

AIORA: An AI-Native Multi-Stakeholder Orchestration Architecture for 6G Continuum

*Original*

AIORA: An AI-Native Multi-Stakeholder Orchestration Architecture for 6G Continuum / Molner, N., Rosa, L., Risso, F., Samdanis, K., Artuñedo, D., Smets, R., Taleb, T., Gomez-Barquero, D.. - In: IEEE NETWORK. - ISSN 0890-8044. - STAMPA. - (2025), pp. 1-1. [10.1109/mnet.2025.3596343]

*Availability:*

This version is available at: 11583/3002352 since: 2025-08-07T06:00:32Z

*Publisher:*

IEEE

*Published*

DOI:10.1109/mnet.2025.3596343

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

# An AI-Native Multi-Stakeholder Orchestration Architecture for 6G Continuum

Nuria Molner<sup>\*††</sup>, Luis Rosa<sup>†</sup>, Fulvio Risso<sup>‡</sup>, Konstantinos Samdanis<sup>§</sup>, David Artuñedo<sup>¶</sup>,  
Rob Smets<sup>||</sup>, Tarik Taleb<sup>\*\*</sup>, David Gomez-Barquero<sup>\*</sup>

<sup>\*</sup> iTEAM Research Institute - Universitat Politècnica de València, Spain <sup>†</sup> One Source Consultoria Informatica LDA, Portugal <sup>‡</sup> Politecnico di Torino, Italy <sup>§</sup> Lenovo Deutschland GMBH, Germany <sup>¶</sup> Telefónica Innovacion Digital SL, Spain <sup>||</sup> Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek TNO, Netherlands <sup>\*\*</sup> Ruhr University Bochum, Germany

**Abstract**—This paper elaborates on a novel AI-native architecture for emerging 6G systems harnessing open APIs, along with supporting mechanisms to empower intelligent and coordinated orchestration of edge-cloud continuum resources. The proposed architecture facilitates a seamless creation, life-cycle management, and exposure of services in multi-segment heterogeneous environments. It integrates new breeds of tools and advanced technologies to enable zero-touch management of an edge-cloud continuum, building on top of the 3GPP Edge Enablement Layer and the respective connectivity models, allowing to cater to the high flexibility, availability, efficiency, reliability, and resilience needs of the future 6G services and applications. Several ongoing industry initiatives—such as ETSI MEC for edge computing platforms, the GSMA Operator Platform for multi-operator service federation, and CAMARA for cross-operator API standardization—demonstrate the growing momentum towards integrated frameworks where edge, cloud, and network resources can be seamlessly orchestrated. Our proposed AI-native orchestration architecture not only aligns with these initiatives but also extends them by leveraging a multi-segment virtual continuum concept and nested AI-driven closed loops for real-time optimization.

**Index Terms**—6G, architecture, edge, cloud, continuum, multi-segment, clustering, AI, AI-native, NPN, NTN

## I. INTRODUCTION

The European Vision [1], [2] for the 6G Network Ecosystem anticipates a landscape built on three foundational pillars: immersive communication, cognition, and digital twinning. These pillars embrace the concept of 6G connected intelligence, ubiquitous computing and connectivity, as elaborated in IMT-2030 [3]. Thus, the realisation of upcoming 6G use cases [4] requires a sustainable, fluid and seamlessly interconnected infrastructure

across multiple system segments across heterogeneous technologies.

Emerging 6G network architectures should be also envisioned to enable interactive and distributed services closer to the user at the network edge. To accomplish this, certain enhancements are introduced to extend the Service Based Architecture (SBA) towards the radio, introducing the notion of Radio Access Network (RAN) and core convergence. RAN-core convergence brings selected features of the core network at the edge, allowing a customized merge of core with radio functionality to enable a flexible 6G architecture in where radio services can directly be consumed by the core network and vice versa [5].

In turn, this calls for a novel framework capable of supporting secure and fully automated service orchestration and multi-objective optimisation on top of a ubiquitous (distributed [6], [7]) edge-to-cloud virtual continuum. Such virtual continuum can be anticipated as an interconnected array of virtual resources, which can either (1) span over a singular segment – e.g., Kubernetes cluster – or (2) be distributed across multiple segments, in which case, the virtual continuum shall also include the essential components that facilitate seamless connectivity across virtual resources in different segments composing such environment, ensuring cohesive operation and resource utilization.

Currently, there is a lack of a holistic approach to solve this problem end-to-end (E2E), as it is composed of multiple smaller and yet complex sub-problems with different actors involved that are being addressed independently. On the one hand, there is a layered approach for software networks and edge locations that do not scale when multiple

<sup>††</sup>corresponding author: numolsiu@iteam.upv.es

stakeholders are involved. For each layer, there are (partial) solutions for a single stakeholder, as each stakeholder has complete control of its resources. However, dynamically allocating resources becomes more complex when more players enter the scene. Business relationships affect privacy, data location, and resource availability, among others, and, thus, create barriers to transparently sharing resources, ultimately impacting the possibility of having a ubiquitous orchestration.

