AI-Native CECC Management architecture: Enabling Scalable and Intelligent Cloud-Edge Computing

Abd Elghani Meliani
Eurecom
meliani@eurecom.fr

Adlen Ksentini
Eurecom
adlen.ksentini@eurecom.fr

mekki mohamed
Eurecom
mohamed.mekki@eurecom.fr

abdelhak kadouma
Fingletek
abdelhak.kadouma@fingletek.com

Dimitrios Amaxilatis
Spark works
d.amaxilatis@sparkworks.net

Amadou Ba
IBM
amadouba@ie.ibm.com

Eduardo Ojeda Coronado
Iquadrat Informatica
j.ojeda@iquadrat.com

John Beredimas
cloud software groupe
john.beredimas@cloud.com

Vrettos Mouloss
University of Piraeus
vrettos@unipi.gr

Souvik Sengupta
Ionos
souvik.sengupta@ionos.com

Dimitris Klonidis
ubitech
dklonidis@ubitech.eu

Christos Verikoukis
ISI/Research Centre ATHENA
veri@isi.gr

***Abstract*—The Cloud-Edge Computing Continuum (CECC) extends IT provisioning beyond centralized data centers, enabling low-latency and high-data-rate applications. This paper introduces an AI-native and cognitive CECC Management (CECCM) architecture, developed within the European project AC3, to enhance scalability, optimize resource allocation, and ensure SLA compliance. CECCM leverages AI, machine learning, and a microservices-based approach to enable efficient deployment, migration, and resource federation. Additionally, a data management Platform-as-a-Service (PaaS) streamlines data handling. CECCM's innovations in life cycle management, infrastructure optimization, and secure resource federation establish it as a key enabler of intelligent and scalable cloud-edge computing.**

***Index Terms*—Application Lifecycle Management, Artificial Intelligence, CECCM, Cloud-Edge Contin- uum, Data Management, Microservices, System Design.**

I. Introduction

Over the past two decades, cloud computing supported by virtualization technologies has been the dominant paradigm, with cloud providers offering a wide range of services—from network management to security, and cloud native engineers operating at top, deploying cloud native applications went well. However, the emergence of latency-sensitive applications such as AR/VR, video streaming, and IoT-based systems with real-time sensor data has highlighted the need to deploy services closer to end users. This shift has led to the emergency of edge computing, and subsequently, far-edge computing. These paradigms involve placing smaller-capacity computing devices geographically near users to reduce latency and improve overall user experience. This has led to the emergence of a new concept known as the cloud-edge continuum, where cloud and edge devices are treated as a single deployment unit. In this model, the cloud serves as the central provider of computing resources, while edge devices support functionalities that require low latency and fast user-to-application data transfer. While the cloud-edge continuum brings significant benefits, it also introduces several challenges: (i) Deploying and managing edge devices across various locations is costly and complex. To address this, industry players are increasingly advocating for resource federation, where application owners do not need to own the infrastructure on which their applications run. While this approach offers flexibility and scalability, it also raises new challenges — notably, application owners must now trust third-party infrastructures to uphold service-level agreements (SLAs) and protect sensitive application data. (ii) The ability to shift execution between cloud and edge environments introduces infrastructure heterogeneity, as applications may run across diverse platforms. For instance, some infrastructures rely on Kubernetes, while others may use OpenShift or alternative container orchestration technologies, adding complexity to the orchestration process. (iii) Traditional cloud-native strategies — such as those for migration, energy optimization, and cost management — become harder to implement in a cloud-edge context due to the variety of input parameters and the dynamic nature of the environment. The Agile and Cognitive Cloud-Edge Continuum Management (AC3) project was initiated to address the challenges inherent in managing the cloud-edge continuum. It introduces a novel AI-native and cognitive Cloud-Edge Continuum Management (CECCM) architecture designed to simplify and optimize the deployment and utilization of cloud-edge resources.

