

6G-INTENSE

Intent-Driven Native Artificial Intelligence Architecture Supporting Network-Compute Abstraction and Sensing at the Deep Edge

Akram Boutouchent¹, Abdelkader Mekrache², Adlen Ksentini³, Gereziher Adhane⁴, João Pedro Celestino da Fonseca⁵, Joseph McNamara⁶, Kostas Ramantas⁷, Marco Palena, Marius Iordache, Renato Lo Cigno⁸, Salwa Mostafa⁹, Swastika Roy, and Christos Verikoukis¹⁰

The 6G ecosystem is anticipated to encompass a diverse array of business actors, including mobile network operators, infrastructure providers, and service providers. Existing network management systems (NMSs), which traditionally map technological and administrative domains on a one-to-one basis, are inadequate for managing such a complex environment. This necessitates innovative distributed architectures with varying levels of abstraction. To address these challenges, we present 6G-INTENSE, a novel network management architecture designed to handle the distributed and heterogeneous nature of 6G through two key abstractions. First, the network-compute fabric (NCF) component facilitates multidomain mapping at the resource level by abstracting resource pools, including deep-edge nodes equipped with sensing capabilities. Second, the domain manager and orchestrator (DMO) serves as a central entity for service orchestration, providing a global perspective on the infrastructure. The proposed architecture leverages the intent-based networking (IBN) paradigm at all architectural layers, emphasizing its critical

role in enabling autonomous networks. Furthermore, it integrates an innovative native artificial intelligence (AI) toolkit, which is pervasive across all layers.

Introduction

While 5G profoundly transformed mobile networks, with its evolved architecture supporting increased capacity, spectral efficiency, and flexibility, the rapid growth of data-hungry, human-centered applications (e.g., emerging social virtual reality and extended reality paradigms, collectively termed the *metaverse*) will soon exceed its capabilities [1]. Consequently, both academia and industry are racing to shape the next-generation communication ecosystem, known as 6G, to drive the ongoing digitization of society. The future 6G smart networks will provide high-performance and energy-efficient infrastructure, enabling the development and deployment of next-generation Internet and other services [1]. The 6G system will foster an industrial revolution and digital transformation, accelerating the creation of smart societies, improving quality of life, and facilitating autonomous systems, haptic communication, and smart health care. Achieving these objectives sustainably requires new approaches to the architecture, federation, and

Digital Object Identifier 10.1109/MVT.2024.3525001

orchestration of telecommunications infrastructures. These new approaches necessitate multistakeholder ecosystems that promote synergies among mobile network operators (MNOs) and owners of computational and networking resources. Such collaboration will share the extraordinary costs of upgrading from 5G to 6G while enabling new business models. These new architectural paradigms introduce unprecedented complexity due to the vast scale and heterogeneity of the orchestration domains involved. This complexity must be matched by equally capable automation capabilities. Therefore, 6G aims for the holy grail of pervasive AI, referred to as *Native AI* [2].

On the other hand, simplified management interactions between stakeholders and the 6G system are crucial in such architectures. Many efforts in research and standardization are currently being made to investigate approaches to achieve this. IBN, for instance, plays a pivotal role in enabling autonomous networks by specifying goals and constraints at a higher level to the NMS [3]. IBN introduces the notion of *intent*, representing an abstract operational goal provided as input to the NMS. The latter generates the necessary low-level configurations to fulfill these intents. Although IBN is a relatively new term and technology, significant efforts have been dedicated to defining and standardizing it [4], [5]. In 6G architectures, IBN will be vital for managing increased complexity and ensuring seamless cooperation among diverse stakeholders.

In this article, we introduce 6G-INTENSE, a novel network management architecture designed to address the distributed nature of 6G networks by adopting the IBN paradigm across all layers of the architecture. The key contributions of this article can be summarized as follows:

- We propose a novel management architecture with a Native AI that facilitates intent declaration, negotiation, and decision automation across autonomous domains (ADs). Moreover, joint communication and sensing is adopted as a key enabler, helping navigate the complexities and lack of reliability of the deep edge.
- The architecture integrates an abstraction layer over resource management. The NCF abstraction framework acts as a unifying entity for all orchestration domains. The abstraction of resources and their separation from service orchestration facilitates a much more scalable DMO.
- Finally, vertical service life-cycle management is delegated to tenants that can be a service provider, a vertical, or a virtual operator. Hence, this separates vertical service management from network service management.

The remaining sections of this article are structured as follows. In the “Related Works” section, we discuss the related works to our contribution. The “The 6G-Intense Innovations” section presents the general 6G-INTENSE high-level architecture and discusses its main novelties

and building blocks. The “Business Perspective” section examines the business model foreseen in 6G-INTENSE, mentioning the set of actors in the value chain. In the “System Design” section, we delve into the internal components of the main building blocks of the proposed architecture and discuss their roles and interactions. Finally, we conclude in the last section by summarizing our work and outlining potential future directions.

Related Works

Several European projects are addressing the challenge of defining the next generation of networks, namely 6G networks. However, each project addresses it from different perspectives. For instance, 6G-BRICKS [6] addresses automation in the radio access network (RAN) by integrating AI-driven zero-touch networks and intent-based automation. The main contributions include delivering a fully decentralized management plane that supports zero-touch orchestration of communication resources. In our contribution, we consider variant types of resources that are provided from different resource pools in administrative domains. In another context, HEXA-X II [7] is a flagship European project that aims to design the 6G platform, defining its use cases requirements, vision, and implementation aspects of the 6G platform. In their recent article [8], the authors discuss a comprehensive methodology for the design of a 6G end-to-end system that includes 10 principles, a blueprint, and a structured design process. The authors included a part related to the management and orchestration of the 6G system, which leverage IBN as a management paradigm. However, they didn’t provide a reference architecture for managing distributed 6G networks, as the main goal of the contribution was to identify key technology enablers for novel management and orchestration systems. Such a challenge can only be addressed by abstracting the underlying layers, facilitating control at a higher level, and creating hierarchical control layers. This is where our proposed architecture steps in, aligning with the vision of sustainable infrastructure sharing to encourage collaboration among all members of the value chain under a unified NCF.