In parallel, the ETSI MEC (Multi-access Edge Computing) framework provides a standardized environment for running edge applications close to the end user, offering service APIs (Application Programming Interface) (e.g., radio network information, location) to facilitate context-aware orchestration. Meanwhile, the GSMA Operator Platform advocates a global model for edge service federation, enabling consistent cross-operator offerings and unified access to edge resources. Although these efforts significantly advance edge-cloud integration, they often focus on single-stakeholder scenarios or require complex, bilateral federation agreements that do not entirely address dynamic, AI-driven orchestration across multiple administrative domains.

On the other hand, several pieces need to be orchestrated seamlessly, integrating different solutions: the management of Edge Computing resources defined by Telco Operator Platform, the APIs for managing Edge Server resources and Applications defined by the 3GPP EDGEAPP Framework, Kubernetes Multi-cluster topologies, User Plane Function (UPF) topologies for accessing Edge services, Network status information, Customers' Location and Requirements defined by 3GPP Connectivity models, etc. The complexity required to combine business relationships, resource availability, network status, and customer requirements in real-time can only be achieved by a novel artificial intelligence (AI) framework that covers all dimensions natively.

To this end, this paper defines a unified 6G AI-native architecture, contributing to identifying and addressing the relevant challenges and current implementations and solutions to be extended for an architecture that works hand in hand with the current and emerging European edge infrastructure to handle appropriately its principles of relevance for the European society [6].

Sec. II describes state of the art of E2E 6G or-

chestration problem, Sec. III elaborates the proposed AI-native multi-stakeholder 6G architecture. Sec. IV describes the business scenarios considering 3GPP EDGEAPP Framework. Sec. V provides an analysis on the capabilities related to the proposed architecture elaborating how to assist the emerging 6G services and applications. Finally, Sec. VI concludes the paper.

## II. STATE OF THE ART

Our architecture is designed to address the E2E 6G orchestration problem, which is composed by multiple smaller yet complex sub-problems with different actors involved, for some of which independent solutions exist.

To satisfy the requirements of 6G, there is a growing need to establish simple, dynamic, automated, and facilitated ways to expose network management services and APIs across operators. In 6G, one common network interface should exist to integrate non-public networks (NPN) management with Operational Technology (OT) networks. Current standards and solutions, such as 3GPP Full Common API Framework (CAPIF [8]), 3GPP Service Enabler Architecture Layer for Verticals (SEAL [9]), CAMARA [10], target the publishing, discovery and usage of APIs which could in principle enable more flexible and dynamic management and monitoring of the network and services to third parties, provide improved control, auditability and ease the management of the exposed APIs to operators. Concerning edge enablement, SA6 (Service and System Aspects) has created the Edge Computing Application Framework (EDGEAPP TS 23.558), integrated with Common API Framework (CAPIF [8]). This framework follows these four Architecture Principles: *(i)* Application Portability: Minimise changes to the Applications in UEs and Application servers; *(ii)* Service Differentiation: capability to enable/disable Edge Computing features; *(iii)* Flexible Deployment: Mobile Network Operators (MNOs) can support multiple Edge Computing Providers; *(iv)* Interworking with 3GPP Networks.

Regarding the management and orchestration of resources, Ligo-based Kubernetes orchestration [11] provides the basic building blocks to handle the life-cycle of the virtual computing continuum, allowing services running in different clusters to interact and communicate as if they were in the same physical infrastructure. ETSI ZSM (Zero-touch network

and Service Management) targets the zero-touch management of resources and provides guidelines for standardized usage with the concept of closed loops for service management [12]. ETSI ENI (Experiential Networked Intelligence) [13] provides a reference and insights into AI and machine learning integration in network management. Additionally, SYLVA [14] Linux Foundation is defining an integrated Software Stack for Telco Edge as a fundamental step to Telco Cloud and Edge homogenization and sustainability.

Some of the existing solutions are open-source, but others are proprietary, and many of these solutions for independent pieces are incompatible. Thus, there is an urgent need for a holistic approach to solve the problem E2E, and the first step is a unified architecture that seamlessly holds the different pieces.

### III. AI-NATIVE MULTI-STAKEHOLDER 6G ARCHITECTURE

Novel real-time and interactive 6G applications need to migrate easily across different edge clouds with assured network connectivity in terms of peak throughput, reliability, and latency while preserving users' data privacy and data location control.

Our architecture illustrated in Fig. 1 addresses scalability, performance, sustainability, and operational issues that will arise from the need to assemble applications, i.e., compose applications with different components stretched across federated multi-segment edge environments. In addition, it integrates principles from ETSI MEC, where a platform manager can manage edge application lifecycles and service APIs, and leverages the GSMA Operator Platform approach to federate resources across multiple operators. By adopting CAMARA-style APIs aligned with 3GPP CAPIF and SEAL frameworks, the solution enables a uniform developer interface for dynamic service deployment and closed-loop automation, while still preserving the distributed (and de-centralized) nature of telco services.