In our previous works, we have addressed key cloud-edge continuum challenges — such as stateful container migration [1], network programmability [2], training models [5] ,security and trust [3], cold and hot data management [4] among others — this paper focuses specifically on presenting the CECCM architecture. The goal is to demonstrate how CECCM can serve as a foundation to standardize future approaches to cloud-edge continuum management. We believe that CECCM will play a critical role in enabling scalability and agility in the management of IT resources within federated infrastructures. To the best of our knowledge, this is the first work of its kind to address the challenges of the cloud-edge continuum in both federated infrastructures and cloud environments. Our approach is built upon a novel infrastructure trust model [3], which empowers application owners to select suitable infrastructures based on trust and compliance with their requirements. The CECCM introduces innovation across three key dimensions to provide the agility required for application deployment and execution, while ensuring efficient IT and network resource management, including energy optimization: (i) Advanced Application Lifecycle Management (LCM): This goes beyond managing the application and its components to also include the associated data sources and the enforcement of agreed Service Level Agreements (SLAs) with application developers. (i) Comprehensive IT and Network Resource Optimization: This includes optimizing computing resources across the entire continuum—from the cloud to the far edge—as well as the underlying network infrastructure. (3) Trusted Resource Federation: This addresses the scarcity of edge resources by enabling shared, federated infrastructure usage. It is especially critical to manage and process large volumes of data (data deluge) efficiently at the edge. The CECCM will leverage state-of-the-art AI/ML algorithms to drive its operations and decision-making processes. It is essential for the CECCM to proactively predict events in order to achieve two main objectives: (i) optimize IT and network resource usage, and (ii) ensure compliance with application SLAs. To support this, the CECCM introduces the concepts of resource and application profiles, which guide decision-making processes related to microservice placement and migration. AI/ML algorithms, combined with a scalable monitoring system, will be employed to continuously analyze the current system load and the volume of data that applications need to process. This enables the dynamic creation and updating of both application and infrastructure profiles, allowing for more intelligent orchestration decisions. Additionally, the CECCM will natively integrate data management procedures within a Platform-as-a-Service (PaaS)-like component. This integration is designed to support application developers in building applications that rely on either stored data or live data streams, such as those collected from sensors or logging modules.

To summarize, the key contributions of this paper are:

- A novel CECCM management framework that handles the life cycle management of microservice-based applications deployed across a federated Cloud-Edge Continuum infrastructure.
- A trusted and secure federated infrastructure, designed to address resource scarcity—particularly at the edge—by enabling secure resource sharing across domains.
- Resource and application profiling mechanisms that classify infrastructures and applications into various types, facilitating SLA-aware orchestration and improved decision-making.
- A Data Management-as-a-Service (DMaaS) component that supports the seamless integration of data (both stored and real-time) throughout the application life cycle.

The remainder of this paper is organized as follows: Section II reviews related work and compares existing approaches to the proposed AC3 CECCM framework. Section III provides a detailed explanation of the CECCM architecture and the interaction between its components. Finnaly, the last section concludes the paper and outlines our future research directions.

II. Related Work

K. Katsaros et al. [6] introduced REASON, a project that tackles technical challenges in future network deployments, including end-to-end (E2E) service orchestration, sustainability, security and trust management, and policy enforcement. The project adopts AI-native principles and integrates multiple access technologies with cloud-native solutions. Although initially designed to address both technical and non-technical challenges of 6G networks, REASON remains relevant to our work. However, our approach offers a broader applicability beyond 6G, as REASON is primarily focused on that domain. Moreover, the proposed system architecture in REASON is overly complex, making real-world deployment difficult, and the paper lacks technical depth in critical components. In contrast, our solution is built on a federated infrastructure, making it more practical and realistic, especially considering the challenges of acquiring and managing edge resources across geographically distributed locations. Nautilus [7] presents an Adaptive Resource-Efficient Microservice Deployment system aimed at ensuring the Quality of Service (QoS) for microservice-based user-facing services while minimizing computational resource usage. It achieves this through SLA-aware scheduling, resource contention-aware allocation, and I/O-sensitive microservice migration. While Nautilus improves certain deployment metrics like latency, it overlooks several critical dimensions, such as deployment cost, energy efficiency, security, and user/application data management—elements that are integral to our CECCM architecture by design. 5G-VIOS [8] proposes an intelligent network service orchestration platform tailored for 5G and beyond, aligned with the Zero-touch Network and Service Management (ZSM) paradigm. It supports automated network slicing and the lifecycle management of network services across heterogeneous administrative and technological domains. The platform employs AI and machine learning techniques to profile services and predict optimal resource configurations, aiming to fulfill performance requirements and Service Level Agreements (SLAs) efficiently.