The 6G-INTENSE European project is one of the stream B projects that address the challenge of defining the next-generation architecture for 6G. Compared to other projects in the same scope, we are the only project that addresses the challenge of managing this novel complex environment while proposing to leverage IBN, as well as a separation between service management and resource management. In terms of standardization activities, TMForum is one of the leading organizations in defining the management plane for next-generation networks. Their proposed architecture [4] adopts IBN operations at all levels. Our proposed architecture follows the same approach by defining intent-management

interfaces at the tenant management platform (TMP), at the DMO, as well as on top of the NCF. In our contribution, we go one step further and propose the TMP, which delegates vertical services to the tenant. We also support the use of natural language for specifying requirements at the tenant level.

Previous research initiatives [9], [10] have contributed innovations that separately address each technological domain, proposing control loop-based automation systems that support on-demand scaling of virtualized infrastructures, dynamic slicing, and optimizing operational costs. Some of these initiatives represent research projects that propose novel contributions to shape the next generation of 6G networks. Authors in [9] and [10] propose new IBN systems powered by machine learning (ML) for network management. However, isolated decision automation for technological domains makes it impossible to provide truly optimal solutions, maintain consistency across layers, or perform adaptations without generating conflicts. Therefore, there is a pressing need to consolidate automation control planes while leveraging declarative, active reconciliation approaches for the entire system.

The 6G-INTENSE Innovations

Infrastructure Abstraction

The multistakeholder infrastructures envisioned in 6G, as per the network of networks concept of federated 6G RANs with dense, massive cells, cell-free access points, and cloud-edge-continuum deployments, will introduce an unprecedented level of network management complexity due to the vast scale and heterogeneity of the

orchestration domains involved [7]. The 6G-INTENSE system aims to abstract and federate all types of computational and communication resources within an Internet-scale framework governed by an intelligent orchestration paradigm. This paradigm is capable of both satisfying service requests from tenants and orchestrating infrastructure resources offered by the underlying heterogeneous, multistakeholder resource pools. The 6G-INTENSE architecture is based on two fundamental abstractions, as represented in Figure 1: 1) the DMO, which represents an operator's platform; and 2) the NCF abstraction framework, which exposes the underlying computational and networking resources of resource pools. The NCF leverages a unified data model to expose information regarding the underlying infrastructure to the DMO. This integration is achieved through various resource orchestrators that control specific resource pools. It is assumed that each DMO has a specific coverage area, which could span a country or metropolitan area while resource pools correspond to the fundamental concept of a region and could represent a single administrative domain, such as a data center, edge site, or even a smart building. Internet-scale operations can be achieved through operator federations facilitated by East/West application programming interfaces (APIs) offered by DMOs. Although efforts in this direction are still in their infancy, telecom operators have been working on global APIs for their platforms that will enable a uniform and consistent exposure of their capabilities to MNO federations.

TMP

The 6G-INTENSE architecture facilitates the recursive deployment of vertical services by enabling tenants to

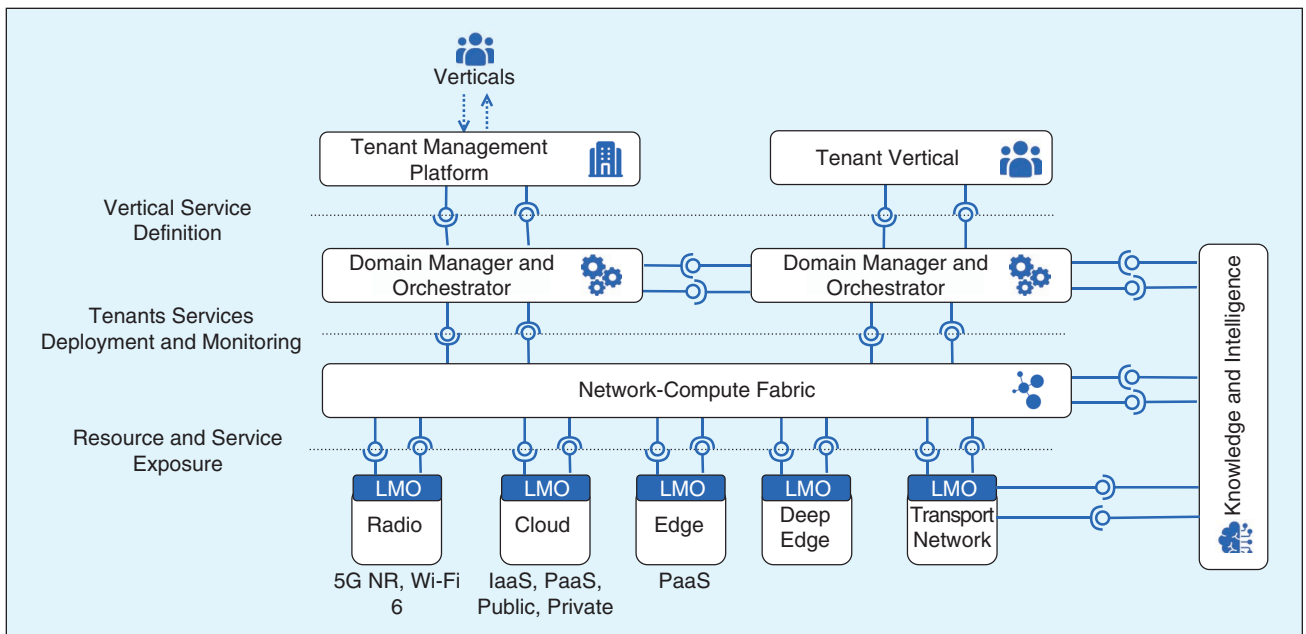


FIGURE 1 The 6G-INTENSE high-level architecture. IaaS: infrastructure as a service; PaaS: platform as a service.

deploy not only their services, but also third-party vertical services. This is achieved through the delegation of certain management functionalities to the tenant by defining the TMP, which has limited visibility over the services it manages. Examples of the network management function possible at the TMP level include the lifecycle management of vertical services and the translation of business intents into service-level intents understandable by the DMO. Additionally, the TMP can implement an application-level closed control loop that uses key performance indicators (KPIs) collected at both the DMO level and from deployed applications, utilizing management utilities deployed with the service to monitor service level agreement (SLA) satisfaction. It is important to note that the TMP has no authority to handle resources directly; all interactions pass through the DMO via intents. By delegating network management functions to the tenant level, the complexity of the DMO is reduced compared to existing network management architectures, allowing for a highly scalable framework. Tenants' verticals can directly interact with the TMP to specify their business requirements in terms of service requests and SLAs. Additionally, tenants can interact with the underlying DMOs to request vertical services by consuming the exposed interfaces.