It is worth mentioning that our architecture is designed to complement –not conflict with– existing and emerging standards such as 3GPP SA6's SEAL framework and CAPIF APIs, Open-RAN (O-RAN)'s RIC architecture, and the International Mobile Telecommunications (IMT)-2030 vision for network softwarization. Rather than duplicating or

replacing these frameworks, our approach integrates them by embedding AI-driven orchestration, automation, and intelligence across layers. It leverages CAPIF for service exposure and discovery, enhances SEAL's service enablement with AI-native optimization, and interfaces with O-RAN's RIC to support adaptive, edge-intelligent operations through higher-level AI agents. Aligned with IMT-2030's goals, our architecture builds on principles of softwarization, intent-based management, and cloud-native design, forming an intelligence fabric that unifies and extends existing capabilities across the 6G ecosystem.

Through these enhancements, we make such applications mobile, to interwork and jointly materialise on the user and business requirements. Scalability issues manifest themselves in (a) unbridled growth of relations between network functions in the domains of the involved MNOs, and applications providers, and (b) the complexity of interaction, i.e., the number of messages that must be exchanged that span numerous network functions, among the involved business partners. Performance issues relate to fulfilling the required Quality of Service (QoS) needed by the application(s) and Quality of Experience (QoE) expected by the end user, allowing the applications to be interconnected over network and computing resources with low latency and ample bandwidth. Sustainability aspects such as energy consumption and carbon footprint requirements must be addressed. Our architecture enables reduced energy consumption and cost through fine-grain control of cloud resources needed by network functions and application servers, and the capability to optimize service placement (with energy/carbon-related constraints) over a wide infrastructure. Finally, this architecture offers operational advantages, supporting the legacy 5G infrastructure and 6G additions to come and ensuring that 6G capabilities lead to operational excellence in monitoring, fault isolation, security, administration, and management, and be compatible with organisational processes.

The proposed architecture addresses the control, relocation and scaling of network management facilities, application assembly (i.e., how applications components are interrelated to composed complex applications), network functions and application servers in and across federated mobile, private and application (cloud) segments. This can assist the assurance of end-user QoE when it becomes belea-

