UNICO I+D Project
Grant No.

6G-EDGEDT-01-E5

System architecture v1

Abstract

This document presents the initial system design of the 6G-EDGEDT-01 architecture. It defines the main entities and building blocks of the system architecture that achieves orchestration of hyper-distributed Edge using emerging concepts of computing, connectivity, and AI in Industry 4.0. The document considers an open edge ecosystem tailored to host decentralized applications operating over the existing IoT-Edge-Cloud computing continuum. This deliverable starts with defining the existing challenges in the state-of-the-art, that are later on used as main pillars in the conceptual design of an architecture for the hyper-distributed edge.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia

Document properties

Document number	6G-EDGEDT-01-E5
Document title	System architecture v1
Document responsible	Antonio de la Oliva (UC3M)
Document editor	Milan Groshev, Antonio de la Oliva (UC3M)
Editorial team	Milan Groshev, Antonio de la Oliva (UC3M)
Target dissemination level	Public
Status of the document	Final
Version	1.0

Production properties

Reviewers	…

Document history

Revision	Date	Issued by	Description
…	…	…	…

Disclaimer

This document has been produced in the context of the 6G-EDGEDT Project. The research leading to these results has received funding from the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO 5G I+D programme.

Contents

List of Figures

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

List of Acronyms

AEN: Atomic Edge Node

ETSI: European Telecommunication Standardization Institution

MEC: Multi-access Edge Computing

IoT: Internet of Thigs

IBM: International Business Machines

AWS: Amazon Web Services

OS: Operating System

I4.0: Fourth Industrial Revolution

HW: Hardware

SW: Software

OPC-UA: Open Platform Communications - Unified Architecture

5G: 5^{th} Generation technology

AI: Artificial Intelligence

ML: Machine Learning

API: Application Programming Interface

VMs: Virtual Machines

vRAN: virtual Radio Access Network

DDL: Decentralized Data Layer

DOL: Distributed Orchestration Layer

CAP: Consistency, Availability, Partition tolerance

E2E: End-to-End

QoS: Quality of Service

KPIs: Key Performance Indicators

DT: Digital Twin

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Executive Summary

This document details the initial system design of the architecture for 6G-EDGEDT-01, which relies on computing, connectivity, and artificial intelligence to manage the hyper-distributed industrial environment. The proposed architecture presents a unique approach towards the integration of IoT, Edge and Cloud into a distributed but inter-connected compute continuum, including integration with traditional Telco edge and cloud systems.

The main results described within the deliverable are:

- the definition of the main existing gaps in the State-of-the-Art;
- the design of an Atomic Edge Node (AEN) so that different types of services and micro-services can run on top of it;
- the design of the Application Layer that encompass system and service-agnostic applications and use case specific service applications;
- the design of an entity that provides a set of essential functionalities and APIs to facilitate the development of these applications;
- the design of the Decentralized Data Layer that unifies the way data moves across the different components of the hyper-distributed system;
- the design of the Distributed Orchestration Layer that provide orchestration of functions/services across the distributed edge;

This subproject is oriented towards the digitalization of industrial process and will relay in the concept of hyper distributed edge. The architecture that is developed in this deliverable will be later on tailored to each of the specific EDGEDT projects under the same coordinated project.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

1. Introduction

Current edge cloud solutions are classified into two main categories: i) Traditional Telco edge solutions based on the ETSI MEC framework featuring hierarchical architectures and complex implementations; ii) Cloud providers' IoT edge solutions such as AWS Greengrass and Wavelength, Microsoft Azure IoT, IBM Watson IoT, or Google Cloud IoT Core, which are typically proprietary to each individual provider. Both approaches are based on a logically centralized architecture and rely on the concept of data threads, where data from IoT sensors are abstracted and fed into the edge for processing, analysis, decision (optimization), actioning and visualization, and then sent back to the sensors for actuation where applicable. In this centralized architecture, the data flows from the IoT devices to the central cloud where data intelligence is used next to enhance applications further up in the central cloud. Considering that i) data is generated by IoT devices, and ii) data is consumed at the edge where data intelligence is produced, in this project, the edge applications and services will be freely and dynamically distributed across the compute and storage continuum from low-capability user-owned devices to medium-capability edge nodes, up next to sophisticated high capability cloud data centers. This edge service paradigm of the "IoT-Edge-Cloud computing continuum" (Software, 2022) provides robustness, flexibility, better resource utilization and an opportunity to optimize the energy consumption for improving cost and environmental impact. The designed solution in this deliverable provides micro-services which use liquid approaches (Cicconetti, 2020) for enabling on-the-fly execution of user functions in a device-independent manner, while optimizing the overall usage of system resources. The designed solution implements the emerging paradigms of pervasive, distributed, and decentralized applications and services where data and related intelligence are produced and consumed collaboratively across various end-devices, edge nodes and the cloud. This represents a technology shift from current big-data and mostly non-real-time based analytics solutions towards a more micro-service and processing close to the origin of the big-data based approach. This, in turn, requires an integrated and continuous development-and-operation lifecycle, guaranteeing the locality and privacy of the data and intelligence. The architectural choices give a high ability and mitigate many potential security risks, both from local attacks and failures, as well as large-scale attacks or events.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

