Towards a Converged Telco Edge Cloud: Architecting the 3C Network for Sustainable Digital Infrastructure

Engin Zeydan*, Abdullah Aydeger^
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Castelldefels, Spain,
^College of Engineering and Science, Florida Institute of Technology
Email: engin.zeydan@cttc.cat, aaydeger@fit.edu

***Abstract*—The convergence of connectivity, cloud and compute, the so-called "3C network", is a turning point that can promote industrial innovation, digital sovereignty and sustainability. In this paper, we present a comprehensive framework for implementing large-scale telco-edge cloud use cases that can integrate heterogeneous computing and communication resources across device, edge and cloud layers. We explore an architecture that incorporates AI-driven orchestration, lightweight virtualisation and privacy protection, tailored to meet stringent latency, energy and mobility requirements. The proposed approach supports open, multi-supplier and interoperable implementations. By considering governance models, security by design and industrial use cases, this paper provides a foundational blueprint for the transition to a federated, scalable digital infrastructure that can support next-generation applications in key vertical sectors.**

***Index Terms*—telco, edge cloud networks, sustainability**

I. Introduction

The digital transformation of Europe's industrial and societal sectors hinges on the seamless convergence of connectivity, cloud, and computing, collectively known as the 3C ("Connected Collaborative Computing") Network [1]. This convergence forms the backbone of a next-generation Telco Edge Cloud ecosystem, enabling real-time, low-latency, and Artificial Intelligence (AI)-driven services at scale. The recent European Commission White Paper on "How to master Europe's digital infrastructure needs?" outlines the urgency of integrating electronic communications networks with distributed computing environments to meet the surging demand for data processing, sovereignty, and sustainability [2].

In the face of increasing global competition and growing dependence on intelligent systems, the 3C Network initiative [3] aims to unify device, edge, network and cloud layers into a holistic platform that not only ensures interoperability and security, but also enables industrial innovation in areas such as manufacturing, energy, automotive, healthcare and agriculture. The focus on open, multi-supplier, and multi-tenant implementations enables the creation of a collaborative and commercially viable Telco Edge Cloud ecosystem. This paper presents a comprehensive vision, architectural framework and research roadmap for the realisation of large-scale pilots of the 3C network. It highlights the key enabling technologies from AI-based orchestration and lightweight virtualisation to privacy-preserving computation and service continuity for mobility. By exploring the integration of security, energy efficiency and regulatory compliance from the outset, the paper creates a foundation for scalable and sustainable digital infrastructures.

The rest of the paper is organized as follows: Section II introduces the unified 3C architecture; Section III explores AI-driven orchestration; Section IV details enabling platform technologies; Section V addresses performance and hardware integration; Section VI discusses standardization and industrial use cases; and Section VII concludes with key insights and future directions.

II. Reference Architecture of the 3C Network

The new 3C network paradigm enables a unified continuum that spans devices, edge and cloud infrastructure. The main goal is to ensure seamless orchestration of data, services and resources across heterogeneous and distributed environments while maintaining sovereignty, reducing latency and improving energy and resource efficiency. To this end, the architecture integrates AI-native control planes, open orchestration frameworks and programmable infrastructure layers. The AI-native Telco Edge Cloud forms the operational heart of this architecture and offers a scalable, secure and energy-efficient platform for the provision of vertical services in areas such as industry, mobility, energy, health and smart cities. It utilises a multi-layered design that includes physical endpoints, programmable edge and core networks, sovereign cloud services and a cross-cutting orchestration and automation layer with AI capabilities.

The architectural vision of the AI-native Telco Edge Cloud presented in this paper is shown in Figure 1. It reflects the convergence of connectivity, compute, and collaboration that form the basis of the 3C network. The architecture is organised into horizontal layers representing the device/user layer, the telco edge, the network structure, sovereign cloud services and a cross-cutting AI-driven orchestration layer. Each layer includes modular innovations such as container-based virtualisation, accelerator-dependent scheduling, predictive QoS management and privacy-preserving analytics. Vertical areas such as industrial automation, cooperative mobility, energy grids, healthcare and smart cities are seamlessly integrated into the infrastructure via open APIs and deployment-as-a-service