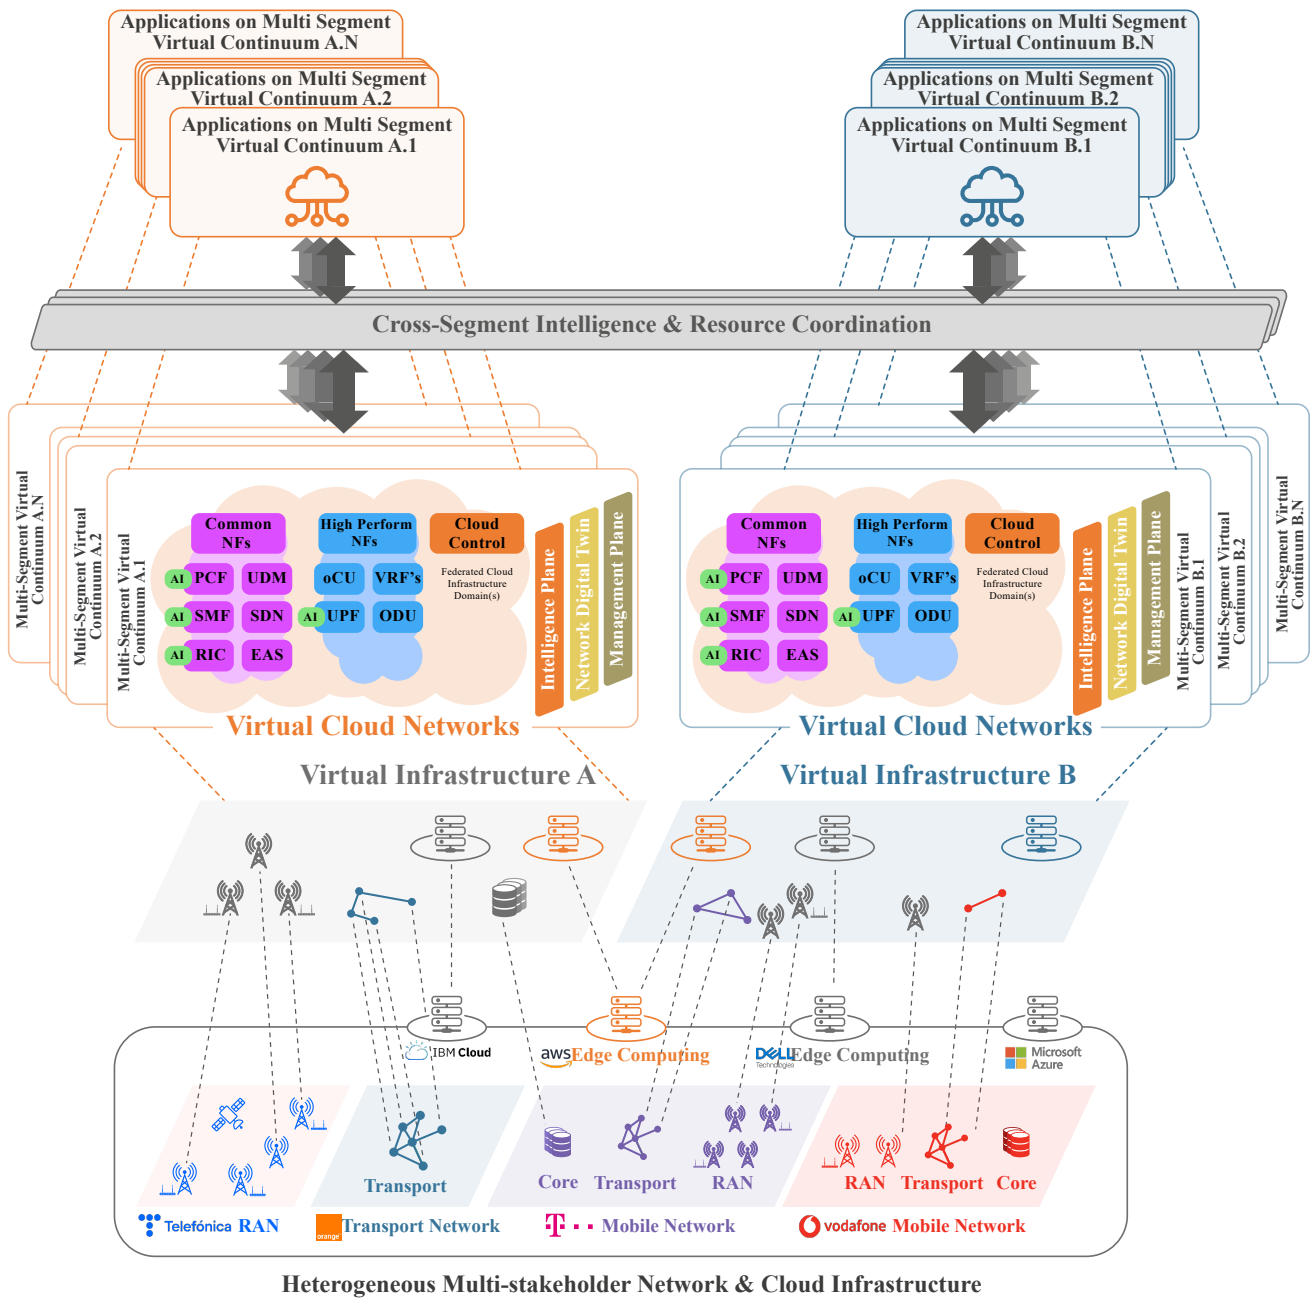


Fig. 1. The envisioned architecture showing the composition of multi-segment virtual continuums based on multi-stakeholder heterogeneous network and cloud infrastructure resources, with each multi-segment virtual continuum being configured with a customised user, control, cloud, management, and intelligence planes. A cross-segment intelligence and resource coordination substrate resolves conflicts among multi-segment virtual continuums, assuring a service targets harmonisation related to the applications enabled on the top.

guered by events, such as but not limited to end-user mobility, server maintenance shutdown, high network loads, and cyber-attack events.

Cross-segment intervening measures and continuous optimisation must ensure the QoE of the end users. This requires a new AI-native control and management framework. The notion of AI shall be embedded on the one hand in the control and management functions, including core network and

radio functions. On the other hand, AI-nativeness shall introduce decision and execution intelligence to address flexibility by continuously considering the different platforms, environments, applications, and end-user behaviour. In other words, AI shall enhance protocol mechanisms and decision-making while simultaneously enabling control closed loops to be created and deployed efficiently, i.e., without conflicts in a cooperative environment, while allow-

ing for nested closed-loop control operations.

Multi-segment virtual continuums must be tailored to the requirements of MNOs, edge computing service providers and application service providers. Each multi-segment virtual continuum may support different applications with similar service characteristics, e.g., geographical area, mobility or QoS parameters. Modifications made to accommodate the needs of a certain multi-segment virtual continuum may impact others that use the same network or cloud resources, especially when modifications occur simultaneously across more than one multi-segment virtual continuum. Cross-segment intelligence and resource coordination is a logical function introduced to resolve conflicts and take decisions based on the virtualised resource information related to the underlying infrastructures and may trigger an interaction with the application providers for negotiating the service quality targets to provide effective coordination and harmonisation.

This system is designed to ensure interoperability across proprietary and open-source orchestration components in multi-stakeholder domains, our AI-native 6G architecture adopts a modular, standards-aligned, and API-centric design. We rely on open, standardized interfaces—such as those defined by 3GPP, ETSI Network Function Virtualization (NFV), O-RAN, and the Linux Foundation—to facilitate plug-and-play integration and reduce the risk of vendor lock-in. By employing intent-based orchestration and AI-driven mediation layers, the architecture can dynamically interpret and translate policies or configurations across heterogeneous systems. Furthermore, the inclusion of interoperability gateways, model-driven architectures (e.g., YANG, TOSCA), and open telemetry standards (e.g., OpenTelemetry, Prometheus) allows both proprietary and open-source components to communicate effectively. Through these design principles, we embed real-world challenges of fragmentation, enabling seamless, policy-compliant orchestration across diverse stakeholders while fostering innovation and openness.