2. SoA analysis and existing challenges

The portability of edge applications across heterogeneous edge environments is fundamental to exploit the potential of the micro-service principles applied at the edge. Edge computing thus requires a completely novel, lightweight, and portable meta operating system (meta-OS) which abstracts the particularity of the various nodes, deployments, and IoT-platforms. The current edge solutions have several drawbacks hindering them from supporting such a fully distributed, multi-domain and multi-stakeholder continuum:

- Telco solutions are designed in a centralized way, requiring complex implementations to realize simple use cases (e.g., ETSI MEC).
- Commercial, generic solutions suffer from vendor or ecosystem lock-in (e.g., Azure IoT, AWS IoT Greengrass).
- These solutions are not fully automated and require human intervention, lacking continuous, multi-stakeholder orchestration from the end-devices to the edge and cloud.
- The way of collecting, processing, and analysing data is not fully customized for the application developers, the data processing is still largely centralized and server-based, and the intelligence is driven by big data and offline non-real time analytics. Partition tolerance1 remains a key challenge in dynamic hyper-distributed2 applications: the lack of supporting interfaces and mechanisms to effectively cope with ephemeral communications within the computing continuum results in intermittent application availability or consistency preventing a fully distributed data and service paradigm.
- There is a lack of automated mesh networking/communication features and orchestration to support workload mobility in a volatile edge-to-cloud continuum fabric.

In this project we aim to to realize an innovative, highly efficient, trustable network and computing ecosystem consisting of interconnected, computing communities (e.g., ad-hoc clouds) and supporting distributed services. By creating efficient mechanisms for leveraging locally generated data, along with the resources available across the IoT-Edge-Cloud continuum, the development and seamless integration of new edge applications and services by third parties will be enabled. Key challenges and advancements that need to address when designing the initial system architecture are:

- *Challenge 1: Liquefying services while addressing the CAP theorem*

The pervasive distribution, portability and reliability of edge applications need to increase for enabling the development of innovative applications. Furthermore, to ensure spatial applications and service availability, multiple instances of such modular functions (or services) must be created at different locations, raising consistency issues. Such issues are greatly exacerbated in dynamic environments with node churning. Last but not least, in a fully distributed system, partition tolerance is the key to ensure the system and applications continue to operate despite arbitrary loss or failure of any part of the large-scale system. However, according to the CAP theorem (Muñoz-Escoí, 2019) only two properties among Consistency, Availability, and Partition tolerance can be ensured at the

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

same time in a system. Identifying the right trade-off among the CAP dimensions and developing algorithms and mechanisms in both DDL and DOL to ensure their performances in a distributed system across heterogeneous edge systems in support of different use case scenarios, conditions and requirements is thus one of the main challenges in the project.

This project tackles the need for distributed applications and services, and the issue captured by the CAP theorem, by leveraging the novel ad-hoc cloud-based architecture and the micro-services approach enabled by the meta-OS developed in this project. By drawing on such innovations, in this project will create novel, data-driven, algorithmic solutions for: i) the definition of the right level of granularity to liquefy applications and services into modular component (e.g., micro-tasks) that can optimally match the resource availability in the IoT-Edge-Cloud continuum; ii) the identification of flexible consistency levels to establish the optimal trade-off among the CAP elements at the data, control, and management planes for application and service provisioning; iii) the development of mechanisms that ensure that data, control, and management planes cohesively work together when different guarantees are provided by each plane; iv) the coordination among ad-hoc cloudlets in the execution of distributed applications and services so as to achieve the targeted CAP trade-off.

