward-compatible Discovery-on-Transit

Backward compatibility is an important stakeholder requirement, given the significant sunk cost in existing communication infrastructure. Thus, while solutions to joint discovery and transfer may well exist, their integration into the existing environments, so called "brown field deployments" is often not a given, thus limiting new communication paradigms to "green field", i.e., new network environments.

State of the Art: Discovery in most communication systems is an explicit function, performing a mapping from (service/information/object) name to an identifier used to reach the communication endpoint at which the desired function can be reached. Most used in the Internet is the DNS, which uses a hierarchical naming and a distributed name resolver architecture to perform the mentioned lookup, returning an IP address for the following communication. Similar approaches are used for, e.g., object broker systems, web service systems, and others, all of which are based on explicit lookups with the caveats of additional latencies, complexity for the additional lookup system, as well as separation of endpoint selection policy from path selection (i.e., routing). Newer communication paradigms like information-centric networking (ICN) attempted to combine these separate functions of lookup and data transfer by migrating to a name-based identifier system over which to either perform path computation or routing; the lacking adoption of those systems is largely due to complexity and lacking backward compatibility with existing (IP-based) infrastructure. The need for backward compatibility has long been recognised, notably in the ICN community. Cisco's attempt to integrate ICN naming into IP addressing, labelled as Hybrid ICN (hICN) [C4-2] was one attempt to achieve ICN deployment in existing IP networks.

What must be researched: Recent work on, e.g., routing-on-service-addresses (ROSA) [C4-3] has demonstrated that discovery steps and data transfer can be integrated through the use of IPv6 extension header piggybacking and processing of the extension headers in a shim overlay. Such solutions, similar to hICN, are backward compatible, even towards transport and application protocols, when preserving existing DNS naming, while opening up the possibility to application-specific naming albeit transported over IP networks. More research is required on *efficient packet processing methods*, extending beyond the existing eBPF-based solutions, while utilising virtualisation capabilities, thus mitigating the need for newly deployed hardware in the network. Also, the *economic roles* of such service routing providers, in extension to IP routing providers, needs a better understanding from a wider ecosystem perspective, particularly in understanding how the joint discovery and delivery may *aid sovereignty* for end users and organisations alike by reducing the reliance on (often global) DNS providers. The impact on data protection, avoiding data leakage to centralised DNS stewards, is another aspect needing greater understanding. Due to the target of these solution being environments with more than one possible service location in the network, driven by virtualisation of SW-based endpoints, the *integration of compute but also energy metrics* in the distribution logic for any incoming service request also needs further understanding, research that would likely also benefit existing DNS-based systems.

4.6.2 Joint Discovery of Endpoint and Network Path

Service provisioning and the packet delivery are mostly governed by separate set of policies, each of which may even be executed by different stakeholders, such as the cloud provider for the former and the transport network provider for the latter. But recent advances in technologies that expose strong path pinning capabilities, foremost in the form of Segment Routing [C4-4], allow for jointly executing endpoint and network path selection under a joint policy framework, thus indicating that this capability is becoming a key enabler for ensuring strongly policy-controlled traffic engineering in future systems.

State of the Art: Traditionally, the selection of a network path, performed either through a distributed routing protocol or a centralised path computation element, has been separated from the selection of the most suitable service endpoint, if more than one exist. The use of DNS names tied to IP anycast allows for some (relatively high-level) control over where traffic is being distributed to, mainly used to direct traffic to regional replicas. CDNs provide similar redirection capabilities as part of the traffic engineering that underpins the distribution to content replicas rather than the original source. However, after the IP address has been determined, the path selection (thus the determination which individual network segment is being used) is handed over to the transport network provider, usually in the form of a local ISP. Segment Routing changed this approach by providing greater control over choosing “segments” that joined together determine the end-to-end path. Exposing suitable interfaces to upper layer functions (including applications) allows for implementing policies, e.g., to avoid delivery through undesirable geographies, to favour higher quality segments for better services, and other. Only recent work, e.g., in the IETF compute-aware traffic steering (CATS) [C4-5] WG, investigated joining compute-centric policies with path selection through, e.g., SRv6.

What must be researched: The definition of ***metrics*** encompassing compute, service, as well as communication constraints started recently in a number of standards developing organizations (SDO) and research efforts, as further elaborated in Section 4.8 for energy efficiency as a first class metric. But the space for policy decisions to drive the end-to-end delivery from client to one of possibly many service endpoints is a large one that not only requires ***standardised ontologies*** to act upon but also the ***decision mechanisms*** that are suitable to be executed at the speed of data transfer, if wanting to implement a vision of joining discovery and transfer into a single operation. Work has shown [C4-6] that simple weighted fair queuing (WFQ) can realised powerful ***compute-centric selections***, which in turn can be combined with routing decisions by dynamically adjusting the routing constraints to the weighted destinations. However, formal approaches, such as ***multi-constraint routing***, that would include not just networking constraints like delay and bandwidth but also service constraints, such as energy consumed, data availability, and others, may provide the basis for joining the service and network policy space while still resulting in a distributed state with a highly efficient routing lookup solution, thus addressing efficiency and policy richness in a deployable system. However, a system that could utilise real world data, observed in near real-time and fed into framework of ***collaborating AI agents at the network ingress decision points*** could likely need to a further increase in dynamicity that can be supported but also join the desired operational efficiency with the observed state of the system, following our vision for new mechanism shown in Figure 4-4-1.

4.6.3 Synopsis of Required Research

Research Theme	*Optimised Capability Discovery*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Backward-compatible Discovery-on-Transit	*Mid-term*	Discovery solution for limited domain type of service communication, compliant with local regulations and governance, particularly on security, privacy, trustworthiness and energy efficiency. **[KPIs]:** (i) reduction of at least 60% of flow completion latency (ii) increase of at least 5x compute resource utilisation in selected realistic use cases, (iii) improve on energy consumption by at least five times	Discovery-on-transite will have strong impact on sovereignty, not just at pan-European but user/organisation level, allowing to optimise concerns on security, privacy, and sustainability, while reducing the dependency on large DNS service providers.
Joint Discovery of Endpoint and Network Path	*Mid-term*	Suite of solution that jointly optimise discovery-on-transit with packet-level path selection, taking into account service and network level requirements, particularly security, privacy, and energy efficiency. **[KPIs]:** improve KPIs of discovery-on-transit only by at least 50%	Aligning service and network usage policies and constraints will allow for improving on sovereignty and sustainability concerns from service level “down to the wire”, thus pushing control into the hands of users/organisations.

4.6.4 Recommendations for Actions/Instruments

Research Theme	*Optimised Capability Discovery*	
Action	**Backward-compatible Discovery-on-Transit**	**Joint Discovery of Endpoint and Network Path**
International Calls		**X** *International calls to drive Europe’s vision into the technical solution space*
International Research	**X**	**X**

	Foster international research to drive Europe's vision into the technical solution space	*Foster international research to drive Europe's vision into the technical solution space*
Open Data		**X** *Foster availability of open operational data to optimise the needed steering mechanisms and replicate results*
Large Trials	**X** *Bring solutions as key input into Euro3C trials as a foundation for Europe's Open Internet Stack*	**X** *Bring solutions as key input into Euro3C trials as a foundation for Europe's Open Internet Stack*
Cross-domain research	**X** Investigate needed security requirements/solutions as well as sustainbility concerns	**X** Investigate needed security requirements/solutions as well as sustainbility concerns

4.7 Ensuring Secure and Compliant Protocols

A crucial question for any (new) protocol is whether it is working as expected, not only as consistent behaviour but also a compliant one against key constraints arising from the stakeholder needs. With the anticipated integration of data and AI training into the innovation cycle, answering this question will require methods taking into account those new factors that define the protocol behaviour. In a network-data-compute continuum, compliance testing must extend beyond individual protocol stacks to encompass the emergent behaviour of jointly optimised networking, computing, and data placement where an AI model steering routing decisions, a Digital Twin predicting congestion, and a transport protocol enforcing determinism interact as a system. Methods for testing compliance and correctness of such tightly coupled, AI-steered protocol ensembles are largely absent and represent a critical research gap.

4.7.1 Transparency, Traceability, and Explainability

Methods must be developed that allow for ***transparency***, ***traceability***, and ***explainability*** of the developed protocols and their deployments so that human developers and operators are still able to understand the current deployment and to perform failure analysis and fixes in case of problems.

State of the Art: There exist multiple approaches from different directions, mainly in the context of network management. For example, a “framework for LLM Agent-Assisted Network Management with Human-in-the-Loop” [C4-7] was proposed within the NMRG of the IRTF that employ an enhanced telemetry module, an LLM agent decision module, and interaction data models between human operators and LLM agent-driven systems, as well as workflows to enforce human oversight. Another approach is Neuro-Symbolic AI [C4-8] that uses Symbolic AI with explicit rules and logic and generates an Answer Set Programming formulation that is rule-based, deterministic, and explainable. Confucius [C4-9] is a framework that integrates multi-agent LLMs with existing management tools and existing network validation methods. It supports human/model interactions, uses domain specific languages, and incorporates its own validation framework to prevent regressions. Agua [C4-10] is an approach that uses human-understandable concepts to support explainability. The LLM-based Net Analyzer rApp [C4-11] is an approach to enable explainable and safe, human-in-the-loop automation for RAN operations. However, none of the described approaches can be directly applied to the context of protocol design and compliance.

What must be researched: Several research questions follow which need to be addressed: which combination of methods is suitable to ensure that the designed protocols are fulfilling the specification and given requirements. This can be a mixture of more constrained models that use domain-specific languages, Symbolic AI, additional telemetry, concepts, validation, or even verification. Moreover, overhead and reaction time that come with these solutions must be considered.

4.7.2 Interoperability

As part of compliance, interoperability is a key requirement in most communication environments. But the emergence of Limited Domains yet interconnected with a global communication fabric poses the question where interoperability is needed, how to handle it best (which SDO, if any), and how to align the anticipated faster evolution of protocols in Figure 4-1 with the organisational processes that come with SDOs. Research on organisational methods in the presence of the growing importance of AI-based methods is needed. A further dimension arises from data sovereignty constraints: in the NDC continuum spanning jurisdictional boundaries, interoperability mechanisms must not only ensure functional compatibility but also verifiable compliance with data locality and governance rules. Interoperability frameworks must be policy-aware, capable of expressing and enforcing placement constraints on both data flows and the computing functions that process them, as part of the protocol behaviour itself rather than as an out-of-band management concern.

State of the Art:

What must be researched:

4.7.3 Synopsis of Required Research

Research Theme	*Ensuring Secure and Compliant Protocols*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Transparency, Traceability, and Explainability	*Mid-term (finished in 5y)*	Methods for testing compliance and correctness, delivering transparency, traceability, and explainability of the developed protocol Systematic compliance testing, human-understandable explanations. **[KPIs]:** At least 75% coverage of protocol specifications and test cases within next 10 years	This research topic contributes to the combination of AI and softwarisation. Transparency and explainability are important values that should be supported between AI and humans.
Interoperability	*Mid-term (finished in 5y)*	Methods to ensure organisational alignment for efficient interoperability testing and assurance. **[KPIs]:** i) At least 50% coverage of protocols within next 5 years ii) re-test within at least 1 hour of deployment of new solution in 75% of all new protocols	Cost reductions and assurance of constraints are well-documented benefits of interop, which increases in importance with AI and softwarisation proliferation in protocols.

4.7.4 Recommendations for Actions/Instruments

Research Theme	*Ensuring Secure and Compliant Protocols*	
Action	**Transparency, Traceability, and Explainability**	**Interoperability**
International Calls		**X** *Ensure the coverage of global SDOs through suitable global partnerships and involvements*

International Research	**X** *Ensure the involvement of international excellence*	**X** *Ensure the involvement of international excellence*
Open Data	**X** *Foster open availability of key operational data to test methods with replicable results*	
Large Trials	**X** *Ensure transfer of results and avoid lock-in lower TRLs*	
Cross-domain research		**X** Ensure involvement of researchers on governance, standardisation

4.8 Energy Efficiency as First-Class Protocol Design Invariant

Energy consumption has historically been treated as a deployment-time concern rather than a property embedded in protocol design itself. The network-data-compute continuum changes this: when networking, computing, and data processing are jointly steered as a single cognitive fabric, the energy cost of a route/transmission decision at a protocol level becomes inseparable from its functional outcome. Protocols must therefore consider energy and greenness metrics as first-class observable and controllable parameters, alongside latency, throughput, and reliability. This requires research on formulising energy efficiency as a cross-layer protocol design invariant, measurable against the sustainability key value indicators (KVIs) central to the 6G framework and the European Green Deal.

4.8.1 Energy Observability and Representation in Protocol Design

Current protocol frameworks have no standardised mechanism for expressing energy cost as a protocol-level parameter. QoS models expose latency, bandwidth, and loss as negotiable properties; energy is absent. Where energy is tracked at all, it is measured externally at the hardware or infrastructure level and is not visible to protocol logic. This means protocol decisions cannot be energy-aware by construction: a routing decision, a retransmission policy, or a congestion control action cannot weigh its energy consequences because no signalling primitive exists to carry that information.

State of the art: Work in green networking has proposed energy-proportional networking and sleep-scheduling for routers. The IETF has explored energy management MIBs (RFC 7326) and within the Routing area several WGs have debated on frameworks for energy aware control planes.

What must be researched: A formal energy metric framework for protocol design is needed, defining: how energy cost is quantified at each protocol layer (transmission energy, processing energy, idle listening energy, in-network compute energy); how these per-layer costs are composed into end-to-end energy path metrics; how energy constraints and capabilities are signalled between protocol peers and across administrative domains; and how energy compliance thresholds — analogous to SLA parameters — can be verified and enforced at runtime. This research must be grounded in real-world operational datasets and requirements, in particular considering the emerging needs in data centers, and the recent trend towards smart data centers and their decentralisation.

4.8.2 Energy compliance testing and certification for protocols

Compliance testing for protocol correctness (functional behaviour, interoperability) is well established. Compliance testing for energy behaviour is not. Energy consumption under realistic operational conditions is highly scenario-dependent: traffic patterns, topology, device heterogeneity, and in-network compute load all interact to produce energy profiles that cannot be predicted from protocol specifications alone. Furthermore, emerging regulatory obligations — the EU Energy Efficiency Directive, CSRD sustainability reporting, and anticipated 6G KVI certification frameworks — will require that protocol deployments demonstrate measurable energy performance, but no standardised methodology or toolchain exists for this.

State of the art: Network benchmarking frameworks (e.g., RFC 2544, Y.1564) define procedures for measuring throughput, latency, and loss but have no energy dimension. Carbon footprint estimation tools for ICT infrastructure exist at the system level (e.g., SCI from the Green Software Foundation) but do not decompose to the protocol layer. Simulation frameworks such as ns-3 model some energy consumption at the MAC/PHY layer in specific wireless scenarios but lack cross-layer, end-to-end energy modelling for heterogeneous protocol stacks.

What must be researched: Standardised benchmarking methodologies for protocol energy compliance, extending existing frameworks with energy measurement procedures applicable across heterogeneous deployments. Open simulation and emulation environments — building on and extending existing tools such as deterministic wireless simulators calibrated to 5G+/6G reference scenarios — that model energy as a cross-layer, scenario-dependent property. Methods for certifying that a deployed protocol configuration meets energy KVI thresholds under representative traffic conditions, including adversarial and anomaly scenarios. Linkage between protocol-level energy benchmarks and system-level sustainability reporting frameworks to enable end-to-end auditability from protocol decision to carbon accounting.

4.8.3 Cross-layer energy-aware scheduling: from MAC to transport

Protocol layer independence is the foundational design principle of the Internet stack. This deliberately hides the internal behaviour of each layer from layers above and below. This abstraction is necessary for modularity and deployability but systematically prevents energy-aware co-optimisation. In IIoT environments.

State of the art: Cross-layer design has been extensively studied in wireless networking, particularly for energy harvesting and battery-constrained devices (e.g., IEEE 802.15.4, LoRaWAN duty cycling). QUIC exposes some transport-layer observability to application logic. TSN and DetNet define cross-layer timing constraints from application through MAC. However, none of these approaches treat energy as a jointly optimised parameter propagated across all layers simultaneously, nor do they address the three-way coupling between networking energy, computing energy (in-network processing), and data placement decisions that characterises the continuum.

What must be researched: Formal models of cross-layer energy coupling in heterogeneous protocol stacks, validated against operational datasets from IIoT deployments (deterministic wireless networks, edge-cloud orchestration environments). Protocol mechanisms for exposing energy state and energy budgets across layer boundaries, analogous to explicit congestion notification (ECN) for congestion, but for energy. enabling upper layers to adapt their behaviour in response to lower-layer energy conditions. Joint optimisation frameworks that treat MAC scheduling, transport flow control, routing, and in-network compute placement as a single energy-constrained optimisation problem, with tractable approximations suitable for real-time protocol operation. Investigation of stability and convergence properties of such cross-layer energy feedback loops, particularly at the timescales relevant to protocol operation (milliseconds to seconds) versus those of AI model retraining (minutes to hours).

4.8.4 AI-driven energy optimisation in protocol operation

AI models trained on operational energy datasets can in principle learn to predict the energy consequences of protocol decisions and to recommend or autonomously select energy-optimal configurations. However, deploying such models within the protocol control loop introduces fundamental tensions: AI inference itself consumes energy, potentially negating the savings it produces; AI models may generalise poorly to unseen traffic or topology conditions, producing unstable or counterproductive energy behaviour; and the timescales at which AI models can be retrained (minutes to hours) are mismatched with the timescales at which protocol parameters must adapt (milliseconds to seconds). Data from IIoT edge deployments shows that AI inference workloads running on the same nodes as protocol processing functions create measurable energy interference, complicating energy budget management.

State of the art: Reinforcement learning has been applied to congestion control (e.g., Aurora, Orca) and network routing, demonstrating that AI can learn protocol policies competitive with hand-crafted ones. However, **none of these approaches optimise for energy:** their reward functions target latency, throughput, or fairness, and the energy cost of the AI inference itself is not accounted for. CODECO has addressed energy-efficient MARL training, but the protocolar and networking aspects have not been worked ipon.

What must be researched: Lightweight AI architectures suitable for energy-constrained protocol control loops, with bounded inference latency and energy cost that can be formally accounted for in the energy budget of the protocol system. Methods for separating slow-timescale AI model training (using operational energy datasets from IIoT and telco-cloud deployments) from fast-timescale protocol parameter adaptation (using model outputs as lookup tables or policy functions), with formal analysis of the stability and optimality properties of this two-timescale architecture. Mechanisms for detecting AI model degradation in energy-optimisation tasks where the model's predictions diverge from actual energy behaviour due to distributional shift and triggering retraining or fallback to energy-safe protocol defaults without disrupting operational continuity.

4.8.5 Synopsis of Required Research

Research Theme	*Energy Efficiency as First-Class Protocol Design Invariant*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
Energy Observability and Representation in Protocol Design	***Mid-term***	A formal, standardised energy metric framework for protocol design, defining per-layer energy cost quantification (transmission, processing, idle listening, in-network compute), composition into end-to-end energy path metrics, and inter-domain signalling primitives . Validated against operational datasets from data centres and IIoT deployments. **[KPIs]:** At least 75% of all protocols exposing necessary energy metric data within next 5 years	Supports 6G KVI sustainability targets and European Green Deal obligations. Enables energy-aware protocol design by construction, a capability absent from all current protocol frameworks. Addresses societal value of reduced ICT energy footprint. SMART: delivers a specifiable, measurable energy metric standard within 5 years.
Energy compliance testing and certification for protocols	*Mid-term*	Standardised benchmarking methodology for protocol energy compliance, e.g., extending RFC 2544/Y.1564 with an energy dimension. Open simulation/emulation environments (e.g., building on ns-3 and 5G+/6G reference scenarios) modelling energy as a cross-layer, scenario-dependent	Directly enables compliance with EU Energy Efficiency Directive and anticipated 6G KVI certification frameworks. Provides industry and regulators with the first auditable, protocol-level energy accountability chain. High societal value: closes the gap between protocol

		property. Certification procedures for deployed protocol configurations under representative and adversarial traffic. Linkage to system-level sustainability reporting (SCI, CSRD). **[KPIs]:** At least 50% of all new protocols tested and certified using benchmarking methodology	behaviour and carbon accounting. SMART: delivers open toolchain and certification methodology within 5 years.
Cross-layer energy-aware scheduling: from MAC to transport	*Long-term*	Formal models of cross-layer energy coupling in heterogeneous stacks (MAC, transport, routing, in-network compute), validated against IIoT and edge-cloud operational datasets. Protocol mechanisms for cross-layer energy state exposure — analogous to ECN for congestion — enabling upper-layer adaptation to lower-layer energy conditions. Joint optimisation frameworks with tractable real-time approximations. Formal analysis of stability and convergence of energy feedback loops at protocol timescales (ms–s). **[KPIs]:** reduction of energy consumption by at least factor 5	Goes beyond SotA (IEEE 802.15.4, QUIC, TSN/DetNet) by treating energy as a jointly optimised parameter across all layers simultaneously, including in-network compute. Supports IIoT sustainability and reliability simultaneously. Addresses efficiency as an ethical obligation in constrained-resource deployments. SMART: delivers validated formal models, protocol mechanisms, and optimisation frameworks within 7 years.
AI-driven energy optimisation in protocol operation	*Long-term*	Lightweight AI architectures with bounded inference latency and energy cost, formally accountable within protocol energy budgets. Two-timescale architecture separating slow AI model training from fast protocol parameter adaptation, with formal stability and optimality analysis. Mechanisms for detecting AI model distributional shift in energy tasks and triggering safe fallback or retraining without operational disruption.	Goes beyond SotA (Aurora, Orca, CODECO) by optimising for energy in the reward function and accounting for the energy cost of AI inference itself. Enables self-optimising, energy-sustainable protocol systems. Societal value: reduces AI-induced energy overhead in network infrastructure. SMART: delivers formally characterised AI-protocol co-design

		Validated on IIoT and telco-cloud operational energy datasets. **[KPIs]:** reduction of energy consumption by at least factor 5	framework within 7 years.

4.8.6 Recommendations for Actions/Instruments

Research Theme	*Energy Efficiency as First-Class Protocol Design Invariant*			
Action	**Energy Observability and Representation in Protocol Design**	**Energy compliance testing and certification for protocols**	**Cross-layer energy-aware scheduling: from MAC to transport**	**AI-driven energy optimisation in protocol operation**
International Calls		**X** Joint calls (e.g. EU-Japan, EU-US) to align energy metric standards across regulatory regimes and protocol standardisation bodies (IETF, ITU-T, 3GPP)		**X** International AI-for-networks calls to share operational energy datasets and co-develop lightweight AI architectures for constrained protocol environments
International Research	**X** Collaboration with IETF working groups (Routing, OPSAWG), ITU-T SG13, and ETSI ENI to develop and vet the formal energy metric framework across jurisdictions		**X** Joint research with IEEE 802, IETF DETNET/TSN, and academia to develop cross-layer coupling models validated on shared IIoT reference datasets	
Open Data	**X** Publication of operational energy measurement datasets from data centres and IIoT deployments to ground the formal metric framework in real-world conditions	**X** Open release of simulation/emulation environments, calibrated reference scenarios, and benchmarking toolchains to enable community-wide compliance testing		**X** Open sharing of IIoT and telco-cloud energy datasets for AI model training, and open publication of trained lightweight model artefacts to support reproducibility
Large Trials			**X** Large-scale IIoT and edge-cloud	**X** Trials in live telco-cloud and

			trials to validate cross-layer energy scheduling models and joint optimisation frameworks under realistic heterogeneous conditions	IIoT environments to evaluate AI-driven energy optimisation under real traffic distributions and measure inference energy overhead
Cross-domain research			**X** Collaboration across networking, power systems, and sustainability research communities to bridge protocol-layer energy models with grid-level and carbon-accounting frameworks	**X** Cross-domain collaboration between AI/ML, protocol engineering, and sustainability science to co-design reward functions that jointly optimise network performance and energy KVIs

4.9 Chapter Summary

Protocols and their enabled communication paradigms will increase in their importance to realise the many possible architectures and overall system view laid out in Chapter 2. But the development of those protocols also undergoes a dramatic change through data-driven AI-based methods, all of which are embedded into the multi-stakeholder, requirement-rich environment in which they operate.

This chapter captured this evolution of protocol development and the required research to drive the evolution towards sustainable, functioning, but also accountable behaviour, encoded in those future protocols. The research topics are wide ranging but appropriately grouped with identified mid-term but also long-term research themes that all contribute to the sustainability, security, privacy, resilience, and performance expected from future communication and computing systems that interconnect the rich world of AI-native and cloud services.

The chapter interacts specifically with:

- **System View and Architecture Orientations:** Protocols implement the interaction models required by distributed and agent-based architectures, enabling coordination across heterogeneous domains and stakeholders.
- **Softwarisation and Intelligence in Networks:** Advanced communication paradigms support AI-driven control loops by enabling efficient exchange of context, knowledge, and intent between distributed components.
- **Physical Layer Innovations:**Protocol design is constrained by physical layer capabilities and, in turn, defines requirements for flexibility, abstraction, and dynamic resource control at the physical layer.

5 Physical Layer Innovations

5.1 Introduction

The physical layer is undergoing a profound transformation as communication systems evolve from connectivity-centric infrastructures into platforms enabling a pervasive cyber-physical continuum. Future networks must simultaneously support extreme capacity growth, ultra-low latency, massive connectivity, precise positioning and sensing, high reliability, and stringent sustainability and security requirements. In this context, the PHY layer becomes a cornerstone of 6G systems, shaping not only achievable performance but also long-term scalability, energy efficiency, resilience, and Europe's digital sovereignty.

Advanced optical communications and networks form the foundation of this transformation. Optical technologies underpin cloud and AI platforms, 6G fronthaul and backhaul, broadband access, and data centres, and non-terrestrial connectivity. As traffic volumes and AI workloads grow exponentially while sustainability constraints tighten, optical networks must evolve from capacity-driven systems into programmable, AI-native, secure, and energy-efficient infrastructures spanning core, metro, access, data-centre, and space–air–ground domains. This chapter therefore addresses key research priorities in optical networking, including ultra-high-capacity transport, advanced fibre technologies, AI-driven network operation, radio-photonics, integrated sensing, converged broadband access, photonic integration, and quantum-enabled secure communications.

On the radio side, 6G builds upon the foundations laid by 5G eMBB, URLLC, and mMTC, but targets a broader set of usage scenarios, including immersive communication, hyper-reliable low-latency services, massive and sporadic connectivity, ubiquitous coverage, AI-native networking, and integrated sensing and communication. The ongoing 3GPP 6GR feasibility study and ITU-R IMT-2030 framework recommendation's research targets define ambitious KPIs such as sub-Tb/s peak data rates, sub-millisecond latency, extreme reliability, centimetre-level positioning accuracy, and unprecedented energy efficiency. Achieving these goals requires continued and strengthened research on radio PHY fundamentals, including waveforms, modulation and coding, non-orthogonal multiple access, massive and cell-free MIMO, mmWave and THz communications, near-field operation, reconfigurable intelligent surfaces, and full-duplex technologies, complemented by AI/ML-empowered PHY design and comprehensive experimentation and validation frameworks.

Non-Terrestrial Networks (NTN) are becoming an integral part of future mobile systems, fundamentally reshaping the PHY design space. The integration of satellites, HAPS, UAVs, and optical wireless links with terrestrial networks calls for interoperable, environment-aware, and secure PHY solutions capable of operating across highly heterogeneous links, frequencies, and dynamics. Similarly, IoT networks introduce contrasting PHY requirements, ranging from massive, ultra-low-power and sporadic communications for billions of devices to deterministic, ultra-reliable links for critical industrial control. Enabling the coexistence of these diverse regimes requires novel PHY/MAC designs, context-aware operation at the edge, and scalable spectrum and resource management.

Across all these domains, security and sustainability emerge as first-order PHY design objectives. Future PHY technologies must embed quantum-safe and physical-layer security mechanisms, support trustworthy AI-driven operation, and provide resilience against interference, spoofing, and infrastructure failures. At the same time, capacity growth must be decoupled from energy consumption, material use, and electronic waste through power-proportional hardware, photonic-electronic convergence, modular architectures, and lifecycle-aware design.

This chapter provides a comprehensive view of PHY layer innovation across optical, radio, non-terrestrial, and IoT systems, identifying the research challenges and opportunities that will enable Europe to lead the development of secure, sustainable, and high-performance 6G infrastructures.

5.2 Chapter Vision: where are we, hot trends, where do we want to be

The physical layer of communication networks stands at a pivotal point. Over the last decade, incremental innovation in modulation, coding, multiple access, antenna systems, and optical transmission has enabled the deployment of 5G systems and the massive growth of cloud and broadband infrastructures. At the same time, communications networks have evolved into critical societal infrastructures, comparable in importance to energy grids and transport systems, underpinning economic activity, public safety, and digital services across Europe.

However, current PHY architectures are increasingly strained. Traffic is no longer driven solely by human-centric services but by artificial intelligence workloads, distributed computing, sensing, automation, and digital twins. These workloads impose requirements that stretch beyond traditional KPIs such as peak throughput or spectral efficiency, demanding ultra-low and predictable latency, extreme reliability, scalable device density, fine-grained positioning and sensing, and strict energy and carbon efficiency. At the same time, network architectures are becoming vastly more heterogeneous, integrating optical core, metro, access, radio access networks, data centres, edge compute, non-terrestrial platforms, and IoT systems into a single, tightly coupled ecosystem.

From a standardisation perspective, the transition from 5G towards 6G has formally started. The 3GPP feasibility studies on 6G radio (6GR), aiming at natively integrated terrestrial/non-terrestrial networks, and the ITU-R IMT-2030 framework recommendation define a new generation of research targets, including sub-Tb/s peak rates, sub-millisecond air-interface latency, centimetre-level positioning accuracy, native support for integrated sensing and communication, and unprecedented levels of energy efficiency and network intelligence. These targets cannot be achieved through incremental PHY evolution alone. They require a fundamental rethinking of PHY design principles across optical, radio, non-terrestrial, and IoT domains, and a much tighter coupling with network architecture, computing, control, and sustainability objectives.

Several converging trends are redefining the role and scope of the PHY layer.

- First, the convergence of communications, computing, and sensing is becoming central. Optical networks are no longer passive transport layers but programmable substrates supporting distributed AI, cloud, and edge computing. Radio PHY technologies are evolving to support not only data transmission but also high-precision positioning, environment mapping, and object detection through integrated sensing and communication. Fibre and radio infrastructures are increasingly repurposed as pervasive sensor networks, enabling proactive resilience, monitoring, and new services.

- Second, AI-native operation is moving into the PHY layer itself. Machine learning is no longer confined to network management but is being explored for waveform design, channel estimation, beam management, coding and decoding, and adaptive resource allocation. At the same time, this raises new challenges around trustworthiness, explainability, robustness, and validation, requiring hybrid solutions that combine physics-based models with data-driven learning.

Third, extreme heterogeneity and multi-layer integration are becoming the norm. Future networks natively integrate terrestrial and non-terrestrial segments, including satellites, HAPS, UAVs, optical wireless links, and underwater communications. They operate across an unprecedented range of frequencies, from sub-GHz and FR3 to mmWave, sub-THz, optical bands, and quantum channels. PHY technologies must therefore be interoperable, scalable, and adaptable across vastly different propagation conditions, device capabilities, and regulatory constraints.

- Fourth, energy efficiency and sustainability are no longer secondary constraints but primary design drivers. AI data centres already operate at gigawatt scale, and 6G deployments risk exacerbating energy consumption if capacity growth is not decoupled from power, material use, and network complexity. This trend drives innovation in power-proportional hardware, photonic integration, optical bypass, multi-hybrid radio architectures, lifecycle-aware design, and carbon-aware control at the PHY and system levels.
- Finally, security, resilience, and trustworthiness are intrinsic PHY concerns. The broadcast nature of wireless media, the reliance on GNSS for timing and positioning, and the increasing scale of automation expose the PHY layer to new attack surfaces. Research trends therefore include physical-layer security, quantum-safe and quantum-native communication, AI-assisted attack detection, GNSS-free positioning, and resilient PHY designs for critical and non-terrestrial infrastructures.

The vision articulated in this chapter is a PHY layer that acts as an intelligent, sustainable, and trustworthy foundation for the 6G system-of-systems.

From a technological standpoint, the future PHY must provide orders-of-magnitude improvements in capacity and flexibility, not only through higher frequencies but through new waveforms, extreme and distributed MIMO, near-field operation, reconfigurable propagation environments, and ultra-high-capacity optical transport. Optical and radio PHY technologies must co-evolve, leveraging photonics for radio processing, fronthaul, and sensing, while ensuring seamless integration across access, metro, core, data-centre, and non-terrestrial networks.

From an architectural perspective, PHY design must be natively aligned with 6G system architecture. This implies support for disaggregation, cloud-native operation, open interfaces, and dynamic functional splits. PHY technologies should expose controllable abstractions to higher layers, enabling end-to-end optimisation across communication, computation, sensing, and control, and enabling new services such as network-as-a-computer and sensing-as-a-service.

From an operational viewpoint, the target is AI-native yet trustworthy PHY operation. Learning-driven PHY components should improve adaptability and efficiency, while remaining explainable, robust, and certifiable. Digital twins, large-scale experimentation, over-the-air validation, and realistic performance evaluation tools must become integral to PHY development and standardisation, reducing risk and accelerating industrial uptake.

From a societal and environmental perspective, the desired endpoint is sustainable and resilient connectivity at scale. Future PHY technologies must decouple performance growth from energy consumption and material use, support circular-economy principles, and enable carbon-aware operation. They must also underpin secure, resilient infrastructures capable of operating under failures, attacks, and environmental disruptions, supporting Europe’s strategic autonomy and public-interest missions.

In summary, the PHY chapter sets out a vision in which the physical layer evolves from a collection of isolated transmission techniques into a holistic, intelligent, and programmable foundation for 6G networks. By aligning optical, radio, non-terrestrial, IoT, security, and sustainability research under a coherent architectural vision, the chapter defines the research directions needed to ensure that Europe not only meets the technical ambitions of 6G but shapes them in line with its broader economic, societal, and environmental goals.

5.3 Advanced Optical Communications and Networks

5.3.1 Optical Networks for AI

The rapid rise of AI, cloud computing and data-intensive services is driving unprecedented traffic growth, at around 100% per year [C5-01]. Photonics already accounts for about 8% of data-centre (DC) power use and its share will increase. Optical networks are therefore becoming critical infrastructure for delivering distributed, AI across Europe. Yet 67% of European firms already use AI while 54% report poor network performance, 77% face connectivity issues, and 81% of telecom operators say customers request AI services their networks cannot yet deliver [C5-02].

Today's AI workloads for model training are concentrated in mega data centres used, configured to act primarily as giant distributed memories. The next phase will push AI inference closer to end-users, into smaller edge data centres, including 6G-ready telecom sites, to support time-critical applications.