The AI-Native implementation of our architecture builds on top of the multi-segment managed infrastructure and its Digital Twin, and it relies on an AI-Native containerised pipeline where AI & Generative AI Models live and are trained using data prepared by real-time dataset generators. This training data is obtained from monitoring, ana-

lytical, orchestration, and actuation functions and will be used to train both AI and Generative AI models. The AI models will enhance monitoring, analytics, orchestration, and actuation operations to provide classification and prediction capabilities to the existing management functions. Meanwhile, Generative AI models will be used to generate the optimal configurations that will enrich monitoring, analytics, orchestration, and actuation sections for continuous optimal enhancement. The AI and Gen AI models, together with infrastructural elements and their digital twin counterparts, form one or more closed loops. The closed-loop coordinator is responsible for initialising and supervising the closed loops to guarantee stable, intent-driven, and, preferably, optimised operation, even when closed loops are nested or intertwined. In this sense, the coordination action shall consider the actuation decision, the managed entities, and the relationship between the closed loops to select the actuation decision to be enacted.

The AI-native closed-loop control system leverages a combination of reinforcement learning, transformers, and GNNs for analytics, with hierarchical orchestration and multi-objective optimization for decision-making. Actuation is handled via standardized APIs from 3GPP, ETSI, and O-RAN frameworks. The AI pipeline supports real-time data ingestion, federated and centralized training, and explainable model deployment through MLOps tools. At scale, challenges like training overhead and latency are addressed through model compression, edge inference, and hybrid pre-trained/continual learning strategies. The architecture is designed to balance intelligence and efficiency, enabling scalable, real-time automation across heterogeneous domains.

This architecture addresses security by using the security-by-design strategy, which relies on a threat analysis of the different architectural alternatives and their operation. Different layers of security services will be integrated to enable confidentiality, integrity, and availability of services. For the secure operation of the system, different security functionalities, ensuring different security levels, are defined and mapped on the different segments to fulfil the security application needs.

#### IV. 6G MULTI-STAKEHOLDER BUSINESS SCENARIOS: A 3GPP EDGEAPP APPROACH

The proposed architecture aims to support different edge data networks following the paradigm of 3GPP SA6 WG, as illustrated in Fig. 2, by enabling the deployment of:

- Edge Application Servers (EAS) resident in the edge data network vertical segment or application side inside the multi-segment virtual continuum business provider.
- Edge Enablement Server (EES) in the multi-segment virtual continuum to assist and configure the edge, and support EAS by offering instantiation, interface invoker and network exposure.
- Edge Configuration Server (ECS) in the virtual infrastructure for provisioning edge configuration to the client, support registration information for EES, and for providing a network interface.

Notably, integrating an ETSI MEC Platform Manager within each multi-segment virtual continuum can further streamline resource orchestration and service discovery, while the GSMA Operator Platform can federate multiple MEC-enabled domains. In scenarios where operators or service providers must expose capabilities to developers, a CAMARA-based API layer ensures a uniform approach to discovering and invoking EES, ECS, and related features.

Three different business scenarios, considering 3GPP SA6 TS 23.558, are supported by the proposed architecture as shown in Fig. 2 (for each scenario, the colour of the components indicates the business entities that offer them and the background colour of the business operator that provides the management). They are as follows:

- Business scenario A: A single business operator, which can be either an MNO or an application service provider, offers and manages the virtual infrastructure (i.e., the multi-segment virtual continuums), and the application server on top of it;
- Business scenario B: A business operator, which can be an MNO or an edge computing service provider, offers the virtual infrastructure (i.e., the multi-segment virtual continuums), while an application service provider offers the application server. The entire system

is then managed by the former MNO or the edge computing service provider, respectively;

- Business scenario C: An MNO offers and manages the virtual infrastructure. At the top, an edge computing service provider offers multi-segment virtual continuums, while an application service provider offers the application server. The former edge computing service provider manages the multi-segment virtual continuums and the respective application server.

The proposed architecture consists of multiple segments creating virtual continuums that leverage the appropriate subset of existing infrastructure-level resources. It assumes that an intermediate multi-segment virtual continuum business provider, i.e., an MNO or a service provider, collects physical resources from other MNOs and edge computing service providers based on service level agreements.

The life-cycle management of multi-segment virtual continuums is a complex issue addressed by introducing the service-enabling framework, as depicted in Fig. 2. It offers a set of mechanisms and tools to enable application developers or service providers to deploy applications that can run in any location of the virtual continuum, with the capability to migrate following the user needs and the dynamics of the underlying infrastructure across heterogeneous resources. Leveraging the benefits of AI/ML, the service-enabling framework introduces automation by employing nested control closed loops. The service-enabling framework aggregates the physical resources and offers a unified view of the virtualised federated cloud and network environment using the following functions:

- Life-cycle manager takes care of the preparation, instantiation, activation, maintenance, modification, and termination of resources.
- Ubiquitous computing manager that ensures the interoperation and management of the cloud resources and corresponding network resources.
- Infrastructure configuration and resource manager is responsible for providing the configuration of federated radio access, cloud, control plane and user plane, i.e., in terms of the radio base stations coverage, composition of network functions, cloud services (e.g., EAS discovery, network capability exposure using CAPIF, service continuity) and control mechanisms.