Intent-Driven Management

To navigate the extreme complexities of the aforementioned Internet-scale ecosystem, 6G-INTENSE will offer full autonomy through actively reconciled intent declaration and translation across all layers of the infrastructure. Intent-driven, continuously reconciling systems represent the state of the art in modern automation architectures [3]. An intent provides knowledge to the autonomous system in the form of expectations and constraints, guiding the system toward achieving business goals [4], [5]. Intent-driven approaches can bridge the gap between application requirements and automated resource provisioning. However, existing intent-based approaches typically target a single automation domain with a one-to-one mapping of declarative intents to predefined policies, or they consider business directives to be imperative on the system, assuming a centralized operation support systems/business support systems layer. This approach has significant limitations regarding scalability and reliability. To address these issues, 6G-INTENSE adopts the TMForum automation architecture paradigm [4], which proposes the federation of ADs operating across different scopes and timescales, from the nonreal-time business domain to the near real-time service operations domain, and finally to the real-time resources domain. The TMForum approach to federation is based on loose coupling via intent APIs, as defined in the recently published TMForum standard [4]. These ADs will be powered by AI/ML agents from the

knowledge and intelligence component, termed *Native AI*. Native AI refers to an architecture where AI/ML is pervasive throughout the different components and technological domains, leveraging federated analytics [2].

Native AI

The 6G-INTENSE architecture delivers the first concrete approach toward this vision by designing and implementing a Native AI toolkit, which relies on a family of data-driven functions broadly applicable in all 6G-INTENSE ADs. The ANT extends the Native AI concept by ensuring coordinated operation both vertically (spanning ADs) and horizontally (within ADs) among the different AI/ML-enabled components. An example of the integration of AI/ML models throughout the architecture is intent handling. Cognitive intent handlers in 6G-INTENSE are utilized across all ADs in the architecture (namely tenant, service, and resource domains). Multiple cutting-edge AI techniques and the latest advances in large language models (LLMs) are leveraged, recognizing the key role of intents in automation architectures. LLMs, known for their ability to process and understand human language, are particularly useful for intent translation, enabling users to specify intents in natural language, as proposed in [10]. This allows intuitive interaction across 6G-INTENSE layers. LLMs can also facilitate natural language reporting, keeping users informed about network events, such as anomalies detected and resolved autonomously, in alignment with the natural language-based intent reporting [11]. Another significant application is resource negotiation, where LLMs interact with users in natural language to modify resource requests when initial intents are unfeasible.

Additionally, 6G-INTENSE will leverage reinforcement learning from human feedback (RLHF), which has proven instrumental in extending AI/ML-based models with human expertise. RLHF will be adopted as the main framework for resolving conflicts in task prioritization at the DMO layer as part of the intent renegotiation process (e.g., when "soft" requirements, priorities, and conflicting stakeholder requirements are involved). The fabric and resource layers can also leverage RLHF to readapt decision-making policies, effectively tuning these policies online with new objectives as intents are modified and renegotiated. Furthermore, 6G-INTENSE proposes the adoption of the hierarchical reinforcement learning (HRL) paradigm to extend the capabilities of RLHF. Specifically, HRL is leveraged to recursively decompose subtasks arising from the intent handler and RLHF selection process into horizontal actions (e.g., at the network-compute orchestration domain) or vertical actions (across ADs) [12]. The propagation of intent and policies across the ADs will typically generate an implicit or explicit hierarchical relationship among the different AI/ML-based functions across the architecture. As human-generated

intents decompose into subtasks, objectives, and decision-making policies, these must be consistent with the intents, objectives, and policies established by the layers above them. In this setting, the framework provided by HRL ensures the downstream propagation of intents from the DMO down to the resource layer, maintaining consistency and compliance with the human-generated intents.

Business Perspective

The distributed nature of the 6G-INTENSE approach promotes the integration of new stakeholders into the market. The separation of service and resource orchestration, in addition to the introduction of the TMP, results in new actors that didn't exist in previous generations. In 6G-INTENSE, we may see stakeholders that own only a DMO or TMP or only resources, which is completely different from the current 5G market, where, generally, network operators manage services using their resources. In Figure 2 we illustrate the interaction between envisioned 6G-INTENSE stakeholders, which include verticals, TMP, DMO, NCF, and resource providers. As stated earlier, the business model proposed by 6G-INTENSE introduces new opportunities to unlock the 6G market by enabling the entry of new actors, such as service owners, DMO owners, and NCF owners. In this model, a TMP and DMO can be a single entity or separate entities, while the NCF owner may also own DMO or 6G resources. In what follows, we will depict each actor and discuss its role and interactions with other value-chain actors.

An infrastructure provider or resource provider is the entity that owns 6G resources. It can possess one or various types of these resources. In 6G-INTENSE, we envision diverse resource types: 1) computing, which includes central cloud, edge cloud, extreme edge, and deep edge; 2) networking, featuring compute interconnection networks, such as software-defined wide area network (SD-WAN); and 3) radio, encompassing 6G radio, Wi-Fi 6, and satellite. The business interaction between infrastructure providers and the NCF is conducted on a business-to-business (B2B) basis. Resource usage can follow the established cloud model, offering pay-as-you-go options or long-term business agreements.

In this context, the NCF is one of the important components of the 6G-INTENSE system, as it contributes to the separation of resource and service management, which is one of the key innovations of 6G-INTENSE. The NCF abstracts the 6G resources to the DMO and is responsible for enforcing a 6G service. The NCF can be an independent entity separate from the DMO and resource providers. It can represent a virtual infrastructure provider or a resource broker that operators interact with to provision resources. It has B2B agreements with DMO(s) and resource providers, and it can serve multiple DMOs while registering and managing the usage of resources from various infrastructure providers.