Across the digital infrastructure, photonics will underpin AI with tailored solutions in each network segment:

- Inside data centres (DCs), low-cost pluggable and co-packaged optics, advanced optical switches, are needed to connect millions of servers efficiently (scale out) while limiting impact of laser failures and energy use, leveraging high radix optical circuit switches for fast reconfiguration, according to optimal DC architectures yet to be refined.
- Across and within campuses of data centres (scale across), hundreds of terabits of data need to be transported in systems at minimum 3.2 Tbit/s per wavelength, while the quasi-synchronous delivery of tokens at minimal packet loss by reworking protocols, containing contention and jitter. Therefore, Europe must rapidly upgrade optical backbones with long-haul systems using high-capacity, coherent optics, novel fibres (multi-core, multimode, hollow-core) by space division multiplexing, and multi-rail amplifiers operated in multiple wavelength bands, some of them unexplored yet, e.g. 1.0µm. Switching will have to be upgraded from packet to band and fibre switching and performed at millisecond to microsecond speed with appropriate control and automation, leveraging integrated full band transceivers. Europe has a global lead in anti-resonant hollow-core fibre – combining low loss, low latency and low non-linearity – which can be vital for time-critical AI and industrial uses. Maintaining and scaling this lead against intense Asian competition requires coordinated investment, standardisation and industrialisation. Hollow core fibre could greater simplify system design by allowing repeater-free, possibly bidirectional optical links at capacities and distances never achieved with standard single mode fibre.
- At the edge, novel orchestrators and switching approaches will be required to ensure the dual use of 6G edge DCs for both communications and data

As fibre networks are gradually transitioning into "the computer" itself and end-devices become mainly sensors, Europe's largely optical connectivity are as fundamental as power grids and transport systems. Photonic systems and networks therefore need sustained policy attention as strategic enablers of Europe's digital sovereignty. In an era of geopolitical tension, extra-territorial regulation, export controls, cyber-threats and climate-related disasters, fibre-network security and resilience must be treated as strategic priorities.

Data centres are already major energy users, with some facilities at gigawatt scale. To keep AI growth compatible with climate goals, Europe must drive large-scale deployment of energy-efficient photonic technologies for both interconnects and processing. Making photonics a core element of industrial, digital and climate policy is a precondition for “sustainable AI”, not an optional add-on.

5.3.2 Optical Networking by AI

Optical networks are becoming increasingly complex due to their growing scale, tighter performance requirements, and the convergence of multiple layers and domains. At the same time, they can now expose real-time telemetry, configuration, and operational data through open interfaces and standardized data models. This situation creates the opportunity for a shift from rule-based reactive network control today towards AI-native, autonomous network operation in the future.

Machine-learning techniques already provide tangible value where sufficient structured data exists and physical behaviour is well understood. Applications such as predictive maintenance, failure detection, root-cause analysis and capacity optimization are increasingly automated and integrated with telemetry pipelines and SDN control loops. In parallel, generative AI and large language models introduce new capabilities for handling unstructured operational data, providing knowledge-base assistance, and offering natural language interaction. They translate human intent into structured workflows and support reasoning in complex operational contexts.

A technology inflection point is now the emergence of agentic AI. Instead of single-step models, agentic AI systems consist of autonomous software agents that can iteratively reason, plan, and act to achieve defined objectives within constrained operational boundaries. In optical networking, such agents can reflect established expert workflows, for example hypothesis-driven troubleshooting or staged remediation planning. Modern agentic architecture combines multiple components: defined agent roles and system contexts, one or more AI models, access to network data and control interfaces, curated domain knowledge, and long-term memory to retain validated experience. Crucially, agentic AI must be tightly coupled with deterministic algorithms, policy constraints, and verification mechanisms to ensure predictable and safe operational behaviour in critical infrastructure environments.

A key enabler for scalable agent-based operation is the abstraction of network capabilities via discoverable tools rather than hard-wired integrations. Model context protocol and similar concepts provide a uniform way for agents to access telemetry, topology, inventory, simulation, planning and control data/functions through well-defined interfaces. This decouples agent operation from vendor-specific interface implementations, simplifies the composition of new agents, and aligns naturally with open SDN and disaggregated transport architectures. The ability to retain and reuse validated operational knowledge in long-term memory allows agents to improve over time, moving from isolated assistance towards continuously learning systems.

Agentic AI bears the potential to revolutionize optical network operations. While first proof-of-concepts look encouraging, substantial research and innovation is needed to move from early demonstrations to hardened and operationally trusted AI-native optical network automation frameworks. Research priorities include the definition of interoperable, agent-ready interfaces for optical networking tools and data; development of hybrid agent architectures that combine AI reasoning with deterministic control, guardrails, and human oversight; definition of domain-specific knowledge bases and retrieval mechanisms tailored to optical network operations; and governed long-term memory frameworks suitable for critical environments. Further emphasis should be placed on the development of digital twins that enable agent-driven experimentation, what-if analysis, and closed-loop validation before deployment, as well as on multi-agent coordination across layers and domains. Finally, robust benchmarking, validation, and certification approaches are required to establish trust by measuring correctness, robustness, and operational impact of AI-driven functions.

5.3.3 Radio Photonics

The evolution towards 6G mobile systems is driving a profound transformation in radio network architectures, with optical technologies emerging as a cornerstone for meeting the escalating demands of capacity, flexibility, and energy efficiency. Two major trends underpin this shift: the relentless increase in bandwidth and data rates, and the growing complexity of radio system architectures, particularly with the proliferation of massive MIMO and advanced beamforming techniques.

Integrated photonics stands out as a key enabler, offering miniaturization, cost-effectiveness, and compatibility with radio frequency (RF) electronics. Silicon photonics offers favourable thermal properties and high bandwidth capabilities making it ideal for next-generation radio systems. The integration of microelectronics and photonics is paving the way for reconfigurable, high-performance antenna frontends capable of versatile operation across diverse frequency bands. As 6G networks evolve, the convergence of high-frequency integrated transceivers, co-packaged optics (CPO), on-board photonic interconnects, and true microwave photonics (MWP) elements will drive significant innovation, addressing the challenges of capacity, energy efficiency, and scalability in future radio units.

Modern radio hardware platforms are transitioning to disaggregated architectures, distributing processing loads across multiple computing units interconnected by high-speed, low-latency links. This paradigm shift is necessitated by data volumes that outpace Moore’s Law, and it introduces new challenges in signal integrity, power consumption, and thermal management. Electrical interconnects, constrained by frequency-dependent losses and crosstalk, are increasingly becoming bottlenecks. Optical interconnects, particularly those enabled by integrated photonics and CPO, offer a compelling solution, providing frequency-independent performance and the ability to transmit signals over much longer distances without degradation.

The placement of the split point between radio units (RU) and distributed units (DU) has deep implications for hardware requirements and the choice of optical technologies. While traditional fronthaul (FH) links relied on bandwidth-hungry split options (e.g., SO 8), the advent of 5G and 6G has seen a shift towards more bandwidth-efficient splits (e.g., SO 7-2), as standardized by the O-RAN alliance. However, even with these optimizations, the capacity requirements for FH and internal interconnects in massive MIMO systems can reach several terabits per second. Pluggable optical transceivers have evolved to support these needs, but their scalability is limited by power dissipation and form factor constraints. CPO modules, with their miniaturized footprint and superior energy efficiency, are poised to replace pluggable optics in future high-bandwidth density radio systems.

To further address the challenges of signal integrity and board-level connectivity, on-board photonic interconnects are gaining traction. Solutions include glass-based devices with embedded waveguides and polymer waveguides integrated into printed circuit boards (PCBs). These technologies enable compact, high-density optical routing, facilitating efficient distribution of RF and baseband signals within radio units. Multi-terabit bi-directional bandwidth can be managed by optical distribution systems, highlighting the potential of these approaches.

Shifting digital processing away from the RU and using analog optical links for RF signal transmission, Radio-over-Fibre (RoF) significantly reduces power consumption and simplifies RU design. Although extensively studied in the past, RoF would represent a disruptive innovation for 6G, where capacity, performance, and integration requirements are more demanding than ever. At the same time, challenges remain in terms of linearity, dynamic range, and the integration of large arrays of optical transceivers for massive MIMO systems. Integrated photonic arrays and wavelength-division multiplexing (WDM) are being explored to address these issues. AI-based baseband processing can further assist in linearizing the system and compensating for propagation impairments. Additionally, microwave photonics (MWP) techniques—such as photonic RF synthesizers, integrated mixers, and optical beamforming—are enabling new capabilities, including high-stability clock generation, wideband signal processing, and advanced antenna architectures.

5.3.4 Optical fibre sensing

Optical fibre sensing has been considered a specialized field relying on dedicated instruments, such as Distributed Acoustic Sensing (DAS) systems, which were usually deployed on dark fibres, or laboratory-grade interferometers, and intrusion-detection devices operating in isolation. Following the seminal article of Giuseppe Marra on Science in 2018, research start discovering new applications concerning earth science topics and the maintenance of the networks. Thanks to the widespread deployment of high-sensitivity coherent transponders in production networks, starting from the initial deployment within submarine cables, is opening an new range of applications. The global sensor fibre plant, which consists of hundreds of millions of km of terrestrial and submarine cables, already deployed and powered, can now function as a pervasive, zero-additional-hardware sensor grid. This “additional network” will have two functions: (i) complementing the existing traditional sensors, by providing coverture of remote areas such as oceans; and (ii) work as a tool for an update accurate knowledge of the cable and the channel characteristics.

This shift from isolated instruments to planetary-scale observatories and network management defines the strategic horizon for the coming decade with the SRIA. Europe is uniquely positioned to lead this transition, leveraging its rich submarine cable infrastructure, dense urban fibre deployment, and strong academic-industry competitiveness in photonics technologies and digital signal processing. All this, jointly to the availability of daily more power AI/ML techniques, will provide an extraordinary tool for optical fibre sensing. Now, the strategic challenge is to convert this structural advantage into research leadership and standardization contributions and leadership.

The most transformative near-term opportunity lies in extracting sensing signals directly from existing coherent WDM networks, i.e., just by using existing live traffic. This goes under the name of Integrated Sensing and Communication (ISAC). Modern coherent transponders continuously compute parameters within the ASIC that are used for equalization purposes, but could be also exploited, e.g., to sense the state of polarization, and thus indirectly sense “movements” close to the fibre. These same quantities encode rich information about polarization-state perturbations induced by temperature, mechanical stress, and seismic activity along the fibre path. The mid-term vision is a “Sensing-as-a-Service” model delivering real-time environmental monitoring at regional scale with zero sensing-specific CAPEX.

The majority of Earth's surface – the oceanic floor – remains essentially blind to seismology, and this submarine infrastructure could be used as a geophysical observatory. Transoceanic cables represent a unique opportunity to address this gap: phase and polarization measurements on submarine links have already demonstrated the ability to detect seismic events, slow-slip earthquakes, ocean infra-gravity waves, and precursor signals ahead of major earthquakes. Key research challenges include improving spatial resolution beyond repeater spacing (e.g., DAS beyond the first span), extending sensitivity into the ultra-low frequency (mHz) regime, and modelling the physical coupling between ocean dynamics and measurable cable perturbations to enable, e.g., tsunami early warning or better understanding of factor related to climate changes.

In this context, no single sensing modality covers the full dynamic range of geophysical and urban events. Consequently, a certain degree of multi-modal data fusion and AI-driven analytics will be needed. For example, DAS provides high spatial density but requires dark fibre; polarimetric sensing operates on lit networks with coarser resolution but it has (if using one single cable) limited resolution in localization; terrestrial seismometer networks are sparse but precisely calibrated. Fusing these heterogeneous data streams into coherent situational awareness – through statistically rigorous frameworks and deep-learning pipelines capable of generalizing across modalities and geographies – is a fundamental research priority. Network-scale / domains triangulation via centralized Optical Network Controllers (ONC) can further constrain event localization and characterization.

State-of-polarization fluctuations encode mechanical cable stress before fibre cuts or outages occur, providing a measurable early-warning window. This branch of research is called proactive network resilience, and it can be applied to both terrestrial and submarine networks, thus helping the network operators to spend large part of their annual budget for infrastructure maintenance. Real-time classifiers can flag perturbation signatures as proactive maintenance alerts, enabling traffic rerouting before service disruption and contributing to a fully autonomous network health layer.

Metropolitan fibre infrastructure can be repurposed as city-wide sensor arrays for traffic monitoring, structural health assessment, and urban seismic mapping - without privacy exposure, since fibre sensing yields aggregate physical data with no individual localization. This concept could enable a better urban sustainability and smart-city geophysics. Across all priorities, critical open problems include metrological calibration of relative measurements against physical quantities, low-latency edge processing for early-warning applications, and the development of open standardized telemetry interfaces and interoperability.

5.3.5 Optical broadband access

Passive Optical Network (PON) technology currently connects approximately one billion customers to wired optical broadband access, marking a significant milestone in global connectivity. The industry is rapidly advancing into a multi-gigabit era, driven by the evolution from GPON to XGS-PON, 25GS-PON, and 50G-PON. This progression is not only enhancing FTTH services but also enabling PON to serve as a foundational technology for diverse converged networking applications.

Beyond traditional residential broadband, PON is increasingly being deployed in smart cities, enterprises, campus networks, industrial factories, for converged mobile-fixed x-hauling, and within buildings.

AI is poised to become a central component in the future of broadband, involving a symbiotic relationship. "AI for broadband" optimizes network operations, while "broadband for AI" provides the necessary high-capacity, low-latency infrastructure for AI workloads and services. Furthermore, "AI and broadband" explores integrated processing and compute solutions. AI will play a crucial role in network management, facilitating fault control, identifying complex traffic patterns, and optimizing resource allocation. This intelligence will also require networks to address a growing demand on more controllable lower-latency, essential for emerging AI-driven applications. The set of advanced applications and use cases for PON impose new requirements at both the physical and transmission convergence layers, demanding more deterministic performance and precise QoS and QoE control. The need for highly reliable and predictable network behaviour is paramount for critical services in these environments.

Moreover, existing and globally deployed fibre optical distribution networks can be utilized for new service considerations, whether for high-capacity point-to-point wavelength overlays or for fibre sensing. Fibre sensing in access networks is gaining interest for monitoring and protecting infrastructure, performing proactive maintenance, and exploring new service opportunities for operators.

Alongside these advancements, an ever-increasing demand for capacity will drive the development of technologies beyond 50G. Simultaneously, there will be a prioritization of power efficiency, robust security, cost-effectiveness, and high availability to meet the complex demands of an intelligent, interconnected access network.

Key open questions for the future of wired optical access networks include the optimal use cases, features, and integration of AI, as well as techniques and applications for fibre sensing within existing deployed fibre infrastructures. Another critical question is which technology, and architectural option will best facilitate the move into higher-end optical access capacities towards 100-400 Gbit/s (be it using Intensity Modulation-Direct Detection, coherent, or a novel approach), while also addressing the stringent demands from non-residential services.

5.3.6 Optical integration 2.0

In the foreseeable future, optical integration will be driven by communication applications showing unprecedented growth in data traffic, AI workloads, and the need for scalable, energy-efficient interconnect technologies. Integrated photonics is a critical platform to overcome electrical bottlenecks and enable next-generation communication systems across datacenters, telecom, wireless, automotive, and space applications. It is a fundamental enabler for all communication and related technologies. Applications driving the development of Photonic Integrated Circuits (PICs) are

Ultra-high-bandwidth Interconnects: Interconnects, such as traditional pluggable modules, Near-Packaged Optics (NPO), Co-Packaged Optics (CPO) and other types of modules, will dominate research as data centers push to 800 Gb/s, 1.6 Tb/s, and beyond. Foundries are accelerating InP and silicon photonics platforms - TSMC's COUPE photonic engine enters volume production in 2026, enabling dense die-to-die optical interfaces for high-performance computing - and platforms based on other materials (e.g. $LiNbO_3$, $BaTiO_3$, $LiTaO_3$, Silicon Organic Hybrids (SOH) and plasmonics) are gaining commercial traction.

Research focus areas include low-loss, low power consumption, high-bandwidth PICs, advanced thermal management for tightly integrated optics and/or optical switching fabrics for AI-driven workloads.

Scalable Photonics Manufacturing and Heterogeneous Integration: While silicon photonics offers tremendous scalability and technological maturity, it is constrained by physical limitations, such as low bandwidth modulators, lack of light sources and comparatively high insertion loss. Other material platforms offer superior performance metrics but are economically less advantageous.

For this reason, research focus areas should include heterogeneous integration of different materials to overcome limitations of monolithic material platforms and the technological advancement of novel materials outside of photonics pilot lines.

Photonic Test, Assembly and Packaging Packaging and test have emerged the highest-cost and highest-risk elements of PIC production. Cost of test in high-volume application is proportional to the time needed for evaluation. With increased complexity of PICs cost of testing is increasing as well, becoming a significant cost factor in photonic systems. Similarly, assembly and packaging has been in the past the largest contributor to costs of photonic modules with often cited 80% of the total BOM. In addition, use in harsh environments (e.g. industrial temperatures, automotive, medical, space) is becoming more prevalent, putting more emphasis on reliability and robustness.

To address these challenges, research areas should include automated, high-precision alignment for mass production, thermal, mechanical, electrical and optical co-design for multi-chip modules as well as radiation-hardened packaging for satellite communications.

Integrated Photonics for Wireless, RF, and 6G: The latest IPSR-I roadmap identifies RF photonics and wireless-optical convergence as essential for future communication infrastructure, including ultra-low-latency fronthaul and high-frequency (mmWave/THz) systems.

Research will focus on integrating photonic frequency synthesizers, modulators, and beamformers.

Quantum Photonic Integration for Secure Communications: Quantum photonics is rapidly transitioning from a conceptual research frontier to a practical enabler of next-generation communication networks. Over the next five years, integrated quantum photonics will be central to building scalable quantum communication systems, secure network infrastructure, and early-stage quantum internet architectures. The pace of progress is accelerating, driven by major breakthroughs in loss reduction, error correction, and manufacturable photonic quantum chips.

Such innovations directly influence communication applications: Low-loss photonic platforms enable longer entanglement distribution distances, more reliable quantum key distribution (QKD), and efficient photon-pair generation, while integrated error correction supports robust operation in noisy environments typical of deployed fibre networks.

Research focus areas may include multi-wavelength quantum-safe channels multiplexed with classical data traffic, stabilized quantum light sources compatible with telecom bands, photonic interfaces for quantum memories and repeaters and/or error-tolerant quantum + classical coexistence on deployed fibre.

5.3.7 Synopsis of Required Research

Research Theme	*Advanced Optical Communications and Networks*		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
RT1: Optical Networks and Switching Architectures for Distributed AI and Cloud Computing	Short–Mid term (2026–2032)	AI-optimized data-centre, campus and metro architectures; fast, programmable optical switching; optical bypass and minimized OEO conversions with deterministic latency and loss. **[KPIs]:**12 Pbps node capacity;	Enables European AI competitiveness and digital sovereignty while reducing energy consumption and network complexity.
RT2: Ultra-High-Capacity Optical Transport and Advanced Fibre Technologies	Mid term (2027–2033)	Tb/s-per-wavelength transmission; multi-band amplification; SDM using multicore, multimode and hollow-core fibres.	Sustains long-term traffic growth while minimizing fibre deployment and maintaining EU

		[KPIs]: 100 THz optical spectrum; 2 Pb/s line capacity; 6.4 Tb/s port speed.	leadership in fibre technologies.
RT3: AI-Native Optical Network Operation with Digital Twins	Short–Mid term (2026–2030)	Agent-based automation, telemetry-driven control loops and physics-aware digital twins enabling closed-loop optimisation and trusted AI operation. **[KPIs]:** Self-optimising operation; autonomous closed-loop control; O(10 s) service provisioning.	Improves scalability, reliability and OPEX efficiency of critical optical infrastructures.
RT4: Radio-Photonics and Optical Support for 6G Fronthaul and Cloud-RAN	Mid term (2027–2033)	Integrated photonic solutions for fronthaul, massive MIMO and beamforming; CPO and board-level optics supporting disaggregated 6G architectures. **[KPIs]:** 6.4 Tb/s optical interfaces; <20 μs access latency; 80% lower power consumption.	Critical enabler for high-capacity, energy-efficient and flexible 6G radio systems.
RT5: Integrated Communication, Fibre Sensing and Proactive Network Resilience (ISAC)	Short–Mid term (2026–2031)	Extraction of sensing information from live coherent networks; AI-based multi-modal data fusion; Sensing-as-a-Service concepts. **[KPIs]:** Network-wide sensing; self-diagnosing operation; cm-level localisation accuracy	Transforms deployed fibre infrastructure into a planetary-scale sensor network for resilience, geophysics and smart cities.
RT6: Next-Generation Converged Optical Broadband Access (Wired and Wireless)	Short–Mid term (2026–2032)	100–400 Gbit/s access technologies beyond 50G-PON; converged fixed–mobile x-haul; AI-driven deterministic QoS/QoE; optical wireless access where fibre is impractical. **[KPIs]:**0.5T-PON readiness; 50 Gb/s residential and 100 Gb/s business services; <100 μs upstream latency.	Sustains Europe's FTTH leadership while enabling converged broadband, industrial and edge-AI services at controlled cost.

RT7: Advanced Photonic Integration and Co-Packaged Optics	Mid term (2027–2034)	Scalable PIC platforms, heterogeneous integration and manufacturable CPO/NPO solutions with reduced power consumption. **[KPIs]:** 6.4 Tb/s optical interfaces; 80% lower power consumption; support for Pbps-class systems.	Enables sustainable scaling of data-centre, telecom and wireless optical hardware.

5.3.8 Recommendations for Actions/Instruments

Research Theme	Advanced Optical Communications and Networks		
Action	**Optical Infrastructure for AI, Cloud and Compute Convergence** (RT1, RT2, RT7)	**AI-Native & Autonomous Optical Networking** (RT3, RT5)	**Converged Optical Connectivity for 6G, Access and Non-Terrestrial Networks** (RT4, RT6)
International Calls	EU–international research on optical transport, data centre interconnects, and advanced fibre technologies for AI workloads	Programmes on AI-driven optical networking, agent-based automation, and digital twins.	Joint calls on optical support for 6G (fronthaul, access, NTN integration) and integrated communication–sensing across RAN and fibre infrastructures
Open Data	Open datasets for optical networks, traffic patterns, and AI-driven workloads	Reusable datasets for AI training, telemetry, and network operation automation	Datasets for converged optical-radio systems, fibre sensing, and integrated RAN–optical sensing scenarios
Large Trials	Trials of multi-band transmission, SDM, and high-capacity optical switching in live networks	Pilots of AI-driven closed-loop control and digital twin-based validation	Trials of optical fronthaul, fibre sensing, and joint communication–sensing capabilities across optical and wireless domains
Cross-domain research	Collaboration across optical, cloud, and AI infrastructure communities	Integration with AI/ML, control systems, and network automation frameworks	Collaboration with radio, NTN, sensing, and access domains for end-to-end 6G integration and multi-layer sensing fabric

5.4 Radio Technologies and Signal Processing

This section briefly addresses the enabling technologies for the 6G radio interface, whose impacts on/connections to 6G KPIs (See Table 5-1) are shown in Figure 5-1. More details about most of these technologies were given in Chapter 6 of SRIA 2024 [C5-03] and the references therein.

Table 5-1: Key KPIs from 5G to 6G radio interface

Target KPIs (Capabilities)	5G (2020~2030) (IMT-2020 by ITU-R M.2083)	6G (2030~2040) (IMT-2030 by ITU-R M.2160)
Bandwidth	<0.5 GHz	<10 GHz
Energy efficiency (NW, UE)	Qualitative	>100% gain vs IMT-2020
Average spectral efficiency	7.8 b/s/Hz (DL), 5.4 b/s/Hz (UL)	1.5x to 3.0x vs IMT-2020
Peak data rate	20 Gb/s	50~200 Gb/s
User data rate	100 Mb/s	300~500 Mb/s
Area traffic capacity	10 Mb/s/m^2	30~50 Mb/s/m^2
Reliability	URLLC: 1-10^{-5}	HRLLC: 1-10^{-7}
Mobility	<500 Km/h	<500~1000 Km/h
U-plane latency	URLLC: 1 ms	HRLLC: 0.1~1 ms
C-plane latency	<20 ms	<20 ms
Connection density	1 device/m^2	1~100 device/m^2
Positioning accuracy	meter-level	1~10 cm

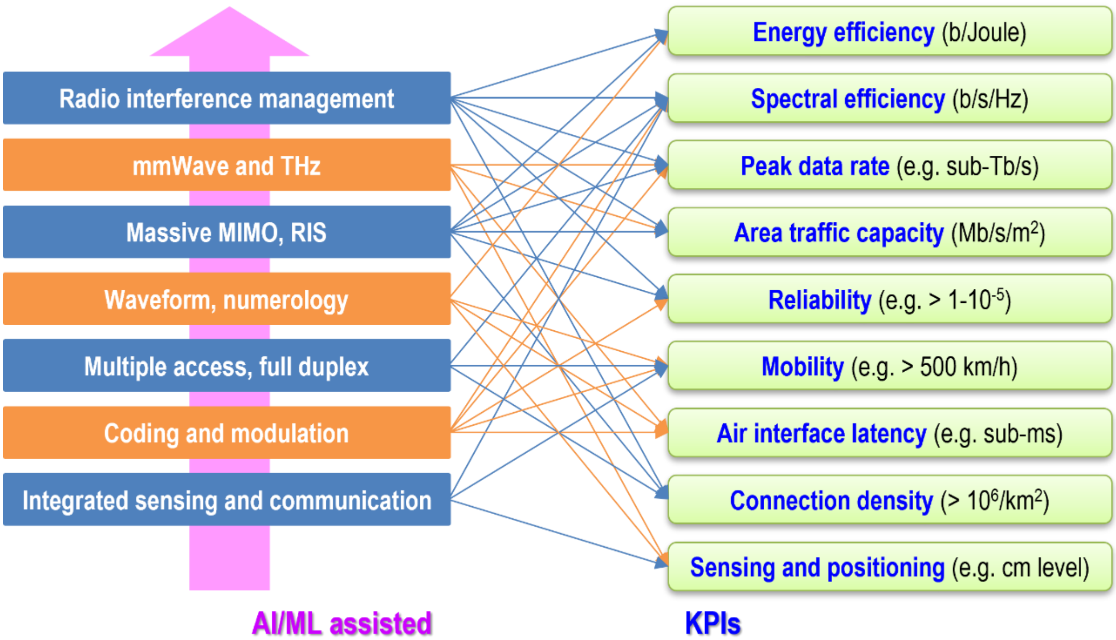

Figure 5-1: Technology enablers strongly impacting on different 6G KPIs, with example ITU-R minimum requirements for IMT-2030 radio interface [C5-04].

5.4.1 Radio Networking and Interference Management

Spectrum re-farming and sharing: Allocated spectrum is one of the main factors that determines the system capacity. It is crucial to reutilize the existing spectrum resources and optimize the access to new frequency bands. Adaptive and dynamic spectrum sharing can be applied to harmonize the joint utilization of both licensed and unlicensed bands. Spectrum sharing can be supported by multi-RAT connectivity, which allows the UE (user equipment) to choose the best RAT depending on the link qualities. Opportunistic spectrum access opens the door to further improvements by means of multi-user diversity, e.g. on a carrier sense multiple access (CSMA)-like basis. In this line, cognitive environment concepts, in which spectrum awareness, e.g. based on a combination of advanced signal design and AI techniques, can be used to ensure coexistence with existing in-band services.

Subnetworks and coexistence: Different scenarios such as factory floor level communications (e.g., within a robot or a production cell, i.e. in-robot and in-production cell communications), sensor/actuator traffic inside a car (in-car communications) and in/on-body communications among smartphones and XR wearables may call for different extreme requirements. To achieve different extreme requirements the so-called subnetworks can be used. The vision is to have an access point that a) controls and serves the needs of these devices next to each other, and b) is, at the same time, a sort of special device connected to an overlay public or private network. In this context, a dedicated air interface (e.g., supporting much shorter symbols), access protocols and diversity mechanisms need to be selected and designed. For life-critical applications, it is essential that mechanisms that allow them running autonomously even when out-of-coverage must be present.

Wireless edge caching: It is widely recognized that a major driver of the wireless data traffic increase is on-demand access to multimedia content. The users' content consumption pattern and the operators' data plans may often be dramatically mismatched. Recent research in information theory and wireless communication has shown that content distribution over a wireless network (e.g., on-demand video streaming) can be made much more efficient than current SotA technology, by caching content at the wireless edge. Treating on-demand content as unicast traffic is highly inefficient, e.g. in video-server services the library of popular movies can be controlled by the service provider, and can be updated at a relatively slow pace. Users' requests arrive to the server in an asynchronous way, such that the probability that many users wish to stream the same file at the same time is basically zero. Coded caching techniques have the ability of turning the unicast traffic (on-demand streaming) into a coded multicast traffic.

5.4.2 Millimeter-Wave and Terahertz Communication

ITU at the 2023 World Radio communication Conference defined several frequency ranges in the sub-THz band for potential future 6G network studies. Now, much has been accomplished in terms of understanding THz propagation, which are affected by molecular absorption loss, spreading loss, and blockage. For communications below one meter, the use of impulse-radio-like communication based on the transmission of one-hundred-femtosecond-long pulses following an on-off keying modulation spread in time has been proposed. For longer communication distances, enabled in part by ultra-massive MIMO, new dynamic bandwidth modulations are needed to not only overcome but even leverage the unique distance-dependent bandwidth created by molecular absorption. As of today, both single-carrier (e.g., M-QAM, M-PSK) and multi-carrier (e.g., OFDM) waveforms for THz systems have been developed. Non-conventional modulation schemes like zero-crossing modulation, have been proposed to mitigate the increased power consumption of the analog frontend. Independent of the modulation, physical-layer synchronization (in time, frequency and phase) remains also a challenge.

Note that the use of large transmit, reflect and receive antenna arrays, lenses and metasurfaces results in THz systems to often operate in the near-field. Conventional beam management strategies which consider far field operation, are no longer valid. New unique types of wavefronts or beams that are only possible in the near field need to be explored. These include, for example, self-healing Bessel beams that can recover from obstacles, or curving airy beams that can bend around corners.

Novel MAC protocols are required for THz networks. The very large available bandwidth almost eliminates the need for nodes to contend for the channel. The transmission of very short signals also minimizes the chances of collisions. However, the need for high-gain directional antennas simultaneously at transmitter and receiver to establish links over realistic distances makes the beam discovery and tracking a major challenge. Innovative neighbour discovery strategies that leverage the full antenna radiation diagram as well as new receiver-initiated MAC protocols have been proposed. Multi-hop relaying and routing strategies may be needed to support mobile ad-hoc THz networks, especially in outdoor applications. Cross-layer solutions that capture the trade-offs between antenna beamwidth, communication distance, available bandwidth, processing overhead and buffer capacity are needed, and the use of different forms of machine learning (ML) can help to operate such networks.

THz phased arrays need to combine agile beamforming and multi-user capabilities. Techniques inspired from lower bands such as hybrid/MIMO beamforming need to be revisited to take the specificities of the wireless channels and hardware implementations. Specifically, the inter-element spacing becomes so small that the antenna array and interconnect structure must be taken into account in the overall system design. Thermal analysis is also critical, given the small dimensions. In addition, new experimental platforms and simulation tools are needed. The majority of the platforms developed to date are channel sounders or physical layer testers that rely on non-real-time DSP, and mostly at sub-THz frequencies, and real-time platforms will become a must for testing systems beyond a point-to-point link.

5.4.3 Next Generation MIMO

Ultra-massive and extreme MIMO: mmWave and THz-band communications suffer from extremely high and frequency-selective path loss (>100 dB over meter-scale distances), necessitating high-gain directional antennas and large-scale arrays such as massive MIMO (mMIMO) with tens to hundreds of elements, where effects like mutual coupling and transmitter nonlinearities addressed via MIMO DPD become critical. At THz frequencies, the reduced antenna spacing enables ultra-massive MIMO (umMIMO), often realized using nanomaterials, metamaterials, and plasmonic nano-antennas to form extremely dense arrays (e.g., 1024 elements/mm² at 1 THz), deployable at Tx/Rx or as reconfigurable intelligent surfaces (RISs) to mitigate path loss and molecular absorption. Operating modes include ultra-massive beamforming with single-feed architectures or multi-band umMIMO exploiting graphene’s tunable response, which motivates new space–time–frequency coding and modulation schemes but faces challenges such as spatially non-stationary channel modelling, large-scale channel estimation and equalization, and cost; holographic MIMO provides related capabilities. Ultrawide bandwidths and large arrays invalidate far-field and narrowband assumptions, causing beam squint, grating lobes, and near-field effects, which require solutions such as true-time-delay beamforming and advanced low-cost algorithms. While advanced beamforming enables parallel links and sensing, conventional phased arrays with many RF chains are costly and energy-inefficient, and although RISs offer low-power alternatives, their linear and reciprocal nature limits multi-user and multifunctional operation. Consequently, research is shifting toward multi-hybrid beamforming, analog wave computing, and the co-design of radiation and processing using multifunctional antennas, metasurfaces, and wave-based processors. In particular, space–time-modulated metasurfaces extend beyond LTI systems to enable frequency conversion, magnetless non-reciprocity, multi-harmonic beamforming, multi-beam and full-duplex-like operation, and joint communication and sensing, though challenges remain in accurate modelling, phase quantization, power consumption, and hardware limits. In parallel, metasurface-based analog wave computing offers real-time, energy-efficient, and parallel linear processing for high-dimensional signal processing and AI, with future work needed to unify microwave analog networks, reconfigurable devices, resistive memory arrays, and nonlinear wave-domain operations to realize hybrid analog–digital communication and sensing architectures.

Distributed and cell-free massive MIMO: Distributed and cell-free mMIMO (CF-mMIMO) combines TDD-based mMIMO with dense, user-centric, distributed access point (AP) deployments, enabling all APs to cooperatively serve users without cell boundaries via fronthaul and backhaul connections to central processing units, yielding significant gains in spectral and energy efficiency, particularly for indoor and hotspot scenarios. However, major challenges remain, including initial access, power control, distributed processing, resource allocation, channel modelling, standardization, and prototyping, with fronthaul capacity emerging as a critical bottleneck since CF-mMIMO can generate signalling traffic far exceeding user payloads at large bandwidths (FR2, FR3, sub-THz). Tight synchronization and calibration across distributed APs/BSs, covering timing and carrier phase, are also crucial, as errors directly degrade EVM and increase multi-user interference. Looking toward 6G, FR3 (7.125–24.25 GHz) is viewed as a key band for high-capacity services, while sub-GHz spectrum (e.g., n71 at 600 MHz) remains essential for coverage, rural connectivity, and resilient IoT, benefiting from favorable propagation and infrastructure reuse. Yet, the large wavelengths at sub-GHz preclude co-localized massive arrays, motivating distributed MIMO architectures that instead face severe FDD channel state information (CSI) feedback overhead, as closed-loop feedback scales linearly with the number of distributed antennas; mitigating this "feedback explosion," especially in bandwidth-limited bands, requires advanced signal processing and CSI compression. Furthermore, extremely large and distributed arrays exhibit non-classical propagation effects, including spatial non-stationarity, near-field spherical wave behaviour, and strong spatial correlation and mutual coupling, motivating non-stationary subarray models, near-field channel models, and evolving cluster models in recent 3GPP extensions, though experimental validation in truly distributed settings is still scarce. Progress ultimately depends on synchronized, multi-frequency distributed measurement platforms to build realistic channel datasets, enabling accurate characterization of spatial consistency, cluster sharing across APs, and robust CSI estimation and channel charting for future CF-mMIMO systems.