978-3-903176-75-1

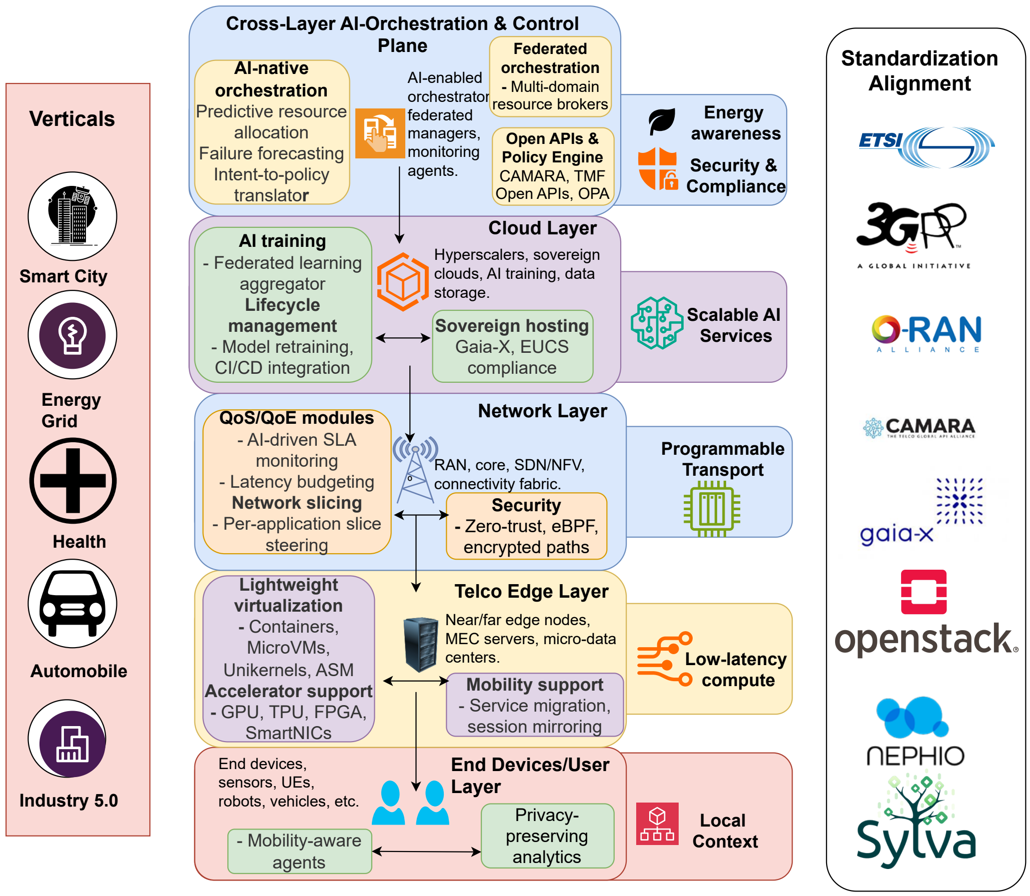

Fig. 1: High-level architecture of the AI-native Telco Edge Cloud for the 3C Network.

models. Cross-layer concerns such as energy awareness, regulatory compliance and alignment with ongoing standardisation efforts (e.g. ETSI ZSM [4], 3GPP, GAIA-X [5], CAMARA [6]) , are embedded by design, ensuring that the platform is future-ready and interoperable. The architecture is modular and composable, facilitating the rapid introduction of new services and technologies.

III. AI-Enabled Orchestration Across Edge-Cloud Continuum

A. *Federated Resource Management and Optimisation*

AI-based orchestration enables intelligent real-time decisions for optimal workload placement, resource allocation and service chaining across heterogeneous infrastructure layers in the edge cloud continuum. This capability is particularly important in multi-domain and multi-operator environments where federated orchestration mechanisms enable shared control and system-wide visibility without sacrificing privacy, autonomy or administrative boundaries [7]. Advanced AI techniques such as reinforcement learning and graph neural networks are used to predict and optimise workload placement and migration, taking into account fluctuating demand, dynamic resource availability and the current state of the network. In parallel, AI agents continuously monitor infrastructure conditions to efficiently allocate bandwidth, compute and spectrum resources, ensure Service Level Agreement (SLA) compliance and enable elastic scaling based on historical trends and real-time telemetry. In addition, multi-objective optimisation strategies- such as Pareto front exploration and constraint-aware evolutionary algorithms- are used to manage complex trade-offs between latency, energy efficiency, throughput and operational costs.

