Bridging Cognitive Computing with Multi-Cloud Integration

Souvik Sengupta
Ali Nikoukar
Maria Barros Weiss
IONOS SE
Germany

Sara Madariaga
ARSYS INTERNET S.L.U.
Spain

Asimakis Lykourgiotis
Athanasios Kordelas
John Beredimas
NetScaler
Greece

***Abstract*—This paper presents a sophisticated architectural framework for Multi-Cloud Federation within the Cloud Edge Computing Continuum (CECC), with a primary focus on the seamless integration and interoperability of multi-cloud and edge computing systems, as delineated in the ongoing Agile and Cognitive Cloud-Edge Continuum Management (AC^3) project. This initiative is instrumental in bridging cognitive computing with multi-cloud integration, thereby optimizing cognitive processing at the network's periphery. This study meticulously addresses pivotal challenges such as data interoperability, latency management, resource allocation, and security, all of which are critical for ensuring scalable, efficient, and secure various cloud operations.**

Our proposed architecture leverages advanced distributed computing paradigms and AI-driven resource management strategies to effectively mitigate these challenges. The framework provides a detailed exposition of its components, their interactions, and the resultant benefits in terms of resilience, efficiency, and context-aware processing. Furthermore, we delve into the AC^3 project's endeavors to achieve robust cloud interoperability within the Cognitive Computing Systems (CCS) platform. Emphasis is placed on a rigorous validation plan comprising both simulation and real-world scenarios, underscoring a phased approach designed to corroborate the feasibility and adaptability of the proposed architecture in meeting contemporary computing demands. Our contributions are poised to significantly advance the state-of-the-art in cloud integration and catalyze the development of innovative applications across various sectors.

***Index Terms*—Multi-Cloud Federation, Cognitive Computing, AC^3, Data Interoperability, Internet of Things (IoT)**

I. Introduction

Cognitive Computing Systems (CCS) represent a significant advancement in artificial intelligence, embodying capabilities that allow them to adapt to situational contexts, understand and interact with environments, and engage collaboratively with humans and a variety of technologies. These systems are instrumental in the engineering of intelligent environments, which aim to enhance operational efficiency, improve comfort, and promote sustainability through the integration of interconnected systems and advanced data analytics. By synthesizing data from multiple sources, CCS facilitate informed decision-making processes and provide personalized, adaptive user experiences. Notable applications include the optimization of energy consumption, predictive maintenance for equipment failure, and the seamless integration of secure human-machine interactions through Natural Language Processing (NLP) and machine learning algorithms.

The initiative focused on Cognitive Computing with Multi-Cloud integration, conducted under the auspices of the Adaptive Cognitive Cloud (AC^3) project, endeavors to harmonize cognitive computing with multi-cloud integration. This synthesis is vital for creating a robust and efficient cognitive computing framework, allowing for seamless interoperability across various technologies, platforms, and services. Ensuring interoperability at both the infrastructure and application levels requires that disparate systems, applications, and services operate cohesively through standardized data formats, communication protocols, middleware solutions, and the connectivity of Internet of Things (IoT) devices. Federation nurtures a cohesive network of autonomously managed systems, enhancing data accessibility and usability through federated data management, resource orchestration, and federated identity and access management.

This level of interoperability and federation is essential for maximizing the potential of cognitive computing, providing scalability, flexibility, efficient resource distribution, load balancing, and fault tolerance, and consequently driving innovation and efficiency across diverse domains. However, the realization of effective multi-cloud federation confronts several challenges, including the complexity of management, data security concerns, establishing interoperability standards, and performance issues related to data transfer and network latency.

The AC^3 project addresses these challenges by proposing an advanced architectural framework that employs distributed computing alongside AI-driven resource management. This framework is designed to elaborate on the components involved, detail their interactions, and outline the consequent benefits in regards to resilience, efficiency, and context-aware processing. Moreover, the AC^3 project is committed to facilitating cloud interoperability within the CCS platform. This commitment is demonstrated through validation plans that span both simulation and real-world scenarios. This strategically phased approach is painstakingly crafted to illustrate the feasibility and adaptability of the proposed solutions in response to the evolving demands of modern computing landscapes, fundamentally advancing the integration of public and private cloud environments and inspiring innovative applications across various sectors.