Reconfigurable intelligent surfaces (RISs): RISs, also known as IRSs or LISs, are composed of many low-cost, tunable reflecting elements that independently control the amplitude and phase of impinging waves to shape propagation conditions. Being nearly passive, lightweight, conformable, and easy to deploy on infrastructure such as walls, ceilings, lamp posts, and vehicles, RISs act as energy-efficient auxiliary devices compatible with legacy user equipment, improving coverage, interference management, capacity, security, positioning, sensing, and wireless power transfer. Variants using many low-complexity antennas connected to a processing unit enable user-adaptive activation, yielding capacity and localization gains, while radio stripe architectures place antennas along linear structures to realize low-cost cell-free systems, though they require efficient signal separation and spatial resource allocation, potentially benefiting from exploitable hardware imperfections. Large RIS apertures and short link distances introduce near-field propagation, making channel estimation challenging due to the high dimensionality of the channel, which motivates parametric estimation and tracking methods aided by positioning information, as well as further study of large-scale fading and spatial correlation. RISs can resolve localization ambiguities in receive or reflect modes, restore coverage in NLoS and near-field conditions, and enable controllable multipath for SLAM and sensing applications such as fingerprinting, received-power maps, multipath processing, and device-free localization. Dynamic multi-RIS coordination can further enhance diversity for detecting and localizing passive objects, especially at mmWave and FR3, but requires new localization algorithms, control strategies, protocols, and experimental validation. Finally, recently proposed physics-compliant RIS channel models and emerging architectures such as beyond-diagonal RISs, stacked multilayer surfaces, wave-domain angle and delay estimation, and refractive smart relays open new research questions on extracting location-dependent information and extending coverage.

Near field communication: Towards 6G, near-field communications are becoming increasingly important due to the combined effects of higher carrier frequencies, which make even moderate-sized antennas electrically large, and distributed mMIMO, which results in physically large apertures. In conventional far-field operation, spherical wavefronts are approximated as planar, allowing beam steering only in the angular domain, whereas below the Fraunhofer distance systems operate in the radiative near-field, where spherical waves enable control in both angle and distance, opening new spatial degrees of freedom. Electrically large antennas further make line-of-sight MIMO feasible for fixed fronthaul and backhaul links, while at distances comparable to the wavelength the reactive near-field emerges, with partially decoupled electric and magnetic fields. Key research directions in near-field communications include the development of channel models that bridge electromagnetic antenna theory and classical communication models by accounting for mutual coupling, polarization, spherical propagation, and realistic environments, along with measurement-based validation for 6G deployment. Fundamental capacity limits derived from electromagnetic information theory provide performance bounds for practical systems, while significant implementation challenges remain in synchronization, RF chain design, antenna technology, and current-shaping materials. Beyond communications, near-field techniques are highly relevant for high-accuracy positioning and sensing, and future systems are expected to jointly integrate near-field MIMO with absorptive RISs and fluid antennas through low-complexity, co-designed architectures.

Advanced antenna system: Fluid antenna systems (FAS) add reconfigurability of the antenna itself to massive MIMO. The term includes liquid antenna arrays and transceivers with radiating elements (e.g., pixel and movable antennas). Liquid antenna arrays can reversibly change shape to modify resonant frequency, polarization, or radiation pattern. They offer solutions to blockages and challenging scenarios. Further work is needed on signal processing and adaptation algorithms consistent with expected behaviours. Channel estimation, feedback, and prediction are key to maintaining low overhead while exploiting reconfigurability. A further challenge is jointly optimizing selected ports and beamforming when FAS combines with MIMO.

5.4.4 Waveform, Multiple Access and Full-Duplex

Waveform: Cyclic prefix orthogonal frequency division multiplexing (CP-OFDM) is widely adopted in wireline and wireless standards, but its reliance on strict synchronization limits performance in scenarios such as machine-type communications, where loss of subcarrier orthogonality occurs, and it is particularly sensitive to phase noise, which becomes more severe at higher carrier frequencies. To address these limitations, alternative multicarrier waveforms such as filter bank multicarrier, generalized frequency division multiplexing or cyclic block filtered multitone, universal filtered multicarrier, and filtered OFDM have been proposed, differing in orthogonality, cyclic prefix usage, and subcarrier filtering to improve frequency localization, with extensive performance comparisons in terms of complexity, spectral efficiency, robustness to multi-user interference, and tolerance to power amplifier nonlinearity. For highly time-varying channels, orthogonal time frequency space modulation maps signals into the delay Doppler domain using long spreading sequences originally developed for radar, while affine frequency division multiplexing offers a chirp-based alternative suitable for high-mobility and integrated sensing and communications. These linear waveforms can be unified through time and frequency domain precoding frameworks, although they typically require advanced equalization to mitigate inter-symbol interference. In contrast, nonlinear waveforms trade spectral efficiency for improved energy efficiency, including constant envelope OFDM, which reduces peak-to-average power ratio via phase modulation with low-complexity receivers, and zero-crossing modulation, which encodes information in timing to support 1-bit ADC receivers and low-rate IoT communications. Given highly variable traffic demands in future energy-efficient networks, the Gearbox PHY concept proposes dynamically switching among waveforms and modulation schemes to adapt performance and energy consumption to instantaneous requirements.

Multiple access: Relaxing strict orthogonality enables more flexible and efficient use of wireless resources, motivating strong interest in non-orthogonal multiple access (NOMA), which can increase achievable uplink and downlink rates while also supporting grant-free access and collision resolution in machine-type communications. Key open challenges in NOMA include optimal user pairing when more than two users share the same resources, power control strategies that balance interference and throughput, including extensions of full-duplex designs to multi-cell scenarios, and physical layer security, as successive interference cancellation allows users to decode other users' signals. Code-domain NOMA schemes such as multiuser shared access, low-density spreading, and particularly sparse code multiple access are promising due to their compatibility with massive MIMO, mmWave systems, and security enhancements, but they require low-complexity code and receiver design. Closely related, massive random access targets scenarios with a very large number of sporadically active devices, where users randomly select non-orthogonal spreading codes, leading to pilot collisions, large timing offsets, and frequency offsets that necessitate advanced blind or data-driven detection exploiting statistical knowledge and CRC information. Alternative solutions focus on improved pilot designs to mitigate collisions, including multiple independent pilots and extremely sparse pilot schemes. Related contention-based approaches such as coded slotted Aloha combine sparse coding with iterative message passing decoding to suppress interference, with further gains achieved by integrating low-rate channel coding and multi-user detection at the physical layer. Activity detection in these settings is closely linked to compressed sensing, where approximate message passing algorithms and multi-measurement vector techniques, potentially combined with coded random access, can asymptotically drive user detection error to zero**.**

Full-duplex: Furthermore, advanced self-interference cancellation techniques can potentially double the spectral efficiency, and enable in-band full-duplex (IBFD) transceivers that offer a wide range of benefits, e.g., for relay, bidirectional communication, cooperative transmission in heterogeneous networks, joint communication and sensing, and cognitive radio applications. However, for the full-duplex technique to be successfully employed in next generation wireless systems, there exist challenges at all layers, ranging from antenna and circuit design (e.g. due to hardware imperfection and nonlinearity, non-ideal frequency response of the circuits, phase noise, etc, especially when taking MIMO and massive MIMO into account), to the development of theoretical foundations for wireless networks with IBFD terminals, and including AI-based algorithms that are capable to perform self-interference cancellation in multiple radio frequency bands. Much work remains to be done, and an inter-disciplinary approach will be essential to meet the numerous challenges ahead.

5.4.5 Coding and Modulation

Channel coding: Channel coding remains one of the most complex components of the baseband processing chain, and while Turbo, LDPC, and Polar codes have been widely adopted thanks to advances in semiconductor technology, significant implementation bottlenecks persist for ultra-high throughput and low-latency applications. For 6G, peak data rates up to 1000 Gbps place stringent demands on decoders, which legacy 5G NR LDPC designs struggle to meet because they are optimized for block-parallel architectures rather than higher-throughput row-parallel or fully parallel decoders. This motivates new LDPC code designs tailored to such architectures, as well as renewed interest in alternative schemes such as Polar codes with successive cancellation decoding. A key challenge is achieving ultra-high throughput with acceptable complexity, flexible rate and length support, and error-correction performance comparable to or better than 5G LDPC codes. For Polar codes, CRC-aided successive cancellation list decoding does not scale well due to its inherently serial nature, which has led to interest in iterative decoding approaches such as multi-trellis belief propagation, along with modified Polar code constructions better suited to iterative decoding. Latency and throughput can be further improved by unfolding iterative decoders using deep neural networks. Looking ahead, unified channel coding schemes are desirable to reduce hardware complexity by enabling Polar codes to perform well at long block lengths or LDPC codes at short ones without sacrificing performance. In parallel, the Guessing random additive noise decoding paradigm offers a code-agnostic, capacity-achieving alternative by decoding via noise guessing, though its practical applicability currently depends on reducing complexity at medium and low code rates, incl. extension to soft-input decoding and concatenated codes.

Modulation: Combining channel coding with higher-order modulation schemes such as QAM can lead to suboptimal performance due to shaping loss caused by uniformly distributed transmitted symbols, resulting in up to 1.53 dB degradation on AWGN channels. To mitigate this loss, constellation shaping techniques have been proposed, including geometric shaping through non-uniform constellations that optimize symbol locations, and probabilistic shaping, where shaping encoders enforce non-uniform symbol probabilities to better match channel capacity. While several shaping encoders support high throughput and low latency, efficient hardware implementations of shaping encoders and decoders remain an open challenge. At the same time, modulation design is evolving from static constellations toward channel-aware and data-driven approaches, as standard AWGN-optimized constellations often perform poorly under realistic propagation conditions and hardware impairments. Machine learning enables adaptive constellations tailored to channel statistics, and specialized formats such as Polar-QAM and circular constellations improve robustness against severe phase noise in high-frequency systems. Delayed bit-interleaved coded modulation extends conventional BICM by introducing controlled bit delays to trade latency against complexity and performance gains. Modulation and coding scheme selection becomes particularly challenging in wideband, frequency-selective channels where aggregate SNR is insufficient, motivating advanced link adaptation based on effective SNR metrics or AI-driven algorithms that learn mappings from high-dimensional channel features to optimal transport block sizes, even in the presence of nonlinear hardware effects and high mobility.

5.4.6 Integrated Sensing and Communication

Integrated sensing and communication (ISAC) with positioning is a key pillar of 6G, aiming to maximize spectral efficiency while minimizing hardware and energy resources by jointly exploiting RF signals for communication, sensing, and localization. The high bandwidth and dense connectivity of 6G enable cooperative data fusion across devices to accurately localize and track passive targets for applications such as traffic monitoring, autonomous vehicle safety, assisted living, industrial automation, IoT, V2X, smart cities, and medical care, with sensing capabilities ranging from object detection to micro Doppler extraction, material characterization, and high precision imaging. A central challenge is channel modelling, where extensions such as 3GPP TR 38.901 introduce stochastic sensing parameters like radar cross section and micro Doppler, though full physical consistency remains an issue, motivating ETSI segment based RCS models and emerging point cloud electromagnetic approaches that combine physical optics with reduced complexity to enable radio object imaging and semantic environment representation. Large scale antenna arrays, especially at short wavelengths, are fundamental enablers for ISAC, supporting spatial multiplexing, high resolution angle estimation, multipath separation, and improved positioning and simultaneous localization and mapping (SLAM) through environmental mapping. Network based sensing can offload SLAM complexity from UEs and operate in bistatic or monostatic modes, but fully exploiting temporal, spatial, and frequency dimensions requires new waveforms and ISAC specific reference signals. Cooperative sensing and positioning significantly enhance accuracy, robustness, and detection probability through spatial diversity and information exchange, although challenges remain in resource allocation, data formats, fusion, and control. Accurate sensing also enables sensing assisted communications and positioning, such as blockage detection and narrow beam selection, with data driven approaches becoming essential in complex propagation environments. Fine grained sensing based on micro Doppler, critical for activity recognition and target classification, is severely impacted by clock asynchrony across network nodes, causing timing and frequency offsets that conventional communication receivers remove as impairments; overcoming this limitation is one of the main obstacles to large scale ISAC deployment, with proposed solutions exploiting correlated static paths as phase references. Achieving sub decimeter positioning accuracy requires several gigahertz of phase coherent bandwidth, feasible through carrier aggregation of non-contiguous bands with fixed phase relations. Finally, future networks can act as large scale radio sensors that continuously map the environment during low traffic periods, enabling clutter suppression and high sensitivity detection, while raising new requirements for physical layer-based protections against unauthorized sensing, as conventional cryptographic approaches are not applicable.

5.4.7 Machine Learning Empowered Physical Layer

Application of artificial intelligence (AI) and machine learning (ML) in communication systems spans a wide range from optimizing a specific function to end-to-end learning of the entire communication system. ML techniques are expected to redefine the classical approach in communication system design to achieve global optimization or performance improvement. A block-structured strategy focuses on the application of a specific ML algorithm for each block of the systems to optimize the individual performance of each block. This methodology inherits the advantages of former models and algorithms. For example, channel estimation can be studied as a regression problem and the selection of modulation and coding scheme (MCS) can be learnt from an exhaustive exploration of the environment. Another perspective addresses end-to-end (E2E) communication and targets at the optimization of the complete communication system. An E2E ML approach may consist in the representation of both blocks as black boxes and the characterization of the transmitter as an autoencoder that can infer MCS through data-driven analysis. ML techniques applied to spectrum/environment awareness may assist in the adoption of both strategies given the scarcity of tools to face 6G challenges such as 3D transmission models, the adoption of new frequency bands (THz), the introduction of new network elements (e.g. RIS), beamforming in ultra-massive MIMO communications, and non-identified interference sources.

The application of ML techniques should be guided by the fact that they do not in general lead to the optimal performance. ML techniques are only advantageous if the considered signal processing problem is nonlinear by nature. Moreover, in case knowledge of the underlying model is available it is always advantageous to use this information and incorporate it into the applied machine learning model. Some of the topics where we expect AI/ML research to be important to improve the physical layer design and performance are listed in the next sub-section (see [C5-03] for more details).

5.4.8 Synopsis of Required Research

Research Theme	Radio Technologies and Signal Processing		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
RT1: Radio Networking & Interference Management	Mid–Long term (2026–2035).	Dynamic spectrum refarming and sharing across licensed / unlicensed / NTN bands; subnetworks for industrial, in-X and life-critical use; edge	Higher spectral and energy efficiency, improved coverage, new wireless-only use cases, better QoE in

	To be specified in 6G standard.	caching with coded multicast and AI-based popularity prediction. **[KPIs]:** 6G energy / spectral efficiency, data rate, area traffic capacity, reliability, connection density (Table 5-1).	dense and heterogeneous environments.
RT2: mmWave & THz Communication	Short–Long term (2028–2035).	Energy-efficient THz transceivers; programmable high-gain arrays and metasurfaces; accurate THz channel models; THz-tailored modulation, coding, beamforming, MAC and routing; experimental real-time platforms. **[KPIs]:** 6G data rate, area traffic capacity, connection density, sensing & positioning (Table 5-1).	Enables sub-Tbps links, short-range ultra-capacity access, and new sensing/ISAC capabilities at extreme frequencies.
RT3: Next-Generation MIMO (mMIMO, umMIMO, cell-free, RIS, advanced antennas)	Mid–Long term (2026–2035). To be specified in 6G standard.	Ultra-massive arrays; near-field-aware beamforming; MIMO DPD; multi-hybrid beamforming and wave-domain computing; distributed and cell-free MIMO with scalable CSI acquisition; RIS; fluid antennas and holographic MIMO. **[KPIs]:** 6G energy / spectral efficiency, data rate, area traffic capacity, reliability, connection density (Table 5-1).	Step-change in area spectral efficiency, energy efficiency, coverage, robustness, and integrated communication-sensing performance.
RT4: Waveforms, Multiple Access & Full-Duplex	Mid–Long term (2026–2035). To be specified in 6G standard.	Robust waveforms beyond CP-OFDM (FBMC, OTFS, AFDM); NOMA; massive random access for mMTC; full-duplex transceivers with advanced SIC; waveform–hardware co-design for energy efficiency. **[KPIs]:** 6G energy / spectral efficiency, mobility, latency, connection density, sensing & positioning (Table 5-1).	Supports high mobility, massive connectivity, low latency, and better spectrum reuse for 6G services.
RT5: Coding & Modulation	Mid-term (2026–2032). To be specified in 6G standard.	Ultra-high-throughput LDPC/Polar codes; unified coding frameworks; iterative and AI-assisted decoding; (geometric and probabilistic) constellation shaping; channel-aware and ML-driven modulation.	Achieves 6G KPIs (100–1000 Gbps, ultra-reliability, low latency) with manageable complexity and energy consumption.

		[KPIs]: 6G spectral efficiency, data rate, reliability, mobility, latency (Table 5-1).	
RT6: Integrated Sensing, Positioning & Communication (ISAC)	Short-Mid term (2026–2035). To be specified in 6G standard.	Joint waveform and reference signal design; physically grounded ISAC channel/target models; multi-band and cooperative sensing; sensing aided communication; communication aided sensing / localization; environment sensing; and clutter suppression. **[KPIs]:** 6G spectral efficiency, mobility, sensing & positioning (Table 5-1).	Enables perceptive networks, cm-level positioning, environment reconstruction, and new safety-critical and industrial applications.
RT7: Machine-Learning-Empowered PHY	Mid–Long term (2026–2035). May be specified in 6G standard.	AI-native PHY blocks; hardware-impairment modelling; CSI and model compression; in-radio network AI computing; radio digital twins and environment-aware learning; distributed source coding for semantic communication. **[KPIs]:** Data rate, area traffic capacity, mobility, connection density, etc (Figure 5-1).	Improves adaptability, reduces overhead and complexity, and enables scalable intelligence in 6G RANs.

5.4.9 Recommendations for Actions/Instruments

Research Theme	Radio Technologies and Signal Processing			
Action	**Radio Networking & Interference Management, mmWave & THz** (RT1, RT2)	**Radio Building Blocks (MIMO, Waveform, Coding & Modulation, etc** (RT3, RT4, RT5)	**Integrated Sensing, Positioning & Communication** (RT6)	**AI-Native Radio Interface** (RT7)
International Calls	Continued EU-international research on radio technologies to lead and contribute to 3GPP Rel. 20+ 6G radio interface standardization.			
Open Data				
Large Trials	Continued trials to prove radio technologies for 3GPP Rel. 20+ 6GR, and verify 6GR pre-commercial standard systems.			
Cross-domain research	Consider NTN, IoT use cases/requirements in radio technology research for 3GPP 6GR standard, to ensure 6GR supporting important NTN and IoT applications. Close collaboration with regulatory & standardisation (ETSI, 3GPP, ITU-R).			

5.5 Non-Terrestrial Networks

5.5.1 Waveform design

OFDM has been serving as the waveform foundation for 4G/5G systems. However, it poses relevant criticalities to NTN: i) in LEO/VLEO constellations, tight time/frequency synchronisation is challenging; ii) large and time-varying propagation delays, combined with large Doppler shifts, degrade the subcarrier orthogonality; iii) large PAPR values lead to power inefficiency and potential distortion in the on-board HPAs. Delay-Doppler (DD) domain waveforms have emerged as the leading research candidates, aiming at a better handling of high mobility and doubly dispersing channels. In addition, ISAC has emerged as a defining feature of 6G, calling for specific adaptations for joint radar-communication operation, [C5-07].

Despite the significant progress, no waveform candidate is capable of fully meeting the requirements for natively integrated TN-NTN in 6G, in terms of: i) *synchronisation and Doppler resilience*, as non-null residual timing/frequency offsets disrupt the sparsity structure of the DD channel and the performance over realistic NTN channels; ii) *receiver complexity and latency*, as iterative detection leads to processing delay and computational overhead that challenge resource-constrained NTN payloads and UEs; iii) *high PAPR*, critical in multicarrier waveforms over NTN where the power amplifier efficiency is pivotal and existing reduction techniques have not yet been adequately validated/tested in realistic scenarios; iv) *limitations in ISAC co-design*, as candidate waveforms exhibit an ambiguity function limiting the sensing accuracy and producing inter-user interference in multi-user scenarios; v) *immaturity of AI-native and neuromorphic approaches*, very promising solutions in the long-term but currently lacking standardisation frameworks and performance validation in more realistic conditions, [C5-08]. A new generation of waveform R&I activities is thus required, as detailed below.

DD waveform optimisation and hybrid architectures. Research activities should advance the theoretical characterisation and practical implementation of OTFS, AFDM, and other waveforms under realistic VLEO/LEO channel models, including non-ideal synchronisation, HW impairments, and multi-beam scenarios. A particularly promising near-term direction is the development of hybrid architectures, *e.g.*, using CP-OFDM for high-throughput data transmission while reserving alternative modulations/access techniques for synchronisation pilots, as a pragmatic path to leveraging the strengths of multiple waveform families.

PAPR reduction. Novel/better performing PAPR reduction techniques, specifically tailored to multi-beam, multi-user NTN payloads, must be developed and validated. These must be jointly evaluated against spectral efficiency, power efficiency, and computational cost at the payload.

ISAC co-design. R&I shall address joint waveform optimisation for communication and sensing, targeting sub-metre range resolution, accurate velocity estimation, and reduced sidelobe levels, without (excessively) sacrificing the communication spectral efficiency. This includes pulse shaping optimisation, ambiguity function engineering, and multi-user ISAC resource management, for monostatic and bistatic/distributed NTN sensing configurations.

AI-assisted/-native PHY design. Medium-term R&I shall explore the use of AI-assisted receivers to handle the complexity of DD equalisation and to dynamically adapt the waveform parameters to the channel conditions. Over the longer term, AI-native waveform paradigms (*e.g.*, neuromorphic communications based on spike-based encoding and impulse radio / ultra-wideband (IR/UWB) signalling) might have matured for multi-user NTN scenarios, with clearer evaluations under realistic channel impairments and theoretical benchmarks.

Cross-domain standardisation and compatibility. Research activities shall systematically address the compatibility of new NTN waveforms with 5G NR and evolving 6G terrestrial standards. This requires close alignment with 3GPP and other standards bodies, and dedicated work on coexistence mechanisms, waveform adaptability across heterogeneous network segments, and the definition of NTN-specific KPIs within standardisation frameworks.

5.5.2 Multi-antenna and federated solutions

Multi-antenna systems can provide high-gain, steerable beams over wide coverage areas, enabling spatial reuse and interference management. Going beyond single-platform solutions, two cooperative architectures are emerging,: i) *opportunistic multi-satellite coordination*, with multiple LEO satellites in visibility forming a Distributed Antenna System in space, [C5-09]; ii) *satellite swarms*, with clusters of small satellites collectively synthesising large apertures and high-resolution beams at a fraction of the cost and mass of a monolithic large-array platforms, [C5-10]. AI is also a key enabling technology for NTN spatial processing solutions.

State-of-the-Art is falling short when aiming at seamless 6G NTN integration, with several critical gaps: i) *distributed beamforming and synchronisation*, as even small residual synchronisation and positioning errors degrade the beamforming directivity increasing the sidelobes levels; ii) *beam management for common signalling and heterogeneous services*, as power limitations on-board prevent the illumination of all beams simultaneously and there is no beam management framework jointly optimising signalling, user data transmission, and sensing/positioning, [C5-11] ; iii) *CSI management and distributed MIMO*, as radiating-element-level CSI knowledge poses acquisition and feedback overhead challenges in addition to lacking of a framing structure; iv) *handover and service continuity*, as current procedures within 3GPP are not yet optimised for federated/cooperative multi-satellite scenarios. In this framework, the following R&I areas shall be prioritised.

Distributed Beamforming for NTN swarms and sparse constellations. R&I activities shall focus on the design of beamforming techniques inherently robust to synchronisation and positioning imperfections, for both swarms and sparse cooperative constellations. In the former, joint optimisation problems over satellite position and precoding vectors under realistic error models (problems expected to be NP-hard) shall be formulated. For the latter, asynchronous transmission strategies and advanced receive algorithms extracting spatial gains despite differential delay and Doppler offsets are essential.

Holistic Beam Management for 6G NTN. A unified beam management framework is needed that simultaneously addresses data, common signalling, and sensing and positioning beams, across single-platform and federated multi-satellite deployments. R&I shall investigate asymmetric UL/DL beam layout schemes, efficient polarisation exploitation, and federated approaches where overlapping fields of view from multiple satellites can be jointly managed to overcome EIRP limitations. Minimising common signalling broadcast time, while preserving user-centric data transmission efficiency, should be a central optimisation objective as well.

Scheduling, CSI, and Standardisation for Distributed MIMO. Distributed MIMO scheduling and procedures shall be designed for efficient acquisition and delivery of ancillary information (*i.e.*, CSI, location) under overhead constraints imposed by NTN links. This includes a critical review and potential modification of current framing structures and CSI reporting procedures. Scheduling solutions should be evaluated under realistic, non-uniform traffic and population distributions, exploiting dissimilarity across users in both channel and geographical domains.

AI-Native Multi-Antenna and Federated Processing. AI should be systematically investigated as an enabler across the full multi-antenna stack for interference detection and mitigation, predictive CSI acquisition to reduce the feedback overhead, beam management, flexible payload reconfiguration, and distributed MIMO. These approaches should be assessed in both standalone and hybrid configurations alongside conventional algorithms. The scalability of AI-native approaches to large constellation sizes and dense user deployments requires particular attention.

ISL Design, Swarm Topology, and Ranging. Cooperative multi-satellite architectures depend critically on the ISL performance for coordination signalling, distributed precoding, and ranging-based positioning. Research advancements are needed on optimal ISL waveforms and protocols for swarm coordination, the relationship between swarm topology and beamforming performance, and ranging techniques capable of delivering the sub-centimetre positioning accuracy required for coherent distributed aperture synthesis.

5.5.3 Next-generation Multiple Access and Resource Management

Multiple access schemes for NTN have historically been built on orthogonal paradigms to separate users. While tractable, these approaches are limited in their ability to scale to massive connectivity and high interference environments in large constellations. Several solutions have been proposed to improve the spectral efficiency via non-orthogonal access schemes or the use of multiple dimensions beyond time, frequency, space, power. On the UL, grant-free Random Access (RA) is emerging for mMTC and IoT-NTN scenarios. As for resource management, AI/ML are being increasingly applied for proactive, data-driven resource orchestration across power, spectrum, and computational dimensions, with cross-layer design tightening PHY adaptations and higher-layer procedures. Current challenges include: i) *scalability and robustness*, as advanced solutions (*e.g.*, RSMA, CD-NOMA) outperform orthogonal solutions in some scenarios, but not in the full range of NTN use cases; ii) *grant-free RA for unsynchronised devices*, as IoT-NTN terminals might operate without reliable synchronisation due to device constraints, large Doppler offsets, and variable propagation delays; iii) *dynamic resource management in highly variable topologies*, with conventional static/semi-static allocation schemes being inadequate for the highly dynamic NTN ecosystem, calling for agile, predictive, and context-aware solutions; iv) *interference and spectrum coordination*, since the heterogeneity of NTN deployments makes interference modelling and spectrum coordination more complex than in TNs; v) *energy efficiency and latency with constrained payloads*, as resource allocation strategies that jointly optimise throughput, energy consumption, and latency have not been adequately developed yet for NTN-specific constraints; vi) *multi-layer and integrated architectures*, as integrating terrestrial and air-/space-borne elements needs advanced and proactive orchestration frameworks operating across the multiple different layers, including protocols and seamless handover.

Advanced MA Design and unified frameworks. Research activities shall enhance the design of RSMA and multi-dimensional MA schemes for specific NTN channels, addressing precoder complexity, robustness to imperfect CSI, and sensitivity to synchronisation and RF imperfections under LEO mobility. A central aspect is the architectural choice between a single generalised MA framework and a service-differentiated approach. Both should be evaluated with NTN-specific KPIs, including UL RA performance with asymmetric delay and Doppler.

Grant-Free RA for asynchronous IoT-NTN. Dedicated studies are needed to extend grant-free RA to fully asynchronous NTN device populations. This calls for new receiver architectures for joint activity detection, channel estimation, and data decoding without prior synchronisation or user scheduling knowledge. The identified techniques shall be assessed over realistic NTN channel models, including very dense IoT user populations.

Predictive and AI-Driven resource orchestration. AI and ML should be systematically applied to NTN resource management, enabling proactive and adaptive allocation of power, spectrum, bandwidth, and computational resources based on historical and real-time network states. R&I priorities shall focus on interference prediction and mitigation, TN-NTN DSS, energy-aware scheduling for power-constrained payloads, and latency-optimised routing across heterogeneous NTN segments. Cross-layer design approaches coupling PHY adaptations with higher-layer scheduling and routing decisions should be pursued.

Integrated communication, computation, and storage resource management. As on-board processing and edge computing capabilities grow in NTN platforms, resource management must go beyond communication resources to jointly optimise computation and storage. R&I shall address offloading strategies, distributed in-network processing, and load balancing under the payload latency and energy constraints. Integrated resource orchestration frameworks for joint communication-computation-storage optimisation problem are needed.

Multi-Layer NTN resource coordination and standardisation. Unified resource management frameworks for integrated satellite, HAPS, UAV, and terrestrial segments must be developed and validated at system level. This includes inter-layer coordination protocols, handover optimisation across heterogeneous segments, and joint end-to-end QoS management. Engagement with 3GPP, ETSI ISG MAT, and ITU standardisation processes is essential to ensure that research outcomes translate into standard-compatible procedures, and to feed NTN-specific requirements into the definition of the future 6G RAN architecture.

5.5.4 Spectrum sharing and coexistence

The emergence of dense NTN constellations operating across multiple orbital regimes has further highlighted the need for clear and robust coexistence solutions with TNs in frequency bands that were historically managed via rigid separation between TN and NTN. This approach is no longer acceptable: in 6G the two domains converge into a unified network and technical, regulatory, and standardisation challenges of spectrum sharing/coexistence will be defining constraints for the entire ecosystem. The separation between IMT and MSS allocations for NTN simplifies interference management, but it limits the NTN available spectrum, particularly for emerging D2D services seeking IMT bands to deliver connectivity directly to standard mobile handsets. Dynamic Spectrum Sharing (DSS) architectures and cognitive technologies have emerged as primary options to such challenges. At the same time, optical ISL and feeder links (OISL/OFL) are gaining momentum to alleviate congestion in RF feeder link bands. In standardisation, 3GPP has identified some MSS bands for NTN use. However, handset adoption for these bands remains limited to a small number of high-end devices, constraining the commercial reach of MSS-based D2D. The broader international regulatory framework for D2D in IMT bands remains unresolved, with WRC-27 that will address new MSS allocations, the harmonisation of D2D conditions in IMT bands, and the extension of MSS into bands currently identified for terrestrial mobile.

In this framework, the following critical issues shall be addressed: i) *obsolete regulatory models and TN-NTN separation*, as regulatory processes were developed under strict TN-NTN separation and SNO constellations are not yet planned for interoperability, leading to potential spectrum fragmentation; ii) *risk of spectrum monopoly*, where a single operator could secure *de facto* exclusive use of contested bands through aggressive filing strategies; iii) *technical coexistence parameters for new/shared bands*, as, without novel well-defined coexistence-oriented KPIs, neither the feasibility of coexistence tools nor the conditions under which it can be guaranteed can be assessed or standardised; iv) *limitations in current DSS*, as multi-radio sharing frameworks do not adequately address NTN and RAT heterogeneity and dynamicity, while AI-assisted solutions are still in their infancy; v) *immature optical and new-band technologies*, as OISL/OFL are still facing challenges for NTN deployment and new bands are subject to regulatory requirements and constrained SWAP whose implications over system design are not yet fully characterised. The following research areas shall be addressed, also targeting a close alignment with the regulatory timeline for WRC-27/-31.

Coexistence Modelling and Technical Parameter Definition. Research is needed to define interference models and techniques for NTN-NTN/TN-NTN coexistence in candidate 6G frequency bands. This work shall feed into 3GPP, ITU-R, and CEPT/ECC standardisation and regulatory processes, and should be structured to provide timely inputs to WRC-27/-31.

Particular attention should be given to NTN-NTN coexistence scenarios involving multiple LEO constellations sharing the same or adjacent bands and tackle the risk of spectrum monopolies.

AI-Native DSS for heterogeneous NTN. R&I shall focus on advanced DSS architectures that go beyond MRSS, leveraging AI to manage highly dynamic spectrum environments in NGSO constellations, heterogeneous channels, and multi-RAT coexistence. This includes AI-driven interference prediction, autonomous TX parameter adaptation, and joint optimisation of terrestrial beamforming, power allocation, and multiple access design for simultaneous communication, sensing, and positioning. Real-time spectrum management interfaces and joint TN/NTN-NTN resource management procedures should be defined at system level.

Regulatory framework innovation. R&I programmes should support research that directly informs and shapes the international regulatory process. This includes developing new methodologies for assessing coexistence in integrated TN-NTN systems, identifying the conditions under which the traditional TN-NTN spectrum separation model can be relaxed without creating risks of harmful interference or spectrum monopolisation, and designing interfaces and procedures to support seamless joint resource management.

D2D technical enablers. Studies shall address the technical enablers for D2D in IMT and MSS spectrum, including interference mitigation techniques for D2D covering both adjacent and co-channel operations with TNs, cross-border interference management between adjacent geographic areas, and handset HW and SW adaptations required to support MSS.

Optical and high-frequency band technologies. Research should progress on characterising optical and Q, V, E/W, D bands and HW constraints for NTN, providing the technical basis for new regulatory allocations covering terahertz frequencies, ESIM, ISL.

5.5.5 Positioning, Navigation, and Timing

Positioning, Navigation, and Timing (PNT) is a feature underpinning NTN operations at every layer of the protocol stack. Currently, GNSS has two roles: i) provide the UE its position to support UL time and Doppler pre-compensation; ii) provide the timing reference for network synchronisation. Such dependence on GNSS, while pragmatic for 5G TN-NTN integration, introduces a structural vulnerability that is incompatible with 6G resilience and security targets. The convergence of communication and sensing creates an opportunity for NTN systems to deliver PNT services close to those of GNSS constellations, without requiring separate infrastructure. The NTN PNT framework defined in 3GPP Rel. 17-19 is based on two complementary pillars: GNSS-derived UE positioning and satellite ephemeris information broadcasting. Combining them at the UE, initial uplink synchronisation and higher-layer procedures are enabled. 3GPP Rel. 20 marks a significant evolution by introducing GNSS-resilient operation, *i.e.*, the ability to maintain efficient network operation even when GNSS signals are temporarily unavailable, degraded, or subject to interference. This requires rethinking UL synchronisation procedures and Random Access (RA). Currently investigated solutions are: i) those in which the UE estimates its own position prior to RA preamble transmission via Positioning Reference Signals (PRS) or DL synchronisation signals, preserving architectural continuity; ii) radically diverse designs involving enhanced RA preambles capable of tolerating the large propagation delays and Doppler shifts of NTN channels, [C5-12] , combined with Timing Advance (TA) mechanisms. Looking further ahead, 6G NTN systems might incorporate dedicated positioning capabilities (*e.g.*, using PRS and associated navigation messages covering multi-satellite geometry, TDoA, and inter-satellite timing information) to deliver a service comparable to standalone GNSS, [C5-13] ; in addition, also opportunistic approaches can be evaluated, aiming at extracting positioning from always-on signals already present in the system for communication purposes (such as SSBs). Enhancing SSB content and structure to support positioning applications could enable GNSS-free operation without the overhead and coordination burden of a full positioning architecture.