- *Challenge 2: Flexible data flow infrastructure*

A multidomain platform such as the one outlined above necessitates a data infrastructure that can facilitate data flows, from the basic sources of data to final data consumers, while supporting dynamic source instantiation and management, consumer lifecycle management, and the possibility of performing different pre-processing operations as requested by the users of the data infrastructure. With increasing volume of data being connected due to multiple devices and systems that a data-driven manufacturer enables, a data storage challenge arises. Powerful cloud storage and computing capability can support big data analytics and optimal decision-making for manufacturing applications with multi-dimensional big data. However, data communication (among devices, the edge and the cloud), synchronization and computation may require considerable time, defying the requirements of time-sensitive applications on the shop floor. A solution would be to store all or large part of the data on premise, which can be costly. This problem has been tackled in literature, e.g., in (Pölöskei, 2021), (Zhang H. , 2020).

This project will create a data collection and management infrastructure that can flexibly handle the data flow regardless of the location of the storage. Batch processing frameworks will be used for data digestion and to parse unstructured data coming from plant floor sensors.

- *Challenge 3: Novel data-based architecture integrating computing, connectivity, and AI*

New-generation network systems combine network, computing and storage resources, and such components are increasingly entangled. Further, AI has emerged as a fundamental component of fully automated network and service orchestration (ETSI, 2022) as well as of user applications (Xiao, 2020). Distributed AI matches the characteristics and potentialities of Edge environments for its ability to leverage local resources and data. A new system architecture is needed to fully exploit the

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

inter-relation among computing, connectivity, and AI. Although initial efforts have been spent towards the realization of data collection and exploitation for ML-driven decision making in 5G systems (see, e.g., the results in (5Growth, 2022)), no working solution exists that can successfully address the above challenges.

This project will develop algorithms, as well as technologies and open, flexible interfaces to: i) optimally federate ad-hoc clouds at the edge which will interact with each other, as well as with the cloud, contributing to the creation of pools of local resources, as well as local functions and services; ii) facilitate the discovery, access and composition of heterogeneous IoT platforms and edge applications through unified APIs; iii) create a monitoring platform that collects system-level and service-level data necessary for IoT applications, with minimal communication overhead thus finding the best trade-off between fine grain monitoring and footprint of monitoring; iv) develop an efficient Decentralized Data Layer capable to stretch from tiny devices up to powerful cloud servers and handling in an eco-friendly manner data and information flows at both application and system level.

- *Challenge 4: Relying on fast automated decision making, IoT-based systems, such as those supporting smart grids, e-health, and smart manufacturing, require processing large amounts of data in real (or near-real) time.*

They must rely on edge computing or on process execution on "far-edge devices" for which the execution of lightweight functions is a perfect match. However, existing IoT approaches (e.g., AWS Greengrass and Azure IoT) make use of short-lived Virtual Machines (VMs) or containers, which are too heavy-weight, and often fail to meet target low latencies (Wang L. , 2018). Also, existing platforms may not be able to allocate sufficient, exclusively reserved, memory and CPUs to functions (AWS, 2022), or provide function isolation, which may lead to inter-function interference, hence performance degradation. The lack of performance guarantees is accentuated by the dependency on heterogeneous hardware at the distributed edge (Zhang Z. , 2015). Even looking at ETSI MEC or proposed conceptual edge network architectures (Xiao, 2020), (Hagenauer, 2019) one cannot yet find a practical, flexible, solution that efficiently fulfills the requirements of existing and emerging applications.

This project will address these challenges by delivering an open, modular software platform with flexible APIs and versatile libraries. Respectively, an Edge architecture and a light-weight meta operating system will be designed and developed which i) integrate, and exploit in a synergic manner, the current available technologies; ii) efficiently coordinate a number of diverse computing communities (namely, ad-hoc clouds), including smart devices, unleashing the full potential of distributed resources; iii) make the IoT-Edge-Cloud continuum concept reality by creating a flexible, fully inter-operable platform overarching the different network segments, which optimally leverages the resources they offer while fulfilling the stringent application requirements.