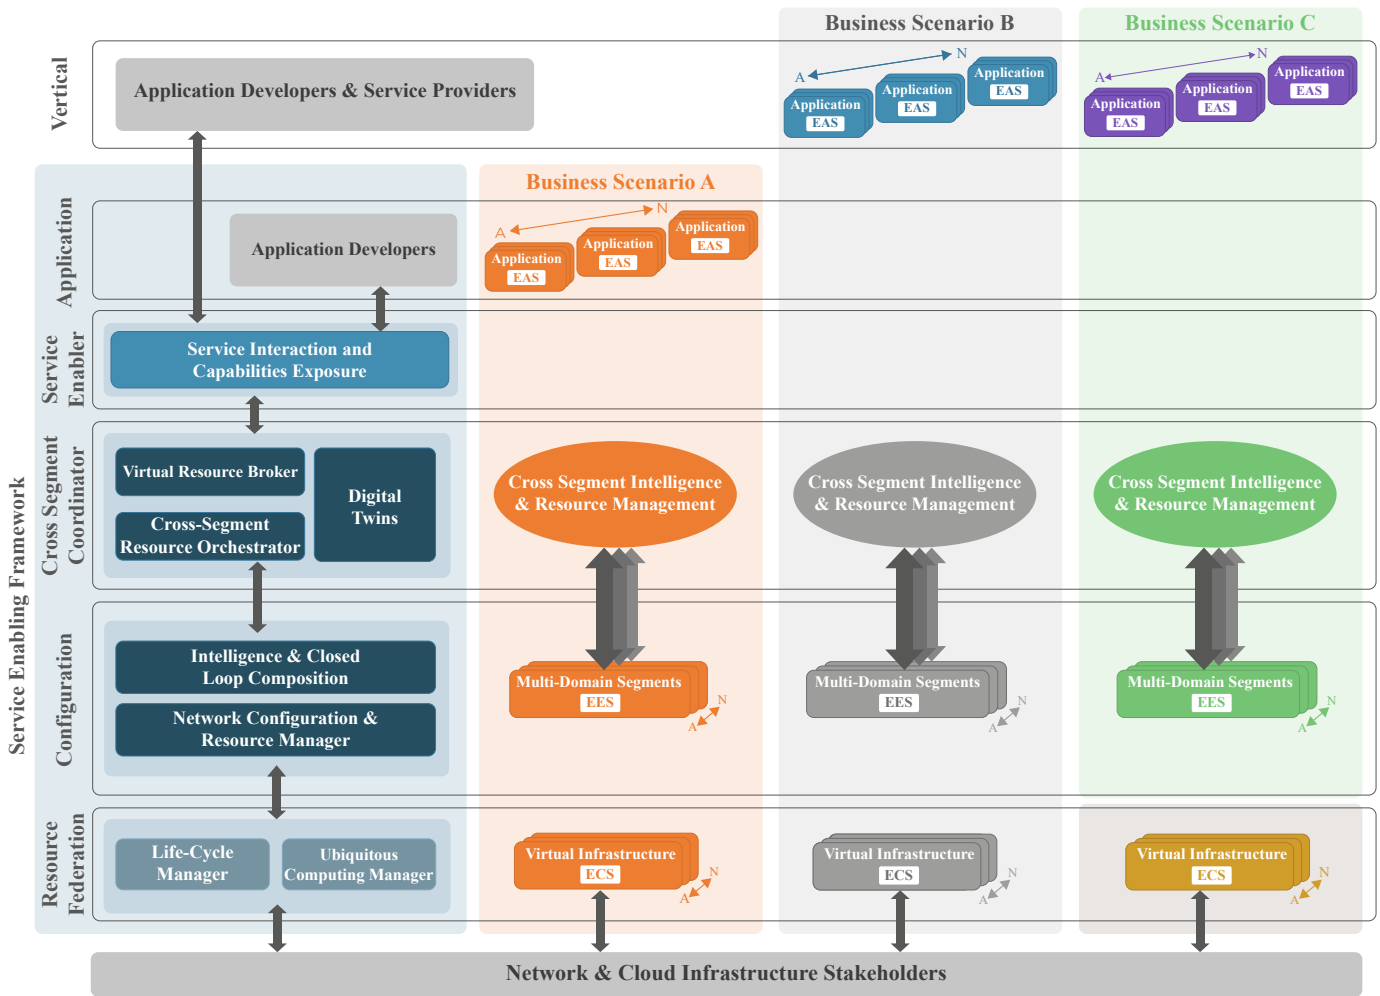


Fig. 2. Service Enabling Framework that offers the tools for facilitating services across a heterogeneous multi-stakeholder cloud and network infrastructure and shows the relation with the 3GPP SA6 WG business scenarios: (i) Business Scenario A: Single business operator that offers and manages all entities of the system, (ii) Business Scenario B: An application provider offers the application servers and an MNO or an edge computing service provider offers the virtual infrastructure and the multi-segment virtual continuums, while a single business operator manages all entities of the system, and (iii) Business Scenario C: An application provider offers the application server and an MNO, or an edge computing service provider offer the virtual infrastructure and the multi-segment virtual continuums, while an MNO is responsible for the management of the virtual infrastructure and service provider for the virtual network on the top including the applications.