Based on the above, the major technical challenges for PNT are: i) *GNSS dependency and structural vulnerability*, as GNSS is susceptible to various attacks well within the reach of a broad range of adversaries, making GNSS availability not guaranteed in all deployments; ii) *NTN channel impairments*, since conventional TN PRS designs degrade significantly when applied to NTN channels due to the time and frequency variability introduced by LEO satellite mobility; iii) *PRS resource grid design*, as in multi-satellite environments the PRS grid shall aggregate the different transmissions also accounting for the differential delay structure across the varying visible satellite geometry; iv) *bandwidth efficiency and communication-positioning trade-off*, since dedicating additional time-frequency resources to PRS in bandwidth-constrained NTN systems necessarily reduces the capacity available for data transmission; v) *procedures for non-GNSS operation*, as currently standardised solution all consider GNSS capabilities as mandatory at the UE. The following priority research areas can be identified.

GNSS-Free procedures. Research activities shall design and assess advanced GNSS-free synchronisation approaches, considering those that have been currently identified for Rel. 20 as a baseline: UE self-positioning via DL reference signals prior to RA and enhanced RA designs with NTN-capable preambles and extended TA mechanisms. Such solutions require new preambles to cope with large delay-Doppler impairments, TA signalling supporting positive and negative offsets, and receiver algorithms capable of achieving reliable synchronisation at the low SNRs encountered in NTN downlinks. In FR2, beam management and/or angle of arrival estimation may play a significant role in UE self-positioning.

Doppler-resilient reference signal design for NTN PNT. Dedicated research is needed to design reference sequences for both PRS and synchronisation signals with cross-correlation properties sufficiently robust to the delay and Doppler uncertainties over NTN channels. This might include Doppler-resilient pseudo-random sequence families suitable for continuous satellite tracking, blind carrier detection algorithms in low-SNR regime, and techniques for superimposing positioning reference signals over data channels using novel multiple access principles that preserve the integrity of both functions without significant bandwidth penalty.

PRS grid design and multi-satellite navigation. Fused positioning paradigms require a comprehensive research programme spanning PRS resource grid design for multi-satellite aggregation under differential delay conditions, navigation message design for NTN geometries, TDoA observable computation and error characterisation, and inter-satellite timing and synchronisation protocols, [C5-14]. This work must assess whether terrestrial PRS architectures can be adapted or fundamentally new designs are required for NTN. Moreover, R&I activities shall define the synchronisation accuracy requirements among the cooperating satellites, with direct implications for ISL and on-board clock design.

Security and resilience of NTN PNT. Over the longer term, R&I endeavours should address the security architecture of NTN-native PNT, developing authentication mechanisms for PRSs, detection and mitigation techniques for jamming and spoofing attacks over NTN channels, and PNT service continuity frameworks that can seamlessly transition between GNSS-aided, GNSS-resilient, and fully GNSS-free operation modes as signal availability changes, [C5-15].

Interference-aware anti-jamming / anti-spoofing techniques. Beyond GNSS-free procedures and robust reference signal design, NTN PNT requires PHY mechanisms that can detect, classify, and mitigate jamming and spoofing events in real time. Research should address receiver and network solutions combining multi-antenna measurements and cooperative UE reporting to identify threatened areas and trigger integrity alerts within application-specific time-to-alert constraints. AI/ML should be investigated as a key enabler for interference classification, fusion of distributed observations, and adaptive mitigation under dynamic jammer behaviour, including adaptive interference suppression, spoofing-aware signal authentication, and network-assisted threat dissemination, [C5-16]. Such techniques are essential because moving jammers create highly dynamic affected regions that static protection approaches cannot address.

5.5.6 Inter-Satellite and Feeder links

ISLs and feeder links are at the heart of NTN architectures, providing high-capacity, reliable, and dynamically managed links among satellites and to/from the ground for mega-constellation networking, distributed space-based service provisioning, and seamless integration with TN. The PHY design of these links involves some of the most demanding engineering trade-offs in NTN: extreme relative velocities, rapidly changing link geometries, tight SWaP constraints for small platforms, and the need to simultaneously maximise data rate, minimise latency, and maintain continuous connectivity. RF ISLs provide robust connectivity with fast link establishment and multi-beam flexibility: phased-array and beamforming techniques allow a single terminal to maintain multiple simultaneous links, well-suited to the complex routing topologies of multi-layer mega-constellations. Q/V band front-ends (37–75 GHz) are advancing toward space qualification, approaching optical bandwidths while retaining RF robustness. Sub-terahertz communications (0.1–1 THz) have also recently emerged as a promising bridge between RF and optical ISLs, offering ultra-wide bandwidths while retaining some RF multi-beam and robustness characteristics. FSO technology remains the primary candidate for very high-capacity ISLs and feeder links; given the breadth of the optical dimension, FSO-specific aspects, including OISLs, OFL, and atmospheric propagation effects, are treated separately in Section 5.7.1. The emerging consensus is that future NTN architectures will deploy hybrid RF/FSO ISL solutions exploiting their complementary strengths.

The major challenges in this context are: i) RF front-end challenges, as extending to Q/V bands and beyond requires components simultaneously meeting power, bandwidth, linearity, and thermal constraints within small-satellite SWaP envelopes; ii) sub-THz synchronisation, as links in this portion of the spectrum face a compounded impairment environment (oscillator phase noise at THz carriers, high inter-satellite Doppler, ultra-wideband timing demands) that conventional pilot-based schemes cannot properly handle [C5-17]; iii) link acquisition and handover, as dense LEO/VLEO constellations require frequent, low-latency handovers that legacy procedures designed for GEO geometries cannot support; iv) hybrid architecture and cross-layer optimisation, as future constellations operating heterogeneous RF/FSO/sub-THz links simultaneously lack dynamic link selection, cross-layer routing optimisation, and standardised fallback and link-quality-reporting protocols for hybrid deployments. In this context, the following research areas shall be addressed for future ISLs and feeder links.

RF beamforming and antenna architectures. Research activities shall advance Q/V band phased-array architectures for fast electronic steering, simultaneous multi-link operation, and SWaP-optimised small satellite implementations, addressing the above trade-offs at mm-wave frequencies alongside calibration techniques for orbital thermal and vibrational environments.

Link acquisition, tracking, and handover. Fast acquisition and handover procedures are needed for RF ISLs in high-dynamics NTN environments, combining beamforming-assisted algorithms and orbital-mechanics-based predictive handover. Fast-PAT procedures for OISLs are discussed in Section 1.7.1.

Sub-THz PHY. Low-complexity, pilot-free synchronisation solutions should be developed and validated under realistic LEO/VLEO sub-THz channel models. Moreover, E2E system-level analyses should define the performance of sub-THz ISLs within hybrid RF/FSO architectures.

Hybrid networking and cross-layer optimisation. Full protocol and networking frameworks for hybrid RF/FSO/sub-THz ISL operation are needed, encompassing dynamic link selection, cross-layer routing optimisation, and standardised interoperability protocols. Intelligent routing strategies for multi-layer LEO/MEO/GEO topologies should be validated against realistic constellation and traffic models, with latency and energy efficiency as primary objectives.

Feeder link capacity and resilience. Q/V band feeder links should be investigated to relieve Ka-band congestion, addressing atmospheric impairment mitigation, hybrid RF/optical ground station architectures, and WRC-31 frequency coordination. Long-term research should integrate feeder link and ISL design within a unified ground-space network architecture to prevent mutual capacity bottlenecks. Optical aspects are detailed in Section 1.7.1.

5.5.7 Synopsis of Required Research

Research Theme	Non-Terrestrial Networks		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
RT1: Waveform design for 6G NTN	Short–Mid term (2026–2032)	1) DD-domain waveforms validation in realistic channels; 2) Multi-beam NTN-tailored PAPR reduction techniques; 3) Joint comm-sensing waveforms with engineered ambiguity function; 4) AI-assisted/-native waveform designs. **[KPIs]:** Peak DL/UL at tens of Mbps for handheld, hundreds of Mbps vehicle/drone, thousands of Mbps aero/maritime; Reliability > 99.999%; GNSS-free waveform operation: PRACH detection >99%, false-alarm <0.01%	Enables robust, power-efficient 6G NTN air interface; underpins TN–NTN integration and ISAC; secures European leadership in 6G waveform design.
RT2: Multi-antenna and federated solutions	Short–Mid term (2026–2033)	1) Federated beamforming and scheduling robust to sync/positioning errors; 2) Holistic beam management (data, signalling, sensing, positioning); 3) AI-native multi-antenna processing at constellation scale; 4) ISL-based ranging for sub-cm coherent swarm operation. **[KPIs]:** Distributed beamforming +50–100% capacity gain and −50% outage probability −50% ; Scheduling energy efficiency +30%; ISL-based ranging for sub-cm coherent swarm; coherent sparse-array sync via ISL	Unlocks LEO multi-satellite and swarm capacity gains; reduces launch mass via distributed apertures; foundation of 6G NTN multi-user performance.
RT3: Next-generation Multiple Access and Resource Management	Short–Mid term (2026–2032)	1) RSMA, multi-dimensional MA, and GNSS-free RA optimised for NTN channels and massive IoT-NTN; 2) Predictive, AI-driven joint spectrum, power, compute orchestration; 3) Integrated comm-	Scalable access for mega-constellations and massive IoT; efficient use of scarce on-board resources; enables multi-layer TN-NTN operation.

		computation-storage resource management. **[KPIs]:** Connection density >1000 per km^2 (vs ≤500 in 5G); Over-the-air latency in line with IMT-2030 TN excl. propagation (vs 5G CP 40 ms / UP 10 ms)	
RT4: Spectrum sharing and coexistence	Short–Long term (2026–2035)	1) Coexistence models and KPIs for WRC-27/-31; 2) AI-native spectrum sharing; 3) Regulatory framework for TN-NTN coexistence; 4) D2D technical enablers in IMT and MSS bands; 5) optical bands characterisation and standardisation. **[KPIs]:** Average & 5th-percentile victim-network throughput loss <5% (3GPP DSS methodology); AI-based interference detection/classification for TN/NTN coexistence	Prevents spectrum monopolies; enables direct-to-handset services; secures European regulatory influence and access to new bands for 6G NTN.
RT5: PNT and GNSS-free solutions	Short–Long term (2026–2035)	1) GNSS-free UE synchronisation and RA; 2) Doppler-resilient reference signal and PRS design; 3) Multi-satellite fused PNT; 4) Authentication and resilience architectures for NTN PNT; 5) AI-driven interference classification and spoofing mitigation. **[KPIs]:** GNSS-free outdoor positioning: ≤50 m H/V @≤1500 km/h; ≤1 m H/V @≤160 km/h; Positioning availability ≥99%, independent of GNSS; Positioning accuracy 1 m, latency below 1 s (vs below 100 s in 5G); Multi-satellite fused PNT (up to 4 sats tracked)	Removes the structural GNSS dependency of 5G NTN; provides sovereign European PNT; enables resilient and secure location services for 6G.
RT6: Inter-Satellite and Feeder Links	Mid–Long term (2028–2035)	1) Q/V-band phased-array multi-link ISLs for small satellites; 2) Fast acquisition, tracking, and handover for RF/FSO; 3) Sub-THz PHY with pilot-free sync validated in orbit; 4) Hybrid RF/FSO/sub-THz networking and cross-layer routing. **[KPIs]:** Terabit-class distributed space networks: approx. 1 Tbps per GEO/LEO satellite; Feeder-link switch-over <100 ms; routing convergence <200 ms; signalling overhead ≤1%; load balancing ≥75%	Breaks the feeder-link bottleneck of mega-constellations; enables terabit/s distributed space networks; reduces ground segment footprint.

5.5.8 Recommendations for Actions

Research Theme	Non-Terrestrial Networks		
Action	**Waveforms, multi-antenna, and MA/RRM** (RT1, RT2, RT3)	**Spectrum sharing, PNT, and regulation** (RT4, RT5)	**ISL and feeder links** (RT6)
International Calls	Joint EU-international research on 6G NTN Air Interface aligned with 3GPP Rel. 20+ timelines.	Coordinated EU-international studies feeding WRC-27/-31 on TN-NTN/NTN-NTN coexistence, D2D, GNSS-independent PNT.	International collaboration on sub-THz and optical ISL/feeder link standardisation, including Q/V-band coordination.
Open Data	Reference NTN channel models and traffic datasets for Air Interface validation and shared AI training benchmarks.	Open TN-NTN interference measurement campaigns and PNT impairment datasets to feed AI-native DSS and anti-jamming research.	Open datasets for RF/FSO/sub-THz ISL propagation, PAT dynamics, and feeder-link atmospheric characterisation.
Large Trials	In-orbit and hardware-in-the-loop testbeds for 6G NTN waveforms, multi-satellite distributed MIMO, and RSMA/grant-free RA at realistic scale.	Large-scale TN-NTN DSS pilots, D2D field trials, and GNSS-free operation demonstrators feeding 3GPP and ITU-R validation.	In-orbit trials of Q/V-band phased-array ISLs, hybrid RF/FSO links, and optical feeder systems integrated with European ground segment.
Cross-domain research	Joint PHY / RRM / AI research bridging NTN and terrestrial 6G; PHY–MAC–application co-design for ISAC and task-oriented services.	Collaboration with regulatory& standardisation (ETSI, 3GPP, ITU-R) on unified TN-NTN spectrum and PNT frameworks.	Cross-domain work with optical, photonics, and space systems communities on hybrid networking & sustainable architectures.

5.6 IoT Networks

5.6.1 A Coverage-Regime Partitioned Analysis

Enabling massive and mission-critical IoT requires addressing spectrum coexistence and interoperability across heterogeneous coverage regimes, each with distinct constraints in terms of scalability, reliability, and operational criticality.

White and underserved areas (NTN-driven coverage). Non-terrestrial networks (NTNs) will be essential to extend connectivity to rural, remote, and disaster-stricken regions where terrestrial infrastructure is unavailable or degraded. A key challenge is the development of unified spectrum management and interoperability frameworks enabling seamless handover between terrestrial networks (TN) and satellite/NTN systems, which currently operate under fragmented regulatory regimes. Ongoing roadmaps aim at harmonizing spectrum allocation and sharing to support operational readiness around 2030, enabling service continuity in mobility scenarios (aerial, maritime) and during terrestrial outages.

Urban and suburban dense areas (macro/micro coverage). Smart cities will face extreme densification of IoT devices, vehicles, and user equipment competing for limited spectrum. Dynamic spectrum allocation and interference management are critical, especially in unlicensed ISM and shared licensed bands. Recent advances in AI-based approaches, including multi-agent reinforcement learning, demonstrate strong potential for real-time resource optimization under imperfect channel state information. These techniques are essential to mitigate cross-technology interference (e.g., Wi-Fi 6/7, 5G NR, and future 6G systems).

Industrial and short-range subnetworks (in-factory, in-building, in-X environments). Industrial and in-X subnetworks support tightly coupled cyber-physical systems where communication directly supports closed-loop control and real-time actuation. They impose stringent requirements on ultra-low latency and deterministic reliability, often at sub-millisecond scales. Alongside control traffic, these environments require dense data acquisition and sensing infrastructures for monitoring and situational awareness. This creates a dual operating regime: deterministic, low-latency links for control, and highly scalable connectivity for distributed sensing. A key challenge is the coexistence and joint orchestration of deterministic control communications with probabilistic, energy-constrained sensing traffic. This requires latency-aware resource allocation and adaptive control of the radio environment. In particular, interference introduced by backscatter and Ambient IoT devices must be carefully managed within the same shared spectrum and physical space.

Networks must ensure interoperability and dynamic spectrum sharing across heterogeneous radio access technologies (RATs), including legacy 5G/6G, Wi-Fi, and emerging radio interfaces. This requires coexistence protocols, AI-driven spectrum management, and adaptive waveform designs operating across licensed, unlicensed, and backscatter bands. ITU-R and 3GPP are actively developing frameworks for multi-RAT spectrum sharing (MRSS) and cross-technology handover toward the mid-2030s.

5.6.2 Device innovations

Achieving ultra-low power consumption is a key requirement for future IoT devices, enabling massive deployments of low-cost nodes with multi-year lifetime and near-zero maintenance. This calls for a paradigm shift in both hardware and communication design: devices must remain asleep most of the time, and connectivity must evolve toward intent-driven and goal-oriented communication, where transmissions occur only when information is relevant. New hardware, harvesting, protocols, and security are needed for decade-long IoT operation across evolving networks.

Embedding lightweight on-device intelligence (TinyML) is a key enabler to filter irrelevant data locally and reduce unnecessary transmissions. In parallel, wake-up radio solutions can drastically reduce idle listening and synchronization costs. Emerging analog and neuromorphic transceiver architectures also offer promising pathways to reduce digital processing overhead and energy consumption.

A major disruptive direction is the development of Zero-Energy Devices (ZED) based on backscatter communications, now considered in wireless standards (e.g., bi-static architectures). Such devices can communicate by reflecting impinging waves while harvesting energy from RF, solar, thermal, or motion sources. Spectrum-agnostic backscatter approaches, capable of operating over very wide bandwidths, provide a route toward long-lived IoT systems resilient to long-term spectrum and technology evolution.

Finally, biodegradable IoT devices represent a critical response to the environmental footprint of large-scale sensing. Using biodegradable substrates, sensors, and antennas (cellulose, silk, bio-polymers), future deployments could enable devices that naturally decompose after a controlled lifetime. This introduces new research challenges where communication performance, energy availability, and network protocols must adapt to nodes with finite lifetime and uncertain degradation dynamics.

5.6.3 Massive access, sporadic communications

Intent-driven and goal-oriented communication paradigms will reshape the design of massive access schemes for ultra-low-power IoT. The core challenge in ultra-dense networks is to dramatically reduce signalling and control overhead while ensuring reliable reception of sporadic, short packets from a vast number of devices. This drives innovation in PHY/MAClayers, including non-orthogonal waveforms, compressive sensing-based signalling, tensor-based multi-user receivers, and short-packet coding, alongside grant-free access to jointly optimize latency, reliability, and energy efficiency.

Over-the-Air Computation (OTA-C) further reduces network overhead by leveraging signal superposition to compute aggregated functions (e.g., sums, averages) directly over the wireless channel, which is critical for large-scale sensing and federated learning. However, robust solutions are needed to address fading, synchronization errors, imperfect CSI, and device heterogeneity, potentially using MIMO and RIS-assisted propagation.

In the downlink, scalable mechanisms are essential to reach low-power devices for control commands, security patches, and firmware updates. This requires energy-efficient wake-up/paging, broadcast/multicast dissemination, coded delivery for intermittent connectivity, and secure update protocols.

Finally, extra-large MIMO (XL-MIMO) arrays enable a new regime for massive IoT access by operating in the electromagnetic near field, where spherical wave propagation creates joint angle–distance–delay channel signatures. This allows distance-aware discrimination and localized energy focusing, enhancing multiplexing and reliability in dense deployments. Realizing these benefits demands scalable signal processing, near-field modelling, calibration, synchronization, and robust operation under mobility and sporadic traffic.

5.6.4 Critical communication

Reliable, resilient, and deterministic communication is a key requirement for industrial automation, robotics, and autonomous platforms, which represent a critical class of cyber-physical systems. Beyond uReliable, resilient, and deterministic communication, future industrial environments will require intelligent resource management, tight integration with sensing and IoT systems, and seamless human–machine interaction. This naturally leads to the convergence of communication, computation, and control, where digital twins of networks, devices, and physical processes are used to enable prediction, optimization, and real-time resource allocation in complex industrial settings.

Key research topics include deterministic wireless networking with integration of TSN principles to guarantee bounded latency and predictable behaviour over heterogeneous infrastructures. A second direction focuses on the joint design of communication and control, ensuring system stability under network constraints through PHY/MAC–control co-design, control-aware scheduling, and explicit characterization of stability–latency–reliability trade-offs in closed-loop systems. A further important area is goal-oriented and semantic communication for industrial applications, where information exchange is driven by control objectives rather than raw data delivery, enabling significant reductions in overhead while preserving safety and performance guarantees. Finally, future systems must support environment- and context-aware ultra-reliable communication, leveraging sensing and digital twin models for predictive resource allocation, as well as addressing extreme mobility and dynamic conditions in advanced robotic and autonomous industrial platforms.

5.6.5 Research Aspect 5 – Context-Aware PHY Control at the Edge

The edge is emerging as a key enabler for context-aware and adaptive physical-layer operation in large-scale IoT systems, bridging heterogeneous access technologies from radio to hybrid optical–radio interfaces. Future deployments require unified timing, scalable synchronization, and coordinated waveform adaptation under mobility, interference, and energy constraints. This motivates an edge-assisted PHY paradigm where sensing and communication are jointly exploited to acquire context at the physical layer and enable predictive resource adaptation. Distributed and federated learning further support collaborative PHY optimization across highly constrained devices and heterogeneous infrastructures, while neuromorphic and event-driven processing offers promising directions for ultra-low-power inference and anomaly detection. Overall, sustainable and resilient massive IoT requires PHY-layer designs that combine sensing-driven context acquisition, scalable synchronization, and learning-based adaptation in an edge-assisted framework.

5.6.6 Research Aspect 6 – infrastructure for experiment-driven research

Building and operating an increasingly complex ecosystem bridging communication, networking, distributed systems, computing, data management and more and more artificial intelligence while reducing the environmental footprint is a challenge. The enormous amount of data generated and exchanged over the network with the rise of Artificial Intelligence (AI) and sensing everywhere, requires to go beyond incremental changes and to investigate radically new infrastructure design and data management approaches. To support such changes, large-scale research infrastructures are needed to test new proposals in experimental but realistic environment and to make results trusted, repeatable, measurable and accessible to the European worldwide community. Deploying and operating a fully remotely programmable research infrastructure covering the entire IoT continuum, from end devices to edge/fog nodes and cloud resources, including integrated monitoring tools, represents a major challenge, but would deliver substantial benefits. First, it would enable cost efficiency through large-scale pooling of resources and capabilities, reducing redundant investments in infrastructure and limiting the need for each stakeholder to develop bespoke expertise from scratch. Second, it would ensure neutrality, as a politically agnostic, goal-independent platform would allow even competing industrial actors to perform evaluations and proof-of-concept validation within a shared and impartial environment. Finally, it would provide long-term sustainability, since while research projects typically span three to four years, a shared and continuously operated infrastructure could persist over decades, enabling sustained experimentation, cumulative impact, and stronger technological continuity.

5.6.7 Synopsis of Required Research

Research Theme	IoT Networks		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
RT1: Multi-RAT IoT Connectivity and Spectrum Coexistence	Short–Mid term (≤7y)	Interoperable IoT operation across ISM, licensed and NTN bands; coexistence frameworks for cellular, Wi-Fi and emerging RATs; AI-driven dynamic spectrum management; coexistence with backscatter devices and cross-technology handover. **[KPIs]:** Coverage (99%, including oceans); Max latency anywhere: 20-50ms, industrial: 1-0.1ms; density: 5M devices/km^3. Reliability (industrial): 99.99999%.	Enables seamless IoT connectivity across all coverage regimes, from white spaces to ultra-dense industrial environments; reduces fragmentation and inefficiencies of isolated RAT ecosystems.

RT2: Massive Access and Sporadic Communications	Mid–Long term (5y–7y+)	Ultra-low duty-cycle connectivity; non-orthogonal and grant-free access; short-packet coding and scheduling; scalable sparse and compressive-sensing receivers; OTA computation; efficient downlink access for massive sleeping devices. **[KPIs]:** device density 5M/m^3, packet length (20 Bytes), signalisation reduction (80%), over-the-air computation accurracy (95%)	Supports reliable, energy-efficient communication for massive IoT populations; reduces signalling overhead and latency; enables large-scale sensing and federated learning.
RT3: Reconfigurable and Near-Field PHY for Dense IoT	Long term (7y+)	Near-field XL-MIMO protocols; scalable processing for high-dimensional channels; distance-aware discrimination and localized energy focusing; sparsity exploitation in angle–distance–delay domains. **[KPIs]:** Energy Focusing Efficiency (80%), Distance Discrimination Resolution (10cm), Near-Field Throughput Density (total throuput per unit area, 20 Gbps/m^2) , Interference Suppression Ratio (30 dB)	Improves reliability and interference mitigation in ultra-dense IoT deployments; unlocks new spatial degrees of freedom for massive access.
RT4: Sustainable and Device-Centric IoT PHY Innovations	Mid–Long term (5y–7y+)	Lifetime-aware PHY/MAC protocols; adaptive modulation and access under device degradation; battery-less, energy-harvesting and biodegradable IoT connectivity; lifecycle-aware deployment and sustainability metrics. **[KPIs]:** Network lifetime (20+ years), consumption (active 20 uW), computation capacity (10 Tops/W) Biodegradability rate (90%), Wake-up latency (5ms)	Extends IoT lifetime toward maintenance-free operation; reduces e-waste and environmental impact; supports ethical, inclusive and sustainable IoT deployments.
RT5: Critical and Deterministic IoT Communications	Mid–Long term (5y–7y+)	Deterministic wireless networking with TSN integration; PHY/MAC–control co-design; stability–latency–reliability guarantees; digital-twin-driven predictive resource allocation; goal-oriented and semantic communication. **[KPIs]:** Reliability (99.99999%), Latency (0.1ms), Jitter (10 us), recovery time (resilience, 50 ms),	Enables wireless replacement of wired industrial networks; ensures cyber-physical stability and determinism; strengthens Europe's leadership in industrial automation.

		accuracy (semantic communication (95%)	
RT6: Context-Aware and Edge-Assisted IoT PHY Control and Experimentation	Mid–Long term (5y–7y+)	Edge-assisted ISAC for context-aware PHY adaptation; learning- and synchronisation-driven coordination across IoT swarms; open, repeatable experimental infrastructure spanning devices, edge and cloud. **[KPIs]:** Context-Aware Adaptation Latency (Time to adapt PHY parameters to real-time context) (10ms), ISAC Sensory Accuracy, Event-Driven Reaction Time, Open Infrastructure Accessibility, Eperiment deployement duration (2H), supported technologies (5+), Infrastructure energy consumption (10 kWh/year/node), Reprocudicbility rate (99%)	Improves resilience and efficiency through proactive PHY operation; enables trustworthy validation and accelerated innovation via shared European testbeds.

5.6.8 Recommendations for Actions

Research Theme	IoT Networks		
Action	**Multi-RAT & Massive IoT Connectivity (RT1, RT2)**	**Advanced IoT PHY & Sustainable Devices (RT3, RT4)**	**Critical, Edge-Assisted & Validated IoT Systems (RT5, RT6)**
International Calls	EU–international research on multi-RAT coexistence, TN–NTN integration, and AI-based spectrum sharing	Programmes on ultra-low-power IoT, backscatter/ZED, energy harvesting, biodegradable devices	Calls on industrial IoT, deterministic communication, and edge-assisted AI-driven PHY
Open Data	Datasets for coexistence, interference, and heterogeneous IoT deployments	Datasets on energy, harvesting, and device lifecycle behaviour	Datasets for industrial/control traffic and edge-assisted PHY validation
Large Trials	Trials of multi-RAT IoT and TN–NTN integration in dense environments	Trials of energy-autonomous IoT (backscatter, biodegradable devices)	Trials of critical IoT, deterministic wireless, and edge/digital twin systems
Cross-domain research	Collaboration across cellular, Wi-Fi, NTN and coexistence frameworks	Collaboration with hardware, materials, and energy domains	Integration with AI, edge/cloud, control systems, and digital twins

5.7 Transversal research themes

5.7.1 Optical satellite FSO

A key need in satellite communications is the rapid capacity increase demanded by new applications and services. This cannot be achieved with present RF technology, whilst resorting to Extremely High Frequency (EHF) is not straightforward (EHF still needs a huge technology development). On the other side, Free-Space Optical (FSO) technology seems the most suitable choice, thanks to the fact that it will be mostly based on existing devices and subsystems (developed and commercial for fibre networks since many years). FSO has long been the primary candidate for high-capacity NTN links, offering data rates and spectral freedom that RF cannot match; several commercial LEO constellations are already implementing optical ISLs, demonstrating terabit-per-second potential. However, FSO deployment faces significant limitations due to the high-precision mechanisms required by Pointing, Acquisition, and Tracking (PAT) between fast-moving spacecraft, the impact of optical path interruptions, and, for feeder links, atmospheric turbulence and fog. This FSO vulnerability translates into critical availability gaps in LEO/VLEO due to PAT precision demands, long re-acquisition times, and atmospheric impairments.

All satellite FSO systems will be strongly dependent on specific link details (type of satellite, orbit, etc.). Yet there are two main categories:

Optical Intersatellite Links (OISL): these will be very popular in future constellations; they can support intra-constellation and inter-constellation communications, with huge speed expected, creating a kind of optical bus with multiple relays in orbit. Presently, only a proprietary development in a private US constellation is known [C5-18] , with bit rate up to 20 Gbps. Other US-based companies are also moving to deploy OISL technology. This is expected to further progress in the near future, reaching hundreds of Gbps, possibly combined with WDM.

Optical Feeder Links (OFL): Optical Feeder Links (OFLs) are a critical enabler for future non-terrestrial networks, providing ultra-high-capacity connections between satellites and the terrestrial fibre network and forming the backbone of integrated 3D space–air–ground architectures. Their development is essential to support seamless, Tbps-scale data pipes across satellites, aerial platforms and ground infrastructure. While recent demonstrations, such as NASA's TBIRD experiment achieving 2×100 Gbps downlink, highlight the potential of OFLs, substantial research challenges remain, notably atmospheric impairments (scintillation, beam wander, and fading), transmission robustness under Doppler and hardware non-linearities, fast and reliable pointing, acquisition and tracking (PAT), and compatibility with terrestrial fibre transceivers. Further work is required on end-to-end network integration, satellite network control, NTN application design, and hybrid RF/FSO architectures enabling dynamic link selection and resilient fallback. Addressing these challenges will unlock Tbps optical satellite links, enabling ubiquitous high-speed connectivity, enhanced resilience to network disruptions, and support for high-priority and emergency services.

5.7.2 Quantum communications and networks

Quantum communication is emerging as a strategic enabler for **secure, resilient, and future-proof 6G infrastructures**. As 6G systems evolve towards distributed, software-defined, and AI-native architectures, the integration of quantum communication technologies will be essential to guarantee long-term security, enable new network functionalities, and support the convergence of classical and quantum information systems. 6G networks are expected to support ultra-critical services across defence, finance, energy, and autonomous systems, where zero-trust security and resilience are fundamental. Quantum communication contributes to this vision through:

Quantum-safe and quantum-native security: Complementing post-quantum cryptography, quantum key distribution (QKD) and entanglement-based schemes provide information-theoretic security for key exchange and network authentication.

Native integration with disaggregated 6G architectures: Quantum links shall be embedded into open, virtualised, and cloud-native network infrastructures, including Open RAN, edge-cloud, and core networks.

Support for distributed intelligence: Entanglement distribution enables synchronisation, coordination, and secure data exchange between distributed AI and compute nodes.

Convergence of terrestrial, aerial, and satellite networks: Quantum communication extends across fibre, free-space optical (FSO), and satellite links, aligning with 6G's NTN vision. NTN-based QKD enables unconditionally secure key distribution over thousands of kilometres via FSO links, directly addressing the vulnerability of classical and post-quantum cryptography to future quantum computer attacks, making it a cornerstone of resilient 6G architectures.

To ensure Europe at the forefront of 6G and quantum convergence, the followings are identified as the key investment areas:

6G-Integrated Quantum Network Infrastructure: Deployment of large-scale, heterogeneous quantum network testbeds tightly integrated with 6G experimental platforms. Interconnection of metro, core, access networks and data centre networks with quantum capabilities, including fibre and wireless optical links. Support for end-to-end use cases, such as quantum-secured Open RAN fronthaul/backhaul and cross-domain orchestration.

Co-existence of Quantum and Classical Traffic: Development of co-propagation techniques enabling quantum and high-capacity classical channels to share the same fibre infrastructure, especially for metro and access networks where fibre is more scarce. Advanced spectral, temporal, and spatial multiplexing schemes aligned with 6G optical transport evolution. Mitigation of noise, crosstalk, and nonlinear effects in realistic network environments, while including new fibre technologies such as Hollow-Core Fibre.

Quantum-Enhanced Network Control and Orchestration: Design of software-defined quantum networking frameworks interoperable with 6G control planes. AI-driven optimisation for entanglement distribution, routing, and resource allocation. Integration with network digital twins to enable predictive management and large-scale validation.

Scalable Quantum Hardware for Telecom Integration: Development of deployable, telecom-compatible quantum systems, including high-rate entangled photon sources, integrated photonic transceivers, field-deployable single-photon detectors, while considering miniaturisation, energy efficiency, and reliability to meet 6G operational requirements.

Interoperability, Standards, and Interfaces: Definition of common interfaces between quantum and classical network layers, aligned with 3GPP, ETSI, and ITU activities. Development of reference architectures and protocol stacks for quantum-enabled 6G networks. Establishment of benchmarking and certification frameworks to support industrial uptake.

QKD on free space NTN: Research activities shall focus on the investigation and implementation of free space communication networks, the design of interfaces integrating quantum systems of both fibre-based and free space based optical networks, the efficient time and frequency multiplexing of data and key transmissions in FSO NTN communications.

5.7.3 Optical wireless communication (OWC)

The continued exponential growth of mobile traffic and the emergence of the cyber-physical continuum, encompassing immersive XR applications, digital twins, and the convergence of communication, computing, sensing, control and robotics, will exceed the long-term capacity limits of the radio-frequency (RF) spectrum alone. Optical Wireless Communication (OWC), exploiting the infrared and visible portions of the electromagnetic spectrum, is therefore emerging as a critical complement to RF technologies (see section 6.3 of [C5-03]). Rather than competing with RF, OWC extends wireless systems into new spectral regions, leveraging largely unregulated spectrum, over three orders of magnitude larger than RF, while providing enhanced intrinsic physical-layer security through highly directional beams. Most current OWC systems rely on intensity modulation and direct detection (IM/DD), while coherent optical wireless transmission has recently attracted significant interest due to higher spectral efficiency and improved robustness. Today's OWC solutions already achieve data rates of tens of gigabits per second and are expected to scale toward hundreds of gigabits per second, making them well suited for future high-capacity terrestrial and aerial wireless connectivity.

In terrestrial networks, Free-Space Optical (FSO) communication is increasingly considered for long-range point-to-point wireless links, particularly as part of X-haul infrastructures complementing millimetre-wave technologies. However, deployment is constrained by atmospheric propagation effects such as turbulence-induced fading, scintillation, phase distortions, and strong sensitivity to weather conditions (e.g. fog). Addressing these challenges requires accurate channel characterisation, resilient transceiver designs, and advanced mitigation techniques. Current research focuses on tightly integrated architectures combining optical hardware innovations, such as adaptive optics and optical phased arrays, with advanced digital signal processing, as well as optical MIMO techniques to exploit spatial diversity and spatial multiplexing.