The DMO is the second component of 6G-INTENSE that enables the separation of service and resource management. This entity manages the life-cycle of 6G services deployed on the 6G infrastructure. Its role involves translating vertical requests and needs, either

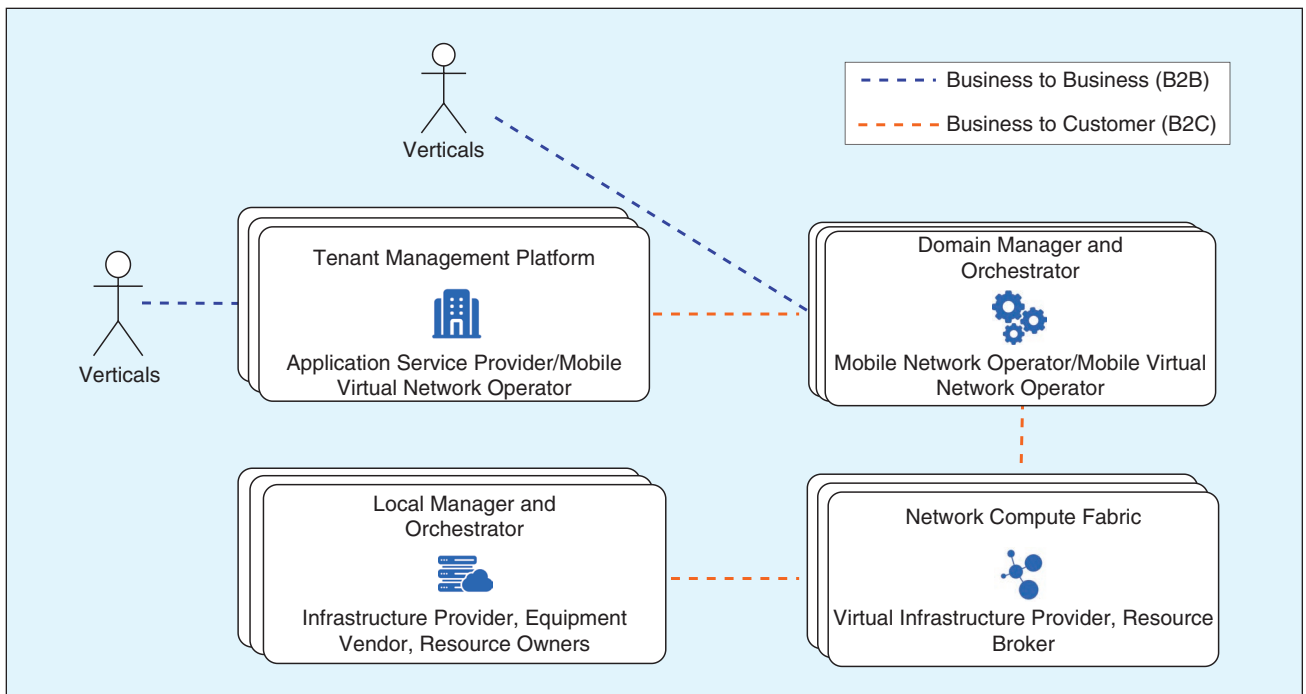


FIGURE 2 The 6G-INTENSE high-level business interactions.

via the TMP or directly, into a service-level mesh. The DMO enforces the 6G service-level mesh using the resources exposed by the NCF. The DMO establishes B2B agreements: 1) with different NCF(s) to utilize 6G infrastructure, and 2) to support various TMP(s). A business-to-consumer (B2C) agreement is established if there is direct interaction with verticals. The DMO can belong to an independent stakeholder that represents an administrative domain, similar to a network operator in 5G.

One of the innovations of 6G-INTENSE is to reduce the complexity of the service management plane by delegating some functions to the vertical service owner. Recognizing that many verticals do not know how to manage network services, the TMP entity has been introduced. The TMP is an entity that sells 6G services to verticals without owning infrastructure or being responsible for enforcing the 6G service on top of the 6G infrastructure, as the DMO does. It establishes B2B agreements with DMO(s), enabling the DMO to delegate service-level management functions. A TMP can be seen as equivalent to the virtual mobile operator actors of 4G and 5G.

Finally, verticals are entities that aim to deploy their applications or services as a 6G service on top of the 6G infrastructure. Vertical services or applications seek to leverage the capabilities of the 6G system to innovate and enhance the final user experience. Examples include autonomous driving, the entertainment industry, smart cities, and Industry 5.0. In 6G-INTENSE, the vertical establishes a business agreement with the TMP or directly with the DMO. This type of agreement covers aspects, such as the required SLA, cost, coverage area, and more. The business model between verticals and a TMP or DMO operates on a B2C basis.

System Design

DMO

While the concepts of *operator federation* and *resource pool* management are envisioned as key aspects of future 6G ecosystems, a system architecture that supports orchestration at the Internet scale is required to ensure operational agility and support the deployment

and assurance of new services in such a complex ecosystem. To fill this gap, 6G-INTENSE proposes a groundbreaking DMO design (illustrated in Figure 3), offering intent APIs toward uniform and consistent capability exposure, and service requests from MNOs, vertical application providers, and B2B customers. The DMO operates an AD with a cognitive intent handler, aiming to facilitate 1) fully autonomous service deployment and assurance in southbound resource pools and 2) assured resource reservation and programmability of the underlying Compute InterConnection networks via SD-WAN overlays translating intents into network policies that will be enforced by the NCF. This would reduce the time-to-market to a minimum, while increasing revenues for all involved stakeholders.

The 6G-INTENSE architecture proposes the following DMO APIs, which are leveraged to exchange intents and nitent reports with tenants, the NCF, and other DMOs:

- *Northbound intent APIs* that receive requests from tenants (via the TMP), requesting services, such as vertical service instantiation in a given coverage area, with certain requirements, and constraints
- *Southbound intent APIs* that send resource intents and receive reports on intent fulfillment from the NCF
- *Westbound APIs* to expose capabilities and send/receive service requests to/from other DMOs (e.g., to request/provide service outside one's coverage area).

For the aforementioned APIs, the API definition is leveraged from TMForum specifications [4].