The organization of this paper is structured as follows: Section II proffers a concise overview of various efforts aimed at achieving cloud interoperability. Section III explores the complexities inherent in multi-cloud federation within the context of cognitive computing, with a focus on guidelines from the National Institute of Standards and Technology (NIST) Cloud Federation and the Gaia-X initiative. Section IV initiates with a discussion on how the AC^3 project addresses the challenges faced by current federation models, followed by a comprehensive presentation of the Architectural Framework for Application Life Cycle Management. This section further explicates how the AC^3 framework ensures Service Level Agreement (SLA) compliance across diverse and potentially unreliable infrastructures for multi-cloud operations, subsequently detailing how NetScaler facilitates the management of stateful microservices across the computing continuum. Lastly, a structured guide is provided for conducting comparative analyses of existing multi-cloud solutions, setting the stage for prospective research and concluding the paper.

II. State of the Art

Conducting a comprehensive review of seminal works in the domains of multi-cloud federation and cognitive edge computing is critical for establishing a robust foundational groundwork. This review commences by examining the National Institute of Standards and Technology (NIST) Cloud Federation reference architecture [1], which outlines best practices and standardized methodologies for incorporating Artificial Intelligence (AI) and Machine Learning (ML) into cybersecurity frameworks. The focus is on model development and testing strategies that enhance security measures. Furthermore, the Gaia-X Framework [2] proposes a federated European data infrastructure, prioritizing interoperability, standardized communication protocols, and stringent data privacy controls. It delineates compliance guidelines to ensure that data exchange processes adhere to national laws and privacy regulations, thereby safeguarding data sovereignty and upholding regulatory standards.

The InterCloud Testbed [3] investigates resource management across multiple cloud providers, concentrating on key areas including interoperability, security, and scalability. This initiative aims to develop performance benchmarks for cloud providers, thereby facilitating seamless service integration across diverse platforms. In their analysis of multi-cloud and cognitive computing, Challita et al. [4] address interoperability challenges by employing frameworks such as FCLOUDS and promoting standardization efforts. Cognitive computing enhances resource allocation, security, and decision-making processes using tools like TensorFlow [5], [6]. Formal models such as Cloud4SOA and PaaSport further ensure seamless integration across cloud service providers [7]. These advancements collectively aim to optimize resource management and fortify secure cognitive computing platforms across various cloud infrastructures [8], [9]. Initiatives like the EUCloudEdgeIoT project[1] and the European Alliance for Industrial Data, Edge, and Cloud[2] contribute to enhancing interoperability and support the digital transformation process. Collectively, these efforts address challenges related to interoperability, security, and resource allocation, leveraging AI and ML to optimize performance and laying a strong foundation for future technological advancements. The subsequent section will explore the limitations inherent in the NIST Cloud Federation model and the Gaia-X Framework, providing insight into the complexities associated with multi-cloud federation.

III. Exploring the Complexities of Multi-Cloud Federation in the Cognitive Computing

Established federation models, such as those proposed by NIST [1] and Gaia-X [2], [10], face several challenges that need to be addressed to build a reliable computing infrastructure. In this section we present the main complexities identified by AC^3.

A. Security and Trust

Security and trust management are critical when we intend to federate on-the-fly resources and data from various domains and devices. In particular, far edge resources introduce unique challenges to the trust model of a federation. The NIST model considers a somewhat controlled membership approach, but does not accommodate quick onboarding of new members to the federation in a manner that still allows for: a minimum yet acceptable level of trust, a quick revocation process, minimal computing and other resources available, short on-the-fly membership timeframes, or wildly diverging APIs and capabilities.

B. Service Level Agreement

Meeting SLAs in these federated environment presents two main challenges. First, NIST has limited provisions for guaranteed SLA when implementing a federated environment. It mentions relevant topics related to SLA, such as low latency, selection of a specific region/location, or resource awareness,

[1] https://eucloudedgeiot.eu/

[2] https://single-market-economy.ec.europa.eu/industry/industrial-alliances/european-alliance-industrial-data-edge-and-cloud_en

but lacks concrete proposals in terms of architecture, deployment process implementation or infrastructure metrics. Second, the introduction of far edge resources poses additional challenges, namely mobility, temporary membership, varying network performance, and diverse hardware resources. These further increase the complexity of the SLA mechanism.