Table I summarizes the comparison between these state-of-the-art solutions and our proposed CECCM framework in areas like microservices management, Network Programmability (NP) End to End Orchestration (E2E) runtime and data management among others.

TABLE I
COMPARISON OF CECCM WITH STATE-OF-THE-ART SOLUTIONS

Solution	Micro	Edge	F-Edge	Energy	security	Feder	Trust Policy	NP	SLAs	Profiling	ML/AI	Monitoring	E2E Orch	RMgmt	DMgmt
REASON	✓	✓	✓	✓	✓				✓	✓	✓	✓	✓	✓	
Nautilus	✓	✓						✓	✓	✓	✓	✓	✓	✓	
5G-VIOS	✓	✓						✓	✓	✓	✓	✓	✓	✓	
CECCM	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

III. AC3 CECCM

The CECCM proposed architecture aims to manage the lifecycle of service-based applications in a Cloud-Edge Continuum Computing infrastructure using advanced AI/ML algorithms. Its primary goals are to ensure application QoS and optimize resource usage. Key features include:

- **Dynamic Application Definition:** Developers can dynamically define applications (components and SLAs) via an Application Gateway with a Northbound Interface (NBI) and GUI.
- **Microservice Composition and Reuse:** Supports composing and reusing microservices using a unified composite catalogue that includes data and service descriptions.
- **Application Descriptor Generation:** Automatically generates application descriptors from the user intents and enforce it with the AI-driven LCM.
- **Zero-Touch Orchestration and Management:** Full automation of the application lifecycle with AI/ML models.
- **Energy-Aware Resource Management:** Optimizes energy efficiency through intelligent resource management.
- **Dynamic Resource Composition:** Dynamically discover and allocates computational, networking, and storage resources to meet application needs.
- **Federation of CECC Resources:** Uses the IEEE SIIF standard to federate CECC resources across infrastructures.

The architecture follows a modular design, resulting in the CECCM, the core component of the AC3 system. CECCM consists of independent modules organized into two planes: the *User Plane and Management Plane* Additionally, the CEC Plane represents the *federated infrastructure layer* where application microservices are deployed the hight level Architecture of the CECCM framework is presented in figure 1.

A. The User Plane

This layer interfaces directly with application developers, primarily through the Application Gateway—a key component of the CECCM. It allows developers to describe, deploy, and manage microservice-based applications. The gateway outputs an application descriptor (in JSON or YAML) detailing each microservice's SLA, resource needs, and network QoS requirements (bandwidth and latency). It also includes metadata like overall SLA and service availability. A dedicated section lets developers define required data sources: hot data (from live sensors, which can be registered via the descriptor) and cold data (pre-existing datasets described in the Service Catalogue).