In addition, it configures the management and orchestration tools needed in the management plane and assembles the closed loops and algorithms for enabling the intelligence plane.

- Intelligence and closed loop composition deals with creating controlled closed loops in a federated resource environment, i.e., selecting the monitoring that shall be used, the analytics to provide the insights, the decision entities, and the orchestrator – actuator that would enforce decisions. It also offers AI/ML algorithms and the respective data preparation and training tools and provisions the desired native AI capabilities to NFs and application services.

Multi-segment virtual continuums may compete

for the same resources or employ intelligent and management mechanisms that impact each other. To resolve these issues, cross-segments intelligence and resource coordination are introduced by adopting the following tools:

- Multi-segment resource orchestrator resolves resource conflicts and other network management-related issues.
- Virtual resource broker collects and maintains federated physical resources (i.e., network, computing, storage) to facilitate service feasibility, check and other related resource operations.
- Network Digital Twin provides a view of how the network resources, services, and applica-

tions perform. It is built using information from the digital twins exposed by external segments and information obtained from the local segment, i.e., from telemetry and performance monitoring data.

Multi-segment virtual continuums are composed by considering the service requirements introduced by the respective applications. The service enabler layer allows the deployment of applications related to specific multi-segment virtual continuums service interaction and capabilities exposure, which is responsible for enabling the composition of a service by analysing the service requirements and checking if the desired service can be supported considering the virtual infrastructure resource availability by interacting with the virtual resource broker. In addition, it enables a set of tools to allow the application and service provider to interact with the multi-segment virtual continuum business provider. The service-enabling framework executes various workflows designed to achieve a desired response from its own technology and external segments. These workflows can be triggered periodically or upon an event, i.e., on-demand, to perform various optimisations.

## V. AI-NATIVE MULTI-STAKEHOLDER 6G ARCHITECTURE: A CAPABILITY ANALYSIS

This novel architecture goes beyond the current 5G architecture, which typically relies on a network slice-level federation of the resources with limited orchestration capabilities. By embedding ETSI MEC's edge orchestration mechanisms and leveraging the GSMA Operator Platform for cross-operator resource sharing, we enable far more granular and dynamic resource management. Moreover, adopting CAMARA-based APIs ensures that developers can consistently access and manage edge functions across heterogeneous domains, while AI-driven closed-loop control provides automated real-time optimizations. Consequently, we enable intelligent, pervasive, fine-grained resource sharing, along with federated access and exposure to real-time monitoring insights.

Building over 5G architectures [15] based on Software Defined Networking (SDN) and application scenarios with mobility management approaches in mobile networks, our proposed architecture exploits the limitations of the previous generation. The cornerstone behind this novel architecture

is the leverage of virtual resources and service groupings on top of the virtualised infrastructure, referred to as multi-segment virtual continuums. These multi-segment virtual continuums enable the creation of sets of network, computing and storage resources, application and services (including network functions), and policies tailored for a specific set of application requirements in terms of network capacity, coverage, latency, etc. In this vision, each multi-segment virtual continuum shall contain its own: (i) Radio access, composed of a set of base stations and radio units; (ii) 6G core network control plane and multipurpose user plane, which could be customised as per the features of the target application(s); (iii) Cloud and (near/far) edge computing resources where the different application functions will be deployed; (iv) Intelligence and management planes implementing the per-virtual continuum Native AI control-loops for radio, network and compute resource management and optimisation; (v) Native AI tools for algorithm lifecycle management (i.e., training and execution); and the virtual continuum-specific management tools enforcing the lifecycle management and optimisation actions.

To maximize efficiency and interoperability, the creation, management, and coordination of multi-segment virtual continuums are orchestrated by a dedicated multi-segment virtual-continuum business provider, which: (i) Provides a high-level and aggregated view of the virtual continuum towards application developers and service providers; (ii) Offers high-level application lifecycle-management primitives; (iii) Offers automated management of the heterogeneous virtualised-infrastructure segments; and (iv) Offering automated coordination and conflict resolutions of closed control loops across multi-segment virtual continuums.

Our architecture represents a transformative evolution from the reactive, policy-driven nature of 5G to a proactive, intent-based framework powered by distributed intelligence. By embedding AI agents across the RAN, core, and edge, it enables closed-loop automation, real-time analytics, and continual learning, which together unlock dynamic service orchestration and fine-grained optimization. This shift allows the network to scale more efficiently across highly diverse use cases and stakeholders, overcoming the centralized bottlenecks typical of 5G. Energy consumption is intelligently managed through predictive workload placement and adaptive

control strategies, leading to significant efficiency gains, especially in IoT-heavy or edge-intensive environments. Meanwhile, latency is reduced by enabling localized, AI-driven decision-making closer to the user, supporting emerging ultra-low-latency applications such as immersive Extended Reality (XR) and tactile communication. Overall, AI-native 6G delivers a more responsive, scalable, and energy-aware infrastructure compared to traditional 5G systems.