The main SLA-related challenges identified by AC^3 are:

- Identification of relevant SLA metrics, including: uptime/availability, network latency (between services and microservice components), environmental restrictions (i.e., 15% of energy coming from green sources), or throughput to far edge devices that generate data.
- Regional constraints introduced by certain use cases, such as agriculture, fire detection, or area security, that may explicitly require microservices deployment on far-edge devices within the area of interest.
- Availability constraints introduced by the unpredictable nature of the far-edge, including: varying network availability and performance, limited computing resources, or reduced feature sets compared to servers available in the near edge and core networks.
- Past/recent SLA trends of the components. For instance, a brown datacenter may be chosen despite a stated preference for green energy if during subsequent cloudy days, the green datacenter has exhibited outages or failures.
- Particular application needs. Ultimately, the application stakeholder determines the requirements, both at application and microservice levels.

C. *Stateful Microservices*

The deployment shift to a container-based microservice model hosting poses additional challenges, particularly when the pods have storage needs. While the NIST framework is broad enough to consider a federation of arbitrary resources, it is mostly oriented towards VMs or higher-level abstractions like DBaaS or PaaS and does not overly consider network and data-intensive microservices.

We provide below the main challenges raised by seamless service migration of stateful microservices.

- When migrating microservices to a different pod, we need to consider complex dependencies between internal volumes/databases and with shared databases in a holistic manner, since the end-to-end service can only be resumed after the complete migration of the microservice and all its dependencies completed managements.
- Network stability, both in terms of availability, latency, and throughput across all the continuum components.
- Energy management of far-edge devices, since any additional edge processing needs to be supported by the energy capacity of the corresponding device.
- Environmental considerations should weigh into any migration decision, including the optimization of energy consumption and a preference for "green" resources.
- Data replication, semantics and data consistency across a variety of storage technologies need to be guaranteed. The system needs "mapping files" acting as an intermediate layer agnostic of vendor-specific implementations.

D. *Lifecycle management across non-uniform infrastructure*

Another challenge not covered by NIST and many of the so-called "hyperscalers", which already provide federated secure data infrastructures, is the integration of edge and far-edge resources belonging to end-users or (very) small organizations. This is a problem of such complexity that providers require end-users and customers to run well-defined software and/or hardware platforms in order to accommodate a common lifecycle management experience across both their public cloud resources and on-premise infrastructure.

The associated challenges are outlined in the previous sections and are related to the unpredictable short/mid-term future state of far-edge components, which is not currently described by NIST.

E. *Data management and federation*

The challenges of data management are huge since multiple heavy operations/jobs, data repositories, and live data streams (e.g., IoT sensors, monitor functions, etc.) need to be handled with robustness. Open-source or enterprise data projects like Apache Spark, Apache Delta-Lake, Apache Dremio, Airbyte, or Gaia-X need to be carefully selected or used as guidelines. We provide below a list of the main identified challenges.

- Secure and authenticated data collection, data governance and sharing mechanisms across business domains. The existence of far-edge devices raises extra challenges since they often connect/disconnect and are not stable members of the federation.
- Web interfaces and API servers, which provide easy access to cloud resources and data, including streaming, transactional, structured and unstructured data as well as online queries. Special handling is needed for the varying latency and connection drops associated to far-edge devices. For example, when a data query is running in a distributed manner, and one of the nodes has a connection drop, then this microservice should be migrated to another node. However, in the case of connection degradation, it should be decided whether to wait for the response or migrate to another node and restart the processing.

- Need to work with existing, legacy, distributed, and monolithic technologies of different operating systems and programming languages. Especially far-edge devices may have fewer hardware resources, run different architectures, or have more restricted runtimes and frameworks.
- Cost and energy reduction, especially for big-data, where the cloud costs can be very high, and for far-edge devices, where autonomy is critical.
- Automated data discovery, data evaluation, meta-data, and feature extraction. The development challenge is to implement multiple clients for the integration of different data solutions, allowing to automatically identify the schema of the data source before proceeding with the evaluation, metadata, and feature extraction.