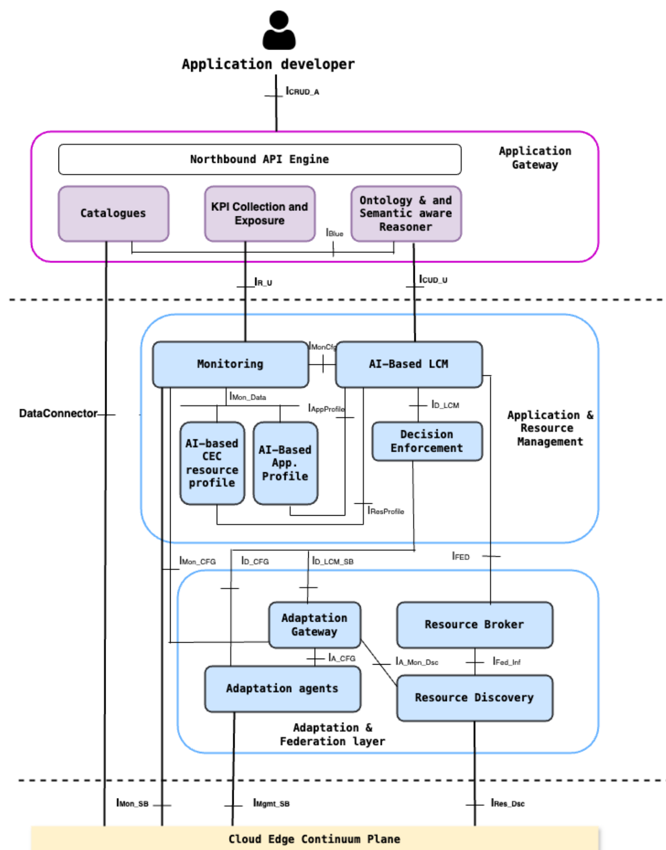

Fig. 1. CECC management architecture

To streamline descriptor creation, the CECCM integrates two key modules: the Ontology Semantic Reasoner (OSR) and the Service/Data Catalogues. Developers express high-level intents through a GUI provided by the OSR, which then translates these into a machine-readable application descriptor. The Service Catalogue offers reusable application blueprints, while the Data Catalogue lists available federated datasets—aligned with Gaia-X, a European initiative for secure data sharing and monetization [1]. OSR supports microservice composition by appending relevant services from the catalogues based on developer needs—for example, automatically adding security and authentication microservices when requested. Similarly, data requirements are handled as detailed in Section V. A KPI collection and exposure module provides real-time insights into microservice-level metrics like CPU, memory usage, and

[1] https://gaia-x.eu/

network traffic.

B. The Management Plane

This is the core of the CECCM, orchestrating the lifecycle of applications and associated data using AI-driven and data-aware strategies, all while accounting for the underlying CECC infrastructure. It consists of two main components:

1) Application and Resource Management (ARM): Responsible for managing microservice-based applications and optimizing the use of CECC resources. It includes several intelligent modules:

- **AI-based LCM:** This module handles the LC of each micro-service composing the application. Its role is to instantiate (i.e., deploy) the micro-services on the CECC infrastructure and manage their dedicated resources during their life cycle. This module receives a request to create an application (i.e., instantiate the application) through the *ICUDU* (sub-interface of lCRUDU introduced in the precedent section) interface in the form of a descriptor (e.g., JSON or YAML). The descriptor will include the software image of the application (i.e., several images, one for each micro-service) along with information on the SLA, such as the needed CPU, memory, networking (latency, bandwidth), availability, etc. The descriptor is issued by the OSR component of the user plane. Besides the application descriptor, the AI-based LCM relies on inputs from AI-based Application Profile and AI-based CEC resource profile modules to decide about the initial placement of the micro-services composing the application. These modules are in charge of providing information on the application profile, if it exists, such as the type of traffic, the pattern of the traffic, dataintensive, etc.; and the status of the CECC resources available for the CECCM, such as computing, type of energy used, networking, etc. Those inputs will be used by AI-based algorithms that derive the optimal decision for placing the micro-services to ensure SLA while optimizing the CECC resources (e.g., energy consumption). Apart from the initial placement, this module is in charge of the runtime management of the application, which consists of tracking the status of the application SLA (i.e., micro-service states) through the AI-based App. Profile and AI-CEC resource profiles to detect any degradation and react accordingly by, for instance, increasing the computing resource dedicated to microservices (running a scale-up process) or migrating part of the micro-service from one computing to another computing node (migrating from central cloud to edge to reduce latency) or updating the network routing configuration (request for more bandwidth). Finally, this module configures the level of monitoring of a newly created application (KPI to collect as well as the frequency to collect the monitoring data) to the monitoring module using *IMonCFG* interface.
- **AI-based Application Profile:** In this module, the AI and ML algorithms will run to create and build the application profiles, relying on the measurement obtained from the monitoring module. The application profile will be first initiated using the application definition provided by the application developer (via the application descriptor) and enriched by the data management module regarding the data sources and the storage policy. Then, the profile will be constantly improved through ML algorithms that monitor the application behavior to answer critical questions such as what type of traffic? When and where is the traffic and data generated? What is the data pattern? The built profiles will be made available to the AIbased LCM, which in conjunction with the SLA and the infrastructure available (or predicated) resources, realizes procedures related to the initial placement of the micro-services and during the run-time management of the application life cycle (such as migration, duplication, resource scaling). The profile will be exposed to the AIbased LCM module using the *I_App_Profile* interface.
- **AI-based CECC Resource Profile:** Similarly to the AIbased application profile, this module creates profiles of the CECC infrastructure used by the CECCM (owned or federated). The resource profile uses the consumed monitoring data (mainly using the resource exposure of each Local Management System (LMS)) to create the profiles, which may include current information on the available computing resources (i.e., CPU, memory, storage), CECC network link status (bandwidth, latency, reliability, etc.), geographical location, type (i.e., cloud, edge or far-edge) and type of used energy (green, brown); but also employ AI/ML algorithms that predict the infrastructure status. As stated earlier, the resource profiles will be combined with the application profile to drive LCM actions to be decided by the AI-based LCM module. The profile will be exposed to the AI-based LCM module using the *I_Res_Profile* interface.
- **Decision Enforcement:** Executes LCM decisions such as deploying, migrating, or removing microservices. It interacts with the abstraction and federation layer to translate high-level decisions into low-level infrastructure commands.