B. *Closed-Loop Automation and Self-Adaptation*

The edge cloud continuum is inherently dynamic and requires continuous adaptation to workload fluctuations, network topology changes and changing user contexts. AI enables closed-loop automation by integrating real-time monitoring, intelligent decision making and automated actuation within a unified feedback cycle [8]. Through telemetry and intent monitoring, AI models ingest real-time metrics from programmable network interfaces along with user-defined intents to detect anomalies, performance bottlenecks and potential SLA violations. Policy-based control mechanisms then guide AI agents to reconfigure services, redistribute workloads or elastically scale resources in response to these triggers, while maintaining alignment with operational and business objectives. In addition, lifecycle management is enhanced through AI-powered orchestration that enables the automated deployment, updating and decommissioning of services- including retraining or migrating AI models themselves when faced with data drift or compute limitations.

C. Predictive Infrastructure Operation

AI significantly enhances orchestration by introducing predictive capabilities that allow systems to anticipate resource demands and network conditions before they impact performance [9]. Predictive scheduling uses historical and contextual data to forecast future workload trends and allocate resources in advance using time series models and Bayesian predictors, reducing latency and avoiding contention during peak utilisation. In addition, AI-driven failure forecasting employs anomaly detection and predictive maintenance models to proactively identify risks such as hardware degradation, network connection instability or impending energy depletion, enabling timely mitigation before a service interruption occurs. To ensure continuity in mobile scenarios, AI algorithms also monitor the mobility patterns of users and devices and preemptively migrate services or data to nearby edge nodes to ensure a seamless user experience.

D. QoS-Aware Orchestration and SLAs

The orchestration layer plays a central role in ensuring Quality-of-Service (QoS) and Quality-of-Experience (QoE) in distributed and heterogeneous infrastructures [10]. By utilising AI capabilities, the system can facilitate dynamic SLA negotiation, continuous monitoring and proactive compliance. Machine learning models can be used to assess SLA compliance in real time, detect potential violations and recommend reconfigurations to maintain Key Performance Indicators (KPIs) under different load conditions. AI-powered controllers also support QoS by selecting optimal communication paths and compute clusters based on metrics such as latency, available bandwidth and network congestion. In addition, orchestration platforms dynamically scale microservices, Virtual Network Functions (VNFs) and AI inference models in response to contextual feedback, such as traffic surges or user mobility patterns, to ensure consistent service quality.

IV. Platform Technologies for Telco Edge Cloud

The realisation of the 3C network depends on robust, scalable and interoperable platform technologies, ranging from cloud data centres to edge nodes. These technologies must support multi-tenancy, automation and portability while meeting the unique requirements of edge environments such as energy efficiency, mobility and hardware diversity.

A. Lightweight Virtualisation and Cloud-Native VNFs

To enable efficient deployment in edge environments, where computing, storage and power resources are much scarcer than in centralised cloud data centres, lightweight virtualisation is becoming a decisive factor. MicroVMs and containers, such as those offered by Firecracker and Kata Containers, provide fast boot times and low-overhead isolation, making them well suited for hosting multi-tenant workloads on edge nodes with limited capacity. In addition, specialised execution environments such as unikernels and WebAssembly (WASM) runtimes support the deployment of minimal single-purpose services with reduced attack surfaces and improved performance, especially for latency-sensitive or security-critical applications. In addition, the transition from traditional VNFs to Cloud Native Functions (CNFs) enables the development of network services as modular microservices enabling fine-grained orchestration, seamless integration into Continuous Integration (CI)/Continuous Deployment (CD) pipelines and elastic scaling in response to dynamic requirements.

B. Service Continuity and Mobility Management

As edge services are deployed closer to mobile users, platform technologies must ensure seamless transitions and uninterrupted services during user mobility. Edge service handover mechanisms such as session replication, pre-fetching and stateful service migration enable application context to be maintained as users move between edge zones, minimising latency spikes and session drops. Context-aware runtime environments further improve adaptability by dynamically adjusting execution policies based on real-time data such as device movement, signal strength or evolving workload conditions. Complementing these mechanisms, mobility-aware network slicing enables the dynamic reconfiguration of network slices in anticipation of user trajectories, optimising resource allocation and reducing the risk of service degradation during mobility events.