At short ranges, optical wireless access based on visible light communication (VLC), commonly referred to as LiFi, enables fully networked, bi-directional connectivity with multi-user access and handover. While early research focused on point-to-point data rates, current efforts address networking, mobility, and interference management. Fundamental differences with RF systems, such as the absence of conventional multipath fading and IM/DD constraints requiring real-valued, non-negative signals, necessitate LiFi-specific PHY and networking solutions. At the same time, the strong spatial confinement of light enables efficient interference control and extremely dense small-cell deployments with very high area spectral densities, reaching up to 88 Gbps/m². Moreover, the short optical wavelengths enable chip-scale detector arrays and massive MIMO structures, opening new opportunities for LiFi-specific MIMO and cooperative transmission. Standardisation is progressing, with LiFi specified in IEEE 802.11bb, complemented by standards for VLC and optical camera communication.

5.7.4 Satellite UAV HAPs

Within the multi-layered architecture envisioned for 6G non-terrestrial networks, Unmanned Aerial Vehicles (UAVs) and High-Altitude Platform Systems (HAPS) occupy a strategically unique position. Operating in the troposphere and lower stratosphere, respectively, they bridge the latency and link-budget gap between terrestrial networks and space-borne nodes, while offering on-demand deployment flexibility not achievable by fixed terrestrial or orbital infrastructures. Over the next decade, research should consolidate UAVs and HAPS as first-class elements of a unified TN/NTN fabric rather than as auxiliary relays.

In the short term, research priorities include native integration of aerial platforms into the 3GPP protocol stack, with particular emphasis on mobility management and handover across aerial–terrestrial and aerial–space boundaries. Owing to their quasi-stationary behaviour, HAPS can act as persistent regenerative nodes hosting gNB-DUs or even full gNB functions, enabling functional-split experimentation under relaxed energy and thermal constraints compared to space-borne platforms. For UAVs, research must address dynamic topology management, coordinated beamforming, and lightweight control-plane procedures capable of coping with high mobility and limited endurance, including the use of UAV corridors for safe and scalable swarm operation (see section 6.5.4 of [C5-03]).

In the medium term, HAPS are expected to become integral components of vertical inter-node links, supporting seamless data relaying between network layers and enhancing resilience through site diversity. AI-driven orchestration frameworks capable of jointly managing UAV and HAPS positioning, trajectories, and radio resources are a key research direction. In parallel, distributed and coordinated multi-antenna solutions exploiting UAV swarms or clustered HAPS can realise virtual large apertures at significantly lower cost than space-borne alternatives, enabling flexible coverage in underserved areas and high-demand scenarios such as large events.

5.7.5 Integrated Communication and Sensing Across Radio and Optical Domains

A transversal research theme emerges from the convergence of Integrated Sensing and Communication (ISAC) in RAN and optical fibre sensing capabilities. Together, they extend networks into pervasive sensing platforms, combining data transmission with continuous environmental awareness.

In RAN, ISAC enables radio signals and infrastructure to support sensing functions such as localization and environment mapping alongside communication. In parallel, optical fibre sensing provides distributed, high-resolution monitoring over wide areas using existing fixed infrastructures. These approaches are complementary: wireless sensing offers flexibility and coverage, while fibre sensing provides precision and continuity.

Their integration creates a multi-layer sensing fabric, supporting applications such as infrastructure monitoring, transportation, environmental observation, and security. It also introduces key research challenges, including cross-domain sensing-data fusion, joint resource management, and integration into real-time network control and orchestration.

From a system perspective, this convergence introduces key research challenges and opportunities, including:

- Joint design of sensing and communication functions across heterogeneous media
- Fusion of sensing data from radio and optical domains for improved accuracy and reliability
- Integration into network control and orchestration frameworks to support real-time decision-making
- Efficient sharing of spectrum and infrastructure resources between communication and sensing tasks

5.7.6 Underwater communications

Maritime and Subsea Communications represent a critical yet underdeveloped component of the emerging 6G connectivity continuum spanning space, aerial, terrestrial, maritime, and underwater domains [C5-05]. While the maritime sector is experiencing rapid technological and economic growth—driven by smart shipping, offshore energy expansion, intelligent ports, and autonomous maritime operations—the subsea segment remains constrained by severe physical-layer limitations and fragmented communication infrastructures. Future 6G architectures are expected to unify these environments into an integrated space–air–ground–sea–underwater network, enabling consistent device identity, management, and service exposure across domains [C5-06]. Achieving this will require new physical-layer innovations at the air–water boundary, enhanced underwater communication methods, and hybrid network architectures capable of supporting real-time monitoring and control of subsea assets.

Emerging technologies such as quantum-enhanced underwater navigation, Integrated Sensing and Communication (ISAC), and Distributed Acoustic Sensing (DAS) using existing subsea fibre cables are expected to play a transformative role in subsea operations. These advances will significantly improve positioning accuracy, situational awareness, and environmental monitoring in high-pressure and low-bandwidth subsea environments. With Europe holding strong industrial leadership in maritime operations, offshore energy, and underwater robotics, dedicated research efforts in maritime and subsea communications will reinforce Europe's strategic position and ensure that these domains are fully represented in the holistic connectivity vision that underpins 6G.

5.7.7 Network performance evaluation tools

The complexity of 6G radio access networks is significantly increased by the introduction of new network elements (e.g. reconfigurable intelligent surfaces and UAV-assisted communications), new frequency bands (e.g. FR3 and THz), new architectures (e.g. cell-free and non-terrestrial networks), and new services such as integrated sensing and communication (ISAC). Mastering this complexity requires novel system-level performance evaluation tools capable of capturing heterogeneous interactions across space, time and network layers (see section 6.2.4 of [C5-03]).

Two complementary approaches currently dominate performance evaluation: discrete-event (system-level) simulation and stochastic geometry (SG). Both must be further advanced for 6G RAN analysis. SG, a probabilistic framework originally developed in materials science, has been increasingly adopted for wireless networks in recent years. By modelling network elements such as base stations and users as spatial point processes, SG enables closed-form characterisation of key metrics including coverage, spectral efficiency, and rate distributions as functions of a limited set of physical parameters. This approach provides statistical-physics-like insights into large-scale network behaviour and has been used for the technical and economic evaluation of cellular, Wi-Fi, vehicular and self-organised networks.

Looking ahead to 6G, there is a strong need to extend SG-based system-level analysis to emerging paradigms such as ISAC networks, cell-free architectures, RIS-assisted systems, and non-terrestrial networks. For RIS-enabled deployments, SG models can represent clouds of RISs associated with base stations and enable quantitative evaluation of their impact on spectral efficiency, including trade-offs between RIS density, element complexity, obstacle density, and base-station densification. Similarly, the integration of non-terrestrial networks into 6G motivates the development of spherical and orbital SG models capable of capturing satellite constellations and their interaction with terrestrial networks. Such tools are essential to analyse coverage, throughput, mobility, interference, and economic coexistence between terrestrial and non-terrestrial systems.

Beyond individual paradigms, future performance evaluation tools must support the systematic analysis of interactions among heterogeneous RAN technologies and architectural options, including levels of centralisation, MEC integration, and joint deployment optimisation. They should also enable the design of real-time and non-real-time controllers, supporting informed orchestration and the identification of joint technical and economic opportunities across the 6G ecosystem.

5.7.8 Over the air test

Within 3GPP, OTA methodologies have evolved from classical radiated conformance toward validating beamformed NR devices under realistic propagation conditions, as seen in RAN4 studies. OTA testing is now a key challenge for 6G validation, as the goal is to assess radiated behaviour in realistic operating conditions. For FR1, FR2, and FR3 beamforming systems, the challenge is to reproduce dynamic, repeatable propagation conditions for beam management, mobility, and multi-user operation under time-varying channels and cell-free mMIMO. Larger apertures and test zones drive the need for 3D multi-probe configurations, but practical constraints arise for large FR3 arrays whose Fraunhofer distances exceed chamber dimensions. Small-device OTA methods do not extend to subarray-based hybrid beamforming or RIS. A shift is also occurring from communication-only OTA toward ISAC/JSAC validation, requiring scene emulation (clutter, target backscatter, multistatic effects) rather than just channel emulation. Future OTA methodologies should serve as a critical bridge between 3GPP-compliant testing and the validation of joint communication, sensing, localization, and massive MIMO in 6G.

5.7.9 Satellite IoT

Realising NTN IoT at scale requires solutions to several technical challenges that remain only partially addressed. Antenna development is a fundamental constraint: devices must remain small and energy-efficient while achieving sufficient gain to close the link with a satellite. The link budget challenge is severe, as LEO path losses far exceed terrestrial link attenuations, critically compressing the margins available for reliable communication on battery-powered devices. Frequency coordination between non-terrestrial and terrestrial systems sharing overlapping bands requires careful management. Perhaps most challenging, the enormous coverage footprint of a single satellite potentially serves tens of thousands of simultaneously transmitting devices, creating unprecedented capacity demands that exceed the capabilities of current multiple access protocols and require advanced interference cancellation techniques tailored to the satellite IoT context.

NTN IoT PHY and Massive Access. Research should jointly optimise antenna, link budget, and battery-life design for LPWAN devices operating over LEO satellite links, closing the margins required for reliable low-power transmission. Multiple access protocols and interference cancellation techniques scaling to the tens-of-thousands-of-devices capacity demands of satellite IoT footprints should be developed and validated. Doppler-based localisation approaches exploiting satellite motion, using single-satellite synthetic apertures or more powerful multi-satellite configurations with inter-satellite links, offer GNSS-free positioning capabilities as a valuable complementary service.

Hybrid TN–NTN IoT Device Architectures. R&I should focus on developing a systematic design framework for selecting between shared-frontend and dual-frontend TN-NTN IoT architectures based on application reliability requirements, cost constraints, and coverage profiles. Simultaneous TN/NTN dual-connectivity solutions for resilience-critical applications should be advanced from proof-of-concept to standardised operational capability, with seamless link switching protocols enabling continuous service as coverage conditions change.

5.7.10 Synopsis of Required Research

Research Theme	Transversal research themes		
Research Topics	**Timeline**	**Key outcomes**	**Contributions/Value**
RT1: Optical Satellite and Free-Space Optical (FSO) Networks	Mid–Long term (2028–2035)	Terabit-class optical inter-satellite and feeder links; fast and robust pointing, acquisition and tracking; atmospheric-impairment mitigation; hybrid RF/FSO	Enables ultra-high-capacity non-terrestrial backbones integrated with terrestrial fibre; improves availability and resilience of global

		architectures with dynamic fallback.	connectivity; reduces RF spectrum pressure.
RT2: Quantum Communications and Quantum-Enabled Networks	Mid–Long term (2028–2035)	Quantum-safe and quantum-native security, including QKD over fibre and free space; co-propagation of quantum and classical channels; software-defined quantum network control.	Future-proofs 6G security against quantum threats; underpins European leadership in secure digital infrastructure and early quantum internet capabilities.
RT3: Optical Wireless Communication (OWC) for Access and X-haul	Mid–Long term (2027–2035)	High-capacity terrestrial FSO for X-haul; robust optical MIMO and adaptive optics; LiFi networking with dense spatial reuse and chip-scale massive MIMO; OWC-specific PHY and networking methods.	Adds ultra-high-capacity spectrum complementing RF; enables dense indoor access and flexible backhaul; enhances physical-layer security and spatial efficiency.
RT4: UAV and HAPS for Integrated Aerial Networking	Mid–Long term (2027–2035)	Native PHY and protocol integration of UAVs/HAPS; coordinated multi-antenna aerial platforms; AI-driven placement, trajectory, and beam management; aerial edge capabilities.	Enables flexible, on-demand coverage and capacity; strengthens TN/NTN integration; supports resilience, rapid deployment, and underserved-area connectivity.
RT5: Integrated Communication and Sensing Across RAN and Optical Infrastructures	Mid–Long term (2027–2035)	Joint ISAC functionality in RAN and distributed optical fibre sensing; cross-domain sensing data fusion; unified frameworks for joint communication–sensing resource allocation; integration of sensing into network control, orchestration, and digital twin platforms.	Establishes networks as pervasive sensing infrastructures combining wireless flexibility with fibre-based precision; enables new services (infrastructure monitoring, transportation, environmental intelligence, security); enhances network situational awareness, resilience, and operational efficiency.
RT6: Underwater and Maritime Communications	Long term (2028–2035)	Hybrid acoustic-optical-RF PHY solutions; air–water interface techniques; integrated sensing and communication for subsea environments; convergence with fibre-based DAS.	Extends the 6G connectivity continuum to maritime and subsea domains, e.g. improved capacity via MIMO and RIS; supports offshore energy, environmental monitoring, and autonomous maritime systems.
RT7: Network Performance Evaluation and Modelling Tools	Short–Mid term (2026–2032)	Advanced system-level simulation; stochastic-geometry-based modelling for RIS, ISAC, cell-free and NTN; joint techno-economic	Provides scalable, physically grounded evaluation of complex 6G architectures; supports informed design,

		evaluation and deployment optimisation tools.	regulation, and investment decisions.
RT8: Over-the-Air (OTA) Testing and Validation	Mid–Long term (2027–2035)	OTA methodologies for massive and distributed MIMO, RIS, ISAC, and near-field operations; realistic channel and scene emulation; standard-ready validation frameworks.	De-risks 6G PHY technologies; bridges laboratory innovation and standard-compliant deployment; accelerates industrial adoption and interoperability.

5.7.11 Recommendations for Actions

Research Theme	Transversal research themes			
Action	**Space–Air–Ground Optical & Wireless Integration (RT1, RT3, RT4)**	**Quantum-Enabled Secure Infrastructures (RT2)**	**Extended Connectivity Domains (RT5)**	**Evaluation, Validation & Experimentation (RT6, RT7)**
International Calls	Joint EU–international calls on hybrid RF/FSO, OWC, and UAV/HAPS architectures for integrated terrestrial–aerial–satellite connectivity.	Coordinated programmes on quantum communications, QKD over fibre and free space, and quantum–classical coexistence in 6G networks.	Targeted actions on maritime and underwater communications, integrating acoustic, optical, and RF technologies.	
Open Data	Open propagation and channel datasets for FSO, OWC, UAV/HAPS, and hybrid RF/optical links.			Open benchmarks for stochastic-geometry modelling, system-level simulation, and OTA testing.
Large Trials		Pilot deployments of quantum-enabled network segments interoperating with classical infrastructures.		OTA and system-level validation of massive/distributed MIMO, RIS, ISAC, and near-field systems.

5.8 Security Considerations

Security is a foundational requirement for 6G systems, where unprecedented connectivity, extensive automation, and heterogeneous infrastructures substantially expand the threat surface. Across optical, radio, non-terrestrial, and IoT domains, future networks must ensure availability, confidentiality, integrity, authentication, and resilience against both conventional cyberattacks and emerging threats enabled by AI, massive connectivity, and quantum computing. Addressing these challenges requires a holistic, multi-layer approach combining cryptography, physical-layer techniques, intelligent management, and trustworthy AI.

Packet-optical networks form the backbone of global communication, cloud, and AI infrastructures and carry massive volumes of mission-critical and sensitive data. Their security is therefore paramount. Future-proof optical networks must simultaneously enhance resilience, protect data confidentiality and integrity, and guarantee authentication. Depending on network attachment points, encryption at different layers is required, with optical-layer (Layer 1) encryption offering stronger security while reducing latency, energy consumption, and bandwidth overhead compared to higher-layer solutions. A key mid-term objective is the transition to quantum-safe security. Although initial post-quantum cryptography standards exist, substantial work remains to harden implementations, optimise performance, and manage migration toward quantum-resistant critical infrastructures by 2030. Photonic technologies can further reinforce security through high-quality quantum random number generators and physical unclonable functions for device authentication, while quantum key distribution provides a complementary quantum-safe mechanism with a fundamentally different attack surface. Security and resilience further depend on intelligent, autonomous network management: actionable situational awareness derived from telemetry, sensing, and continuous monitoring is essential to counter espionage, information leakage, natural hazards, and sabotage. AI-driven, intent-based management enables automated anomaly detection and rapid recovery, while availability-driven resource allocation reduces downtimes. In parallel, optical timing distribution offers a secure alternative to GNSS-based timing, inherently resistant to spoofing and jamming.

In the radio domain, ubiquitous access and the massive deployment of low-cost devices, key features of 6G, significantly increase exposure to security threats. Emerging 6G PHY layer innovations, particularly ISAC, introduce new security, privacy, and trustworthiness challenges. Since communication and sensing functions share waveforms, spectrum, and infrastructure, future research must address secure sensing, protection against sensing manipulation and inference attacks, resilient beamforming and waveform design, and trustworthy AI driven PHY layer operation. Due to the broadcast nature of wireless channels, transmissions are inherently vulnerable to attacks such as jamming or passive eavesdropping. Beyond higher-layer security mechanisms, physical layer security exploits physical channel characteristics to mitigate these threats. While traditional approaches rely on artificial noise and diversity, emerging solutions leverage AI/ML together with massive MIMO, full-duplex, and mmWave technologies. ML-based spectral and signal analysis can further enable the detection of jamming, eavesdropping, and rogue base stations. Physical-layer key generation, based on channel reciprocity and uniqueness, can particularly benefit from AI/ML in complex environments such as multi-hop mMIMO and high-frequency systems where accurate channel models remain limited (see section 6.10 of [C5-03]).

The increasing reliance on AI/ML across 6G also raises critical trustworthiness concerns. Key challenges include explainability, robustness to inputs outside the training distribution, and resilience to spoofing and adversarial attacks. Unlike model-based techniques, many AI/ML algorithms behave as black boxes, motivating research in explainable AI, the integration of physical models into learning, and mechanisms to detect and handle out-of-distribution operation. Ensuring robustness against both unintentional perturbations and deliberate manipulation is essential as AI/ML becomes integral to radio access, resource control, and network automation.

Security in NTN spans the entire system chain, from eavesdropping protection to resilience against GNSS jamming and spoofing. Integrated terrestrial–non-terrestrial systems require end-to-end security frameworks combining post-quantum cryptography, quantum-safe key distribution, blockchain-based distributed trust, and physical-layer security. Such frameworks must support high secure-key generation rates, robustness to multi-hop impairments, secure key handling under constellation-scale handovers, and protection against eavesdropping, man-in-the-middle attacks, and data poisoning. AI-based interference classification for GNSS and hybrid GNSS/6G positioning is critical to address the structural vulnerability of satellite navigation, requiring generalisable and fast-adapting models, federated learning, and robust domain-adaptation techniques.

Finally, security for IoT devices is a critical concern as future deployments rely on massive numbers of constrained, autonomous, and often unattended nodes. Lightweight authentication and access-control mechanisms are essential under strict energy and processing constraints. Physical-layer security offers complementary protection by leveraging channel properties to enhance confidentiality without heavy computational overhead. As IoT systems increasingly incorporate local intelligence and autonomous decision-making, establishing trust in devices, data, and AI-based training at the edge becomes central, especially in dynamic and adversarial environments where compromised data or models can directly impact system behaviour and reliability.

Research Domain / Theme	Security assets concerned	Main security needs	Key mechanisms & services	Security management & assumptions
Optical networks (transport, DC, timing)	Critical systems (optical infrastructure, control planes); Critical data (high-value traffic)	High availability and resilience; confidentiality, integrity and authentication at very high data rates; long-term quantum resistance	Optical-layer (Layer-1) encryption; post-quantum cryptography; quantum key distribution; autonomous monitoring and recovery; optical timing distribution	Automated quantum-safe key lifecycle management; trusted physical infrastructure and photonic devices; secure control planes and telemetry
Radio communications (RAN, spectrum, access, sensing)	Critical systems (gNBs, mMIMO, RIS, ISAC infrastructure); Personal, sensing and critical user data	Protection against jamming, eavesdropping, spoofing and sensing manipulation; scalable security for massive access; trustworthy AI-driven control; privacy preserving sensing	Physical layer security; AI-based attack detection; PHY layer key generation; secure waveform and beamforming design; PHY layer authentication; explainable and robust AI/ML	Dynamic reconfiguration of security mechanisms; continuous model updates; assumption of channel reciprocity and partial environmental trust; secure coordination between sensing and communication

Non-terrestrial networks (NTN, GNSS)	Critical systems (satellites, ISLs, PNT infrastructure); Critical positioning and timing data	End-to-end security across TN/NTN; resilience to GNSS spoofing/jamming; distributed trust in multi-stakeholder environments	Multi-layer security combining PQC, QKD, blockchain-based trust, PHY security; AI-based interference classification and detection	Federated security management across space, aerial and terrestrial segments; distributed consensus; limited trust in RF environment
IoT devices and edge systems	Critical systems (constrained devices); Personal and safety-critical data	Lightweight authentication and confidentiality; resilience with minimal energy and compute overhead; trustworthy edge intelligence	Lightweight cryptography; physical-layer security; secure boot and identity; protected AI-based edge training and inference	Energy-aware security updates; minimal trust in physical access; secure but intermittent connectivity
AI/ML across domains	Critical systems (models, training data, controllers)	Explainability, robustness, and resilience to adversarial and spoofing attacks; predictable behaviour	Explainable AI; robustness to out-of-distribution inputs; secure and auditable learning pipelines	Trusted training data pipelines; secure model updates; human-in-the-loop and auditability where required

5.9 Sustainability Considerations

Sustainability is a critical cross-cutting requirement for future communication systems, as networks already account for a non-negligible and growing share of global electricity consumption and material use. This trend is expected to accelerate with the large-scale deployment of AI workloads, 6G infrastructures, immersive applications, and edge intelligence. Future network design must therefore decouple capacity growth and service performance from energy consumption, operational complexity, and material waste, while enabling long-term scalability across terrestrial and non-terrestrial domains.

In optical networks—which form the backbone of metro, metro-access, fronthaul, and AI data-centre connectivity—sustainability hinges on energy-efficient high-capacity transmission and reduced network complexity. Coherent transceivers are a key enabler, with research priorities including low-power operation, point-to-multipoint architectures that reduce the number of active elements, and SDN-controlled optimisation supporting multi-vendor interoperability and flexible functional splits. Power-proportional transceivers, dynamically adapting modulation, symbol rate, and DSP complexity, help avoid chronic over-provisioning. In parallel, co-packaged optics, silicon photonics, and selective photonic offloading of DSP functions can significantly reduce per-bit energy consumption in dense AI data-centre fabrics. Sustainability is further enhanced by keeping traffic in the photonic domain end-to-end through optical bypass, ROADMs, optical circuit switching, passive splitting, and wideband transmission, thereby minimising optical-electrical-optical conversions, cooling demand, and operational overhead. AI-driven control planes additionally enable carbon-aware routing based on real-time grid conditions, while modular, backward-compatible platforms and circular-economy principles address long-term material sustainability.

In radio communications, sustainability must become a first-order design objective alongside performance. A key trajectory is the transition from RF-chain-intensive architectures toward multi-hybrid transceivers that shift part of the processing into analogue and wave-propagation domains, reducing reliance on power-hungry RF chains, converters, and digital processing. Beyond reducing operational energy, sustainable radio design also requires modular, upgradable hardware that extends lifetime, lowers embodied energy, and limits electronic waste. Life Cycle Assessment should therefore evolve from an ex-post evaluation tool into a decision framework guiding eco-design, hardware selection, and deployment strategies throughout the 6G lifecycle.

For non-terrestrial networks, sustainability considerations are tightly coupled to efficient use of limited space and energy resources. Key research directions include energy benchmarking of candidate waveforms under realistic satellite power-amplifier constraints, launch-mass-to-capacity optimisation for federated and swarm-based architectures, energy-aware multiple access and scheduling exploiting orbital predictability and renewable energy availability at ground stations, and improved spectrum sharing to reduce redundant infrastructure. Enhancing the efficiency of inter-satellite and feeder links and reducing reliance on GNSS through authenticated assistance data further contribute to lowering energy consumption and emissions. Across all domains, the development of common sustainability metrics—such as bits-per-Joule, CO_2-per-service-hour, and lifecycle impact—integrated into system-level design trade-offs is essential to align 6G evolution with the EU Green Deal and emerging ICT environmental reporting standards.

Research theme / aspect	**Sustainability impact areas (Q1–Q6)**	**Primary LCA phase**	**Concrete sustainability impact (SMART)**	**Relation to sustainability**
Energy-efficient coherent optics (low-power, power-proportional, P2MP)	Q1 Technology (ICT); Q2 Economy (ICT); Q4 Technology (enabled by ICT)	Use / Manufacturing	Reduces energy per transmitted bit and the number of active optical elements in metro, access, and fronthaul networks; measurable via W/bit and transceivers per endpoint.	Dedicated to sustainability
Photonic integration and optical bypass (CPO, silicon photonics, ROADMs, OCS)	Q1 Technology (ICT); Q2 Economy (ICT); Q4 Technology (enabled by ICT)	Manufacturing / Use	Cuts electrical I/O losses and avoids unnecessary O-E-O conversions, reducing power consumption and cooling demand; measurable via pJ/bit and regeneration-	Dedicated to sustainability

			stage reduction.	
Wideband fibre utilisation and asset reuse (L-band, modular optics)	Q1 Technology (ICT); Q2 Economy (ICT); Q4 Technology (enabled by ICT)	Lifetime extension	Increases fibre capacity without new deployments, extending infrastructure lifetime and avoiding new material extraction; measurable via added Tb/s per installed fibre.	Helps with sustainability
AI-driven carbon-aware network control (optical and radio)	Q1 Technology (ICT); Q2 Economy (ICT); Q3 Society (ICT); Q6 Society (enabled by ICT)	Use / Operation	Dynamically steers traffic based on real-time grid carbon intensity, reducing CO_2 per service hour; measurable via carbon-aware KPIs.	Dedicated to sustainability
Multi-hybrid radio transceiver architectures (analogue/wave-domain)	Q1 Technology (ICT); Q2 Economy (ICT); Q4 Technology (enabled by ICT)	Manufacturing / Use	Reduces RF chains, converters, and digital processing blocks, lowering operational energy and embodied footprint; measurable via energy/bit.	Dedicated to sustainability
Modular and upgradable radio and optical hardware	Q1 Technology (ICT); Q2 Economy (ICT); Q4 Technology (enabled by ICT)	Lifetime extension / End-of-life	Extends hardware lifetime and reduces e-waste by enabling upgrades rather than replacement; measurable via extended service years.	Dedicated to sustainability

Energy-aware NTN architectures (waveforms, scheduling, spectrum sharing)	Q1 Technology (ICT); Q2 Economy (ICT); Q5 Economy (enabled by ICT)	Use / Deployment	Improves satellite energy efficiency via energy-aware waveform selection and resource use; measurable via bits/J and reduced infrastructure redundancy.	Dedicated to sustainability
Sustainability metrics and LCA-driven eco-design (cross-domain)	Q5 Economy (enabled by ICT); Q6 Society (enabled by ICT)	Cross-lifecycle	Enables measurable and comparable sustainability governance (bits/J, CO_2/service, lifecycle impact) embedded into design and operation decisions.	Dedicated to sustainability

5.10 Device Considerations

Future 6G systems will rely on an unprecedented diversity of devices, spanning smartphones, IoT nodes, industrial controllers, vehicles, wearables, aerial platforms, and satellite payloads. At the physical layer, this diversity translates into highly heterogeneous constraints and capabilities, making device-aware PHY design a core requirement rather than an afterthought. Compared to previous generations, 6G devices must simultaneously meet extreme performance targets, tight energy budgets, long operational lifetimes, and growing requirements for security, trustworthiness, and sustainability.

In the radio domain, user and infrastructure devices must operate across a much broader range of frequencies, bandwidths, and antenna configurations than in 5G, including FR3, mmWave, and sub-THz bands, near-field regimes, and large or distributed antenna arrays. This imposes strong constraints on RF front ends, antenna integration, thermal management, and processing complexity. A key trend is the shift away from RF-chain-intensive designs toward multi-hybrid and wave-domain processing, reducing the number of power-hungry RF chains, converters, and digital blocks. These approaches are essential to keep energy consumption, cost, and form factor compatible with large-scale deployment, particularly for massive MIMO, RIS-assisted systems, and cell-free architectures. Device design must also accommodate tighter integration with optical technologies, such as radio-photonics, co-packaged optics, and optical interconnects within radio units and distributed cloud-RAN architectures.

IoT devices represent an even more extreme case of heterogeneity. Massive IoT targets billions of ultra-low-cost, energy-constrained devices expected to operate autonomously for years, often with sporadic transmissions, minimal synchronization, and limited hardware capabilities. In contrast, critical IoT devices for industrial automation, robotics, and control systems require deterministic, bidirectional, ultra-reliable, and low-latency communication with precise timing. PHY-layer device design must therefore support both ultra-low-power operation—leveraging wake-up radios, backscatter, energy harvesting, and lightweight processing—and highly reliable, tightly synchronized operation for safety-critical use cases. This duality drives the need for adaptable PHY modes that can switch between energy-centric and reliability-centric regimes as requirements evolve.

Non-terrestrial and aerial devices, including satellites, UAVs, and HAPS, further amplify device-level challenges. These platforms are constrained by size, weight, power, and thermal limits while operating under extreme mobility, Doppler, and propagation dynamics. PHY device design for NTN must prioritise power-efficient waveforms, robust synchronisation, flexible multi-band and multi-beam front ends, and close integration of communication, positioning, and sensing. Federated and cooperative architectures shift part of the complexity to coordinated system behaviour but still impose stringent requirements on on-board processing, timing accuracy, and interoperability.

Across all domains, sustainability and security increasingly shape device considerations. Modular, upgradable architectures are needed to extend device lifetime and reduce e-waste, while power-proportional operation and lifecycle-aware design help decouple performance growth from energy consumption. At the same time, devices must embed physical-layer security, quantum-safe mechanisms, and support trustworthy AI-assisted operation without excessive overhead. The overarching challenge for the PHY layer is therefore to enable device designs that are performant, energy-efficient, secure, and adaptable across highly heterogeneous 6G deployment scenarios, while remaining economically viable at scale.

5.11 Chapter Summary

The chapter addressed the evolution of fundamental communication technologies including advances in optical and radio systems, integrated sensing and communication, non-terrestrial networks, and device technologies. Research directions focus on extending capacity, improving energy efficiency, enabling new functionalities such as sensing, and supporting extreme heterogeneity in devices and environments. Identified research topics reflect the need for tight integration with higher layers and for alignment with cross-cutting objectives such as sustainability and resilience. The Physical Layer (PHY) Innovation Layer is inherently interdependent with all other SRIA chapters, reflecting the system-level nature of future network infrastructures where performance, efficiency, and functionality emerge from cross-layer integration.

- **System View and Architecture Orientations:** PHY innovations define the fundamental capabilities (e.g., capacity, latency, reliability, sensing) that constrain and enable architectural designs, while system-level requirements drive the need for adaptive, reconfigurable, and energy-efficient physical layer solutions.
- **Softwarisation and Intelligence in Networks:** The PHY layer both benefits from and enables AI-driven control and Softwarisation, requiring programmability, real-time observability, and exposure of low-level capabilities to support closed-loop automation and intelligent network operation.
- **Protocols and Communication Paradigms:** Advances in PHY technologies directly shape feasible communication paradigms (e.g., deterministic, semantic, and joint communication-sensing), while protocol requirements demand increased flexibility, abstraction, and dynamic resource control at the physical layer.
- **Transversal and Disruptive Technologies:** The PHY layer serves as the primary integration point for emerging technologies such as photonics, quantum communication, sub-THz systems, sensing, and non-terrestrial networks, enabling their incorporation into scalable, end-to-end infrastructures.

6 Transversal and Disruptive Technologies

This chapter captures transversal and disruptive, high-impact technologies not yet mainstreamed in frameworks like the ETSI technology radar **Error! Reference source not found.** or Gartner hype cycle. It adopts a storytelling format to explore user-centric scenarios, emphasising European innovation, holistic societal impact, and alignment with global challenges (e.g., UN SDGs). Unlike prior chapters, it bridges fragmented topics into cohesive visions, highlighting unresolved ethical, technical, and behavioural hurdles. Future Emerging Technologies (FETs) are presented in a storytelling fashion of user scenarios, including aspects as:

- A description of the user scenario
- A description of the technology (or technologies)
- A view on the potential impact on the UN Sustainable Development Goals (SDGs)

In addition to the storytelling fashion that cover the above points, several identified technologies (or set of technologies) are accompanied by:

- An estimation of the TRL. We should expect that FETs are at most at TRL2 at the time of publication of this SRIA, so this point will mostly be omitted.
- Unresolved issues (technical and non-technical, social acceptance, change of human behaviour, ethics...).

6.1 The Future Fabric of Computing and Storage

As Europe advances toward its 2030+ sustainability, sovereignty, and digital transition goals, computing is approaching a fundamental inflection point. For decades, progress in information processing has relied on Moore's Law and digital logic miniaturisation. However, the physical and energetic limits of CMOS scaling are now converging with global constraints on energy production and material availability. The Decadal Plan for Semiconductors identified five seismic shifts **Error! Reference source not found.**, including the need for new compute paradigms, that must be addressed to sustain innovation and economic growth. Simultaneously, initiatives such as Ecologic Computing **Error! Reference source not found.** emphasise that information processing must become symbiotic with its energy and environmental context — computing that is not only efficient, but ecologically scalable for the long term.

This chapter explores how extending computation beyond the digital domain can create a resilient, low-carbon, and high-performance technological fabric for Europe's post-2030 innovation agenda. It also discusses how post-Shannon communication models could complement these paradigms, enabling holistic matter-to-information ecosystems.

6.1.1 The Necessity of Extending Computing Paradigms

Energy consumption from ICT already represents 8–10% of global electricity use and is projected to rise steeply with the proliferation of AI and immersive services **Error! Reference source not found.**. The Decadal Plan warns of an "energy–information singularity" where computing energy demand could surpass the world's total generation capacity if current digital architectures persist. For the European Union, this represents both a technological and a strategic challenge: achieving the Green Deal objectives, digital sovereignty, and global competitiveness requires decoupling computational growth from proportional energy growth.

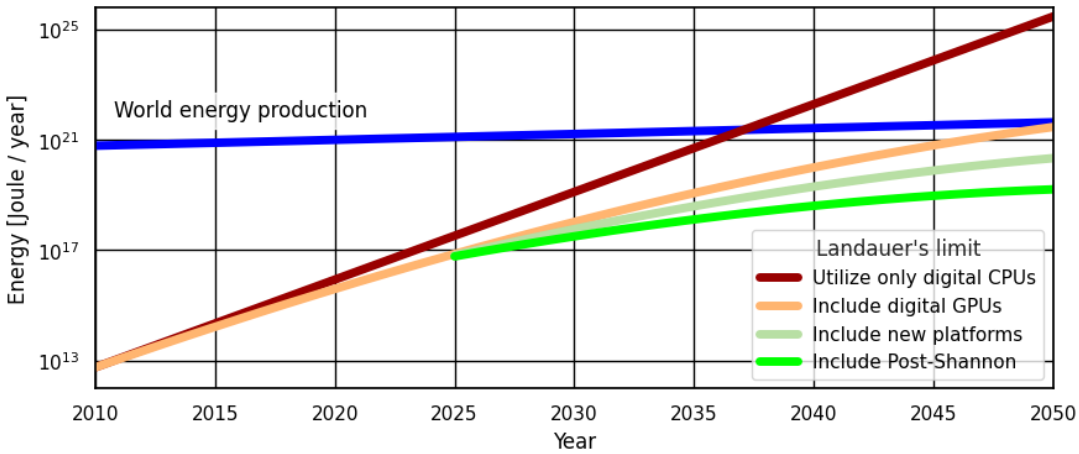

Figure 6-6-1 - Energy-computation singularity (Source Error! Reference source not found.**,** Error! Reference source not found.**).**

Simply improving transistor efficiency or scaling digital logic will not suffice. Every bit processed carries an intrinsic energy cost bounded by Landauer's limit **Error! Reference source not found.**, and incremental architectural optimisation can no longer offset exponential demand growth. Extending the computing spectrum toward physical, quantum, and biological substrates offers a way to avoid digital inefficiencies rather than merely optimise them, enabling ICT growth compatible with the EU's 2050 climate neutrality targets.