IV. Multi-Cloud Federation: An Architectural Proposal

A. How AC^3 Addresses the Challenges Faced by Current Federation Models

The Adaptive Cognitive Cloud Collaboration Model (Cloud-Edge-Computing Continuum Manager (CECCM)) introduced by the AC^3 project offers a pioneering approach to surmount the myriad challenges delineated in Section III. This architectural framework is quintessential for bolstering scalability and agility in managing heterogeneous Information Technology (IT) resources - specifically computing, memory, and storage—required by burgeoning applications throughout their life cycles. The framework meticulously addresses three fundamental dimensions to effectively execute and deploy applications while optimizing IT and network resources for enhanced resource efficiency and reduced energy consumption.

- **Sophisticated Application Life Cycle Management (Life Cycle Management (LCM)):** This incorporates comprehensive oversight of an application's life cycle and its components, including the data sources it depends on and the Service Level Agreement (Service Level Agreement (SLA)) established with the application developer.
- **CECC IT Components:** This encompasses optimization of far-edge and networking resources.
- **Trusted Resource Federation:** This addresses the scarcity of resources, particularly at the edge, essential for the optimal management of data influx. The CECCM integrates cutting-edge Artificial Intelligence (AI), Machine Learning (ML), and semantic and context-awareness algorithms to proficiently manage the application LCM, alongside cloud, edge, far-edge, and networking resources. Anticipating events and exhibiting proactivity are imperative for the CECCM to optimize IT and networking resources, ensuring adherence to application SLAs. Artificial Intelligence (AI) algorithms, in conjunction with a scalable monitoring system, evaluate current and forecasted loads and data requirements, enabling accurate prediction of application resource demands.

B. Architectural Framework for Application Life Cycle Management

Figure 1 illustrates a sophisticated, multi-layered architecture devised to streamline application life cycle management within a cognitive computing continuum framework. This framework employs advanced AI techniques, alongside NetScaler Global Server Load Balancing (Global Server Load Balancing (GSLB)) [11], and load balancing capabilities within IONOS and Arsys cloud infrastructures for effective resource allocation and management. These components enhance interoperability and support multi-cloud operations, fostering a seamless and dynamic operational milieu. The architecture harnesses AI for strategic decision-making and optimization processes, ensuring proficient management of applications across a federated CECC infrastructure.

The architecture seamlessly integrates a multitude of components, intercommunicating through a rigorous API across three planes: the User Plane, the Management Plane, and the CECC Plane. A notable innovation within the User Plane is the introduction of an Orchestration Semantics Engine (Ontology and Semantic Engine (OSE)), comprising an application descriptor that systematically aggregates required resources, edge deployment considerations, and SLAs. This descriptor is subsequently translated into a machine-readable format, relaying information to the Management Plane.

The Management Plane is pivotal to the CECCM, amalgamating essential management and orchestration functionalities. It utilizes data-driven and AI-based strategies to orchestrate both the life cycle management (LCM) of applications and their corresponding data, in consideration of the dynamic state of CECC infrastructure resources. AC^3 envisions a federated CECC infrastructure constituted by public and private cloud resources, edge resources, far-edge resources, and networking resources, which may not necessarily be owned by the CECCM proprietor but could be sourced externally through federated partnerships.

Within the User Plane, the **Northbound API** serves as the interface through which application developers and external systems engage with the CECCM. It facilitates create, read, update, delete (CRUD) operations on applications and resources, translating high-level application intents into actionable descriptors. The **Key Metrics Collection and Presentation** component aggregates performance metrics from diverse sources to offer insights into application performance and resource utilization, integral for AI-based decision-making processes.

The **OSE**, exploiting semantic technologies, interprets high-level intents into policies for stakehold-