- *Challenge 5: Distributed and energy-efficient ML training and inference*

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

ML has become ubiquitous (Xiao, 2020), (Peltonen, 2020); it is essential to resource and network services orchestration (e.g., vRAN management) as well as to IoT-based applications (e.g., automation in Industry 4.0), and, as mentioned, the distributed ML paradigm (like Federated Learning (Wang S. , 2019) or D-SGD (Gemulla, 2011)) well matches the capabilities of the distributed Edge (Chen, 2020), (Tu, 2020). Distributed learning also implies keeping data local; hence it contributes to providing data privacy. It is however critical to properly map the diverse ML tasks onto the system components with diverse capabilities and data availability. Additionally, awareness of the need for energy-efficient ML is quickly raising, thus making ML sustainable is a vital emerging challenge (Schwartz, 2020), (Deng, 2020). Still, fundamental issues concerning the design, use, and performance of ML in highly distributed, volatile systems remain open and urgently need a solution.

In this project will i) design new, energy efficiency-aware distributed learning approaches that extend Federated Learning or D-SGD allowing more flexible, reconfigurable distributed paradigms; ii) develop algorithms that, for each ML-driven application, determine the quantity of data, and select high quality data, to be used for learning so that an optimal trade-off between learning accuracy, latency, and energy consumption is achieved; iii) create algorithms to optimal map ML tasks onto the IoT-Edge-Cloud continuum; iv) devise mechanisms to cope with intermittent learning nodes availability caused by dynamic connectivity conditions and topology.

3. General system architecture

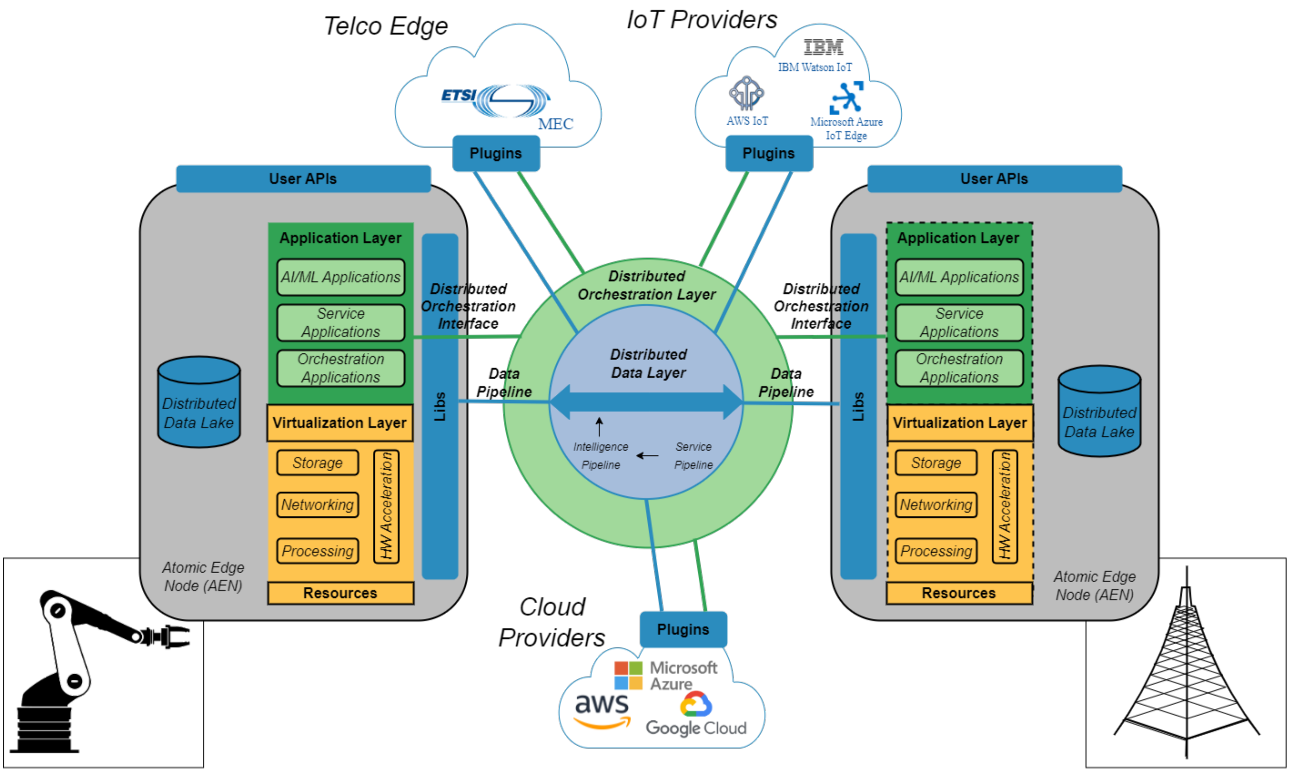

FIGURE 1: GENERAL SYSTEM ARCHITECURE

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

In this section we describe a general system architecture to exploit the usage of distributed data and distributed orchestration in industrial environments.