In Figure 4, the two main 6G-INTENSE DMO workflows for autonomous service deployment and assurance are detailed. Both workflows are actuated by the cognitive intent manager DMO workflow for service creation and assurance. The main control loop of the DMO is responsible for topology-aware placement and optimization of services to the optimal resource pool(s) given the declarative intents. The DMO leverages AI/ML models from Native AI components to analyze the requirements and constraints received through the intent. This intent, in the form of ontology, contains requirements regarding the services to be provisioned as well as the KPIs to be respected. Based on this input,

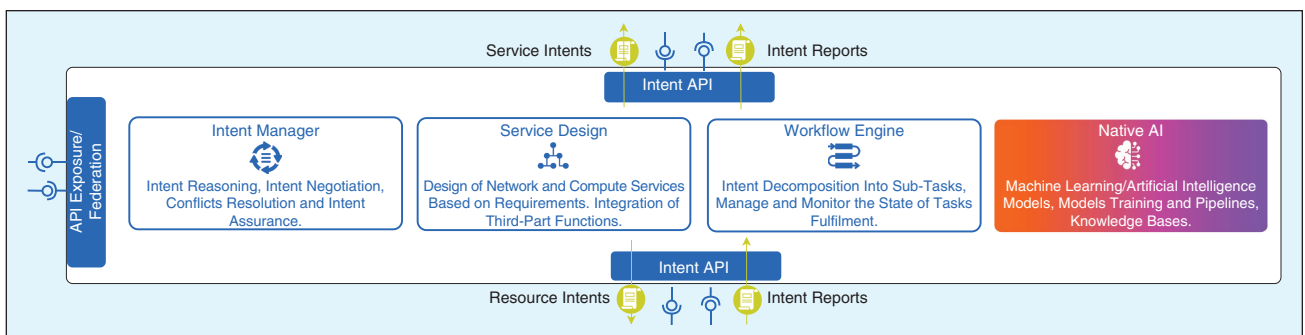


FIGURE 3 System design of the DMO.

the DMO performs feasibility checks and conflict resolution to validate that the intent can be fulfilled given the current infrastructure state. In case of nonfeasibility, an intent negotiation procedure is initiated with the intent owner (tenants or verticals). This procedure begins by assessing the available resources exposed by the local management orchestrators (LMOs). Based on this assessment, agentic LLMs can be employed to engage in a negotiation with the user using natural language. The LLM negotiates by proposing alternative solutions that align with the available resources, allowing users to either accept or decline the proposals based on their needs. When conflicting intents are detected—for example, when two or more intents request incompatible resources or service parameters—the DMO evaluates the conflict by identifying the specific resource constraints or KPI violations. The conflict resolution process involves either adjusting the resource allocation or proposing alternative configurations to ensure that all feasible intents are satisfied.

In some cases, users may be prompted to prioritize certain intents over others, depending on the relative importance of the services or KPIs involved.

The fulfillment of service intents received from tenants and verticals requires the design of services composed of various network functions, applications as microservices, and other management functions. This responsibility is handled by the service design component, which maps the received requirements to a service model that includes entities and connections. In this context, the 6G-INTENSE DMO adopts the generalized concept of a service mesh [13], leveraging cloud-native technologies for workload packaging. This encompasses cloud-native network functions, vertical microservices, and cyberphysical systems services (e.g., for sensing). The DMO considers all pathways that link events and responses internal to the service mesh, optimally distributing microservice components across the distributed, multistakeholder cloud-edge computing continuum. The Native