2) Abstraction and Federation Layer: The abstraction and federation layer enables seamless interaction between CECCM and the federated infrastructures, each managed by its own LMS, such as Kubernetes, OpenShift, or SDN controllers. Since each LMS uses a specific Northbound Interface (NBI) to expose its CRUD APIs and resource information, this layer serves as an intermediary that abstracts the heterogeneity of these systems and unifies resource access. The key components of this layer are:

- **Adaptation Agents:** : they have the role of interacting with the infrastructure using the *lMgmtSB* interface (NBI interfaces of LMS). These agents depend on the LMS technology used by the infrastructure owner. For each technology (such as Kubernetes, Openshift, K3S, as well as SDN controllers), an adaptation agent is available at the adaptation and federation layer. The adaptation agents are configured by the adaptation layer when a new resource from the federation will be used using *(ID_CFG)* or after discovering new LMS's NBI through *ID_Mon_Dsc*. An adaptation agent directly interacts with the decision enforcement module of the management layer using the IDSB interface. An adaptation agent translates the LCM decisions of the decision enforcement module to LMS-specific requests using the NBI exposed by the corresponding LMS. The adaptation agent is an API end-point for a specific infrastructure(LMS) used by the management layer.

- **Resource Discovery Module:** This module gathers detailed information from each LMS via their resource exposure APIs (e.g., resource types, location, availability, energy source, QoS, monitoring KPIs, and API access details). It standardizes the data into a common model and exposes it to other modules through *l_inf* and *l_Mon_Dsc* interfaces. **Resource Broker Module:** When multiple infrastructures can satisfy a resource request, this module uses selection algorithms to choose the optimal one. It acts as a decision point for complementing CECCM resources with those from the federation. **Adaptation Gateway:** it has the role to associate an adaptation agent with a LMS of an infrastructure, which will be used later to enforce LCM decisions issued by the management layer (i.e., the decision enforcement module). The agent is configured using the NBI endpoint of the LMS obtained via the discovery module. Per the request of the decision enforcement module (using *I_CFG* interface) indicating the LMS of the infrastructure, the adaptation gateway provides the endpoint to use, which is typically an adaptation agent address. Finally, the adaptation gateway also exposes to the monitoring module the NBI endpoint exposed by LMS discovered via the resource discovery module.

C. Federated CECC Infrastructure Plane

AC3 envisions a federated CECC infrastructure that enables the CECCM to deploy microservice-based applications across a diverse range of infrastructure providers. This federation includes public/private cloud, edge, far-edge, and networking resources. Resource federation in CECC can support two scenarios:

- The CECCM owns some infrastructure and uses federation to complement local shortages (e.g., no edge resources in a specific region).
- The CECCM does not own any infrastructure and fully relies on federated resources.