C. Edge Discovery and Deployment-as-a-Service

Decentralised on-demand discovery and deployment mechanisms are critical for the flexible and scalable provisioning of edge services in heterogeneous and distributed infrastructures. Edge discovery protocols can enable platforms to dynamically locate and select appropriate compute resources based on real-time application requirements, geographic location and system capabilities, using registries or distributed discovery mechanisms. Zero-touch provisioning can further enhance automation by utilising infrastructure-as-code tools such as Ansible, Terraform and Nephio to seamlessly deploy, configure and update edge services without manual intervention. In addition, Deployment-as-a-Service interfaces can expose abstracted Application Programming Interfaces (APIs) that enable third-party developers and vertical service providers to dynamically onboard and instantiate services, regardless of the underlying hardware or network domain.

D. Hardware Acceleration and Specialised Infrastructure Support

To meet the stringent performance and energy efficiency requirements of next-generation services, platform technologies must effectively utilise the heterogeneity of the underlying hardware resources. Hardware abstraction layers represent an important middleware capability that enables transparent access to GPUs, FPGAs and ASICs through orchestration systems and AI frameworks, simplifying integration and improving portability. Intelligent resource scheduling mechanisms, often supported by AI, dynamically allocate accelerator resources based on workload criticality, application profiles and real-time system telemetry to maximise utilisation and responsiveness. In addition, energy-aware execution strategies are embedded in edge platforms through techniques such as power profiling and Dynamic Voltage and Frequency Scaling (DVFS), which enable the system to adapt performance levels to current demand while minimising energy consumption.

E. Multi-Domain and Multi-Tenant Isolation

To support commercial-grade deployments at scale, telco edge cloud platforms must enforce robust isolation between tenants and administrative domains to ensure security, performance stability and regulatory compliance. Network isolation is achieved through technologies such as SR-IOV [11], eBPF [12] and software-defined network (SDN) slicing, which provide secure, deterministic network behaviour tailored to the needs of individual tenants. At the compute and storage level, isolation is maintained by Kubernetes-native primitives such as namespaces, resource quotas and taints, as well as service mesh frameworks such as Istio, which enable granular control over inter-service communication and resource utilisation. In addition, policy-based access control mechanisms based on attributes such as roles, identities and usage contexts enable operators to enforce fine-grained authorisation rules that ensure tenants access only permitted resources while meeting SLAs and compliance requirements.

V. QoS/QoE Assurance and Specialised Hardware Integration

QoS and QoE are critical performance metrics for next generation services deployed over the Telco Edge Cloud. As applications become more latency-sensitive, data-intensive and mobility-aware - especially with the proliferation of industrial automation, extended reality (XR) and AI inference workloads- ensuring deterministic, high-performance delivery becomes increasingly important. This section discusses the technologies and mechanisms that support end-to-end QoS/QoE assurance and the integration of specialised hardware to achieve demanding service level objectives.

A. End-to-End QoS and QoE Mechanisms

To meet the diverse and stringent requirements of vertical sectors, the telco edge cloud must provide predictable and measurable performance guarantees across the entire edge cloud continuum. Network slicing enabled by 5G and 6G architectures facilitates traffic differentiation based on application profiles, while AI-powered traffic steering dynamically routes data flows to meet specific latency, throughput or reliability targets. QoS-aware service placement further improves performance by using constraints such as delay, jitter and bandwidth to guide orchestration decisions, ensuring latency-sensitive workloads are deployed at the most appropriate edge location. Real-time monitoring and adaptation mechanisms continuously collect telemetry data and apply inline analytics to detect SLA violations and automatically trigger service reconfiguration or scaling actions. Complementing these mechanisms, QoE prediction models use Machine Learning (ML) to estimate the end-user experience based on underlying system metrics. This enables proactive service optimisation, such as adjusting video bitrate or reallocating edge resources before user satisfaction degrades.