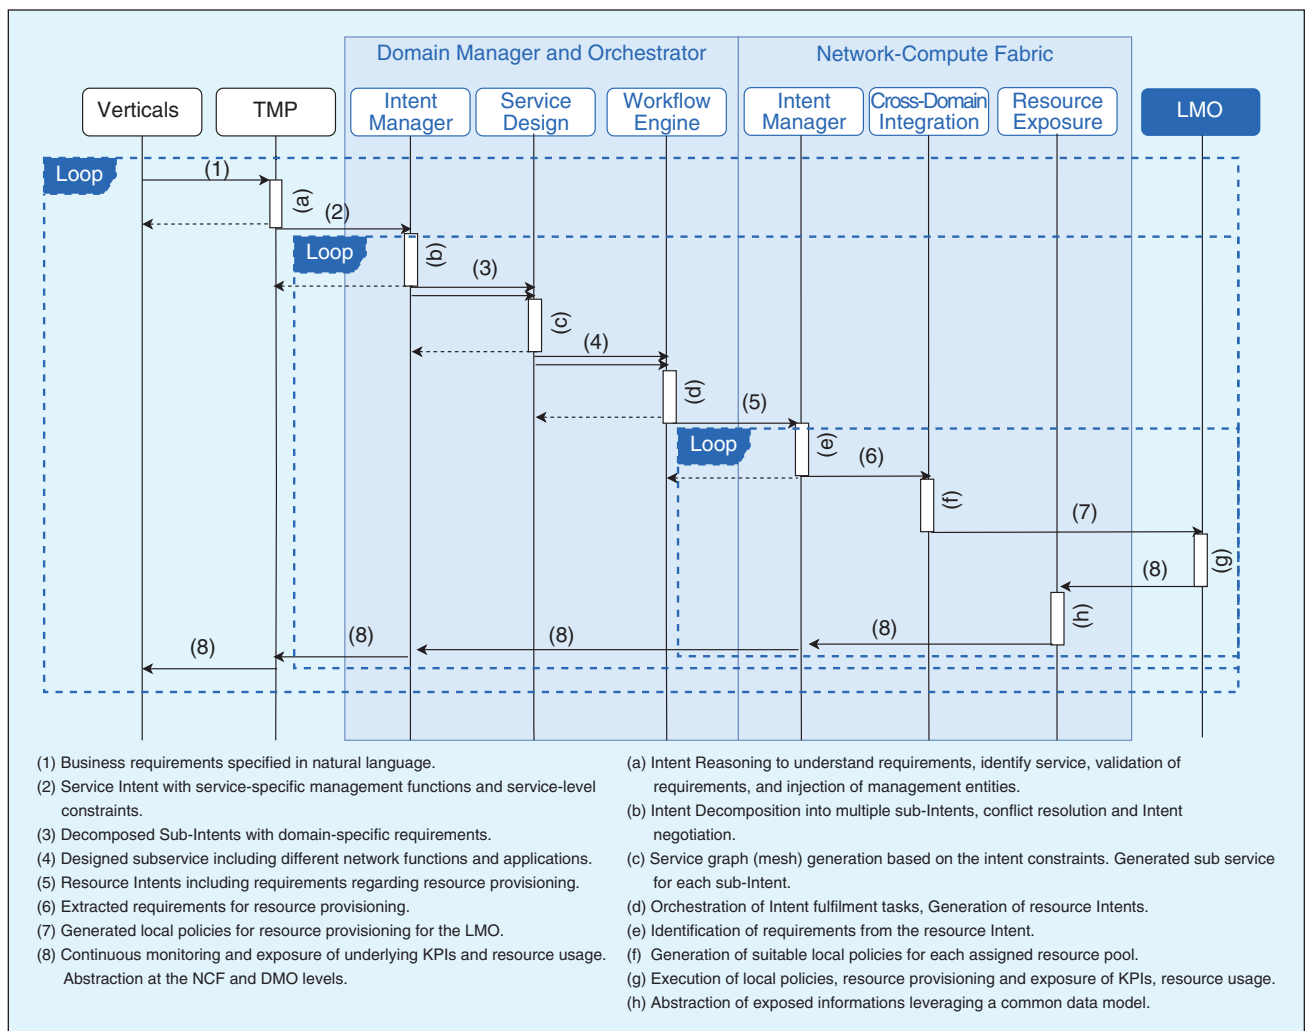


FIGURE 4 The 6G-INTENSE internal workflows. LMO: local management orchestrator.

AI concept proposed is instrumental for achieving the tasks undertaken by the DMO, as the Native AI is integrated into the main intent manager DMO loop, which propagates intents and objectives downstream. It is worth noting that the DMO decomposes the received intents into multiple subintents to reduce the complexity of the requirements and constraints embedded in the intent. Service design can address each subintent and realize topology-aware placement and optimization. Additionally, the workflow engine further decomposes the fulfillment of these intents and subintents into smaller tasks. Some tasks are propagated to the NCF as resource intents, while others leverage the East/West APIs for inter-DMO coordination. The workflow engine is responsible for orchestrating the fulfillment of the intent and continuously monitoring the state of its task executions.

On the other side, the “DMO loop” continuously performs cost/benefit analysis, conducting topology-aware optimization while considering changes in context and behavioral constraints, such as tolerable latency, utilization, and cost, assuring both new and existing services. This process involves collecting intent reports from the NCF, which include the current state of the service, including KPIs and resource usage. Additionally, the DMO leverages the resource exposure feature of the NCF to make decisions based on the exposed abstract model of the infrastructure. The intent manager closes the control loop using Native AI components for anomaly detection and makes proactive decisions to fulfill the intent. These decisions made by the intent manager

influence all other components and trigger a new intent activation pipeline.

NCF

The vision of 6G service meshes disaggregated across multistakeholder resource pools adds an unprecedented level of complexity due to the heterogeneity of the orchestration domains involved that must coordinate to fulfill the tenant intent. Failures in a single resource pool, network link, or a single function of a service mesh can unexpectedly propagate, affecting services in all adjacent domains. This is a well-known failure mode of hyperscaler platforms and would be immensely pronounced in 6G, where resource pools can aggregate multiple types of sometimes transient resources (e.g., RAN, transport, and compute). To address this, the NCF is proposed as a key future 6G enabler, acting as a single entity providing access to all types of resources (illustrated in Figure 5). The 6G-INTENSE architecture goes a step further by incorporating generative AI techniques to simplify interaction with the NMS using natural language [10], aiming to provide a unified intent-based resource management API for all technological domains (cloud, edge, deep edge, RAN). In this context, LLMs are intended to be used and adapted through in-context learning techniques or parameter-efficient fine-tuning to translate natural language intents into LMO-defined structures (e.g., the European Telecommunication Standard’s Institute’s network services descriptor for the cloud/edge domain). Generative AI in the NCF layer improves interaction by allowing users, regardless of technical expertise, to

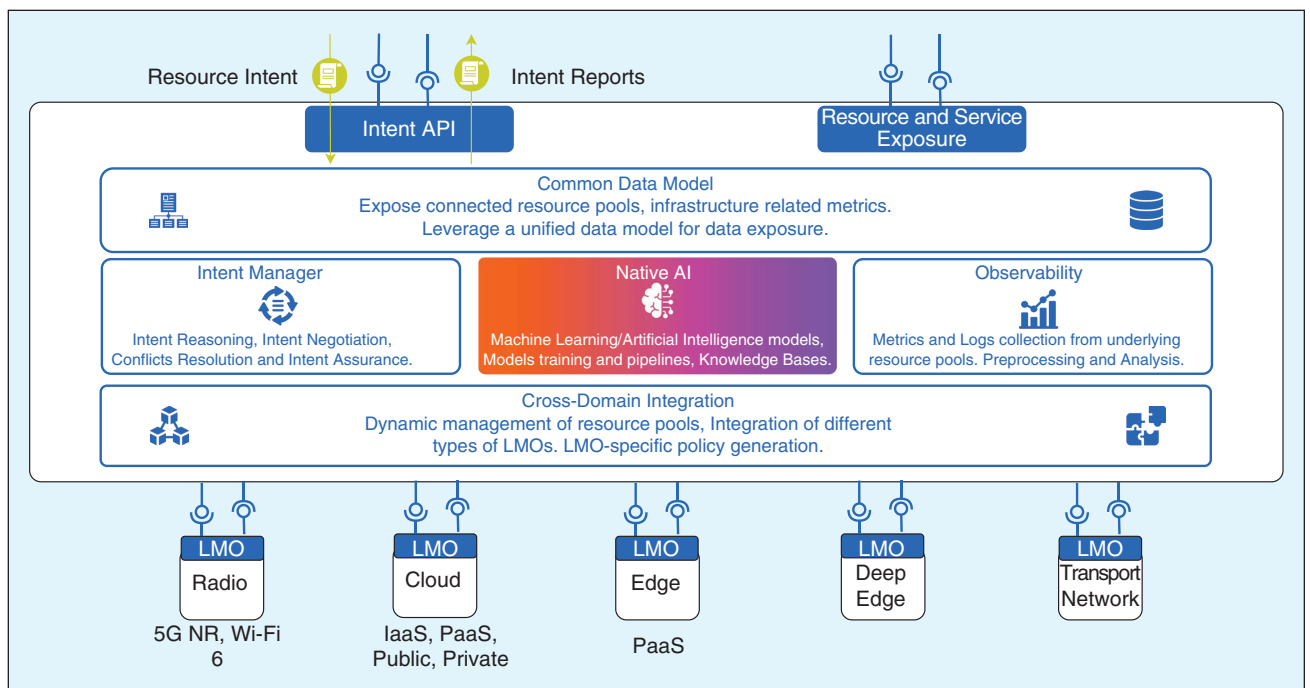


FIGURE 5 NCF architecture design.

specify high-level intents in natural language, which the LLMs can then interpret and map into structured requests recognized by the underlying infrastructure. Conversely, the NCF layer uses LLMs to translate structured, technical responses from the LMO back into natural language for natural language-based intent reporting, making it accessible to NCF consumers.