The federation model follows the NIST reference architecture and integrates with the IEEE Standard for Intercloud Interoperability and Federation (SIIF). This is achieved through the Resource Discovery module, which serves as the point of contact with the Federation Hosting Service (FHS) to identify and retrieve services available in the federation. To participate in this federated environment:

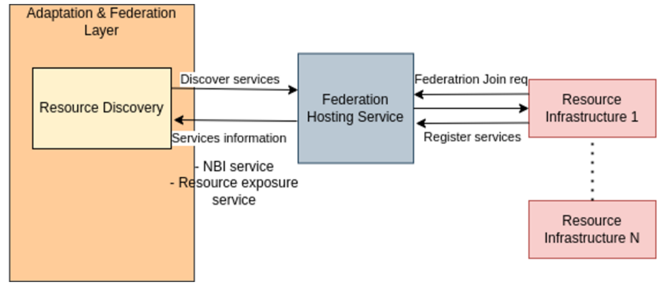

Fig. 2. CECCM interaction with the CEC federation

1) The CECCM must be registered as a federation member, with the necessary permissions to discover and consume services.
2) Any infrastructure willing to contribute resources must join the federation at the FHS level with a service/resource owner membership.
3) Upon joining, the infrastructure must register its key services—specifically, the LMS's Northbound Interface (NBI) and the Resource Exposure API—with the FHS.

Once registered, the Resource Discovery module retrieves this information and passes it to the Adaptation Gateway within the Adaptation and Federation layer, enabling integration and management within the CECCM. Importantly, the Federation Hosting Service (FHS) is considered external to the CECCM framework and is operated as an independent administrative entity. Figure 2 illustrates the CECCM's interaction with the CECC federation, aligned with the IEEE SIIF standard.

IV. Data Management

The proposed CECCM framework enables the application developer to deploy not only the application microservices but also configure and attach the data source ingested by the application, considering the case of IoT and other sensor-based applications, Since this is a key component of the CECCM framework, we address it in a dedicated section. Additionally, we have a separate publication that provides an in-depth analysis of how the AC3 manages data [4]. Indeed, the proposed framework eases the integration of data sources and automate the connection of the application to the data sources to ingest data. To achieve this objective, the following components are introduced:

- Data Catalogue (component of the user plane) that provides a comprehensive inventory of available data, including information about formats, types (live or stored), locations, and access methods. It offers all the necessary procedures and interfaces to register, federate, manage, and provide access to data spaces of application developers or third-party IoT infrastructure providers.
- Data Sources (Hot and Cold) refer to the locations where the actual data available to the applications of AC3 reside. These are categorized based on their access frequency, usage patterns, and lifecycle. HOT data is frequently accessed and requires fast retrieval times; often stored

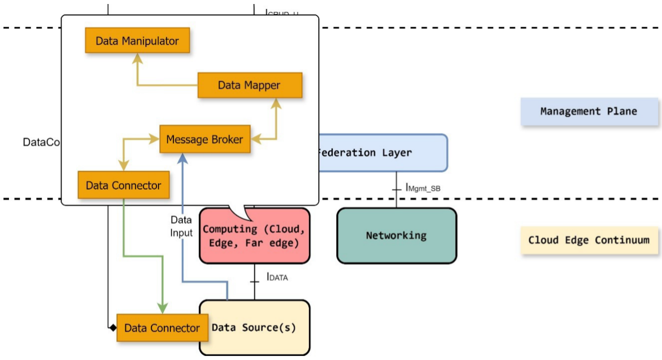

Fig. 3. Data management as a service

in high-performance storage for immediate use, with a live feed nature, where each new measurement is added to the available dataset once produced. Time is crucial for this type of data, and it needs to be processed as fast as possible. COLD data is accessed less frequently, stored in cost-efficient storage solutions, and may require more time to retrieve when needed based on its size and format.