B. Specialised Hardware for Performance Acceleration

The inclusion of domain-specific accelerators is critical to meet the stringent performance and energy efficiency requirements of modern applications, especially in edge environments with limited compute and power resources. For AI and ML workloads, edge nodes are increasingly equipped with GPUs and Tensor Processing Units (TPUs), which provide the necessary computing capacity for low-latency inference and high-throughput data processing. In parallel, programmable FPGAs are being used to offload specialised network functions such as encryption, packet classification and deep packet inspection, reducing CPU overhead and improving system responsiveness. In addition, Application-Specific Integrated Circuits (ASICs) and SmartNICs provide line-rate packet processing and virtual switch acceleration directly at the network interface, enabling the efficient execution of programmable network functions at the data plane. To maximise performance per watt, AI-enabled orchestration frameworks perform energy-aware scheduling and dynamically allocate workloads to the most energy-efficient hardware resources based on real-time telemetry, power profiles and historical usage data.

C. Resource-Aware Orchestration and Hardware Abstraction

The integration of heterogeneous hardware at distributed edge sites necessitates advanced orchestration and abstraction mechanisms to ensure efficiency, portability and performance consistency. Hardware-agnostic APIs enabled by technologies such as Kubernetes device plugins, SR-IOV and DPDK- allow orchestrators to schedule workloads across different compute and acceleration resources without being locked into specific vendor implementations. To optimise performance, topology-aware scheduling takes into account physical hardware layouts, including Non-uniform memory access (NUMA) boundaries, PCIe connectivity and thermal constraints, reducing latency and avoiding interconnect bottlenecks. In addition, orchestration agents in federated multi-domain environments must be coordinated across administrative boundaries to balance the use of hardware resources while adhering to policies, isolation requirements and data protection guarantees.

D. Performance Assurance for Mobility and Dynamic Workloads

The Telco Edge Cloud must be able to deliver real-time performance guarantees in highly dynamic operational contexts, including scenarios with mobile users, fluctuating workloads and hybrid public-private network topologies. To achieve this, predictive QoS management techniques are used where AI models analyse usage patterns, mobility trajectories, and network congestion indicators to predict potential SLA violations and trigger proactive measures such as resource scaling or service migration. Mobility-aware resource reservation further improves service continuity by anticipating user movements and allocating compute and network resources at the appropriate edge zones. In addition, latency budgeting and enforcement mechanisms ensure that allowable delay thresholds are precisely distributed across the various segments of the network so that end-to-end performance targets can be consistently met, even under volatile conditions.

VI. Open Orchestration, Standardisation Alignment and Industrial Use Cases

To achieve interoperability, scalability and long-term sustainability in the 3C Network, open orchestration frameworks and a strong focus on global standardisation initiatives are required. As the Telco Edge Cloud spans across different infrastructure providers, operators and technology areas, open interfaces, reusable components and convergence with established open source ecosystems are essential for broad adoption and seamless integration.

A. Open Orchestration Platforms and Interfaces

Open orchestration enables dynamic service lifecycle management, intelligent workload placement and coordinated resource allocation in multi-domain, multi-vendor and multi-cloud environments. Cloud-native orchestration frameworks, primarily built on Kubernetes, are extended to support network-centric workloads through projects such as Nephio [13], which enables declarative automation for network functions, and Sylva [14], which provides a standardised abstraction for edge infrastructures. Service meshes and policy engines such as Istio and Open Policy Agent (OPA) [15] enable granular control of service behaviour, including policy enforcement, observability and multi-tenancy in complex, heterogeneous deployments. In addition, open orchestration supports the federation and portability of workloads across clouds and edge environments through the use of standard APIs defined by initiatives such as ETSI NFV MANO, TOSCA and TM Forum Open APIs [16]. Finally, the shift to declarative and intent-based management enables operators and applications to define high-level objectives that AI-powered orchestrators then interpret and execute as actionable configurations, promoting automation, adaptability and resilience across the continuum.

B. Alignment with Standards and Open-Source Initiatives

Harmonisation with existing standardisation efforts is crucial to ensure technology convergence, market compatibility and the long-term viability of Telco Edge Cloud solutions. Integration with established frameworks such as ETSI NFV/MANO, 3GPP specifications for network slicing and service management, and O-RAN interfaces ensures compliance with evolving 5G and new 6G standards. The adoption of CAMARA APIs [17] promotes consistent and harmonised service exposure to application developers and vertical sectors and facilitates dynamic, cross-network service provisioning. Collaboration with complementary initiatives - e.g. the Cloud Edge IoT framework [18] - can ensure coherence in the orchestration infrastructure, service layers and application ecosystems, minimises fragmentation and promotes interoperable, open innovation.