The abstraction framework interacts with each resource pool using interfaces exposed by the LMOs. This significantly expands opportunities for future expansion by dynamically adding or removing resources (both long-term, such as cloud resources, and volatile, such as deep-edge devices) without requiring detailed bookkeeping from the DMO, which would otherwise create bottlenecks. While the LMO handles the details of managing and orchestrating the resources, the abstraction framework is responsible for fully automating FCAPS (fault, configuration, accounting, performance, and security) in the region, resolving conflicts across LMOs from different technological domains (RAN, transport, compute). The NCF workflow is presented in Figure 4. To this end, the “fabric loop” continuously reconciles intents received from the DMO with the state of all (network-compute) southbound platforms via a local deployment of the Native AI toolkit. More specifically, the NCF abstraction framework offers:

- *Autoconfiguration*: Leveraging generative AI to create on-the-fly runtime configurations required by the vast ecosystem of LMOs (e.g., Nephio,¹ Open-RAN,² EdgeX Foundry,³ etc.)
- *Self-healing*: Acting as an escalation point to handle failures and capacity exhaustion by enacting adaptation primitives (e.g., service migration, service operation in “degraded” quality-of-service mode, or service termination and escalation to the DMO via intent reports generated by the Native AI HRL mechanisms)
- *Self-optimization*: Detecting and addressing conflicts (e.g., due to conflicting service requirements for the multiple underlying platforms) and processing intent reports where potential performance degradation (e.g., due to workload variations) or deviation from the target SLAs is reported
- *Abstraction*: Leveraging common data models understandable by the DMO to abstract the underlying multistakeholder complex infrastructure.

The NCF will continuously expose information regarding resource usage and KPIs, allowing the DMOs to make high-level decisions without considering low-level details handled by the NCF. Following this approach, Internet-scale orchestration can be achieved, and scalability issues related to one-to-one mapping of current state-of-the-art solutions can be resolved.

¹<https://nephio.org/>

²<https://www.o-ran.org/>

³<https://www.edgexfoundry.org/>

Conclusion

The evolution toward next-generation networks, namely 6G, is impacted by the high expectations and constraints imposed on the infrastructure. It is expected to achieve unprecedented performance and enable novel applications. Realizing such ambitions will require rethinking how network and computing infrastructures are managed. The 6G-INTENSE architecture takes the first step toward this ambition by proposing a novel architecture that pushes the boundaries and limitations of the current management architecture to a new level. In this article, we discussed the innovations brought by the proposed architecture, as well as our vision of how it can impact new businesses with the rise of novel actors and stakeholders. This is the first step toward decentralized NMSs that enable the novel 6G ecosystem.

Acknowledgment

This work is supported by the European Union’s Horizon Program under the 6G-Intense project (Grant 101139266).

Author Information



Akram Boutouchent (akram.boutouchent@eurecom.fr) is a Ph.D. candidate in the Communication Systems Department at EURECOM, CS 50193-06904 Biot Sophia Antipolis cedex, France. He received his engineering degree and M.Sc. in computer science from the Ecole Nationale Supérieure d’Informatique in Algiers, Algeria, in September 2022. His research interests include artificial intelligence, network automation, and intent-based networking for next-generation wireless networks. He is a Member of IEEE.



Abdelkader Mekrache (abdelkader.mekrache@eurecom.fr) is a Ph.D. candidate in the Communication Systems Department at EURECOM, CS 50193-06904 Biot Sophia Antipolis cedex, France. His primary focus is advanced network management frameworks in next-generation wireless networks. He is an active participant in collaborative research, and notably contributes to the OpenAirInterface project. His research interests include next-generation networking, and artificial intelligence for beyond 5G networks. He is a Member of IEEE.



Adlen Ksentini (adlen.ksentini@eurecom.fr) is a professor leading the Network Softwarization group in the Communication Systems Department at EURECOM, CS 50193-06904 Biot Sophia Antipolis cedex, France. His research interests include network virtualization, software-defined networking, and edge computing for 5G/6G networks. He has participated in several H2020 and Horizon Europe projects and is the technical manager for 6G-Intense and AC3.


He is a member of the OpenAirInterface Board of Directors and a Senior Member of IEEE.



Gereziher Adhane (g.adhane@iquadrat.com) is a researcher at Iquadrat, 08006 Barcelona, Spain. He received a Ph.D. degree in computer science from the Universitat Oberta de Catalunya. His research interests include artificial intelligence and computer vision, with a special emphasis on large language models, explainable artificial intelligence, and reinforcement learning. His work explores the application and implementation of artificial intelligence techniques in various applications.



João Pedro Celestino da Fonseca (j.fonseca@iquadrat.com) is a researcher at Iquadrat, 08006 Barcelona, Spain, and a computer science Ph.D. candidate pursuing a joint degree from the Universities of Minho, Aveiro, and Porto, hosted by Instituto de Telecomunicações. He actively promotes software freedom and advocates for the Linux kernel as a member of GLUA-Grupo de Linux da Universidade de Aveiro. He is a Graduate Student Member of IEEE.



Joseph McNamara (joseph.mcnamara@ericsson.com) is a senior researcher with Portfolio Architecture & Research, Ericsson, Westmeath N37 PV44 Athlone, Ireland. Joining Ericsson in 2018 and pursuing a Ph.D. degree has led to experience working on both current network management solution products and future innovation and technologies. Initially starting with rules-based management, his primary research interest is in network management automation through a mixture of Intent-based management and GenAI.