6.1.2 Emerging Computing and Storage Paradigms

Analogue computing

Analog computation, once overshadowed by digital logic, is experiencing a renaissance in the AI and sensor era. Rather than discretising physical signals into binary form, analogue systems compute directly in the physical domain using continuous variables — voltage, current, phase, or wave interference — as computational substrates. This provides a structural energy advantage: it avoids the overhead of analogue-to-digital conversion, clock synchronisation, and high-precision binary representation where such precision is not required, matching computational resolution to application needs.

Key implementations include in-memory analogue processing for neuromorphic inference, where resistive memory crossbars (RRAM, FeFET) perform matrix multiplications natively within memory arrays **Error! Reference source not found.**, dramatically reducing data movement — one of the dominant contributors to energy consumption in modern architectures. Wave-based analogue computing uses optical or acoustic interference to perform transforms such as convolution or Fourier operations at the speed of physics with minimal switching energy. Within the Ecologic Computing framework **Error! Reference source not found.**, these architectures are positioned as energy-adaptive substrates that process only what the environment requires. Analog computing is a transversal enabler for sustainable edge intelligence, real-time sensing-to-action systems, and ultra-low-power AI acceleration, with a realistic deployment horizon of five to ten years.

Quantum computing

Quantum computing **Error! Reference source not found.**, [C6-09] leverages superposition and entanglement to perform exponentially complex operations with potentially lower thermodynamic cost per logical operation. It is addressed in depth in chapter 0.

Biological computing and storage

Biological systems process information with extraordinary energy efficiency. Cells perform sensing, decision-making, memory, and actuation through biochemical pathways that are inherently parallel, self-repairing, and thermodynamically balanced with their environment. Translating these principles into ICT architectures opens a radically new frontier for sustainable computing **Error! Reference source not found.**.

DNA-based storage illustrates the transformative potential. DNA molecules store information at densities orders of magnitude beyond conventional media — approximately 100 trillion GB per gram compared to 1 TB per square inch for solid-state storage— and require no active energy input at rest. Legacy storage media has an average lifespan of around seven years; DNA can survive, even under non-ideal conditions, for more than 100,000 years **Error! Reference source not found. Error! Reference source not found.**. Projects such as DNA-Fountain have demonstrated the practical viability of this approach.

Beyond storage, synthetic biology enables programmable genetic circuits capable of logical operations within living cells. Enzymatic reactions implement logic gates **Error! Reference source not found.**; cellular networks respond dynamically to temperature, pressure, or chemical signals; and emerging research into mycelium-based bio-electrical networks demonstrates that biological materials can function as distributed, decentralised, and adaptive computational substrates — offering a glimpse of how future networks might become self-healing and self-organising. Bio-neuromorphic systems, integrating living neuronal networks with silicon interfaces, extend this further toward hybrid architectures.

For Europe, investment in biological computing aligns naturally with sustainability and circular economy principles. Biological substrates are biodegradable, potentially regenerative, and intrinsically energy-proportional — directly addressing ICT's environmental footprint while opening new avenues for long-term archival storage and bio-hybrid sensing platforms.

Together, these paradigms could enable ICT growth that decouples from exponential energy consumption, essential to achieving the EU Green Deal’s 2050 climate neutrality targets.

6.1.3 Post-Shannon Communications

Shannon's theory defined the limits of reliable digital communication through discrete bits and entropy. Future hybrid computing fabrics — where information may exist in analogue, quantum, or biochemical forms — demand post-Shannon frameworks that go beyond this foundation.

Post-Shannon communications explore three directions. Semantic and goal-oriented transmission, where only meaningful or actionable information is communicated, aligns with the Decadal Plan's concept of 100,000:1 data-to-information compression and is addressed in depth in following section. Quantum communication — including Quantum Key Distribution (QKD) and entanglement-assisted networking — provides both security and reduced signalling energy and is developed further in chapter 6.2. Analog and biological signalling channels, such as chemical or optical gradients in intra-body or environmental networks, are the natural communication substrate for the biological computing paradigm described above. By co-designing communication and computation under unified thermodynamic and informational principles, post-Shannon architectures can substantially reduce global data movement energy.

Semantic Communications

Shannon explicitly chose to address only the technical problem of communication — how to transmit bits reliably — setting aside the semantic problem of whether transmitted symbols convey meaning, and the effectiveness problem of whether communication achieves its intended purpose. With the rise of machine-to-machine communication and the stringent bandwidth and latency demands of emerging applications, the time to address those deferred problems has arrived.

Goal-oriented communication frameworks transmit only what is task-relevant, requiring fundamental extensions to classical notions of entropy and mutual information to incorporate semantic content. The full realisation of semantic communication across all layers of a network could have deep transformational effects on smart infrastructure operation. The central unsolved challenge is bootstrapping a shared knowledge base across communicating parties — a problem that is as much about distributed knowledge representation as it is about protocol design, and one that current research has not yet addressed with the multidisciplinary depth it requires, particularly for application contexts such as XR and multi-sensory communications.

6.1.4 Conclusions

Extending computing beyond CMOS and digital logic is not an option but an imperative. Analogue, quantum, and biological paradigms are complementary responses to the dual crises of energy and information scalability, each operating closer to fundamental physical and thermodynamic efficiency limits than conventional digital architectures can approach. By investing in multi-paradigm computing and post-Shannon communication infrastructures, Europe can build a computing continuum that mirrors the resilience and energy proportionality of natural systems — paving the way toward sustainable intelligence for 2030 and beyond.

6.2 Quantum technologies

Quantum computing **Error! Reference source not found.**, **Error! Reference source not found.**, introduces a computational fabric based on qubits, whose states are described by probability amplitudes rather than binary logic. Quantum superposition and entanglement allow certain classes of problems (i.e., molecular modelling, optimization, and cryptography) to be solved exponentially faster than by digital machines. However, the true value of quantum computing within the SRIA 2030+ vision lies not only in speed, but in energy-information density: a single cryogenic quantum processor could replace thousands of high-power digital clusters for specific workloads. European programs like Quantum Flagship or OpenSuperQ demonstrate this shift toward hybrid quantum-classical platforms, where energy efficiency, scalability, and fault-tolerance remain the key research frontiers [C6-07]. Emerging *post-digital hybrids* integrate quantum accelerators with neuromorphic and analogue units **Error! Reference source not found.**, forming energy-heterogeneous compute fabrics optimized for task-based orchestration.

The next generation of digital transformation will be driven by the convergence of multiple technological domains with quantum technologies. Together with advances in AI, photonics, sensing, space communications technologies, etc. they are expected to reshape the foundations of global digital infrastructure over the coming decades. Managing infrastructures composed of classical networks, AI systems, quantum devices, cloud platforms, edge nodes, satellites, and sensing systems requires entirely new operational approaches **Error! Reference source not found.**, **Error! Reference source not found.**.

6.2.1 Quantum Technologies and the Quantum Internet

Quantum technologies will become one of the most transformative technological fields of the coming decades. Their importance extends far beyond quantum computing itself and increasingly includes communications, networking, sensing, and cybersecurity applications. The first large-scale deployments are focusing on secure communications through Quantum Key Distribution (QKD). These systems are operating over point-to-point fibre links or trusted-node infrastructures and post-quantum cryptographic mechanisms designed to protect infrastructures against future quantum-computing threats. The longer-term objective is the evolution toward fully entanglement-based quantum networks using quantum repeaters and entanglement swapping to eliminate the need for trusted intermediary nodes and enable large-scale quantum networking applications.

The longer-term vision, however, is considerably more ambitious. Future quantum networks may evolve toward the creation of a Quantum Internet capable of distributing quantum entanglement across geographically separated infrastructures. In contrast to today's internet, where digital packets are exchanged between devices, quantum networks may exchange quantum states themselves, enabling entirely new forms of communication and computation **Error! Reference source not found.**.

This evolution could support ultra-secure communications, distributed quantum computing, blind quantum computing services, ultra-precise synchronisation, and collaborative sensing applications. Entanglement may become a core network resource managed dynamically across infrastructures in much the same way that bandwidth and routing are managed today.

Achieving this vision requires substantial technological advances in enabling hardware systems such as quantum repeaters, quantum memories, entangled-photon sources, superconducting devices, and photonic integrated circuits. These technologies aim to overcome current limitations related to communication distance, signal fidelity, scalability, and interoperability. Future quantum infrastructures may gradually transition from isolated experimental links toward metropolitan, and eventually global quantum networking systems integrated with terrestrial and satellite communication platforms.

Beyond hardware limitations, future quantum networks will also require entirely new networking protocols and control architectures. Because quantum information cannot be copied or routed using conventional networking principles, future quantum communication systems will require dedicated quantum protocol stacks, entanglement-routing mechanisms, and quantum-aware network management frameworks.

6.2.2 AI-Native and Quantum-Enabled Network Infrastructure

AI-native networks are expected to evolve alongside quantum computing as complementary technologies for future communication infrastructures. As 6G and post-6G systems become increasingly autonomous and complex, AI will provide the intelligence required for real-time orchestration, optimisation, and management of highly distributed networks. However, many of these optimisation tasks, such as traffic engineering, spectrum allocation, resilience management, large-scale resource scheduling, etc. involve computational problems that may become too complex for classical computing alone.

Quantum computing may therefore become a potential accelerator for specific network optimisation and simulation tasks. Future infrastructures may rely on hybrid quantum-classical architectures in which AI systems manage overall operational intelligence while quantum processors accelerate selected computational workloads. This convergence could support more efficient network management, large-scale digital twins, advanced cybersecurity, and distributed cloud-edge optimisation, while also enabling future “Quantum Computing as a Service” capabilities integrated directly into communication infrastructures.

Quantum Machine Learning (QML) may further strengthen the convergence between AI and quantum technologies. Hybrid quantum-classical learning systems could accelerate selected machine learning and optimisation tasks relevant to future communication networks, including traffic prediction, anomaly detection, network digital twins, and distributed resource orchestration. Although practical quantum advantage remains uncertain for many AI applications, QML is increasingly viewed as a promising long-term research direction for highly complex network environments **Error! Reference source not found.**, **Error! Reference source not found.**, **Error! Reference source not found.**.

6.2.3 Hybrid Quantum-Classical Computing

The future of computing will evolve toward highly distributed and heterogeneous architectures combining classical processors, AI accelerators, high-performance computing platforms, edge systems, and quantum processing units. Rather than replacing classical computing, quantum technologies are more likely to augment existing infrastructures by accelerating highly specialised computational tasks.

Hybrid quantum-classical systems may become particularly valuable for solving optimisation problems, simulations, cryptographic operations, and AI workloads that are computationally difficult for traditional architectures. Quantum processors could operate as specialised accelerators integrated into broader computing infrastructures managed through cloud and edge orchestration platforms.

This convergence may also redefine how computing resources are delivered and consumed. Future networks could expose “Quantum Computing as a Service” capabilities in which users access remote quantum resources dynamically through cloud-based infrastructures. Distributed quantum computing architectures may eventually allow multiple quantum processors to cooperate across networked environments through entanglement-based communication systems.

At the same time, distributed computing may shift away from centralised hyperscale data centres toward highly decentralised infrastructures. Processing tasks may increasingly be distributed dynamically across edge nodes, cloud systems, mobile devices, satellites, and quantum accelerators according to latency requirements, energy constraints, and workload characteristics.

6.2.4 Advanced Photonics and Optical Technologies

Advanced photonic systems offer significant advantages compared with conventional electronic approaches, particularly in terms of bandwidth, signal integrity, latency, and power efficiency. As a result, silicon photonics, optical interconnects, integrated quantum photonics, and optical switching technologies will play central roles in future communication networks, data centres, AI infrastructures, and quantum systems.

Photonics are becoming critical for scaling quantum technologies themselves. Future quantum communication systems rely heavily on photonic devices capable of generating, manipulating, transmitting, and detecting quantum states with extremely high precision. Integrated photonic circuits may eventually provide compact, scalable, and energy-efficient platforms capable of supporting both classical optical communications and quantum networking functions within unified architectures.

The convergence between photonics, microelectronics, and quantum systems may therefore become one of the most strategically important industrial and research domains of future digital infrastructure development.

6.2.5 Quantum Sensing and Precision Technologies

Quantum sensing technologies may deliver major advances in measurement precision, timing, and environmental awareness. By exploiting quantum mechanical effects, future sensing systems may achieve levels of sensitivity and accuracy that exceed current classical technologies.

One of the most important applications lies in ultra-precise timing and synchronisation. Quantum-enhanced synchronisation systems may support future communication infrastructures, financial systems, autonomous transportation networks, and industrial automation platforms that require very accurate coordination across distributed systems. Future quantum sensor networks may also become part of 6G infrastructures, enabling collaborative sensing across highly distributed environments.

Quantum sensing may further support navigation systems that operate independently of traditional GPS or GNSS infrastructures. Such capabilities are strategically important for resilience, security, and defence. Future infrastructures may therefore combine satellite-based navigation with terrestrial quantum positioning approaches to improve robustness and operational continuity.

Another emerging frontier is the integration of communication and sensing into unified infrastructures. Future wireless systems may simultaneously support data transmission, localisation, environmental monitoring, and electromagnetic sensing. In this context, Rydberg atom-based receivers illustrate how quantum sensing concepts could contribute to highly sensitive radio-frequency detection while also supporting communication-related functions. Their significance lies not only in their sensing performance, but also in the way they point toward future systems where sensing and communication become more tightly coupled **Error! Reference source not found.**, **Error! Reference source not found.**, **Error! Reference source not found.**.

The broader challenge is to translate these promising quantum sensing concepts into reliable capabilities for real infrastructures. This requires research not only on the sensing principles themselves, but also on how such technologies can fit into future communication, navigation, and distributed intelligence systems.

6.2.6 Space-Based and Hybrid Infrastructures

Future hybrid terrestrial-space architectures may support intercontinental quantum communication links, enabling secure global connectivity, resilient synchronisation services, and the integration of terrestrial and satellite-based quantum networks.

Optical inter-satellite communications and future space-based quantum networking systems may also support secure global-scale communication infrastructures capable of operating independently of vulnerable terrestrial pathways **Error! Reference source not found.**. The integration of terrestrial and space assets is strategically important for technological sovereignty and resilience.

6.2.7 Standardisation, Security, and Technological Sovereignty

Standardisation will be particularly critical for emerging fields such as quantum networking, AI orchestration, photonic systems, and hybrid cloud-edge infrastructures. Without common interoperability frameworks, the large-scale deployment of quantum and hybrid-quantum technologies may become fragmented and economically inefficient.

At the same time future infrastructures combining quantum systems, AI, distributed computing, and autonomous operations will introduce entirely new categories of vulnerabilities, including side-channel attacks, quantum-era cryptographic threats, and AI manipulation risks.

In the context of quantum-enabled telecommunication systems, these developments are closely linked to broader geopolitical and technological sovereignty concerns. Control over quantum-aware hardware, photonic components, semiconductors, and secure communication infrastructures is increasingly associated with resilience, security, and strategic autonomy. Europe should therefore strengthen investments in quantum-aware supply chains, strategic quantum research programmes, and advanced manufacturing ecosystems to reduce technological dependencies and reinforce industrial competitiveness in future communication infrastructures.

6.2.8 Conclusion

Quantum technologies are expected to play an increasingly important role in the evolution of future telecommunication networks. The convergence of quantum communications, AI-native network management, advanced photonics, distributed computing, and terrestrial-space infrastructures will redefine how future communication systems are secured, operated, and integrated.

Future quantum-enabled telecommunication networks may evolve beyond conventional data transport infrastructures toward intelligent and adaptive platforms capable of supporting quantum-safe communications, entanglement distribution, distributed quantum computing services, ultra-precise synchronisation, and advanced sensing capabilities. In this context, quantum technologies will increasingly interact with AI-driven orchestration systems, cloud-edge infrastructures, and photonic communication platforms within highly distributed 6G and post-6G ecosystems.

Although many enabling technologies, including quantum repeaters, quantum memories, photonic components, and large-scale quantum networking architectures, remain at early stages of maturity, their progressive integration into telecommunication systems will drive a major transformation of future digital infrastructures over the coming decades.

6.3 Non-Terrestrial Networks

Non-Terrestrial Networks (NTN), encompassing Low Earth Orbit (LEO) and Geostationary (GEO) satellites, High Altitude Platform Stations (HAPS), and unmanned aerial vehicles, extend connectivity to underserved regions, provide resilient coverage for critical services, and introduce new capabilities in sensing, positioning, and distributed computing that terrestrial networks alone cannot deliver. Six research aspects define the long-term NTN research needs. They address the intelligence, integration, security, and efficiency challenges that must be solved before NTN can function as a trusted, sovereign, and scalable component of future 6G infrastructure. The aspects are deeply interconnected: AI underpins nearly all of them;

security cuts across every layer; and the compute–communication co-design thread runs from the physical layer through to cloud orchestration.

6.3.1 Integrated Infrastructure: ISAC + PNT + SatCom

Converged ISAC, PNT (Positioning, Navigation, and Timing), and SatCom improve coverage, automation, resilience, and European sovereignty, requiring cross-layer waveform design, synchronisation, and data fusion under high Doppler and heterogeneous geometries.

NTN geometries increase Line-of-Sight diversity for positioning and sensing. OTFS may outperform OFDM in high-Doppler environments, while beam-overlap interference and PRS multiplexing become key optimisation challenges. Towards trustworthy PNT and resilient timing, hybrid NTN–terrestrial integrity mechanisms and distributed synchronisation are needed for graceful degradation and trusted fallback, building on 3GPP integrity metrics and cross-domain validation.

UEs can act as distributed sensors reporting anomalies to the network, enabling threat localisation and integrity alerts through data fusion, enabling cooperative threat detection at scale. Key challenges include multi-antenna processing, low-cost monitoring, and AI-assisted interference classification.

Integrating PNT into future NTN architectures reduces reliance on externally owned GNSS. Resilience requires integrity-aware, network-assisted architectures combining authenticated assistance data, satellite health monitoring, receiver integrity checks, and terrestrial cross-validation.

NTN-ISAC is a joint problem spanning waveform design, spatial geometry, networking, and intelligence. Key enablers include spectrum and hardware sharing, 3D integration, mobility-aware sensing, and sensing-assisted control. Major challenges include 3D beamforming complexity, synchronisation, ambiguity suppression, and communication–sensing trade-offs.

6.3.2 Semantic Communications for NTN

Semantic communications (introduced in Section 0) are a natural fit for NTN environments, where bandwidth is scarce, links are highly variable, and many services are inherently task-oriented. The NTN-specific challenge is adapting generic semantic architectures to conditions that existing systems handle poorly: high Doppler, large propagation delays, rapidly switching beams, and strict onboard Size, Weight, and Power (SWaP) constraints that make standard deep learning inference impractical. Privacy and security risks introduced by semantic encoders are also unresolved in NTN contexts, where wide beam footprints increase exposure to adversarial interception.

The NTN R&I agenda therefore focuses on what general semantic communication theory does not yet address: channel-adaptive architectures based on transfer and meta-learning for rapid adaptation across varying satellite link conditions; joint semantic–channel coding co-designed with delay-Doppler waveforms; lightweight and neuromorphic inference implementations for space-grade hardware; and secure semantic frameworks incorporating federated learning, differential privacy, adversarial robustness, and semantic watermarking under NTN-specific threat models.

6.3.3 AI for NTN

The scale, heterogeneity, and dynamism of future NTN deployments exceed the capabilities of rule-based control. AI and ML can exploit orbital dynamics, traffic evolution, and propagation variability to enable proactive, adaptive, and trustworthy operation across the physical layer, resource management, positioning, security, and orchestration.

Current research already spans these domains. Multi-agent reinforcement learning supports scheduling, beam management, and TN–NTN traffic offloading, leveraging orbital predictability for energy-efficient control. Deep learning mitigates NTN Doppler effects through improved channel estimation and equalisation, while strict onboard Size, Weight, and Power (SWaP) constraints drive research into edge AI and neuromorphic computing. AI also enables

GNSS–5G–NTN data fusion, anomaly and spoofing detection, and XAI-based auditability for safety-critical systems.

However, major gaps remain. Real NTN datasets and open benchmarks are scarce, limiting model generalisation. Multi-agent RRM has yet to demonstrate stable operation under dynamic topologies and asymmetric NTN latencies. XAI for NTN remains immature, with no certification framework for safety-critical multi-agent decisions, while AI-driven security lacks integrated solutions for moving jammers, GNSS spoofing, and constrained onboard authentication.

The research challenges include distributed learning for RRM and orbital prediction, AI-assisted multiple access, lightweight neuromorphic architectures, XAI and open benchmarking aligned with 3GPP, and proactive end-to-end security frameworks for interference classification and signal authentication.

6.3.4 End-to-End Security for Unified TN/NTN Networks

Unified 6G-NTN security is challenging by the coexistence of three distinct frameworks: 3GPP procedures for 5G/6G access, IETF protocols (IPsec, TLS) for networking, and CCSDS standards (including the Space Data Link Security protocol) for TT&C links. Space-time-variant NTN topology means security contexts cannot be kept consistent over time. A coherent framework must align with best practices from all three bodies while addressing gaps at the protocol and architecture level. The ongoing evolution to post-quantum cryptography (PQC) adds urgency, requiring integration of new primitives into existing protocols. The research challenges include extending 3GPP and IETF security frameworks to address NTN-specific challenges, as well as the integration of post-quantum cryptography at both link and service levels.

6.3.5 Cross-Layer Compute–Communication Optimisation

System-level analysis based on 3D stochastic geometry is needed to model multi-layer TN–NTN coexistence, quantify interference and coordination gains, and provide rigorous foundations for future 6G NTN architecture design. Future 6G systems will integrate 2D terrestrial cellular networks with 3D NTN infrastructures, leading to a “network of networks” composed of heterogeneous LEO, MEO, GEO, HAPS, UAV, and multi-tier terrestrial systems sharing spectrum and spatial resources. This integration is further complicated by the introduction of RIS-enhanced environments, ISAC/ICAS functionalities, and increasingly dynamic resource orchestration mechanisms.

In such settings, stochastic geometry emerges as one of the few tractable analytical frameworks capable of evaluating and predicting the statistical properties of large-scale 3D NTN/TN systems. Beyond classical interference analysis, these models can capture the organised interplay between terrestrial and non-terrestrial infrastructures, including traffic offloading, coordinated resource sharing, and distributed computing interactions. Such a framework is essential not only for performance evaluation, but also for the design of optimal architectures, resource allocation strategies, and techno-economic models for future NTN-enabled 6G systems.

A major challenge specific to NTN environments is the high mobility of network elements, including satellites, UAVs, and HAPS. This mobility impacts multiple layers simultaneously, including Doppler-sensitive physical-layer communications, association and handover dynamics, routing stability, distributed AI coordination, and cross-layer orchestration. Research is therefore needed on dynamic stochastic geometry and mobility-aware system models capable of jointly analysing communication, sensing, computing, and networking behaviour in highly dynamic 3D architectures.

6.3.6 Distributed and Federated Processing

LEO constellations are inherently distributed systems, and this distribution is increasingly a structural asset rather than a constraint to be managed. Satellite Mobile Edge Computing (SMEC) transforms satellites into active compute nodes, enabling onboard analytics and autonomous decisions that are infeasible if all processing is deferred to the ground. As sensing capabilities grow, satellites generate massive volumes of heterogeneous data — imagery, continuous sensor streams — that must be processed efficiently to extract timely, actionable insights without overloading communication links.

Distributed and federated learning naturally align with this paradigm: data and computation are spread across the constellation, enabling collaborative model training without centralising data. Orbital dynamics can be exploited to enable cyclical training processes as satellites sequentially contribute during visibility windows, turning mobility into a scalability enabler. Split learning further reduces communication overhead by exchanging intermediate representations rather than raw data.

The research challenges include distributed AI architectures tailored to the partial and locally constrained views of satellite constellations, enabling coordinated learning and inference with minimal inter-satellite communication. Strict onboard power limitations also require energy-efficient AI through model compression, edge inference, and runtime optimisation, closely aligned with edge AI and neuromorphic computing efforts. In parallel, future satellites must evolve toward autonomous and adaptive onboard operation, using online and incremental learning to support real-time, context-aware decision-making with reduced dependence on ground infrastructure.

6.3.7 Quantum Technologies

While the 6G standardization is in progress, it is relevant to ask: what are the contours of the wireless generation that will be deployed in the 2040s, after 6G? Instead of linearly extrapolating the enhanced wireless systems into “7G”, a viable option is that the generation beyond 6G should start to integrate elements of quantum communication, computation, and sensing. It is therefore termed 1Q: first quantum wireless generation, generalizing the traditional concept of cellular communication to include quantum information processing. Considering the fact that communication networks are the key ingredient of European technological sovereignty, as well as the fact that the impact of quantum information processing can have on security and computation, it is the right time to start to build a large private-public partnership that will define the quantum-classical networks that will evolve within the next decade. The development of quantum information processing is characterized by both optimism and anxiety, but most of all by uncertainty. Yet, Europe must not wait for quantum technology and especially quantum hardware to mature. Instead, it must use architectural design as an accelerator, creating demand for quantum hardware, integrating even small and near-term quantum devices into a scalable network vision.

The essence of 1Q is that it expands the notion of wireless coverage to the quantum realm. A classical Base Station (BS) offers classical wireless communication through radio waves within a given geographical area denoted as cell. A Quantum Base Station (QBS) uses Free Space Optical (FSO) transmission to distribute entanglement within a quantum cell. As shown by quantum teleportation, the distribution of entanglement plus classical communication is sufficient to support transfer of quantum information. The QBS acts as a quantum orchestrator: generating and distributing entangled pairs to Quantum User Equipment (QUE), enabling teleportation-based communication, performing entanglement swapping for long-range quantum links, and supporting distributed quantum computation within its quantum cell. Satellites play a critical role, as FSO links from Low Earth Orbit can distribute entanglement over hundreds of kilometres and 1Q can advantageously use the fact that satellites will already be part of the 6G deployments.

6.4 Drone Corridors and 3D Radio Connectivity

Uncrewed aerial vehicles (UAVs), commonly known as drones, are poised to drive significant economic growth and societal change. Due to their low cost and high mobility, UAVs are expected to become crucial for goods delivery, surveillance, search and rescue, and monitoring wildfires, crowds, and assets. As urbanization strains ground transportation, electrical vertical take-off and landing vehicles (eVTOLs) could revolutionize urban mobility by serving as air taxis or ambulances, enhancing the speed, safety, and interconnectedness of transportation systems. Autonomous levitating pods, once considered science fiction, are now being tested and could transform commuting, influencing where we live and work [C6-22], [C6-23].

Use case and technology requirements: As UAVs become more widespread, their flights may be restricted to specific aerial corridors defined by air traffic regulation authorities. Within these corridors, UAVs will need to exchange vast amounts of real-time data with the network, necessitating ultra-reliable wireless connectivity [C10-60]. This connectivity must support safe UAV operations through low-latency control and mission-specific data, encouraging legislators to relax regulations on civilian pilotless flights and paving the way for autonomous UAVs and related vertical markets.

Research challenges: Traditional cellular base stations are optimized for 2D ground connectivity, which limits UAV reach to upper antenna sidelobes and causes signal fluctuations with UAV movement. UAVs flying above buildings may also encounter line-of-sight interference from multiple cells, complicating the decoding of critical command and control messages. Achieving 3D connectivity will require re-engineering ground-based deployments. Traditional trial-and-error cell shaping is not scalable, prompting the need for automated optimization techniques using data-driven models. Alternatively, UAV connectivity could be supported by non-terrestrial networks (NTNs), such as Low Earth Orbit (LEO) satellite constellations, with terrestrial network operators leasing spectrum and infrastructure from NTN providers.

6.5 High Frequency Integrated Photonics and Communication Systems

Background: The telecommunications sector is shifting its focus towards next-generation networks, namely B5G/6G. A key technological advancement in this area is the transition to the upper millimetre wave (100-300 GHz) and THz (0.1 -10 THz) spectrum. The larger bandwidth and shorter wavelengths offer the benefits of higher data rates and greater miniaturisation. This also allows for denser antenna arrays on integrated platforms, with the potential for integrated directional beamforming.

Integrated Sensing and Communications (ISAC): Concurrently, the high photon energies of THz radiation can enhance the resolution and accuracy of radar and imaging systems, as well as enable new sensing capabilities. Notably, specific frequencies in the THz band are significantly affected by molecular absorption. While traditional communication systems have avoided these frequencies, a shift in perspective reveals that the inherent characteristics of the THz range make this spectrum ideal for joint communication and sensing (JCAS) systems.

Challenges: Communications in the THz band face some remarkable challenges. The very high attenuation significantly limits the coverage, implying that new techniques must be developed for coverage extension and energy efficiency enhancement. New waveforms with multiple numerologies must be developed to achieve full integration of the sensing and communication subsystems. At the same time, receiver impairments such as phase noise play an important role in THz frequencies and new approaches to minimize their effect must be investigated. At present, there is a scarcity of integrated systems capable of operating in frequencies up to THz for both communication and sensing applications. High frequency components face technological constraints and integration challenges within photonic circuits. These limitations include high power consumption, complex system design, and difficulties in achieving high efficiency and miniaturisation for THz devices. Among these, the most significant is the effectiveness and efficiency of current THz technologies, which impede their practical application and widespread use. Firstly, the demand for ultrafast broadband connectivity and precise environmental monitoring in various applications (such as safety, industrial control, and infrastructure monitoring) cannot be fulfilled with existing solutions. Secondly, the absence of miniaturised, integrated devices that can operate efficiently at such high frequencies restricts the potential for new applications and technologies that could leverage its capabilities. However, by combining advancements in high frequency components with photonic circuit integration, these limitations could be overcome, leading to the creation of compact, accurate, and cost-effective devices for mm and THz communication and sensing. With complementary technologies operating at these high frequencies, high frequency

Integrated Photonics holds the promise of revolutionising ultrafast broadband connectivity and real-time environmental monitoring, thereby unlocking new applications and enhancing existing ones.

6.5.1 Analog Signal Photonic Processing

Analogue signal photonic processing is a key enabling technology for future communication, sensing, and computing systems, allowing RF, microwave, millimetre-wave, and optical signals to be processed directly in the photonic domain. Functions such as frequency conversion, wideband filtering, beamforming, equalisation, interference cancellation, and signal classification can be performed with significantly lower latency and energy consumption than conventional electronic processing while supporting ultra-wide bandwidths. These capabilities are relevant to 6G networks, satellite and space communications, radar systems, fibre-wireless convergence, quantum technologies, and emerging photonic AI platforms.

Applications include photonic front-ends for communication networks and datacentres, wideband sensing and electronic warfare systems, simplified satellite payloads, advanced industrial monitoring, and low-noise interfaces for quantum sensing and computing. As a transversal technology, analogue photonic processing also supports advances in integrated photonics, heterogeneous packaging, and photonic-electronic co-design.

Key challenges include thermal stability, large-scale integration and packaging, limited design and modelling tools, and achieving sufficient linearity, noise performance, and dynamic range for RF-optical conversion. Reliability and radiation tolerance remain important requirements for aerospace and defence applications. Representative targets include instantaneous bandwidths above 100 GHz, sub-picosecond processing latency, dynamic ranges exceeding 100 dB/Hz, and energy consumption below 1 pJ per operation.

Research priorities include low-loss and programmable photonic platforms, high-linearity modulators and detectors, photonic-electronic co-design methodologies, advanced packaging technologies, integrated stabilisation mechanisms, and radiation-tolerant photonic subsystems. Current technologies are typically at TRL 3–5, with progression towards TRL 6–7 by 2030 and wider deployment by 2035.

Analogue photonic processing can contribute significantly to European objectives in energy-efficient computing and communications, technological sovereignty, secure connectivity, advanced space infrastructures, and future AI and quantum systems.

6.5.2 Photonic-based AI Interference

Photonic-based AI inference enables the execution of neural-network operations directly in the optical domain, offering ultra-low latency and significantly reduced energy consumption compared with conventional electronic accelerators. By exploiting the inherent parallelism of photonic systems, it can support real-time perception, classification, and decision-making in edge devices, autonomous systems, aerospace platforms, medical instruments, robotics, and other resource-constrained environments.

Applications include distributed AI in 6G and edge networks, onboard processing for satellites and aerospace systems, low-power medical diagnostics, autonomous mobility, and security and defence systems. By reducing dependence on cloud resources and power-hungry processors, photonic inference can enable advanced AI capabilities in environments with stringent latency, energy, or reliability requirements.

Key challenges include limited precision and noise accumulation in photonic computations, the lack of compact optical nonlinearities and memory elements, scalability of programmable photonic architectures, and integration with electronic control and memory systems. Additional challenges relate to calibration, verification, and certification, particularly for safety-critical applications. Representative targets include energy consumption below 100 fJ/MAC, sub-nanosecond inference latency, and compute efficiencies exceeding 100 TOPS/W.

Research priorities include scalable photonic computing architectures, hybrid photonic-electronic integration, optical nonlinearities and memory technologies, low-power stabilisation mechanisms, and methodologies for validation and certification. Current solutions are typically at TRL 2–4, with progression towards TRL 5–6 by 2030 and TRL 7–8 by 2035.

Photonic AI inference supports European objectives in energy-efficient computing, technological sovereignty, and trustworthy AI. It offers an alternative computing paradigm for future AI infrastructures while contributing to strategic priorities in communications, space, defence, healthcare, and secure digital systems.

6.5.3 Spectrum-Aware Energy Autonomy for Near-Zero-Power Communication

The increasing densification of wireless networks is creating growing RF spectrum congestion and generates significant amounts of ambient electromagnetic energy that can be harvested and reused. Energy present in non-allocated bands, idle channels, guard bands, and neighbouring communications can be converted into electrical power, enabling reduced energy consumption and potentially near-zero-power operation of communication devices.

Realising this vision requires spectrum-awareness capabilities that can identify spatial, spectral, and temporal opportunities for energy harvesting in highly dynamic environments. Combined with integrated RF energy harvesting, such capabilities can support autonomous operation of electronic and photonic subsystems in 5G/6G networks, fibre-wireless systems, satellite terminals, industrial IoT deployments, and spectrum-monitoring infrastructures. At the network level, spectrum-aware energy autonomy can improve resource utilisation, reduce maintenance requirements, and enable self-powered communication platforms.

Key challenges include improving RF-to-DC conversion efficiency, developing ultra-low-power energy management circuits, integrating harvesting functions without degrading communication performance, and advancing broadband spectrum-awareness techniques for reliable energy discovery. Research is also needed on long-term reliability, resilience, and autonomous operation of RF-powered communication systems.

6.6 Digital Twinning

Digital twinning is emerging as a foundational capability for future communication networks, evolving from static modelling tools into continuously synchronised cyber-physical systems. In the context of future 6G and post-6G infrastructures, digital twins are expected to become integral to network design, operation, experimentation, and lifecycle management, particularly as networks expand to include AI-native control systems, distributed cloud-edge resources, and emerging quantum and photonic technologies.