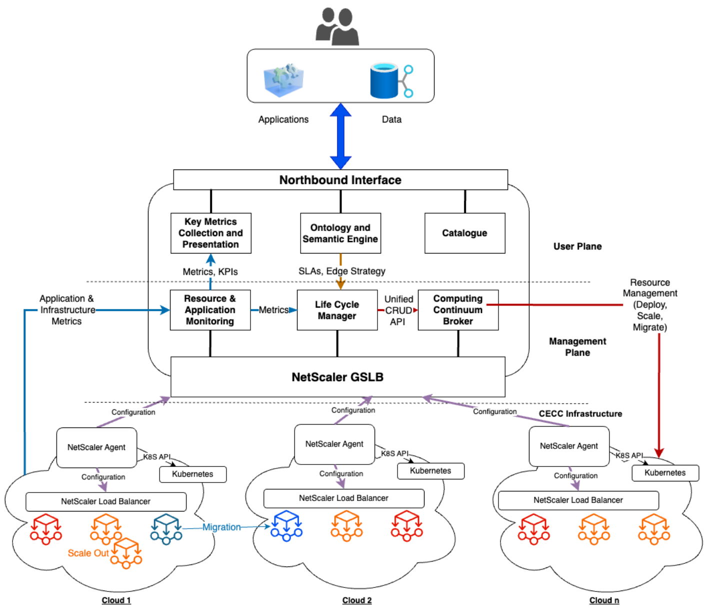

Fig. 1. High-level architectural framework proposed by AC^3.

ers—such as data source providers and application developers—ensuring consistent understanding and implementation across the CECCM ecosystem. A **Catalogue** acts as a repository for application blueprints and metadata, managing service information and resource dependencies to support efficient application deployment and management, offering extensibility through adaptive design.

Within the Management Plane, **Resource & Application Monitoring** plays a crucial role in vigilant oversight of application status and infrastructure performance, collating resource usage and behavior data to inform decision-making.

The **LCM** oversees the life cycle of microservice-based applications, managing their instantiation, scaling, migration, and termination, optimizing resource allocation while ensuring SLA compliance. It incorporates predictive functionality to anticipate application demands using AI-based algorithms, enabling proactive scaling across federated infrastructures to avert SLA violations proactively. The **Cognitive Computing Broker (Computing Continuum Broker (CCB))** facilitates resource federation and brokering, managing provider interactions for optimal federated resource utilization.

Additionally, the system integrates **NetScaler GSLB** for high availability and optimal application performance, distributing traffic across multiple locations by continuously monitoring operational health and navigating traffic based on resource availability and performance.

In the CECC Plane, **NetScaler Load Balancers**, operating within IONOS and Arsys infrastructure, manage localized load balancing, ensuring high availability and consistency with latency and throughput SLAs. **NetScaler Agents** decouple Load Balancers from other application lifecycle components, integrating with Kubernetes control plane for real-time application scaling and migration updates, enhancing system reliability and performance while simplifying LCM operations.

1) Advanced AI Techniques for Optimizing Application Life-Cycle Management:

2) Description of Architectural Components: The architecture is meticulously designed to facilitate interoperability and multi-cloud operations via diverse infrastructure providers and technologies. It abstracts the underlying infrastructure through standard Application Programming Interface (API)s and interfaces, supporting smooth application operations across heterogeneous environments—a crucial aspect for integrating computational and storage resources from assorted providers. The CCB component fosters federated resource integration, enabling optimization in application deployment and management.

Interoperability is further enhanced by the semantic engine, which translates high-level intents into executable tasks across platforms, ensuring consistency in application behavior. The architecture adeptly supports multi-cloud deployments, efficiently amalgamating resources from various cloud providers—achieved through resource federation incorporating IONOS, Arsys, and other cloud resources to optimize deployment and management.

Advanced AI techniques in LCM bolster application lifecycle management by employing machine learning models to predict application behavior and resource needs, dynamically adjusting resource allocations to meet SLA targets while optimizing performance. AI-driven functionalities encompass load pattern prediction and proactive adjustments, ensuring

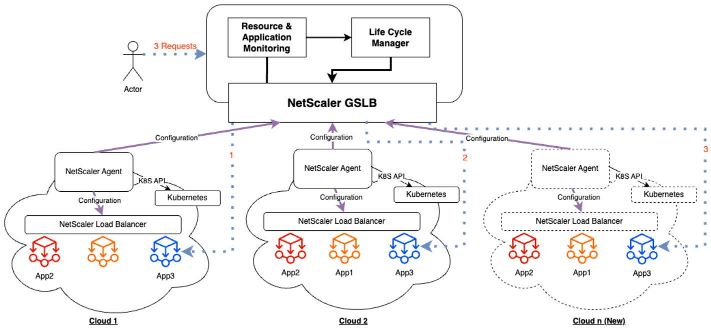

Fig. 2. NetScaler Multi-Cloud Federation within AC^3.

optimal resource utilization. AI-driven decisions are autonomously implemented, optimizing responses to evolving conditions, as continuous monitoring provides real-time data to the AI-based LCM, facilitating adjustments that ensure decision enforcement is effective—be it deploying instances, scaling resources, or migrating services. Resource federation further allows infrastructure adaptability to changing demands by integrating diverse provider resources.