Kostas Ramantas (kramantas@iquadrat.com) is a senior researcher at Iquadrat, 08006 Barcelona, Spain. He received a Diploma of Computer Engineering, M.Sc. degree in computer science, and Ph.D. degree from the University of Patras, Greece. He has been the recipient of two national scholarships and has published more than 35 journal and conference papers.



Marco Palena (marco.palena@cnit.it) is a researcher at Consorzio Nazionale Interuniversitario per le Telecomunicazioni, 10129 Turin, Italy. His research interests include automated reasoning, distributed machine learning, and explainable artificial intelligence. He received his Ph.D. degree in computer and control engineering from Politecnico di Torino in 2017.

Marius Iordache (marius.iordache@orange.com) is a Telco Cloud and Service Platforms manager at Orange, 010665 Bucharest, Romania. He received a degree in engineering and telecommunications from University



Polytechnics Bucharest in 2005. He is currently involved in several research projects related to 5G and 6G areas, focus of interest for networks, AI/ML, automation, orchestration, and network programmability in general. He co-authored of more than 30 papers in journals and conferences.



Renato Lo Cigno (renato.locigno@unibs.it) is a full professor of communication networks at the University of Brescia, 25121 Brescia, Italy. He received his Electronic Engineering degree from Politecnico di Torino and was a visiting scholar at the University of California, Los Angeles. He is area editor for Elsevier Computer Communications and chair of the Wireless On-demand Network Systems and Services Conference Steering Committee. His research interests include network performance, vehicular networks, and systems beyond 5G. He is a Senior Member of IEEE.



Salwa Mostafa (salwa.mostafa@oulu.fi) is a postdoctoral researcher at the CWC-Radio Technologies, Information Technology and Electrical Engineering, University of Oulu, 90570 Oulu, Finland. She received her Ph.D. degree from the Department of Electrical Engineering, City University of Hong Kong. Her research interests include wireless communication, deep learning, radio frequency, and nonorthogonal multiple access. She is an Affiliate of IEEE.



Swastika Roy (s.roy@iquadrat.com) is a research engineer at Iquadrat Informatica SL, 08006 Barcelona, Spain. She received her Ph.D. degree in signal theory and communication engineering from the Technical University of Catalonia, Barcelona, Spain. Previously, she worked as an early stage researcher at the Centre Tecnològic de Telecomunicacions de Catalunya, Barcelona, Spain, as part of the H-2020 Marie Skłodowska Curie ITN SEMANTIC project.



Christos Verikoukis (cveri@isi.gr) is a professor with the University of Patras, 265 04 Patras, Greece, and a collaborating faculty member with the Industrial Systems Institute/Athena Research Center, 265 04 Patras, Greece. He received his Ph.D. degree from the Technical University of Catalonia, Barcelona, Spain. He has authored 170 journal articles and more than 250 conference papers. He is currently the editor-in-chief of the *IEEE Networking Letters*. He is a Senior Member of IEEE.

References

- [1] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Netw.*, vol. 34, no. 3, pp. 134–142, May/Jun. 2020, doi: 10.1109/MNET.001.1900287.

- [2] W. Wu et al., "AI-native network slicing for 6G networks," *IEEE Wireless Commun.*, vol. 29, no. 1, pp. 96–103, Feb. 2022, doi: 10.1109/MWC.001.2100338.
- [3] A. Leivadreas and M. Falkner, "A survey on intent-based networking," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 1, pp. 625–655, 1st Quart. 2023, doi: 10.1109/COMST.2022.3215919.
- [4] D. A. Cooperson, "Intent in autonomous networks," TMForum, IG 1253, version 1.3.0., Aug. 2022. <https://inform.tmforum.org/research-and-analysis/reports/intent-in-autonomous-networks>
- [5] A. Clemm, L. Ciavaglia, L. Z. Granville, and J. Tantsura, "Intent-based networking - Concepts and definitions," RFC 9315, Oct. 2022. <https://datatracker.ietf.org/doc/rfc9315/>
- [6] K. Ramantas et al., "6G-bricks: Building reusable testbed infrastructures for cloud-to-device breakthrough technologies," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2023, pp. 751–756, doi: 10.1109/GCWkshps58843.2023.10465223.
- [7] M. A. Uusitalo et al., "6G vision, value, use cases and technologies from European 6G flagship project Hexa-X," *IEEE Access*, vol. 9, pp. 160,004–160,020, 2021, doi: 10.1109/ACCESS.2021.3130030.
- [8] S. Kerboeuf et al., "Design methodology for 6G end-to-end system: Hexa-X-II perspective," *IEEE Open J. Commun. Soc.*, vol. 5, pp. 3368–3394, 2024, doi: 10.1109/OJCOMS.2024.3398504.
- [9] A. Boutouchent et al., "AMANOS: An intent-driven management and orchestration system for next-generation cloud-native networks," *IEEE Commun. Mag.*, vol. 62, no. 6, pp. 42–49, Jun. 2024, doi: 10.1109/MCOM.003.2300367.
- [10] A. Mekrache and A. Ksentini, "LLM-enabled intent-driven service configuration for next generation networks," in *Proc. 10th IEEE Int. Conf. Netw. Softwarization (NetSoft)*, 2024, pp. 253–257.
- [11] A. Mekrache, M. Mekki, A. Ksentini, B. Brik, and C. Verikoukis, "On combining XAI and LLMs for trustworthy zero-touch network and service management in 6G," *IEEE Commun. Mag.*, early access, Nov. 4, 2024, doi: 10.1109/MCOM.002.2400276.
- [12] M. A. Habib et al., "Intent-driven intelligent control and orchestration in O-RAN via hierarchical reinforcement learning," in *Proc. IEEE 20th Int. Conf. Mobile Ad Hoc Smart Syst. (MASS)*, 2023, pp. 55–61, doi: 10.1109/MASS58611.2023.00015.
- [13] W. Li, Y. Lemieux, J. Gao, Z. Zhao, and Y. Han, "Service mesh: Challenges, state of the art, and future research opportunities," in *Proc. IEEE Int. Conf. Service-Oriented Syst. Eng. (SOSE)*, 2019, pp. 122–1225, doi: 10.1109/SOSE.2019.00026.

VT