The proposed system architecture presents a unique approach towards the integration of IoT, Edge and Cloud into a distributed but inter-connected compute continuum, including integration with traditional Telco edge and cloud systems, as shown in Figure 1. The proposed architecture is conceived as a fully decentralized and distributed system where all nodes are equal, and they interact and cooperate in a flat, modular, and scalable architecture. It considers the liquification of all the functionalities running on the compute continuum, as a way of dealing with the dynamicity and volatility of the resources composing the underlying infrastructure. As such, the main components may run, e.g., from Edge terminals or robots (far-edge) up to the Edge located at an IoT Gateway deployed in specialized HW, and progressively extended towards the Telco Edge (near-edge) and even the Cloud.

The basic element of the architecture is the Atomic Edge Node (AEN) representing an edge node, as shown in Figure 1. The AEN abstraction will be based on extensions (an integration) to open-source projects led by partners of the consortium, such as Open Nebula. This abstraction of computing, networking and storage resources is virtualized so that applications to implement different types of services or micro-services can run on top of it, conforming the Application layer of the AEN. Examples of such applications encompass both generalized system and service-agnostic applications (e.g., orchestration Apps, AI/ML Apps, and data collection or data pre-processing Apps) and use case specific service applications (e.g., IoT Apps running directly on edge nodes processing computation tasks, with or without networking capabilities, or storage). The applications are built on top of the Core technologies, namely the Libs, which provides a set of essential functionalities and APIs to facilitate the development of these applications. The atomic process can be seen as a virtualized micro-service and all communication for the data storage, distribution, exchange and sharing as well as for the distributed orchestration (i.e., federation) between the AENs and the other edge and cloud systems is performed through Decentralized Data Layer (DDL) and Distributed Orchestration Layer (DOL).

In the following we present the main components of the proposed architecture:

3.1. Libs (Elibs)

The Libs sets the basis of the distributed and (micro-)services-based architecture. The Libs include a set of minimum functionalities that a developer can use - through well-defined APIs - to build a wide functional spectrum of applications and services. The key characteristic of the library is the easiness developers experience when operating the network, compute, and storage resources at a reasonable level of abstraction, hiding the technical details of the underlying technologies composing the infrastructure. Hence, developers can focus on implementing the application logic of their services without incurring in development, configuration, and deployment overhead. The functionality provided by the Libs will focus on providing the means to easily distribute an application or service and their associated data, allocating an ad-hoc cloudlet of resources, and managing: i) exceptions

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

and errors, ii) data consistency across the different AENs, iii) partition reliability including volatility of resources, iv) access to lightweight distributed AI resources, v) security and privacy, and vi) integration with other IoT platforms.

3.2. Decentralized Data Layer (DDL)

The DDL unifies the way data moves across the different components of the distributed architecture (Figure 1). All exchange of information between the AENs will use the DDL. It is more than just a

communication protocol for unifying data and computation, since it also provides a unified paradigm for data-in-motion, data-in-use, data-at-rest, and computations in the IoT-Edge-Cloud continuum. The DDL does this by carefully interconnecting heterogeneous sets of resources and applications that operate in diverse environments and networks. It will blend traditional publish/subscribe technologies with geo-distributed storage, queries, and computations in a consolidated and location-transparent API.

The DDL includes: i) the design of breakthrough protocols and data storage/caching mechanisms to build a distributed data pipeline for effective and robust data sharing across the decentralized edges for different network topologies, different types of data with diverse nature, communication patterns, lifetime and location, data producers/consumers, including constrained devices and networks; ii) novel resource management and partitioning mechanisms to address the required "CAP" guarantees, with exceptions and security handling while supporting heterogeneous and distributed use case scenarios such as IoT, vehicles, robots, smart grids, and datacenters; iii) data consistency mechanisms for managing fully decentralized and heterogeneous edge systems, aiming for deployments a factor 100x more devices than nowadays; iv) data distribution mechanisms to be used by any AI or intelligence consuming/producing entity, which, albeit transported by the DDL, may have its own specific requirements and storage logical domains (information security and authorization levels).