C. Support for Multi-Vendor and Multi-Supplier Ecosystems

The open orchestration approach promotes both competition and collaboration within the European industrial ecosystem and enables the flexible and scalable integration of innovations from SMEs, start-ups and established vendors alike. At its core, the Telco Edge Cloud utilises modular and composable reference architectures that support plug-and-play integration and lifecycle management across different components and technology stacks. The adoption of common data models and standardised APIs- such as YANG, OpenConfig and TOSCA- ensures seamless interoperability between orchestrators, resource managers and service platforms, reducing dependency on individual vendors and increasing system flexibility. In addition, by integrating orchestration frameworks with testbeds and CI/CD pipelines, the ecosystem enables real-time validation, iterative testing and rapid deployment of new features and services across the continuum.

D. Governance and Evolution Pathways

Sustained collaboration between stakeholders is essential to ensure long-term alignment and to evolve orchestration capabilities in response to new technical and operational requirements. Open governance models based on community-driven structures and public-private partnerships enable inclusive development, promote the sustainability of open source solutions and ensure transparent, consensus-based decision-making. Active participation in standardisation bodies such as ETSI, ITU-T, 3GPP and ISO/IEC enables project partners to contribute requirements, validate research results and influence the development of technical specifications. In addition, open orchestration frameworks provide clear pathways for the transition from legacy systems to AI-native, cloud-native infrastructures through gradual migration strategies that maintain backward compatibility and minimise service disruption.

E. Industrial Use Cases

The deployment of 3C Network infrastructure through large-scale pilots offers a unique opportunity to validate advanced digital capabilities across all industrial and societal sectors. Table I finally provides an overview of vertical use cases and potential pilot demonstrators in the 3C Network.

VII. Conclusion and Future Directions

This paper presents a comprehensive vision, a reference architecture and a technology blueprint for the realisation

TABLE I: Summary of Vertical Use Cases and Potential Pilot Demonstrators in the 3C Network

Vertical Domain	Key Use Case	Potential Pilot Demonstrator
Manufacturing and Industry 5.0	Synchronized multi-site factory operations and predictive maintenance powered by edge-hosted AI models for real-time process optimization.	Edge clusters integrated with factory floor devices, sensors, and mobile robots to support seamless orchestration and service continuity.
Smart Energy Systems	AI-driven demand forecasting and anomaly detection in decentralized energy grids for dynamic load balancing.	Secure edge gateways deployed at substations and renewable energy nodes, enabling closed-loop control with central cloud coordination.
Automotive and Mobility	Roadside edge nodes enable V2X communication, collective perception, and real-time traffic optimization for connected autonomous vehicles.	MEC-enabled corridors supporting low-latency inference and validating mobility-aware orchestration with seamless vehicular handovers.
Healthcare and Emergency Services	Edge-based analytics for real-time medical diagnostics, patient telemetry, and emergency incident decision support.	Portable edge devices in ambulances or rural clinics securely interfacing with hospital cloud systems under privacy and bandwidth constraints.
Smart Cities and Communities	Real-time orchestration of video analytics, traffic signals, and public safety alerts using multi-sensor fusion and edge AI.	City-wide edge nodes scaling based on population dynamics and time-of-day, with policies enforcing service prioritization.

of the 3C network, a converged, collaborative and compute-integrated telco edge cloud infrastructure. By integrating connectivity, edge intelligence and AI-powered orchestration across heterogeneous infrastructures, the 3C network enables a sovereign, secure and sustainable digital ecosystem that can support next-generation industrial and societal applications. Key enabling technologies such as lightweight virtualisation, distributed AI, privacy-preserving data processing and open orchestration frameworks are central to achieving the goals of latency reduction, service continuity and energy efficiency. The focus on open standards, ongoing initiatives and multi-stakeholder ecosystems ensures interoperability, scalability and long-term viability.

Looking ahead, several strategic priorities are emerging that will shape the next phase of research and development within the 3C Network ecosystem. Integration with 6G-native architectures and Web 4.0 paradigms will be critical, enabling support for semantic communication, distributed AI agents and immersive virtual environments. Advances in autonomous infrastructure management- through self-organising, intent-driven orchestration and AI-powered observability- will be critical to ensure scalability and operational resilience. In parallel, green digital transformation must be prioritised by embedding sustainability into platform design, including energy-aware workload scheduling, carbon-conscious orchestration and circular hardware strategies.