Load balancing is pivotal in maintaining application reliability and performance, with NetScaler's GSLB and Application Delivery Controller (Application Delivery Controller (ADC)) capabilities distributing traffic across cloud providers to enhance redundancy and disaster recovery. Figure 2 demonstrates a scenario where requests to an application are dynamically distributed by NetScaler GSLB across federations. Beyond round-robin, intelligent traffic management directs client requests based on metrics—such as proximity and server load—thereby enriching user experience, enhancing performance, and accommodating system changes dynamically. With resilience anchored in health monitoring, NetScaler's capabilities ensure high availability and seamless integration into a federated multi-cloud environment.

C. Ensuring SLA Compliance Across Diverse and Unreliable Infrastructures

NetScaler GSLB and Load Balancers are critical for ensuring seamless scalability and SLA compliance under unforeseen conditions leading to performance degradation. The NetScaler Agent bridges cloud infrastructure loading balance and GSLB by intelligently distributing traffic, considering factors such as server capacity and response times, ensuring even load distribution to avert server overloads.

Scalability and migration actions are executed, utilizing NetScaler's savvy load balancing and traffic management, optimizing performance, and minimizing disruptions. Detailed execution involves:

- **Scaling Out**: Adding additional instances to shoulder increased loads by traffic distribution across servers or locations is facilitated through dynamic traffic allocation, strategically utilizing GSLB for efficient scale-out across data centers, preventing excessive scaling.
- **Migration**: Shifting services or data to different locations ensuring uninterrupted service. The GSLB handles dependencies management ensuring migration orchestration and DNS configurations are updated timely to route requests appropriately, with sticky sessions and high session reliability maintained amid transitions.
- **Scaling Up**: Enhancing existing instance capacities by allocating additional resources—such as CPU, memory—requiring backend configuration for adjustments and GSLB-based traffic distribution enhancements, ensuring seamless exploitation of upgraded resources in managing high loads efficiently.

D. How NetScaler Enables Stateful Microservices

Managing stateful microservices necessitates sophisticated techniques for maintaining session persistence and data consistency during migrations or fluctuations. NetScaler GSLB distributes traffic strategically, enhancing availability and address challenges encountered in stateful microservices deployment, particularly within containerized environments. The following delineates how NetScaler GSLB addresses challenges in stateful microservices migration:

- **Sticky Sessions:** Maintain consistent user sessions by directing traffic to the appropriate federation member, ensuring uninterrupted state throughout migrations.
- **Proximity-Based Routing:** Optimizes latency and throughput by routing traffic to the nearest or most responsive centers, mitigated by continuous health monitoring ensuring only viable services are leveraged.
- **Load Balancing Policies:** Incorporate energy capacity and consumption parameters of far-edge devices to optimize energy use and extend operational lifespan, complemented by prioritizing nodes with renewable energy for sustainability.
- **Data Synchronization:** Works in tandem with replication mechanisms to ensure consistent multi-location data synchronization, integrating

mapping file techniques to maintain seamless access across heterogenous back-end implementations.

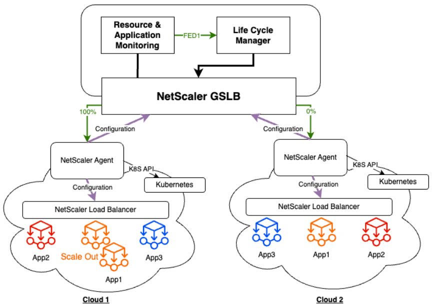

Fig. 3. NetScaler Multi-Cloud Federation with Stateful Microservices.

Figure 3 demonstrates leveraging LCM alongside NetScaler's programmable capabilities to dynamically adjust operations within a multi-cloud federation:

1) The resource monitor identifies an energy shift to brown sources in Federation 2 (FED2).
2) It is reported that Federation 1 (FED1) possesses sufficient capacity to accommodate existing and forecasted application loads.
3) The LCM, via the CCB, scales the application in FED1, transferring load progressively.
4) Upon reconfiguration, NetScaler GSLB redirects the load exclusively to FED1, maintaining the application in FED2 for disaster recovery purposes.