6.6.1 Digital Twin Applied in Communication Networks

Digital twin technology is increasingly recognised as a core enabler of future intelligent network infrastructures. It provides a virtual representation of physical communication systems and processes, continuously updated through real-time and historical data streams, with the goal of supporting prediction, optimisation, and autonomous decision-making. In future telecommunication systems, digital twins will evolve beyond traditional modelling approaches by incorporating AI-native analytics, real-time telemetry, and potentially quantum-enhanced simulation capabilities for highly complex optimisation problems. In this context, digital twins are no longer limited to monitoring and visualisation but become active components of network intelligence. They enable closed-loop interaction between physical and virtual systems, supporting adaptive control of communication networks, resource allocation, resilience management, and service assurance. This evolution is closely linked to the development of autonomous 6G infrastructures, where networks must continuously self-optimise across highly dynamic and heterogeneous environments.

Research challenges include:

- High-fidelity virtual representation of networked systems
- Real-time synchronisation and multi-modal data integration
- AI-native and predictive modelling capabilities
- Support for distributed and large-scale system dynamics

- Integration with emerging computing paradigms

6.6.2 Digital Twinning for Network Innovation and Experimentation

Digital twinning is increasingly used as a key methodology for accelerating innovation in communication networks by enabling large-scale virtual experimentation environments. These environments provide high-fidelity replicas of complex network systems, allowing researchers and operators to test architectures, protocols, and services in a controlled yet realistic digital space.

In future network ecosystems, digital twins will evolve into comprehensive experimentation platforms that integrate real-time operational data with large-scale simulation capabilities. These platforms will support continuous validation of network functions, service orchestration strategies, and infrastructure evolution scenarios, significantly reducing reliance on costly and time-consuming physical testbeds. A key advantage of this approach lies in its scalability and flexibility. Digital twin-based environments can be extended dynamically, allowing new technologies, network components, and service models to be incorporated without disrupting existing experimental setups.

Future digital twin platforms will also play an important role in closed-loop network optimisation, enabling continuous feedback between operational networks and their virtual counterparts. This supports predictive maintenance, resilience planning, energy optimisation, and automated service assurance, contributing to the broader evolution toward autonomous communication infrastructures **Error! Reference source not found.**.

Research challenges include:

- Scalable emulation of heterogeneous network ecosystems
- Seamless integration with physical testbeds and operational networks
- Closed-loop optimisation and autonomous control
- High-fidelity scenario generation and reproducibility
- Support for future network paradigms

Overall, digital twinning is expected to evolve from a supporting engineering tool into a central pillar of future communication network design, operation, and innovation, enabling faster, safer, and more scalable experimentation for next-generation digital infrastructures.

6.7 Cyber-Physical Intelligence

This section sets out a unified research vision across three interconnected themes. The Digital Continuum transforms communication networks into closed-loop cyber-physical systems, jointly orchestrating sensing, computing, and communication as a single programmable fabric. The Electromagnetic Continuum engineers the physical layer of that fabric, turning programmable EM fields and surfaces into shared substrates for communication, sensing, and computation. Physical and Embodied AI connects both to the real world, addressing autonomous agents that perceive, reason, and act under physical constraints. Together, the three agendas converge on one long-term ambition: networks that do not merely transport data, but perceive, infer, and act — closing the loop between the digital and physical worlds.

6.7.1 The Digital Continuum

Future communication systems are increasingly shaped by the integration of communication, computation, and storage resources. Emerging architectures—spanning terrestrial and non-terrestrial networks—are moving toward systems where distributed computing capabilities, advanced signal processing, and intelligent resource orchestration jointly determine performance.

Network nodes will progressively play a dual role: not only communicating data but also supporting distributed computation. This evolution is particularly visible in edge and cloud–

edge systems, where computation is increasingly coordinated with communication rather than treated as a separate layer.

At the physical layer, this evolution opens new research directions, including communication schemes that exploit computational capabilities at terminals and network edges, advanced multi-antenna and multi-user transmission techniques, and mechanisms that jointly optimize spectral efficiency, processing load, and energy consumption. In this perspective, terminals are no longer passive receivers but active participants in distributed communication-and-computation processes.

Future networks must evolve from data transport systems into distributed intelligence platforms — geo-distributed cyber-physical systems where sensing, communication, computing, and storage are jointly programmable resources of a single continuum-scale fabric. The network becomes a closed-loop control system: it senses the physical world, communicates state, computes decisions, and acts — all under constraints of latency, energy, sovereignty, and first order environmental effects.

Goal-driven orchestration: Classical KPI-centric design (throughput, latency, packet loss) is insufficient for cyber-physical applications. The Digital Continuum replaces it with intent-driven orchestration: high-level objectives are dynamically translated into coordinated resource allocation across sensing, communication, and computing assets, adapting continuously as conditions change.

Compute–communication co-design: In distributed AI systems, communication is the dominant bottleneck. Information flow must be determined by the structure of computation, not network topology. This motivates task-oriented communication, over-the-air computation, and compute-aware protocols that compress or aggregate data based on its semantic role rather than its header fields.

Network sensing plane: Sensing becomes a first-class network primitive. Beyond telemetry, the network continuously infers physical environment state, mobility patterns, propagation conditions, blockage events, enabling the construction of live digital twins and supporting predictive, proactive orchestration rather than reactive control.

Agentic AI-native control: Centralised control cannot scale to continental networks with sub-millisecond timing requirements. A hierarchy of cooperating agents across device, edge, and cloud replaces it, operating under partial observability, coordinating through negotiation and policy constraints, and optimising collective objectives without any single point of authority.

Lifelong and federated learning: Communication environments are intrinsically non-stationary. Models and policies must adapt continuously as conditions drift, without forgetting prior knowledge, and without centralising data across administrative boundaries. Federated learning under privacy and heterogeneity constraints is the enabling framework.

Volatility-aware and risk-sensitive control: Tail events, rare but consequential failures caused by mobility bursts, interference, or workload spikes, determine whether a system is fit for purpose. Risk-aware control minimises variance and bounds failure probability, integrating stochastic control, model predictive control, Lyapunov methods, and safe reinforcement learning into a coherent framework for autonomous network management.

6.7.2 The Electromagnetic Continuum

The EM Continuum reframes the wireless physical layer: rather than a noisy medium to be compensated for, it becomes a programmable fabric to be engineered. Driven by mmWave and sub-THz operation and the emergence of electrically large reconfigurable surfaces, future systems manipulate EM fields directly — jointly realising communication, sensing, and computation through wave control.

From channels to programmable EM fields: Classical channel abstractions are inadequate for near-field, large-aperture systems. Holographic MIMO (HMIMO) and Large Intelligent Surfaces (LIS) operate in a field-centric regime where beam steering gives way to beam focusing, far-field plane-wave assumptions give way to spherical wavefronts, and spatial degrees of freedom are determined by geometry and wavelength rather than fading statistics.

Fundamental limits and near-field modelling: A rigorous scientific foundation is needed: EM-consistent models for near-field propagation, mode coupling, and continuous apertures; capacity and spatial DoF limits derived from first principles; and tractable system-level approximations that bridge EM physics and network-level analysis. Random matrix and random field theory are key mathematical tools.

Wave-domain computing: Computation performed directly through wave propagation, such as linear transformations via interference, spatial filtering via diffraction, task-specific processing via engineered material responses, offers ultra-low-latency and ultra-low-energy processing at the physical layer. Programmable WDC (pWDC) extends this with runtime reconfiguration. Key challenges include limited expressivity of linear wave operations, model–reality mismatch, and scalable control of dense programmable structures.

Integrated sensing, communication, and computing (ISAC+): The same waveforms and surfaces used for data transmission can simultaneously probe the environment: estimating positions and dynamics, inferring propagation conditions, and building real-time spatial scene representations. This ISAC+ paradigm enables EM-aware digital twins — continuously updated propagation models that support predictive beamforming and context-aware resource allocation.

Programmable wave-matter systems: Dense meta surfaces, RIS architectures, multi-layer EM structures, and space-time modulated systems enable direct control of the EM transfer function of the medium itself. Engineering challenges include nonlinear wave behaviour, sensitivity to perturbations, and control algorithms that operate within channel coherence times.

6.7.3 Physical AI and Embodied Intelligence

Physical AI addresses autonomous agents that do not merely observe the world but act within it — navigating uncertain environments, manipulating objects, coordinating with other agents and humans, and doing so with the real-time computational and communication resources available at the network edge. It is the primary consumer of the goal-driven, latency-guaranteed, compute-aware services that the Digital and EM Continua aim to provide.

The embodiment gap: Recent AI progress has concentrated in symbolic and perceptual domains like language, images, strategy. The embodiment gap separates this from physical competence: real environments are continuous, partially observable, and subject to dynamics that resist precise modelling. Sensors are noisy; actuators are imprecise; timing is real. Bridging this gap requires agents whose internal representations are grounded in the causal structure of the physical world, not merely in statistical patterns of text or pixels.

World models and sim-to-real transfer: Training embodied agents in the physical world is slow and expensive. Simulation is necessary but introduces a reality gap: simulated environments are imperfect models, and agents trained in them often fail in deployment. Internal generative models of physical dynamics enable planning under uncertainty and provide a pathway to robust sim-to-real transfer. Building accurate, tractable world models with calibrated uncertainty over unobservable parameters remains a central open problem.

Multi-agent physical systems: Many applications involve populations of agents sharing a physical space. Their actions interact physically, and their communication is subject to the latency and bandwidth constraints of the underlying network. Coordination protocols must be scalable, safe, and robust to communication failures. The co-design of agent policies and network resource allocation is a direct connection point with the Digital Continuum.

Safety, robustness, and human–machine interaction: Physical AI systems operate alongside humans. Failure consequences are physical. Safety requires both formal verification, and empirical evaluation across the full distribution of deployment conditions. Agents must also be legible: their intentions and likely actions must be interpretable by human collaborators, and they must be capable of adjusting behaviour in response to human feedback and instruction.

6.7.4 Cross-Cutting Themes

The physical–digital loop: The three research areas described in this sub-section, address different segments of the same loop: EM systems perceive and communicate physical state; the Digital Continuum orchestrates resources and decisions; Physical AI agents act in the world; their actions are in turn perceived. Closing this loop with the latency, reliability, and accuracy required by real cyber-physical applications is the overarching engineering objective.

Shared mathematical foundations: Stochastic control and MDPs (Markov Decision Process), Gaussian processes and Bayesian non-parametric, distributed optimisation, and random matrix theory appear across all three themes. Investing in these shared foundations, enables methodological transfer and economies of scale.

Sustainability by design: Wave-domain computation reduces digital processing energy at the physical layer. Task-oriented communication transmits only semantically relevant information. Efficient physical agents achieve objectives with fewer actuations. Carbon-aware orchestration must span all three layers simultaneously.

Open interoperable infrastructure: The scale of this agenda requires shared infrastructure: open-source software stacks, common APIs for intent, context, policy, and resource control, and shared simulation environments. Open infrastructure is a first-class research output, and a prerequisite for ecosystem-wide progress.

6.8 Robotics, Autonomous Systems and Embodied AI

Future robotic systems will operate as networked entities rather than isolated machines, collaborating with humans, digital infrastructures, and other autonomous agents across physical and virtual environments enabled by advances in AI, distributed computing sensing technologies, and communication networks. Applications ranging from advanced manufacturing and logistics to healthcare, mobility, environmental monitoring, and critical infrastructure management will require communication systems that provide unprecedented levels of reliability, latency, scalability, security, and intelligence. Consequently, robotics provides a compelling framework for identifying future network requirements and validating emerging technologies such as AI-native networking, semantic communications, integrated communication-computing infrastructures, digital twins, and autonomous network management. In this context, robotics is a beneficiary of future network innovations and a key driver shaping the design, architecture, and capabilities of communication systems beyond 6G.

6.8.1 Networked Robotics and Robot Swarms

Future robotic ecosystems will comprise large numbers of interconnected autonomous agents operating cooperatively in dynamic environments. Such swarms may span terrestrial robots, drones, autonomous vehicles, sensors, and intelligent infrastructure elements, collectively executing tasks that exceed the capabilities of individual agents. These systems demand highly scalable communication architectures supporting real-time coordination, distributed decision-making, cooperative perception, and collective learning. Key open research areas include network-assisted swarm intelligence, resilient multi-agent communications, semantic information exchange, and distributed consensus mechanisms robust to highly dynamic conditions. The convergence of communication networks and swarm intelligence may ultimately enable large-scale autonomous systems operating across industrial facilities, smart cities, transportation networks, and environmental monitoring infrastructures **Error! Reference source not found.**.

6.8.2 Human-Robot Collaboration

Future robotic systems will increasingly operate in close collaboration with humans across industrial, healthcare, service, and domestic environments. Rather than replacing human operators, robots are expected to augment human capabilities through continuous interaction, shared situational awareness, and adaptive cooperation. Communication networks will be critical to enabling safe and trustworthy human-robot collaboration. Future systems must support multimodal interactions spanning speech, vision, gestures, haptic feedback, physiological signals, and contextual information. These requirements extend beyond traditional quality-of-service metrics toward human-centric measures such as trustworthiness, responsiveness, explainability, and safety. Key research directions include multimodal communication frameworks, privacy-preserving interaction mechanisms, context-aware networking, and architectures supporting cognitive collaboration between humans and autonomous systems **Error! Reference source not found.**.

6.8.3 Teleoperation and Tactile Internet

Teleoperated robotic systems are becoming important in sectors such as healthcare, industrial maintenance, hazardous-environment operations, offshore energy, and space exploration. These applications require human operators to interact remotely with physical systems while maintaining high levels of situational awareness and control fidelity. The concept of the Tactile Internet extends communication beyond audio and visual information to include real-time transmission of touch, force, and motion information. Such applications impose extremely demanding requirements on communication infrastructures, including ultra-low latency, deterministic behaviour, ultra-high reliability, and resilient operation under adverse conditions. Future research should address end-to-end latency optimisation, network resilience, edge-assisted control loops, predictive communication techniques, and integrated sensing and communication mechanisms supporting remote physical interaction. Despite its transformative potential, the Tactile Internet has seen limited real-world deployment since the concept was formalised in 2014. Progress has been constrained by the difficulty of simultaneously achieving the required ultra-low latency, ultra-high reliability, and deterministic behaviour across end-to-end network paths

6.8.4 Distributed Intelligence and Edge Robotics

The evolution of robotic systems is driving a transition from centralized cloud-based intelligence towards distributed intelligence operating across a continuum of devices, robots, edge platforms, and cloud infrastructures. Future robotic applications will increasingly rely on dynamic placement of sensing, perception, planning, and learning functions according to latency, reliability, energy-efficiency, and privacy requirements. Rather than treating communication networks as passive transport infrastructures, future systems will integrate communication, computing, storage, and artificial intelligence into a unified digital continuum. Key research challenges include distributed AI orchestration, collaborative edge-cloud learning, federated and continual learning, semantic information exchange, and autonomous management of heterogeneous computing resources. Such capabilities will be essential to support large-scale deployments of autonomous robots, collaborative robotic systems, and embodied AI operating in dynamic real-world environments.

6.8.5 Robotics Digital Twins

Robotics digital twins enable continuous synchronization between physical robots and their virtual counterparts. Beyond monitoring and simulation, future robot digital twins will support real-time control, predictive maintenance, autonomous adaptation, and AI-driven optimization throughout the robot lifecycle. Recent research points toward Robot Digital Twins (RDTs) as a distinct technological domain, integrating robotics, artificial intelligence, sensing, communication networks, and edge-cloud infrastructures. Future developments will increasingly focus on multi-robot digital twins, human–robot collaboration, generative AI-enabled simulation environments, and digital-twin ecosystems operating across the digital continuum. These trends will place significant demands on communication networks regarding latency, synchronization, reliability, scalability, and trustworthy data exchange **Error! Reference source not found.**.

6.9 AI-Native Autonomous Networks and Intent-Based Operations

6.9.1 Natural-Language Intent in Autonomous Networks

Future communication infrastructures are increasingly evolving toward AI-native and intent-driven operational models in which operators may interact with networks through high-level objectives expressed in natural language. Rather than manually configuring network functions and policies, future systems may increasingly rely on autonomous orchestration platforms capable of translating operational intent into dynamic network behaviour. This evolution introduces semantic reliability as a major technological challenge. Traditional telecommunication standards have historically depended on highly constrained and rigorously defined language to eliminate ambiguity and ensure deterministic interoperability across vendors, operators, and jurisdictions. In contrast, natural language is inherently contextual, ambiguous, and linguistically variable.

As autonomous networks begin interpreting human-defined operational intent directly, ambiguity may introduce operational instability, inconsistent behaviour across systems, unintended policy execution, and new cybersecurity risks. The challenge becomes particularly significant in multi-vendor and multilingual environments where equivalent intent expressions may not always produce identical interpretations or actions. Future intent-based networking systems will therefore likely require layered semantic architectures in which natural language serves primarily as the human interaction interface, while execution relies on formally constrained semantic representations, deterministic policy models, and machine-verifiable intent abstractions. In this context, AI systems may increasingly operate as semantic translation and validation layers between human operators and autonomous infrastructures.

The convergence of AI-native orchestration, autonomous operations, and intent-based networking may consequently require entirely new approaches to standardisation, semantic governance, explainability, verification, and operational trust within future 6G and post-6G communication systems.

6.9.2 From Linguistic Intent to Deterministic Network Execution

The challenge extends beyond linguistic ambiguity itself and increasingly concerns whether independently developed AI systems can produce sufficiently consistent operational outcomes from equivalent high-level intent expressions. Future autonomous networks may not require identical semantic interpretation internally, but they will require sufficiently deterministic behavioural outcomes to guarantee interoperability, stability, and operational trust.

This problem resembles long-standing challenges in distributed computing and protocol implementation. Even when systems comply with the same specification, implementation-specific optimisation strategies, internal decision processes, and vendor-specific algorithms may still produce divergent behaviours under identical operating conditions. Future AI-native network orchestration systems may amplify this effect significantly.

As a result, the objective of future intent-based networking may not be strict semantic equivalence across platforms, but bounded operational equivalence within constrained execution domains. The narrower and more formally constrained the operational intent space becomes, the greater the probability that heterogeneous AI systems can produce sufficiently aligned outcomes. This suggests that future autonomous telecommunication systems may increasingly rely on constrained semantic frameworks, domain-specific intent languages, formal policy abstractions, verification layers, and deterministic orchestration boundaries rather than unrestricted conversational interaction. Natural language may remain useful at the human interaction layer, but executable intent will likely require transformation into constrained machine-verifiable representations before operational deployment.

The problem also highlights the emerging research challenge for future 6G and post-6G systems to define measurable semantic distance and operational deviation metrics for AI-driven network control systems. Similar to how communication theory models noise, distortion, and error tolerance in transmission systems, future intent-based infrastructures may require formal mechanisms capable of quantifying acceptable divergence between intended and executed network behaviour. In this context, future standardisation needs to evolve beyond protocol interoperability, toward semantic interoperability and behavioural assurance frameworks for autonomous infrastructures.

6.10 Security, Privacy, and Data Sovereignty Technologies

6.10.1 Confidential computing and AI/ML security

Confidential computing encompasses a set of cryptographic and architectural techniques that together address how computation can be performed — and trusted — without exposing sensitive data or system internals.

Homomorphic encryption allows computations to run directly on encrypted data, so that results remain protected throughout processing. The appeal for privacy-sensitive cloud workloads is clear, but the technology still faces a hard scalability barrier: performance degrades severely at practical scale. Hardware-assisted approaches using Trusted Execution Environments (TEEs) and Intel SGX enclaves offer a partial workaround, but introduce their own vulnerability to side-channel attacks. The core research challenge is finding alternative paths to scalability that do not inherit these weaknesses.

System-level trustworthiness cannot be achieved by securing individual components in isolation. The 6G threat landscape — defined by architectural disaggregation, open interfaces, and multiple competing stakeholders — requires a holistic approach combining privacy-preserving technologies, hardware and cloud trust anchors, quantum-safe cryptography, physical layer security, distributed ledger technologies, and automated closed-loop security operations **Error! Reference source not found.**.

Zero-Knowledge Proofs (ZKPs) allow a party to prove a statement is true without revealing any underlying information, making them attractive for privacy-preserving verification in multi-stakeholder environments. The main barriers are computational cost, scalability with data size and complexity, trusted setup requirements, and integration with existing cryptographic systems — all active research areas.

AI/ML security is a multi-layered problem spanning both attacks on AI systems and the use of AI for defence. On the attack side, the principal threats across the AI pipeline are data and feature poisoning during training, model stealing at inference time, prompt injection in large language models, agent exploitation, and — most concerningly — deceptive backdoors that survive standard safety fine-tuning and become harder to detect as models scale. On the data side, sharing is constrained by privacy and business concerns, while standard anonymisation techniques destroy the feature correlations that make data useful for training. Federated learning is the most promising response, enabling collaborative model training without centralising data.

Large language models are emerging as active security tools in their own right — processing logs and traffic at scale, automating threat response, analysing phishing and social engineering, and adapting continuously to new attack patterns. Their limitations mirror broader AI/ML challenges: privacy, robustness, and scalability.

Broader AI/ML challenges relevant to 6G include the energy cost of training at scale (where the overall lifecycle benefit is not yet established), the need for context-specific rather than generalised models (addressed in part by transfer learning), data bias and fairness, and the design of robust reinforcement learning policies that can operate under 6G's dynamic, heterogeneous topology without violating system constraints. Generative AI adds a further dimension: it can automate network configuration, generate digital twins, and mediate between non-expert intent and network commands — but it is also resource-intensive, requiring tight integration of communication, compute, and control infrastructure to deploy efficiently. The question of whether any of this trajectory leads toward general artificial intelligence remains open and contested **Error! Reference source not found.**.

6.10.2 Data Protection and European Sovereignty

Data protection has traditionally been approached along two main directions. On the one hand, normative frameworks such as the General Data Protection Regulation (GDPR) establish rules governing the handling of sensitive personal data. On the other hand, proprietary digital infrastructures and specialised protocols aim to ensure the protection of confidential information, whether private or public in nature.

However, experience increasingly suggests that regulatory measures alone may not be sufficient to guarantee effective privacy protection in practice. At the same time, infrastructure-based approaches often rely on proprietary solutions that require a significant degree of trust, frequently without offering users full transparency or control.

In this context, the OECD report **Error! Reference source not found.** identifies fourteen categories of Privacy-Enhancing Technologies (PETs) as promising tools for strengthening data protection. These approaches are generally based on non-proprietary infrastructures, yet they still depend on trusted environments. The report also highlights that the current maturity of PETs remains limited, indicating that further technological progress is needed before they can be widely deployed at scale.

An alternative perspective emerges from the data usage control paradigm, originally introduced in the early 2000s and so far applied primarily in niche contexts **Error! Reference source not found.**. This paradigm is based on associating data with explicit usage policies defined by the data owner, with enforcement mechanisms ensuring that data processing complies with these policies throughout its lifecycle. Rather than focusing solely on access control, this approach extends governance to the actual use of data.

Within the broader discussion on European data protection and sovereignty, these two directions—privacy-enhancing infrastructures and data usage control **Error! Reference source not found.**—can be seen as complementary technological trajectories. Both remain relatively underexplored in large-scale deployments and have yet to receive sustained attention across research, industrial ecosystems, and policy frameworks.

In particular, the data usage control paradigm resonates strongly with the European vision that individuals should retain meaningful control over their personal data **Error! Reference source not found.**. At the same time, it opens perspectives for new hardware–software co-design approaches, potentially influencing future computing architectures. In this sense, developments in this area may also intersect with strategic considerations around technological sovereignty, including opportunities for European leadership in emerging semiconductor and system design domains.

6.11 Energy-Aware Systems and Sustainable ICT Ecosystems

Energy has emerged as a first-order design constraint and catalyst for innovation in future telecommunication and digital infrastructures. As networks evolve into AI-native, hyper-distributed systems—integrating cloud, edge, sensing, and quantum-enhanced components—energy transitions from a peripheral concern to a foundational system parameter. Future 6G and post-6G networks must be reimagined as energy-aware, adaptive, and sustainable ICT ecosystems, where communication, computation, and energy dynamics are co-designed with environmental, economic, and societal resilience in mind. This paradigm shift demands architectures that not only optimize energy efficiency but also embrace renewable energy harvesting, lifecycle-aware resource management, and policy frameworks for net-zero operations. By embedding sustainability into the DNA of network design, these systems will enable a circular economy for ICT while balancing growing demands for performance, scalability, and global connectivity.

6.11.1 Energy harvesting devices

Energy harvesting is emerging as a key enabler for ultra-distributed, low-maintenance communication and sensing infrastructures. In future digital ecosystems, massive numbers of devices are expected to operate in environments where conventional power supply or battery replacement is impractical, including environmental monitoring systems, industrial sensing fabrics, agricultural networks, oceanic deployments, and smart infrastructure embedded in the physical world.

In this context, energy harvesting technologies allow devices to operate by extracting energy from ambient sources such as vibration, light, thermal gradients, motion, and radio frequency signals **Error! Reference source not found.**. These capabilities are particularly relevant for future large-scale sensing and communication systems that extend beyond traditional network boundaries.

However, harvested energy levels remain extremely constrained, often in the micro-watt range, while communication and processing demand remain significantly higher. This fundamental imbalance drives a shift toward ultra-low-power system design, event-driven communication models, and intermittent connectivity paradigms. Devices should be able to operate in highly duty-cycled modes, activating only when sufficient energy is available or when externally triggered.

Future systems will therefore require radical redesign of communication protocols, including ultra-lightweight security mechanisms, energy-aware identity management, and new physical-layer approaches that minimise transmission overhead. Conventional mobility and continuous connectivity models will become increasingly incompatible with such energy-constrained devices, leading to more opportunistic and context-driven networking paradigms.

6.11.2 Energy efficiency with impact on standardisation and policy

Energy efficiency is a first-class design and regulatory parameter for current and future telecommunication systems. In emerging 6G and AI-native infrastructures, networks will need to expose, manage, and optimise energy usage as a measurable service attribute.

Future communication systems must be able to provide means for monitoring, predicting and exposing energy related metrics as part of service-level information. This enables new operational models in which users or applications can define energy budgets, trade-offs between performance and energy consumption, and constraints on service execution.

From a standardisation perspective, this requires the development of unified, technology-agnostic metrics for energy consumption across highly distributed and virtualised environments, including cloud-edge infrastructures, AI accelerators, and future quantum-enabled components. Such metrics must support real-time monitoring, cross-domain comparability, and lifecycle-aware energy accounting.

Key developments include:

- Definition of energy consumption models for distributed digital infrastructures
- Standard interfaces for energy telemetry and exposure
- Integration of energy constraints into service orchestration and network slicing
- Coupling of energy metrics with performance indicators such as latency, reliability, and throughput
- Development of “energy-aware QoS” and “energy-aware service contracts”

In parallel, regulatory and policy frameworks may evolve toward requiring explicit disclosure of energy consumption for digital services, similar to existing requirements in transport or consumer goods sectors. This would mark a significant shift toward accountable and measurable digital sustainability.

6.11.3 Sustainable ICT

Sustainability in ICT is evolving beyond a narrow focus on carbon accounting toward a systemic, multi-dimensional assessment of impacts across the full lifecycle of digital infrastructures. Future communication systems must be evaluated holistically across environmental, economic, and social dimensions, reflecting the broader scope of emerging frameworks for sustainability.

From an environmental perspective, the impact of digital systems spans the entire lifecycle of ICT infrastructures, including resource extraction, manufacturing, deployment, operation, and end-of-life recovery. First-order effects relate to direct environmental effects such as energy consumption, materials usage, and emissions associated with infrastructure lifecycle stages. Second-order effects arise from the enabling role of ICT in optimising other sectors such as transport, energy systems, and industrial processes. Third-order effects capture systemic social transformations, including behavioural change, digital dependency, and rebound effects that may ultimately increase overall resource consumption despite efficiency gains.

Economic sustainability introduces an additional pillar of complexity, as future communication infrastructures must remain economically viable. This includes considerations such as total cost of ownership of distributed digital systems, investment efficiency across cloud-edge-network infrastructures, and the economic implications of resource constraints, including spectrum. In this context, sustainability also requires balancing innovation with affordability.

Social sustainability further extends this framework by considering the impact of future communication systems on inclusion, accessibility, digital equity, and social resilience. The social implications of increased automation, data dependency, and human–machine interaction must be considered as integral parts of sustainable system design.

A key cross-cutting challenge is that efficiency improvements at the system level do not necessarily translate into reduced global resource consumption. As digital services become more efficient, they tend to scale more rapidly, leading to increased total demand, a phenomenon consistent with Jevons' paradox. This effect is particularly relevant in the context of AI-intensive applications, cloud-scale infrastructures, and immersive communication services.

Future sustainability frameworks must therefore move beyond single-metric optimisation and incorporate a broader set of dimensions, including:

- Material lifecycle assessment of ICT systems
- Critical raw material availability and circular economy considerations
- Water and chemical usage in semiconductor and photonic manufacturing
- Recyclability and recovery rates of advanced electronic components
- System-wide rebound effects driven by digital adoption
- Economic viability and total cost of ownership of large-scale digital infrastructures
- Social inclusion, accessibility, and digital equity considerations

In this context, sustainable ICT emerges as a multi-dimensional optimisation challenge that must jointly balance environmental impact, economic feasibility, and social value creation.

6.11.4 Environmentally sustainable communication networks beyond 6G

Without fundamental architectural changes, the growth of telecommunication demands, risks driving unsustainable energy consumption at global scale. Next-generation networks must therefore evolve toward energy-adaptive and energy-autonomous architectures, where energy consumption is dynamically managed as a core system property. This includes tight integration between communication systems, computing infrastructure, and energy provisioning mechanisms.

An important emerging direction is the development of energy-neutral and self-powered network elements, particularly for rural, remote, or infrastructure-poor regions. These systems rely on renewable energy sources such as solar, wind, kinetic, and ambient RF harvesting, combined with high-efficiency storage and adaptive load management. Future base stations and edge nodes may operate under energy-neutral constraints, dynamically adjusting service quality, connectivity density, and computational load based on available energy. This introduces a new dimension of network control where energy availability becomes a primary driver of system behaviour. At a system level, future communication infrastructures will require:

- Joint optimisation of data and energy networks
- Energy-aware orchestration across cloud-edge-radio domains
- Integration with smart grids and distributed energy systems
- Policy frameworks enforcing energy accountability in network operation
- Cross-layer coordination between AI workloads and network resource allocation

Ultimately, future communication networks must evolve into integrated data–energy infrastructures, where communication, computation, and energy systems are co-designed rather than independently optimised. This co-evolution is essential to ensuring that future digital societies remain both scalable and environmentally sustainable.

6.11.5 Energy efficient computing for large scale distributed systems

Future communication systems will increasingly rely on large-scale distributed system architectures, therefore advances in edge intelligence, massive MIMO, and future sensing-integrated communication systems will require orders-of-magnitude improvements in computational efficiency.

Conventional computing architectures based on classical von-Neumann designs are reaching fundamental limits in energy efficiency due to memory bottlenecks, data movement costs, and limited parallelism scalability **Error! Reference source not found.**. As a result, future systems will require new computing paradigms that prioritise energy efficiency at architectural level. Key directions include:

- Highly parallel, data-flow-based processing architectures
- Near-memory and in-memory computing systems to reduce data movement
- Heterogeneous accelerators combining CPU, GPU, AI, and domain-specific processors
- Analogue and mixed-signal computing approaches for ultra-low-power operations
- Task-specialised architectures for signal processing and network optimisation

The evolution towards large-scale distributed systems is closely linked to emerging hybrid computing paradigms, where classical, AI, and potentially quantum accelerators are integrated into unified computing fabrics. Energy-efficient computing will therefore become a foundational pillar of future communication systems, enabling the scalability of AI-native and ultra-dense network architectures.

6.12 Human and Physical Interfaces

6.12.1 Human-Centric Multimodal Communication

6G will enable communication that goes far beyond audio and video. Haptic, tactile, and emotional modalities open applications ranging from remote surgery and industrial teleoperation to immersive gaming, athletics training, and accessibility tools for people with sensory impairments — where touch-based alerts exploit the faster neural response to tactile stimuli compared to sight or sound **Error! Reference source not found.**, **Error! Reference source not found.**, **Error! Reference source not found.**.

The technology demands are significant: XR and teleoperation require very high throughput combined with ultra-low latency for haptic feedback, drawing on a mix of existing 5G bands and new sub-THz spectrum.

Holographic sense: Current AR/VR is limited to audio-visual experience. True holographic telepresence adds touch, and eventually taste and smell. Touch alone requires roughly 1 Gbps per hand; taste and smell would depend on replicating chemical reactions at both ends of a link — technically uncharted territory. Compression and coding latency add to transmission latency, making end-to-end delay a central design constraint **Error! Reference source not found.**.

Augmented cognition: Retinal implants — electrode arrays that electrically stimulate the optic nerve — can already restore rudimentary vision, and the technology is advancing toward higher fidelity. The logical extrapolation is using the same pathway to overlay digital information directly onto perception, making displays obsolete. The core technical challenges are signal precision and context-switching; the harder challenges are ethical and regulatory. Security questions are acute: how do you protect a neural interface from malware or unauthorised transmission? Societal acceptance will vary, and legislation will inevitably shape what is permissible.

Entangled personality: As holographic presence and augmented cognition mature, the boundary between a person's physical and virtual identity becomes increasingly fluid. Avatars and metaverse representations evolve from novelty into primary social actors — a scenario that raises profound questions about identity, agency, and the governance of virtual–physical interaction.

The disappearance of the smartphone: The smartphone may eventually dissolve as a physical object, with its functions distributed across the ambient interfaces surrounding a user — smart displays, voice systems, in-ear devices, vehicle infrastructure — while the device persists as a virtual service point in the cloud **Error! Reference source not found.**. The concept is technically plausible but faces hard barriers in trust, confidentiality, and economic viability. Realising it would require a complete redesign of security architecture and significant regulatory change.

6.12.2 Nano, Bio-/Molecular Technologies and Communications

At the nanoscale, computation, sensing, and actuation become possible within environments — blood vessels, cellular structures, chemical systems — that are entirely inaccessible to conventional devices. Envisioned since Feynman's 1959 proposal to build machines at the 10–100 atom scale, nanobots small enough to travel through the bloodstream and be guided by magnetic fields are now moving from theoretical concept toward experimental reality **Error! Reference source not found.**. Key application areas include targeted cancer cell detection and destruction, environmental chemical sensing, and — in the longer term — swallowable surgical agents that operate autonomously inside the body.

Individual nano-machines are limited to very simple tasks. The research challenge is enabling them to coordinate: sharing and fusing information to perform complex collective behaviours that no single device could achieve alone. This requires understanding and building the interfaces and gateways that connect the nanoscale world to macro-scale networks — a problem that is fundamentally different from conventional wireless communication, since electromagnetic behaviour at sub-100 nm scales is governed by quantum-mechanical effects. One promising foundation is the observation that electro-mechanically resonating nano-materials can generate and process electromagnetic waves at this scale, providing a potential signalling primitive.

The core research agenda spans three areas: the design of macro/nano interfaces that can reliably observe, control, and retrieve information from sub-microscopic systems; the programmability of nano-machine properties and behaviour; and the security of these interfaces, which is critical given that applications operating directly within biological systems or environmental contexts carry significant risk if compromised or misbehaving. As with implant-based augmented cognition, the ethical and regulatory dimensions are at least as demanding as the technical ones.