The DDL is the main pillar towards the special objectives on CAP challenges. The design of the architecture does not make assumption regarding these three characteristics of the system but will provide mechanisms making use of the Elibs and the DDL integrated communication and storage, as well as the DOL distributed orchestration and resource management algorithms, to cope with that, enabling application developers to choose the best suitable characteristics and the best trade-off for their use cases.

3.3. Distributed Orchestration Layer (DOL)

The other main pillar of the project is the DOL, which includes several breakthrough technologies to provide orchestration of functions/services across the distributed edge.

First, novel orchestration algorithms for the federation of distributed (micro-)services and ad hoc resources will be developed, leveraging data-driven approaches, digital ledger technology, and integration of AI/ML techniques. These algorithms will be used to provide a fully automated distributed orchestration to enable automated provisioning and dynamic adaptation of virtualized

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

services, and coordinate the computing, networking, and storage tasks running within the architecture, while coping with the ecosystem diversity and volatility.

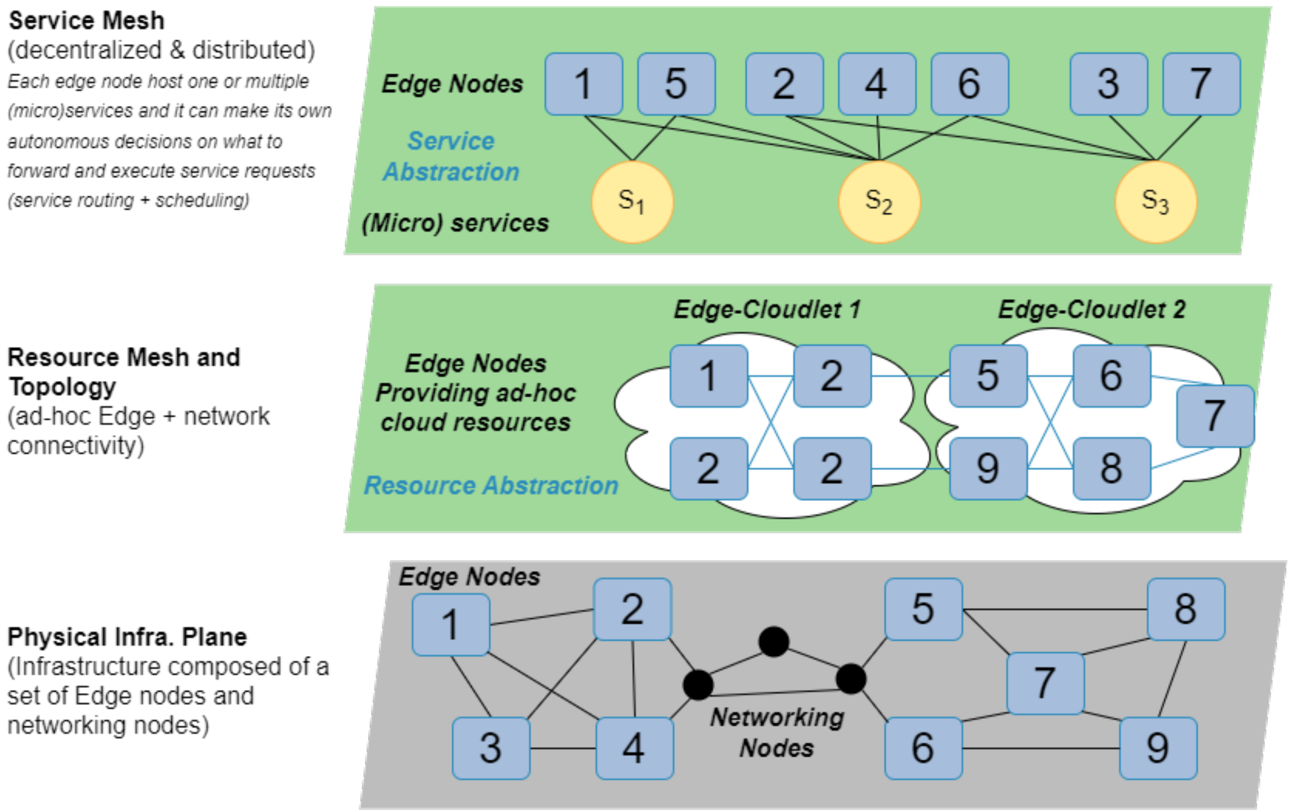

FIGURE 2: DISTRIBUTED ORCHESTRATION CONCEPT

The distributed orchestration concept is illustrated in Figure 2. Based on the abstraction of the Resource Mesh and topology provided by the DDL and Service Mesh provided by the Application Layer, the design of the DOL will include novel data-driven algorithms for i) decomposing services into micro-tasks (micro-services) consumed by the edge nodes and making autonomous decisions on where to forward and execute these micro-services, i.e., routing and scheduling of micro-services requests; and ii) for dynamic aggregation of edge resources into service specific ad-hoc edge cloudlets while allowing automatic data driven scaling. This approach, combined with service liquification, will lead to an effective, full exploitation of the system resources which would otherwise be underutilized or totally unused.

Second, these algorithms should be able to run on constrained devices, improving learning latency and accuracy and with a reduced energy footprint. These new approaches will enable distributed learning, inference, and intelligence placed along the ad-hoc cloudlets of resources composing the distributed edge working over trade-offs between learning performance and energy efficiency. This new distributed intelligence will work over distributed decision making by optimally combining decisions made by different entities and with different confidence levels. Volatility of computing resources will be considered while developing the consensus algorithms and the data distribution and storage features of the system, to be able to cope with the dynamic nature of the envisioned IoT-Edge-Cloud compute continuum. Third, new monitoring mechanisms and extensions suitable for

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