This comprehensive approach by AC^3 epitomizes a convergence of inter-disciplinary methodologies, exploiting advancements in AI, cloud technology, and network engineering to ensure optimized, reliable, and sustainable technological frameworks adaptable to modern computational imperatives.

V. Future Direction & Concluding Remarks

Along this paper we showed how the combination of cognitive computing with multi-cloud environments significantly enhances operational efficiency and resilience. The architecture proposed by AC^3, which leverages AI-driven resource management and robust frameworks, effectively addresses critical challenges such as data interoperability, latency, and security.

Next steps will involve implementing the interfaces between the final models and the different components to ensure seamless integration and provide the expected functionality and performance. We will also execute field trials to recreate a realistic federated CECC infrastructure, deploying a centralized cloud to run the CECCM components and three locations representing the edge, far edge, and end-devices. Each location is associated to a use case with its own requirements (including SLA), but consuming resources from a unified API that provides access to the respective local infrastructure.

Real-time data will therefore be gathered by the monitoring systems, facilitating continuous updates and refinement of models. Advanced model development will enhance explainability and accuracy, ensuring these improvements do not compromise performance. In particular, the LCM will be able to dynamically adapt application profiles based on ML predictions, taking into account traffic patterns, resource utilization, and user activity. For resource management, we will train the models on the data collected by both the resource discovery mechanism and the monitoring process, also adapting the algorithms to become energy-aware.

The use case validation will further substantiate the practical benefits and scalability of AC^3's solution, thereby advancing the fields of multi-cloud federation and cognitive computing. This research paves the way for innovative applications and more efficient different private and public cloud integration, pushing the state-of-the-art forward in this domain.

Acknowledgment

The AC^3 project has received funding from the European Union's Horizon Europe Research and Innovation programme under grant agreement No. 101093129.

References

[1] R. Bohn, C. Lee, and M. Michel. (2020, February) The NIST Cloud Federation Reference Architecture. [Online]. Available: https://www.nist.gov/publications/nist-cloud-federation-reference-architecture

[2] P. Kraemer, C. Niebel, and A. Reiberg, "Gaia-X and Business Models: Types and Examples," 2023, Gaia-X Hub Germany. White Paper 1/2023.

[3] D. Bernstein and Y. Demchenko, "The IEEE Intercloud Testbed–Creating the Global Cloud of Clouds," in *2013 IEEE 5th International Conference on Cloud Computing Technology and Science*, vol. 2. IEEE, 2013, pp. 45–50.

[4] S. Challita, F. Paraiso, and P. Merle, "Towards formal-based semantic interoperability in multi-clouds: the fclouds framework," in *2017 IEEE 10th International Conference on Cloud Computing (CLOUD)*. IEEE, 2017, pp. 710–713.

[5] G. Dhayanidhi, "Research on IoT threats & implementation of ai/ml to address emerging cybersecurity issues in IoT with cloud computing," Master's thesis, University of Alberta, Edmonton, Canada, March 2022, available at https://era.library.ualberta.ca.

[6] C. Ramalingam and P. Mohan, "Addressing semantics standards for cloud portability and interoperability in multi cloud environment," *Symmetry*, vol. 13, no. 2, p. 317, 2021.

[7] D. Saxena, R. Gupta, and A. K. Singh, "A survey and comparative study on multi-cloud architectures: emerging issues and challenges for cloud federation," arXiv:2108.12831.

[8] S. Cuñat, M. Julian, A. Belsa, C. I. Valero, M. Esteve, and C. E. Palau, "Secure, Trusted, Privacy-Protected Data Exchange in an Edge-Cloud continuum environment," in *IoT Edge Intelligence*. Springer, 2024, pp. 201–231.

[9] N. Singh, R. Buyya, and H. Kim, "Securing Cloud-Based Internet of Things: Challenges and Mitigations," 2024, arXiv:2402.00356.

[10] Gaia-X European Association for Data and Cloud. (2023, October) Gaia-X Architecture Document - 23.10 Release. [Online]. Available: https://docs.gaia-x.eu/technical-committee/architecture-document/23.10

[11] NetScaler, "Global Server Load Balancing," May 2023. [Online]. Available: https://docs.netscaler.com/en-us/citrix-adc/current-release/global-server-load-balancing.html