6.12.3 Advances in Material Sciences

Biodegradable and ingestible materials: Conventional electronics present a growing environmental liability, particularly in massive machine-type deployments where sensors are placed in inaccessible locations — oceans, forests, sewage systems — and cannot practically be recovered at end of life. Biodegradable materials offer a path to sensors that simply dissolve after use, eliminating the recovery problem entirely. The medical case is equally compelling: implantable devices made from transient materials degrade naturally in the body, removing the need for surgical extraction and reducing both inflammatory responses and medical waste **Error! Reference source not found.**.

The research challenges are precise control of degradation rates — ideally programmable to match the intended device lifetime — and reliable wireless data and energy transfer for implanted or ingestible devices. Self-powered solutions are strongly preferred: harvesting energy from body movement, heat, or biofluids avoids the battery replacement problem that makes conventional implants burdensome. Ingestible smart pills are an active frontier, with applications in continuous clinical monitoring, targeted drug delivery, and endoscopic diagnosis — all areas where dissolving after use is a feature rather than a limitation **Error! Reference source not found.**, **Error! Reference source not found.**.

Bioreactors for energy production: Where biodegradable materials address the waste problem, bioreactors address the power problem. Microbial Fuel Cells (MFCs) generate electricity directly from organic matter through microbial metabolism: microorganisms oxidise organic compounds, releasing electrons that can be harvested as usable current **Error! Reference source not found.**. At micro scale, MFCs become viable power sources for embedded devices deployed in environments where battery replacement is impossible — a sewage inspection robot, for instance, that harvests energy from the organic material in its environment to sustain movement, sensing, and communication indefinitely without maintenance. This closes the loop between deployment context and energy supply in a way that no conventional battery chemistry can match, and opens a class of truly autonomous long-life machine-type devices for infrastructure monitoring, environmental sensing, and similar applications.

6.13 Hidden Communication Systems and Weird Networks

Digital surveillance, censorship, and the growing need to detect and disrupt covert malicious activity have made hidden communication a two-sided research challenge: building channels that evade detection and detecting channels that should not exist. Hidden Communication Systems (HCS) and Weird Networks sit at this intersection, addressing both the design of covert, resilient, anonymous communication and the analytical and forensic tools needed to expose and counter it.

The field builds on a long lineage. The Tor network, whose principles were developed by the US Navy in the mid-1990s to protect government communications, remains the most widely deployed example [C6-42]. Like Tor, most current solutions are implemented as overlay networks on top of existing internet protocols — hidden in plain sight within normal traffic rather than requiring separate physical infrastructure. What is changing is the growing interest in solutions that operate below the overlay layer, exploiting unconventional physical channels that fall entirely outside conventional network specifications.

Weird networks are the conceptual heart of this. A weird network is one that emerges spontaneously within an existing infrastructure, using properties of that infrastructure in ways never intended by its designers. They are use-case agnostic — the same principles that support internet freedom efforts in censored environments can protect critical infrastructure communications or enable covert sensing. What distinguishes them from conventional secure networks is that their existence is not merely encrypted but structurally concealed.

6.13.1 Research challenges and directions

The core research agenda spans four areas. The first is formal modelling and detection: developing domain-specific languages and analysis tools to characterise the properties of weird networks, and methods to reveal steganographic channels embedded in normal traffic. The second is resilient anonymous communication: advanced cryptography, protocol steganography, endpoint obfuscation, and post-quantum cryptographic solutions that remain robust as adversaries gain computational power. The third is adaptable network architecture: software-defined approaches and AI-based configuration that allow hidden networks to reconfigure dynamically in response to detection attempts. Emerging approaches include the use of generative AI to synthesise adaptive communication patterns, traffic signatures, and protocol behaviours that blend into ambient network dynamics, as well as synthetic or ephemeral protocol constructs that exist only transiently within host infrastructures. The fourth is the fundamental performance trade-off: concealment degrades throughput and scalability, and principled methods for navigating this trade-off with mathematical guarantees rather than empirical heuristics remain an open problem.

The most distinctive long-term research direction is the exploration of entirely unconventional physical channels. These include the Internet of Sound, using ambient noise or acoustic signals for covert data transmission; Light Communications, exploiting visible light, Li-Fi, or infrared outside conventional spectrum allocations; Biological Systems, leveraging extracellular vesicles and intercellular communication mechanisms already studied in biosensing and medical contexts; and Primary Particle Communications, investigating muons for deep underground or underwater channels where radio frequency propagation is impractical. These are speculative but not fanciful — muon tomography already has industrial applications, and molecular communication is an active research field. Their value for HCS lies precisely in their operating in domains where conventional surveillance infrastructure has no visibility.

6.13.2 Research roadmap for Hidden Communication Systems and Weird Networks

Near-term priorities are analytical and infrastructural. The immediate need is to characterise existing overlay networks and obfuscation methods, build experimentation and emulation platforms, and apply AI and machine learning to traffic anomaly detection and steganographic analysis — drawing on the AI/ML security frameworks developed in 6.8 rather than treating detection as a separate problem. Protocol enhancements integrating 6G capabilities, LiFi, and acoustic channels into the HCS toolkit belong in the same horizon, with 6G architectural enablers addressed in the NTN and cyber-physical chapters (6.6, 6.3). This also includes the emerging use of generative AI for synthesising covert traffic patterns and adaptive protocol behaviours for both offensive and detection purposes.

In the medium term, quantum key distribution and post-quantum cryptography move from proof-of-concept to validation in realistic deployment scenarios. The foundational quantum infrastructure and PQC primitives are developed in 6.2; the work here is their specific application to hidden and resilient communication contexts, including hybrid quantum-classical approaches for hidden traffic analysis. The confidential computing techniques covered in 6.8 — zero-knowledge proofs, homomorphic encryption — are also directly relevant to this horizon, providing cryptographic building blocks for anonymous and verifiable communication without centralised trust.

The longer-term directions — biological channels, quantum internet-scale infrastructure, and exotic particle communications — are speculative but grounded in active adjacent research. Bio-inspired and molecular communication channels connect directly to the nano/bio work in 6.11, where extracellular vesicles and intercellular communication are already studied for sensing and medical applications; the covert communication angle is a natural extension of that physical substrate. The quantum internet horizon — large-scale QKD infrastructure, quantum memory, and repeater solutions — is the long-term trajectory already described in 6.2, and the contribution here is specifically the application to distributed, sovereign, and censorship-resistant communication. Muon and primary particle communications stand somewhat apart, with no current home elsewhere in the chapter, and represent a genuinely novel physical layer research direction requiring dedicated foundational study.

6.14 Mastering complexity

Tractable modelling and optimization tools are needed which allow mastering the joint complexity of the many new radio paradigms to be introduced in 6G. This complexity stems from the simultaneous introduction of several new network elements and functionalities, including Reconfigurable Intelligent Surfaces (RIS), Non-Terrestrial Networks (NTN) such as UAVs, HAPs, and LEO/MEO constellations, the use of radically different frequency bands ranging from mmWave to THz communications, as well as novel paradigms such as Joint Communication and Sensing (JCAS), cell-free architectures, and massively RIS-enhanced and 3D networks. While each of these paradigms is currently being studied independently, understanding their interactions remains an open challenge. Questions such as the comparative efficiency of densifying RIS relays versus terrestrial base stations, the mutual impact of dense NTN deployments and terrestrial infrastructure sharing radio resources, or the optimal allocation of radio resources between sensing and communication functionalities illustrate the combinatorial growth of 6G system complexity.

However, more global network architecture questions, like e.g. their hierarchical organization, their level of centralization/decentralization, the level of virtualisation, or more concretely the balance between core and edge computing, also contribute significantly to the complexity of 6G. These questions and their impact on network performance metrics are most often addressed independently using queuing network or resource sharing theory. Yet, the joint complexity resulting from the interaction between architectural design choices and the new radio paradigms alluded to above is not sufficiently captured today, despite being essential for 6G to reach its promised performance targets.

This joint analysis can be conducted either through large-scale discrete event simulation or through the development of unified stochastic geometry frameworks capable of incorporating the new 6G RAN components and their interactions, as developed in recent literature **Error! Reference source not found.**, **Error! Reference source not found.**, **Error! Reference source not found.**. Such unified stochastic models would make it possible to evaluate, in a mathematically tractable manner, the interplay between radio access innovations and architectural choices such as virtualization, resource sharing, edge computing, and hierarchical control. These approaches could therefore provide a principled way to master the unprecedented complexity of future 6G systems.

References

[C1-01] M. Draghi (2024), “The future of European competitiveness: A competitiveness strategy for Europe, (Part A:strategy + Part B: analysis&recommendations)” European Commission, Available: https://commission.europa.eu/topics/competitiveness/draghi-report_en (Accessed: 5.6.2026)

[C1-02] Letta, E. (2024), “Much more than a market: Speed, security, solidarity – Empowering the Single Market to deliver a sustainable future and prosperity for all EU citizens,” European Union. Available: https://www.consilium.europa.eu/media/ny3j24sm/much-more-than-a-market-report-by-enrico-letta.pdf (Accessed: 5.6.2026)

[C1-03] M. Heitor, C. Aerts, A. Breedenord, M.C. Carrozza, D. Dimitrova, H. Fassmann, M. Feguson, A. Kiopa, K. Palm, A. Piotrowski, I. Ryl, G. Schutte, S. Schwaag Serger, A. Vasara (2024), “Align, Act, Accelerate: Research, Technology and Innovation to boost European Competitiveness,” Publications Office of the European Union, Luxemburg, https://doi.org/10.2777/9106236

[C1-04] European Commission, “Factsheet: One Year of The Draghi Report,” Available: https://ec.europa.eu/commission/presscorner/detail/en/fs_25_2104 (Accessed: 5.6.2026)

[C1-05] GSMA Intelligence, “The Mobile Economy” (regional EU figures, 2026 edition); as reported in Mobile World Live, “GSMA reveals growing EU mobile GDP contribution,” 12 June 2026, https://www.mobileworldlive.com/gsma/gsma-reveals-growing-eu-mobile-gdp-contribution/ and https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-economy/europe/ (Accessed:12.6.2026)

[C1-06] GSMA Intelligence, “Mobile investment needs in Europe,” 6 May 2026, https://www.gsmaintelligence.com/research/mobile-investment-needs-in-europe ; see also GSMA, “€475 billion required for Europe to complete its 5G journey and regain digital leadership,” Brussels, 6 May 2026 (Accessed:12.6.2026)

[C2-01] European Commission, The Tech Sovereignty Package, available online, https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty, (Accessed: 23.6.2026).

[C2-02] Networld Europe SRIA 2024, available online, https://www.networldeurope.eu/wp-content/uploads/2025/05/ta-sria-2024-final-published-pdf.pdf (Accessed:23.6.2026).

[C2-03] IETF RFC 8799, “Limited Domains and Internet Protocols”, July 2020.

[C2-04] IETF SFC WG. Available Online: https://datatracker.ietf.org/wg/sfc/about/ (Accessed:23.6. 2026).

[C3-01] https://one6g.org/download/6873/?tmstv=1774425872

[C4-01] B. Carpenter, B. Liu, “Limited Domains and Internet Protocols”, RFC 8799, Internet Society, 2020, available at https://datatracker.ietf.org/doc/html/rfc8799

[C4-02] L. Muscariello, “Cisco Advances Open-Source Hybrid Information-Centric Networking for 5G”, accessed 2026, available at https://blogs.cisco.com/innovation/cisco-advances-open-source-hybrid-information-centric-networking-for-5g

[C4-03] D. Trossen, L. Contreras, J. Finkhaeuser, P. Mendes, “Routing on Service Addresses”, IETF draft, 2023, available at https://www.ietf.org/archive/id/draft-trossen-rtgwg-rosa-02.html

[C4-04] “Segment Routing”, accessed 2026, available at https://www.segment-routing.net/ietf

[C4-05] Charter for Computing-Aware Traffic Steering (cats) WG, IETF, 2026, available at https://datatracker.ietf.org/group/cats/about/

[C4-06] K. S. Khandaker, D. Trossen, R. Khalili, Z. Despotovic, A. Hecker and G. Carle, "CArDS: Dealing a New Hand in Reducing Service Request Completion Times," 2022 IFIP Networking Conference (IFIP Networking), Catania, Italy, 2022, pp. 1-9, doi: 10.23919/IFIPNetworking55013.2022.9829778.

[C4-07] Y. Cui , M. Xing , L. Zhang, “A Framework for LLM Agent-Assisted Network Management with Human-in-the-Loop”, IETF draft, 2026, available at https://datatracker.ietf.org/doc/draft-cui-nmrg-llm-nm/

[C4-08] L. Colombi, S. Cavicchi, F. Poltronieri, M. Tortonesi, C. Stefanelli, P. Varga, "Investigating Neurosymbolic AI for Intent-based Service Management", 21st International Conference on Network and Service Management (CNSM), 2025, available at https://opendl.ifip-tc6.org/db/conf/cnsm/cnsm2025/1571190267.pdf

[C4-09] Z. Wang, S. Lin, G. Yan, S. Ghorbani, M. Yu, J. Zhou, N. Hu, L. Baruah, S. Peters, S. Kamath, J. Yang, and Y. Zhang. 2025. Intent-Driven Network Management with Multi-Agent LLMs: The Confucius Framework. In Proceedings of the ACM SIGCOMM 2025 Conference (SIGCOMM '25). Association for Computing Machinery, New York, NY, USA, 347–362. https://doi.org/10.1145/3718958.3750537

[C4-10] S. Patel, D. Han, N. Narodytska, and S. A. Jyothi. 2025. Agua: A Concept-Based Explainer for Learning-Enabled Systems. In Proceedings of the ACM SIGCOMM 2025 Conference (SIGCOMM '25). Association for Computing Machinery, New York, NY, USA, 329–346. https://doi.org/10.1145/3718958.3754341

[C4-11] T. V. Ngo, M. V. Ngo, B. Chen, T. Q. S. Quek, T. Kumari, M. Nekovee, "LLM-Based Net Analyzer rApp for Explainable and Safe Automation in O-RAN Non-RT RIC", available at https://arxiv.org/abs/2603.13775

[C4-12] D. Guzman, D. Trossen, J. Ott, Distributed Consensus through Network Support, The International Federation for Information Processing (IFIP) Networking 2024 Conference, June 2024

[C4-13] D. Guzman, D. Trossen, T. V. Doan, J. Ott, Proliferation of the Service-centric Distributed Consensus Model and its Impact on Ethereum, 2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), May 2024

[C4-14] D. Guzman, D. Trossen, J. Ott, If Iterative Diffusion Is The Answer, What Was The Question?, Proceedings of the ACM SIGCOMM Workshop on Future of Internet Routing & Addressing at ACM SIGCOMM 2023 Conference, August 2023

[C5-01] Meta at ECOC 2024 conference

[C5-02] AI is too big for the European Internet", Nokia white paper, 2025, survey of 1,024 senior decision-makers across Europe, carried out online by Dynata between July and August 2025

[C5-03] Networld Europe SRIA 2024. https://www.networldeurope.eu/wp-content/uploads/2025/05/ta-sria-2024-final-published-pdf.pdf

[C5-04] Recommendation ITU-R M.2160. Framework and overall objectives of the future development of IMT for 2030 and beyond. 2023.

[C5-05] N.-N. Dao et al., "Neglected Infrastructures for 6G—Underwater Communications: How Mature Are They?" Journal of Network and Computer Applications, vol. 213, 2023

[C5-06] A. Giannopoulos et al., "From 6G to SeaX-G: Integrated 6G TN/NTN for AI-Assisted Maritime Communications," Journal of Marine Science and Engineering, 2025.

[C5-07] R. De Gaudenzi et al., "Beyond 5G non terrestrial networks for direct-to-device joint communication and positioning services provision: Part I—system scenarios and architectures," Wirel. Technol. 2, 16 (2026), doi: 10.1038/s44459-026-00047-w

[C5-08] P. Savazzi et al., "Comparison of Neuromorphic Coding for Distributed Wireless Spiking Neural Networks Based on Mutual Information and Energy Efficiency," in IEEE Journal of Radio Frequency Identification, vol. 9, pp. 658-668, 2025, doi: 10.1109/JRFID.2025.3600048.

[C5-09] A. Guidotti et al, "Federated Cell-Free MIMO in Nonterrestrial Networks: Architectures and Performance," IEEE Transactions on Aerospace and Electronic SystemsOpen source preview, 2024, 60(3), pp. 3319–3347

[C5-10] X. Artiga et al., "Random Satellite Swarms for LEO Direct-to-Handheld in Interference-Limited Regimes," in Proc. International Communications Satellite Systems Conference 2025 (ICSSC), 2025

[C5-11] X. Artiga et al., "Pragmatic Earth-Fixed Beam Management for 3GPP NTN Common Signaling in LEO Satellites," arXiv preprint arXiv:2512.14368, Dec. 2025

[C5-12] M. Caus et al., "GNSS Independent Random Access Schemes for Beam Hopping Satellite Systems," European Wireless 2024, 2024, pp. 210-215

[C5-13] Shahid, Husnain, et al. "Emerging advancements in 6G NTN radio access technologies: An overview." 2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit). IEEE, 2024.

[C5-14] O. M. Picchi et al., "Fused PNT on Future Wideband European Non-Terrestrial Network Infrastructure," 2025 IEEE/ION Position, Location and Navigation Symposium (PLANS), Salt Lake City, UT, USA, 2025, pp. 417-428

[C5-15] I. Gallardo-Duval et al., "Towards GNSS-Free Synchronization: An Improved PSS-Based Frequency Alignment Technique for NTN," in IEEE Open Journal of the Communications Society, vol. 7, pp. 461-479, 2026, doi: 10.1109/OJCOMS.2026.3651826

[C5-16] J. R. van der Merwe et al., “Low-cost COTS GNSS interference monitoring, detection, and classification system,” Sensors, 2023.

[C5-17] P. Savazzi et al., "Experimental Evaluation of Joint Clock Recovery and Equalization for Sub-Terahertz Links," 2025 IEEE WiSEE, doi: 10.1109/WiSEE57913.2025.11229864

[C5-18] Travis R. Brashears, "Achieving 99% link uptime on a fleet of 100G space laser inter-satellite links in LEO," Proc. SPIE 12877, Free-Space Laser Communications XXXVI, 1287702 (13 March 2024

[C6-01] Networld Europe, “Technical Annex to Strategic Research and Innovation Agenda 2022-27”, Eds. Rui Aguiar, Ari Pouttu, Jyrki Huusko, Artur Hecker, Luis Miguel Contreras Murillo, Emmanuel Dotaro, Josef Urban, Wen Xu, Raul Muñoz, Tomaso de Cola, Andre Bourdoux, Anastasius Gavras, 2024, DOI: https://doi.org/10.5281/zenodo.15165400

[C6-02] Semiconductor Research Corporation (SRC) & Semiconductor Industry Association (SIA), Decadal Plan for Semiconductors, January 2021

[C6-03] Ecologic Computing, Product Overview, Ecologic Computing, 2024. Available at: https://www.ecologic-computing.com/product/

[C6-04] Lambson, B., Carlton, D. and Bokor, J., 2011. Exploring the Thermodynamic Limits of Computation in Integrated Systems: Magnetic Memory, Nanomagnetic Logic, and the Landauer Limit. Physical review letters, 107(1), p.010604.

[C6-05] Ulmann, B., 2022. Analog computing. Walter de Gruyter GmbH & Co KG.

[C6-06] Xu, J., 2025. Biological computing (p. 329). Springer Nature.

[C6-07] European Commission, Quantum Flagship Strategic Research Agenda 2023–2030, Directorate-General for Communications Networks, Content and Technology (DG CONNECT), 2023. Available at: https://qt.eu/media/pdf/Strategic-Reseach-and-Industry-Agenda-2030.pdf

[C6-08] The Quantum Moment: A Global Report; Outcomes of the International Year of Quantum Science and Technology, 2026, https://doi.org/10.54677/NTNB6257

[C6-09] EPO/OECD, Mapping the global quantum ecosystem: A comprehensive analysis based on innovation, firm, investment, skills, trade and policy data, 2025, EPO, Munich/OECD Publishing, Paris, https://doi.org/10.65216/20251217-0001.

[C6-10] Buko, T., Tuczko, N., & Ishikawa, T. (2023). DNA Data Storage. Biotech (Basel Switzerland)), 12(2), 44. https://doi.org/10.3390/biotech12020044

[C6-11] Tavella, F., Giaretta, A., Dooley-Cullinane, T.M., Conti, M., Coffey, L. and Balasubramaniam, S., 2019. DNA molecular storage system: Transferring digitally encoded information through bacterial nanonetworks. IEEE Transactions on Emerging Topics in Computing, 9(3), pp.1566-1580.

[C6-12] Milenkovic, O. et al. (2021) Exabytes in a test tube: The case for DNA Data Storage, IEEE Spectrum. Available at: https://spectrum.ieee.org/exabytes-in-a-test-tube-the-case-for-dna-data-storage (Ac-cessed: 22 October 2025).

[C6-13] Kozlowski, Wojciech and Stephanie Wehner. “Towards Large-Scale Quantum Networks.” Proceedings of the Sixth Annual ACM International Conference on Nanoscale Computing and Communication (2019).

[C6-14] Juan Yin et al., “Satellite-based entanglement distribution over 1200 kilometers”. Science356,11401144(2017). DOI:10.1126/science.aan3211 [C6-15] F. Phillipson. Quantum Machine Learning: Benefits and Practical Examples. QANSWER, 2020, pp. 51–56.

[C6-15] Jessica Illiano, Marcello Caleffi, Antonio Manzalini, and Angela Sara Cacciapuoti. “Quantum Internet protocol stack: A comprehensive survey”. Comput. Netw. 213, C (Aug 2022). https://doi.org/10.1016/j.comnet.2022.109092

[C6-16] V. Dunjko and H. J. Briegel. Machine learning & Artificial Intelligence in the quantum domain: a review of recent progress. Reports on Progress in Physics, vol. 81, no. 7, p. 074001, jun 2018, doi: 10.1088/1361-6633/AAB406

[C6-17] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean. Power of data in quantum machine learning. Nature Communications 2021 12:1, vol. 12, no. 1, pp. 1–9, may 2021, doi: 10.1038/s41467-021-22539-9

[C6-18] P. Rebentrost, M. Mohseni, and S. Lloyd. Quantum Support Vector Machine for Big Data Classification. Physical Review Letters, vol. 113, no. 13, Sep 2014, doi: 10.1103/physrevlett.113.130503

[C6-19] M. T. Simons, A. H. Haddab, J. A. Gordon, et C. L. Holloway, « A Rydberg atom-based mixer: Measuring the phase of a radio frequency wave », Applied Physics Letters, vol. 114, no 11, p. 114101, mars 2019, doi: 10.1063/1.5088821.

[C6-20] S. Borówka, U. Pylypenko, M. Mazelanik, et M. Parniak, « Continuous wideband microwave-to-optical converter based on room-temperature Rydberg atoms », Nat. Photon., vol. 18, no 1, p. 32-38, janv. 2024, doi: 10.1038/s41566-023-01295-w.

[C6-21] F. Zhang et al., « Quantum Wireless Sensing: Principle, Design and Implementation », in Proceedings of the 29th Annual International Conference on Mobile Computing and Networking, Madrid Spain: ACM, oct. 2023, p. 1-15. doi: 10.1145/3570361.3613258

[C6-22] G. Geraci, et al., “What will the future of UAV cellular communications be? A flight from 5G to 6G,” IEEE Communications Surveys and Tutorials, 2022

[C6-23] Y. Zeng, et al. “UAV communications for 5G and beyond”, Wiley – IEEE Press, Dec. 2020.

[C6-24] Fredrik Heintz, Nabil Belbachir, and Edward Curry, “Strategic Research, Innovation, and Deployment Agenda, 2025-2027”, January 2026, Brussels – Adra

[C6-25] The IEEE Global Initiative 2.0 on Ethics of Autonomous and Intelligent Systems, https://standards.ieee.org/industry-connections/activities/ieee-global-initiative/

[C6-26] Qiang Qin, Zhihao Liu, Ruirui Zhong, Xi Vincent Wang, Lihui Wang, Magnus Wiktorsson, Wei Wang, Robot digital twin systems in manufacturing: Technologies, applications, trends and challenges, Robotics and Computer-Integrated Manufacturing, Volume 97, 2026, 103103, ISSN 0736-5845, https://doi.org/10.1016/j.rcim.2025.103103.

[C6-27] V. Ziegler, P. Schneider, H. Viswanathan, M. Montag, S. Kanugovi and A. Rezaki. Security and Trust in the 6G Era. IEEE Access, vol. 9, pp. 142314-142327, 2021, doi: 10.1109/ACCESS.2021.3120143

[C6-28] OECD, "Emerging privacy-enhancing technologies: Current regulatory and policy approaches", OECD Digital Economy Papers, No.351, OECD Publishing, Paris, (2023), https://doi.org/10.1787/bf121be4-en

[C6-29] https://internationaldataspaces.org/data-sovereignty-updated-position-paper-ondata-usage-control-in-the-ids/

[C6-30] E. Del Re, Technologies of Data Protection and Institutional Decisions for Data Sovereignty. Information 2024, 15, 444. https://doi.org/10.3390/info15080444

[C6-31] European Commission. IoT Privacy, Data Protection, Information Security. 2013. Available online: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679

[C6-32] E. Hubinger et al. Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training. arXiv e-prints. Jan 2024. arXiv:2401.05566. doi: 10.48550/arXiv.2401.05566

[C6-33] Ericsson blog - Zero-energy devices – a new opportunity in 6G, retrieved June 2022, https://www.ericsson.com/en/blog/2021/9/zero-energy-devices-opportunity-6g

[C6-34] G. P. Fettweis, "The Tactile Internet: Applications and Challenges," in IEEE Vehicular Technology Magazine, vol. 9, no. 1, pp. 64-70, March 2014, doi: 10.1109/MVT.2013.2295069

[C6-35] ITU-T Focus Group on the Tactile Internet, https://www.itu.int/en/ITU-T/techwatch/pages/tactile-internet.aspx

[C6-36] Rahim Tafazolli, online https://www.itu.int/en/ITU-T/Workshops-and-Seminars/20190218/Documents/Rahim_Tafazolli_Presentation.pdf

[C6-37] Richard Feynmann, “There is plenty of room at the bottom”, Pasadena, Dec. 1959, https://web.archive.org/web/20170105015142/http://www.its.caltech.edu/~feynman/plenty.html

[C6-38] Flavio Meneses, Daniel Corujo, Rui Aguiar. Virtualization of Customer Equipment: Challenges and Opportunities. ITL internet technology letters Vol. 3, Nº 6, May, 2020

[C6-39] Hosseini, Ensieh S., Dervin, Saoirse, Ganguly, Priyanka, Dahiya, Ravinder. Biodegradable Materials for Sustainable Health Monitoring Devices. ACS Applied Bio Materials 2021 4 (1), 163-194. doi: 10.1021/acsabm.0c01139

[C6-40] Litvinova O, Klager E, Tzvetkov NT, Kimberger O, Kletecka-Pulker M, Willschke H, Atanasov AG. Digital Pills with Ingestible Sensors: Patent Landscape Analysis. Pharmaceuticals (Basel). 2022 Aug 19;15(8):1025. doi: 10.3390/ph15081025. PMID: 36015173; PMCID: PMC9415622.

[C6-41] Yang SY, Sencadas V, You SS, Jia NZ, Srinivasan SS, Huang HW, Ahmed AE, Liang JY, Traverso G. Powering Implantable and Ingestible Electronics. Adv Funct Mater. 2021 Oct 26;31(44):2009289. doi: 10.1002/adfm.202009289. Epub 2021 Feb 4. PMID: 34720792; PMCID: PMC8553224.

[C6-42] The Tor -project, https://www.torproject.org/about/history/ [online: accessed on 10.10.2024]

[C6-43] Stochastic geometry and architecture of communication network, F. Baccelli, M. Klein, M. Lebourges and S. Zuyev, Telecommunications Systems, 7, pp.~209--227, 1997

[C6-44] Modeling and Analysis of Stochastic Mobile-Edge Computing Wireless Networks, Yixiao Gu; Yao Yao; Cheng Li; Bin Xia; Dingjie Xu; Chaoxian Zhang, 2021 https://ieeexplore.ieee.org/abstract/document/9385418

[C6-45] Wireless Networks for Mobile Edge Computing: Spatial Modeling and Latency Analysis (Extended version) Seung-Woo Ko, Kaifeng Han, Kaibin Huang, https://ieeexplore.ieee.org/document/8370831

List of Contributors (to be updated)

Name	Affiliation
Jyrki Huusko (coordination)	VTT Technical Research Centre of Finland
Rui Luis Aguiar (coordination)	Instituto de Telecomunicações - IT Aveiro
Adrian Gallego-Sanchez	Deutsche Telekom AG, T-Systems GmbH
Adrian Parez-Portero	i2CAT
Alessandro Bassi	Eurescom GmbH
Alessandro Guidotti (ed.)	University of Bologna
Alessandro Vanelli-Coralli	University of Bologna
Alexandros Kaloxylos	6G Industry Association – 6G-IA
Ali Hamidian	Huawei
Amedeo Giuliani	CTTC
Ana Aguiar	University of Porto
Ana G. Armada	Universidad Carlos III de Mardid
Anastasius Gavras (ed.)	Eurescom
Andre Bourdoux	IMEC
Andrea Giorgetti	University of Bologna
Andrés Meseguer Valenzuela	ITI
Andrew Lord	British Telecom
Anna Vizzielo	University of Pavia
Anne Lönnqvist	VTT Technical Research Centre of Finland
Antonio Napoli	Nokia
Arman Shojaeifard	InterDigital, UK
Ari Pouttu	University of Oulu
Artur Hecker (ed.)	Huawei
Bahare Masood Khorsandi	Nokia
Bengt Holter	SINTEF
Benjamin Wohlfeil	ADTRAN
Bruno De Filippo	University of Bologna
Carla Amatetti	University of Bologna
Catherine Bouillard	IMT Atlantique
Chiara Lombardo	University of Genoa
Christian Rode	Fraunhofer IIS
Claudio Cicconetti	Consiglio Nazionale delle Ricerche - CNR
Colja Schubert	Fraunhofer HHI
Damir Filipovic	AIOTI
Daniel Nunes Corujo	Instituto de Telecomunicações - IT Aveiro
Deria Malak	EURECOM
Didier Arnaud Bourse (ed.)	Nokia
Dinh-Thuy Phan-Huy	Orange
Dirk Trossen (ed.)	DaPaDOT Tech/Datacom Industry Assoc.
Djamal Zeghlache	Telecom SudParis
Elisa Rojas	University of Alcala
Emmanuel Dotaro (ed.)	Thales
Enrico Del Re	Università degli Studi di Firenze and CNIT
Ernestina Cianca	University of Rome
Ernesto Ciaramella	Scuola Superiore Sant'Anna
Eva Lagunas Targarona	University of Luxemburg
Fabio Cavaliere	Ericsson
Flávio Oliveira Silva	University of Minho
Franco Davoli	University of Genoa
François Baccelli	INRIA - École Normale Supérieure
Frank Schaich	Nokia
Frédéric Munoz	Commissariat à l'énergie atomique - CEA
Georgios Karagiannis	Huawei
Gerhard Fettweis	TU Dresden
Gilberto Berardinelli	Aalborg University
Giovanni Geraci	Telefónica
Giuseppe Caire	TU Berlin
Giuseppe Gramegna (ed.)	IMEC
Golsa Ghiaasi Hazefi	Technical University of Wien
Harald Haas	University of Cambridge
Hatim Chergui	i2CAT
Helmut Grießer	ADVA
Henk Wymeersch	Chalmers University of Technology
Hugo Tullberg	Ericsson
Hui Song	Sintef
Husnain Shahid	CTTC
Ian F. Akyldiz	Georgia Institure of Technology
Ijaz Ahmad	VTT Technical Research Centre of Finland
Ilora Maity	University of Luxemburg
Joan A. Ruiz de Azua Ortega (ed.)	i2CAT
Joan F. Munoz-Martin	i2CAT
Joaquin Escudero Garzas	Gradiant
Joerg Widmer	IMDEA Networks
Jorg-Peter Elbers	ADTRAN

Jose F. Monserrat	Universitat Politècnica de València - UPV
José Francisco Cabaça Carvalho (ed)	Instituto de Telecomunicações - IT Aveiro
Josep M. Fàbrega	CTTC
Josep M. Jornet	Northeastern University
Jouaber Badii	Telecom SudParis
Konstantin Mikhaylov	University of Oulu
Konstantinos Ntontin	University of Luxemburg
Kostas Katsaros	Digital Catapult
Laurent Clavier (ed.)	INRIA
Lorenzo Favalli	University of Pavia
Luigi Iannone	Huawei
Luis Miquel Contreras Murillo (ed.)	Telefonica
Luis Pedro Santos	UbiWhere
Luis Perez Roca	Gradiant
Luis Perez-Freire	Gradiant/AIOTI
Magnus Skjegstad	University of Oslo
Malisa Vucinic	INRIA
Malte Schellmann	Huawei
Marc Martínez	CTTC
Marja Matinmikko-Blue (ed.)	University of Oulu
Marta Castro López	Gradiant
Maxime Guillaud	INRIA
Michael Welzl	University of Oslo
Miguel Ángel Vázquez	CTTC
Miguel Ferreira Cao	Gradiant
Mika Rantakokko	VTT Technical Research Centre of Finland
Nandana Rajatheva	University of Oulu
Nathatlie Mitton	INRIA
Nikolaous Bartzoudis	CTTC
Paulo Mendes	Airbus
Pawani Porambage	VTT Technical Research Centre of Finland
Petros Elia	EURECOM
Petros S. Bithas	National and Kapodistrian University of Athens
Philippe Martins Goncalves	IMT Atlantique
Pietro Cassara	ISTI-CNR
Pietro Savazzi	University of Pavia
Pouria Sayyad Khodashenas (ed.)	i2CAT
Raffaele Bolla	University of Genoa
Raffaele Bruno	Consiglio Nazionale delle Ricerche - CNR
Raffaele D’errico	Commissariat à l'énergie atomique - CEA
Raheeb Muzaffar	Silicon Austria Labs - SAL
Ramin Khalili	Huawei Technologies
Ramon Casellas	CTTC
Raul Muñoz (ed.)	CTTC
Raymond Knopp	EURECOM
Rene Bonk	Nokia
Ricard Vilalta	CTTC
Roberto Bruschi	University of Genoa
Roland Bless	Karlsruhe Institute of Technology - KIT
Rui Dinis	Instituto de Telecomunicações/Nova University of Lisbon
Rui Wang	University of Bristol
Rute Sofia	Fortiss
Sahana Raghunandan	Fraunhofer IIS
Samuel Marchal	VTT Technical Research Centre of Finland
Sebastian Bigo	Nokia
Sebastian Robitzsch	Interdigital
Sebastian Troia	Politecnico di Milano
Silvia Liberata Ullo	Università degli Studi del Sannio di Benevento
Steve Taylor	University of Southampton
Thang Xuan Vu	Univ. of Luxembourg
Tim Hentschel	Barkhausen Institut
Tomaso De Cola	German Aerospace Center
Tommy Svensson	Chalmers University of Technology
Valeria Loscri	INRIA
Wen Xu (ed.)	Huawei
Xavier Artiga	CTTC
Xavier Jordan	i2CAT
Yvette Koza	ZTE