such a distributed environment will be developed, along with End-to-End Quality of Service (E2E QoS) management. Monitoring E2E QoS in context of the target KPIs for different services deployed across different (ad-hoc) cloudlets is essential.

At last, novel mechanisms will be designed for Digital Twinning of the running system and services, based on the monitoring of data collected across different edge nodes over the time, to validate the orchestration and configuration decisions before their actual enforcement. The DOL can leverage the DDL, therefore making use of its security and privacy mechanisms to access information.

3.4. Plugins ensuring compatibility and inter-operability with existing fog, edge, and cloud SDOs

One of the key problems of currently available approaches is that they tend to tie users to specific proprietary implementations. This is exacerbated on current solutions for the IoT-Edge (e.g., Azure IoT, Amazon GG) which completely lock users into an ecosystem of selected functionalities and partners. This project aims at creating an open edge ecosystem which needs to be able to interconnect and interoperate with the de-facto industry standards, such as AWS Greengrass or Microsoft Azure IoT, as well as traditional private or public Clouds, fog, edge systems. Following past experiences on the integration of such systems (e.g., Kubernetes), this project will use the concept of plugins to integrate these existing systems into the designed architecture. The architecture specific plugins, which need to be developed for each type of external system, will conceptually be similar to a user space driver in the GNU Linux world. Each plugin will implement the translation and mapping functionalities to allow the external system to integrate with the DDL for the exchange of data with the native applications. Moreover, it will provide a system specific implementation of the Libs APIs, to enable the external system to be seamlessly used, as part of the AEN concept. In this way, the compatibility and inter-operability of the technology is offered under any external fog, edge, and cloud systems.

3.5. DevOps evolution towards EdgeOps

The proposed architecture design is based on the concept of DevOps (Development and Operation) into the Edge, a trend that is known as EdgeOps. DevOps builds on top of the micro-services concept, where each component of a DevOps toolchain is independent and can be replaced by a different one, as long as the interfaces remain common. This fact increases modularity, reliability, and flexibility of SW solutions, departing from the traditional monolithic approaches. EdgeOps [28] applies DevOps practices to the Edges, taking into consideration new requirements and challenges for the edge applications with respect to the traditional Cloud. Each vertical using an Edge deployment has its own associated characteristics, including specific standards and certifications, hardware, types of sensors, communication protocols or integration of diverse technologies. Consequently, the portability of edge applications across heterogeneous edge environments, which is fundamental to exploit the full potential of the micro-services principles applied at the edge, requires a meta-OS which supports EdgeOps for allowing a continuous development and integration of services within the ecosystem.