- Data Source Connectors serve as the vital links that facilitate the flow of data between the continuum and external data sources. These connectors are engineered to accommodate a diverse array of data types, formats, and communication protocols, ensuring that data can be ingested from and dispatched to a wide range of sources, including IoT devices, cloud repositories, and edge computing platforms.
- Data Management Application Add-ons are components that live and operate inside the application's premises with a set of roles to access, retrieve, and process the data stored in the available Data Sources. As part of the application, the proposed OSR component appends the following components to the application descriptor facilitating the data managemnet procedure. Figure 3 illustrates the integration of the data management components to the CEC architecture. The add-on components of the CECCM framework enhance data handling and integration across the architecture. *Data consumer connectors* ensure efficient delivery of data to end-users or downstream systems by adapting data flow to consumption patterns. *Data mappers* transform data formats and structures to facilitate interoperability and seamless integration between heterogeneous systems. *The message broker* acts as a central communication hub, managing the flow of data requested by applications and routing it to the appropriate destinations. Finally, *the data manipulator* refines raw data through cleaning, transformation, and aggregation processes, making it more suitable for analysis, decision-making, or other targeted operations within the CECCM.

V. Conclusion

In this paper, we presented a novel AI-native and cognitive Cloud-Edge Continuum Computing Management (CECCM) architecture. CECCM improves scalability and agility in managing IT resources—such as computing, memory, and storage—while optimizing the deployment of emerging applications across cloud-edge infrastructures. The architecture introduces innovations across three main areas: (i) advanced application Life Cycle Management (LCM), which integrates data sources and service-level agreements (SLAs) for efficient deployment; (ii) optimized IT and network resource management, including far-edge computing; and (iii) a trusted resource federation to overcome edge resource limitations and handle data surges effectively. Our future work will focus on demonstrating the applicability of the CECCM architecture in real-world scenarios, ranging from IoT environments and smart monitoring systems to large-scale use cases involving the processing of hundreds of terabytes of astronomical data. These results, along with the detailed design and technical implementation of each CECCM component, will be presented in forthcoming publications.

Acknowledgment

This work was partially supported by the European Union's Horizon Research and Innovation program under AC3 project and grant agreement No 101093129.

References

[1] Mekki, M., & Ksentini, A. (2025). Resiliency focused proactive lifecycle management for stateful microservices in multi-cluster containerized environments. Computer Communications, 236, 108111.

[2] A. Mokhtari and A. Ksentini, "SD-WAN for Cloud Edge Computing Continuum interconnection," GLOBECOM 2024 - 2024 IEEE Global Communications Conference, Cape Town, South Africa, 2024, pp. 2533-2538, doi: 10.1109/GLOBECOM52923.2024.10901751.

[3] S. Messaoudi, A. -E. Meliani, A. Mokhtari and A. Ksentini, "Security and Trust Management in Cloud Edge Continuum: AC3 Project Approach," 2024 IEEE 29th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), Athens, Greece, 2024, pp. 1-6, doi: 10.1109/CAMAD62243.2024.10942825.

[4] D. Amaxilatis et al., "Data Management Over Cloud Edge Continuum in the Context of AC3," 2024 IEEE 29th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), Athens, Greece, 2024, pp. 1-7, doi: 10.1109/CAMAD62243.2024.10942767.

[5] Abd Elghani, M., Arora, S., Ksentini, A., & Knopp, R. (2025, June 8–12). TimeTrack: A dataset for exploring temporal patterns and predictive insights into OpenAirInterface (OAI) CI/CD cluster. In IEEE International Conference on Communications 2025 (ICC 2025) (Montreal, Canada). IEEE. © 2025 IEEE.

[6] K. Katsaros et al., "AI-Native Multi-Access Future Networks—The REASON Architecture," in IEEE Access, vol. 12, pp. 178586-178622, 2024, doi: 10.1109/ACCESS.2024.3507186.

[7] K. Fu, W. Zhang, Q. Chen, D. Zeng and M. Guo, "Adaptive Resource Efficient Microservice Deployment in Cloud-Edge Continuum," in IEEE Transactions on Parallel and Distributed Systems, vol. 33, no. 8, pp. 1825-1840, 1 Aug. 2022, doi: 10.1109/TPDS.2021.3128037

[8] Moazzeni, S., Katsaros, K., Ferdosian, N., Antonakoglou, K., Rouse, M., Kaleshi, D., ... & Simeonidou, D. (2024). 5g-vios: Towards next generation intelligent inter-domain network service orchestration and resource optimisation. Computer Networks, 241, 110202.