This architecture not only advances technological innovation but also drives economic sustainability. By optimizing resource utilization and energy consumption, it significantly reduces both deployment and operational costs. Furthermore, it fosters an open, interoperable ecosystem that encourages collaboration, accelerates market growth, and lowers entry barriers for new players (newly offered cloud infrastructure can now be absorbed as a resource and become part a virtual continuum). Its inherent flexibility ensures seamless adaptation to emerging technologies, protecting long-term investments and paving the way for future innovation. Ultimately, this approach stimulates technological progress, fosters new value chains, and unlocks novel business models, catalyzing innovation and safeguarding long-term investments.

## VI. CONCLUSION

This paper presents a novel AI-native, distributed, multi-segment, and multi-stakeholder architecture tailored for 6G applications, which are expected to support real-time interactivity, with the requirement to migrate easily across different edge clouds with assured network connectivity quality in terms of peak-throughput, reliability, and latency while preserving users-data privacy and data location control.

Our architecture is centered on the usage of groupings of virtual cloud and network infrastructure resources and services of multiple system segments (e.g., public and private, terrestrial and satellite, and inter-operation among MNOs, service and application providers) on top of a virtualised infrastructure, referred to as multi-segment virtual continuums. It enables a more pervasive and fine-grained sharing of the resources of the current 5G architecture, federated at the network slice level, and federated access and exposure of the monitoring information.

In summary, our AI-native orchestration framework — by embracing ETSI MEC for edge service enablement, the GSMA Operator Platform for multi-operator federation, and CAMARA APIs for standardized network capability exposure — can expedite the industry’s move toward open, interoperable, and continuously optimized 6G ecosystems. The proposed multi-segment virtual continuum concept and its supporting intelligence layer provide a unified, future-proof foundation capable of accommodating emerging technologies and diverse stakeholder requirements.

## REFERENCES

- [1] The 5G Infrastructure Association, “European Vision for the 6G Network Ecosystem,” <https://5g-ppp.eu/wp-content/uploads/2021/06/WhitePaper-6G-Europe.pdf>, Accessed: 2024-09-11.
- [2] 6G Smart Networks and Services Industry Association (6G-IA), “Key Strategies for 6G Smart Networks and Services,” [https://6g-ia.eu/wp-content/uploads/2023/10/6g-ia-position-paper\\_2023\\_final.pdf](https://6g-ia.eu/wp-content/uploads/2023/10/6g-ia-position-paper_2023_final.pdf), Accessed: 2024-09-11.
- [3] International Telecommunication Union, “Framework and overall objectives of the future development of imt for 2030 and beyond,” 2023. [Online]. Available: <https://www.itu.int/rec/R-REC-M.2160-0>
- [4] M.A. Uusitalo, et.al., “6G Vision, Value, Use Cases and Technologies From European 6G Flagship Project Hexa-X,” *IEEE Access*, Vol.9, Dec. 2021.
- [5] T. Chen, et.al., “Service-oriented Architecture Evolution towards 6G Networks,” *IEEE CSCN*, Munich 2023.
- [6] 6G Smart Networks and Services Industry Association (6G-IA), “EU-US Beyond 5G/6G Roadmap,” <https://6g-ia.eu/wp-content/uploads/2024/01/eu-us-aligned-6g-roadmap-joint-paper.pdf>, Accessed: 2024-09-11.
- [7] European Commission, “White Paper - How to master Europe’s digital infrastructure needs?” <https://digital-strategy.ec.europa.eu/en/library/white-paper-how-master-europes-digital-infrastructure-needs>, Accessed: 2024-09-11.
- [8] 3GPP, “Common API Framework for 3GPP Northbound APIs,” [https://www.etsi.org/deliver/etsi\\_ts/123200\\_123299/123222/15.03.00\\_60/ts\\_123222v150300p.pdf](https://www.etsi.org/deliver/etsi_ts/123200_123299/123222/15.03.00_60/ts_123222v150300p.pdf), Accessed: 2024-09-11.
- [9] —, “Service Enabler Architecture Layer for Verticals (SEAL); Functional architecture and information flows,” <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3587>, Accessed: 2024-09-11.
- [10] CAMARA project, “APIs enabling seamless access to Telco network capabilities,” <https://camaraproject.org/>, Accessed: 2024-09-11.
- [11] M. Iorio, F. Risso, A. Palesandro, L. Camiciotti, and A. Manzalini, “Computing without borders: The way towards liquid computing,” *IEEE Transactions on Cloud Computing*, vol. 11, no. 3, pp. 2820–2838, 2023.
- [12] ETSI, “Zero touch network & Service Management (ZSM),” <https://www.etsi.org/technologies/zero-touch-network-service-management>, Accessed: 2024-09-11.

- [13] ETSI, “ETSI ENI (Experiential Networked Intelligence);” <https://www.etsi.org/technologies/experiential-networked-intelligence>, Accessed: 2024-09-24.
- [14] SYLVA, “SYLVA: A fundamental step to Telco Cloud & Edge homogenization and sustainability;” <https://sylvaproject.org/>, Accessed: 2024-09-11.
- [15] Q. Long, Y. Chen, H. Zhang, and X. Lei, “Software defined 5g and 6g networks: A survey,” *Mobile networks and applications*, vol. 27, no. 5, pp. 1792–1812, 2022.