4. Summary and Conclusions

This deliverable specifies the first version of the overall system architecture proposed to perform orchestration in hyper-distributed environment. First, the deliverable specifies the exiting gaps in the State-of-the-Art, defining a set of challenges that this project will address. Next, the deliverable specifies the overall architecture considering the IoT-Edge-Cloud continuum to gather, process and orchestrate information. The main component of the specified system architecture is the Atomic Edge Node (AEN) that provides abstraction of computing, networking, and storage resources in virtualized environment so that applications that implement different types of services or micro-services can run on top of it. Each AEN has its own Virtualization Layer Application Layer and Libs to ease the development and execution system and use case specific services and applications.

The proposed architecture provides a unified paradigm for data-in-motion, data-in-use, data-at-rest, and computations in the IoT-Edge-Cloud continuum through the Decentralized Data Layer (DLL). Additionally, the architecture supports orchestration of functions/services across the IoT-Edge-Cloud continuum with the help of the Distributed Orchestration Layer (DOL).

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

5. References

5Growth. (2022, December). *5G-enabled Growth in Vertical Industries*. Retrieved from H2020 5Growth project: https://5growth.eu

AWS. (2022, December). *Choosing the Optimal Memory Size*. Retrieved from docs.aws.amazon.com: https://docs.aws.amazon.com/whitepapers/latest/serverless-architectures-lambda/choosing-the-optimal-memory-size.html

Chen, M. (2020). A Joint Learning and Communications Framework for Federated Learning over Wireless Networks. *IEEE Transactions on Wireless Communications*.

Cicconetti, C. (2020). Toward Distributed Computing Environments with Serverless Solutions in Edge Systems. *IEEE Communication Magazine*.

Deng, S. (2020). Edge Intelligence: The confluence of edge computing and artificial intelligence. *IEEE Internet of Things Journal*.

ETSI. (2022, December). *technologies*. Retrieved from ETSI: https://www.etsi.org/technologies/zero-touch-network-service-management

Gemulla, R. (2011). Large-scale Matrix Factorization with Distributed Stochastic Gradient Descent. *ACM SIGKDD*, ACM.

Hagenauer, F. (2019). Efficient Data Handling in Vehicular Micro Clouds. *Ad Hoc Networks*.

Kavakl, E. (2018). An Architecture for Disruption Management in Smart Manufacturing. *IEEE International Conference on Smart Computing (SMARTCOMP 2018)* (pp. 279-281). IEEE.

Muñoz-Escoí, F. D. (2019). CAP Theorem: Revision of Its Related Consistency Models. *The Computer Journal*.

Peltonen, E. (2020). 6G White Paper on Edge Intelligence. *arXiv*.

Pölöskei, I. (2021). Data engineering case-study in digitalized manufacturing. *IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI)* (pp. 491-494). IEEE.

Schwartz, R. (2020). Green AI. *Communications of the ACM*.

Software, O. A. (2022, December). *Open Automation Software*. Retrieved from Open Automation Software: https://openautomationsoftware.com/blog/iiot-edge-computing-vs-cloud-computing/

Tountopoulos, V. (2018). Towards a Cloud-Based Controller for Data-Driven Service Orchestration in Smart Manufacturing. *Sixth International Conference on Enterprise Systems (ES)*, (pp. 96-99).

Tu, Y. (2020). Network-Aware Optimization of Distributed Learning for Fog Computing. *IEEE INFOCOM*. IEEE.

Wang, L. (2018). Peeking Behind the Curtains of Serverless Platforms. *USENIX ATC*. USENIX ATC.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Wang, S. (2019). Adaptive Federated Learning in Resource Constrained Edge Computing Systems. *IEEE JSAC*.

Wei, W. (2019). Data-driven Manufacturing Service Optimization Model in Smart Factory. *IEEE 15th International Conference on Automation Science and Engineering (CASE)* (pp. 362-367). IEEE.

Xiao, Y. (2020). Toward Self-Learning Edge Intelligence in 6G. *IEEE Communications Magazine*.

Zhang, H. (2020). Design of Data Acquisition Platform for Industrial Internet of Things. *IEEE 3rd International Conference on Information Systems and Computer Aided Education (ICISCAE)* (pp. 613-618). IEEE.

Zhang, Z. (2015). Platform Selection for MapReduce Processing in the Cloud. *ICCAC*. ICCAC.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D