References

[1] Y. Ji, J. Zhang, X. Wang, and H. Yu, “Towards converged, collaborative and co-automatic (3c) optical networks,” *Science China information sciences*, vol. 61, no. 12, p. 121301, 2018.

[2] Digital Infrastructure Association of Europe, “How to master europe’s digital infrastructure needs,” Digital Infrastructure Association of Europe, Tech. Rep., Oct. 2023, white Paper, Accessed: 1 August 2025. [Online]. Available: https://digitalinfrastructure.eu

[3] T. Norp, M. Raftopoulou, P.-Y. Danet, C. De Majo, P. Sachdeva, K. Trichias, and P. Mohnani, “Research and innovation in europe on cloud for 6g networks,” in *2025 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, 2025, pp. 375–380.

[4] ETSI ZSM Industry Specification Group, “Etsi zero-touch network and service management (zsm),” https://www.etsi.org/technologies/zero-touch-network-service-management, 2023, accessed: 2025-09-21.

[5] Gaia-X AISBL, “Gaia-X: A federated, secure data infrastructure for europe,” https://gaia-x.eu/, 2025, accessed: 2025-09-21.

[6] Linux Foundation CAMARA Project, “Camara: Open telco network apis for service enablement,” https://camaraproject.org/, 2025, accessed: 2025-09-21.

[7] S. Ghosh, H. Asgari, D. Hond, K. Antonakoglou, I. Mavromatis, N. Butler, R. Hussain, S. Zeb, K. Katsaros, S. Moazzeni *et al.*, “Future open networks cross-domain cognitive orchestration: A novel design paradigm,” *IEEE Access*, 2025.

[8] L. Blanco, E. Zeydan, S. Barrachina-Muñoz, F. Rezazadeh, L. Vettori, and J. Mangues-Bafalluy, “A novel approach for scalable and sustainable 6g networks,” *IEEE Open Journal of the Communications Society*, vol. 5, pp. 1673–1692, 2024.

[9] S. Sharma, N. Kumar, Y. Dash, A. Dubey, and K. Devi, “Intelligent multi-cloud orchestration for ai workloads: enhancing performance and reliability,” in *2024 7th International Conference on Contemporary Computing and Informatics (IC3I)*, vol. 7. IEEE, 2024, pp. 1421–1426.

[10] M. Escheikh and W. Taktak, “Online qos/qoe-driven sfc orchestration leveraging a drl approach in sdn/nfv enabled networks,” *Wireless Personal Communications*, vol. 137, no. 3, pp. 1511–1538, 2024.

[11] Y. Dong, X. Yang, J. Li, G. Liao, K. Tian, and H. Guan, “High performance network virtualization with sr-iov,” *Journal of Parallel and Distributed Computing*, vol. 72, no. 11, pp. 1471–1480, 2012.

[12] D. Soldani, P. Nahi, H. Bour, S. Jafarizadeh, M. F. Soliman, L. Di Giovanna, F. Monaco, G. Ognibene, and F. Risso, “ebpf: A new approach to cloud-native observability, networking and security for current (5g) and future mobile networks (6g and beyond),” *IEEE Access*, vol. 11, pp. 57 174–57 202, 2023.

[13] Nephio Project (Linux Foundation Networking), “Nephio: Cloud-native network automation,” https://nephio.org/, accessed: 11 August 2025.

[14] Sylva Project (Linux Foundation Europe), “Sylva: A cloud software framework for telco and edge,” https://sylvaproject.org/, accessed: 11 August 2025.

[15] Open Policy Agent Project, “Open policy agent (opa): Policy-based control for cloud native environments,” https://www.openpolicyagent.org/, 2024, accessed: 2025-09-21.

[16] TM Forum, “Tm forum open apis,” https://www.tmforum.org/oda/open-apis/, 2023, accessed: 2025-09-21.

[17] CAMARA Project (Linux Foundation), “Camara: Global telco api alliance,” https://camaraproject.org/, accessed: 11 August 2025.

[18] EUCloudEdgeIoT Initiative, “Building the european cloud, edge & iot continuum,” https://eucloudedgeiot.eu/, accessed: 11 August 2025.