Ref. Ares(2025)777747 - 31/01/2025

HEDGE-IoT

Holistic approach towards Empowerment of the Digitalization of the Energy Ecosystem through adoption of IoT solutions

D3.3
HEDGE-IoT Technological Enablers (First Release)

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

DOCUMENT CONTROL SHEET

PROJECT INFORMATION

Project Number	101136216		
Project Acronym	HEDGE-IoT		
Project Full title	Holistic Approach towards Empowerment of the Digitalization of the Energy Ecosystem through adoption of IoT solutions		
Project Start Date	01 October 2024		
Project Duration	48 months		
Funding Instrument	Horizon Europe Framework Programme	Type of action	HORIZON-IA HORIZON Innovation Actions
Call	HORIZON-CL5-2023-D3-01-15		
Topic	Supporting the green and digital transformation of the energy ecosystem and enhancing its resilience through the development and piloting of AI-IoT Edge-cloud and platform solutions		
Coordinator	European Dynamics Luxembourg SA		

DELIVERABLE INFORMATION

Deliverable No.	D3.3
Deliverable Title	HEDGE-IoT Technological Enablers (First Release)
Work-Package No.	WP3
Work-Package Title	Technological Enablers Specification, Design and Development
Lead Beneficiary	INESC TEC
Main Author	Vasco Campos (INESC) Paulo Monteiro (INESC) Fábio Coelho (INESC)
Other Authors	Josipa Stegić (KONC) Nika Mlinarič (JSI) Aleksandra Rashkovska (JSI) Gregor Kosec (JSI) Tomislav Antić (UNIZG) Mattia Alfieri (APIO) Mateo Cardenas (NESTER) Inoussa Laouali (NESTER) Vagelis Karakolis (ICCS) Giannis Papias (ICCS)

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

	Sotirios Filippou (HENEX) Faheem Muhammad (TAU) Mehdi Attar (TAU) Kari Maki (VTT) Anna Kulmala (ABB) Xander Wilcke (VU) Tudor Cioara (TUC) Ionut Anghel (TUC)						
Due date	31/01/2025						
Deliverable Type	X	Document, Report (R)		Data management plan (DMP)		Websites, press & media action (DEC)	X	Other
Dissemination Level	X	Public (PU)		Sensitive (SEN)		Classified		
	PU: Public, fully open SEN: Sensitive, limited under the conditions of the Grant Agreement Classified R-UE/EU-R – EU RESTRICTED under the Commission Decision No2015/444 Classified C-UE/EU-C – EU CONFIDENTIAL under the Commission Decision No2015/444 Classified S-UE/EU-S – EU SECRET under the Commission Decision No2015/444						

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Document Revision History

Version	Date	Description of change	List of contributor(s)
V0.1	11/11/2024	Table of Contents (ToC) draft	INESC
V0.2	22/11/2024	ToC revisions	INESC, ICCS
V0.3	05/12/2024	Final ToC, structure and guidelines	INESC
V0.4	19/12/2024	First contributions	All authors
V0.5	03/01/2024	Consolidation	INESC
V0.6	14/01/2024	Final partner's contributions	All authors
V0.7	20/01/2024	Consolidation and final touches	INESC
V0.8	24/01/2025	Quality review	NESTER, TUC
V0.9	28/01/2025	Draft for submission	INESC
V1.0	31/01/2025	Ready for submission	ED

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

PARTNERS

Participant number	Participant organisation name	Short name	Country
1	EUROPEAN DYNAMICS LUXEMBOURG SA	ED	LU
2	RHEINISCH-WESTFAELISCHE TECHNISCHE HOCHSCHULE AACHEN	RWTH	DE
~~3~~	~~ENGINEERING - INGEGNERIA INFORMATICA SPA~~	~~ENG~~	~~IT~~
4	EREVNITIKO PANEPISTIMIAKO INSTITOUTO SYSTIMATON EPIKOINONION KAI YPOLOGISTON	ICCS	EL
5	INESC TEC - INSTITUTO DE ENGENHARIADE SISTEMAS E COMPUTADORES, TECNOLOGIA E CIENCIA	INESC	PT
6	NEDERLANDSE ORGANISATIE VOOR TOEGEPAST NATUURWETENSCHAPPELIJK ONDERZOEK TNO	TNO	NL
7	TAMPEREEN KORKEAKOULUSAATIO SR	TAU	FI
8	TEKNOLOGIAN TUTKIMUSKESKUS VTT OY	VTT	FI
9	TRIALOG	TRIALOG	FR
~~10~~	~~CYBERETHICS LAB SRLS~~	~~CEL~~	~~IT~~
11	CENTRO DE INVESTIGACAO EM ENERGIA REN - STATE GRID SA	NESTER	PT
12	INTERNATIONAL DATA SPACES EV	IDSA	DE
13	ELIA TRANSMISSION BELGIUM	ETB	BE
14	HRVATSKI OPERATOR PRIJENOSNOG SUSTAVA D.D.	HOPS	HR
15	UNIVERSITATEA TEHNICA CLUJ-NAPOCA	TUC	RO
16	CLUSTER VIOOIKONOMIAS KAI PERIVALLONTOS DYTIKIS MAKEDONIAS	CLUBE	EL
17	F6S NETWORK IRELAND LIMITED	F6S	IE

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

18	SOCIAL OPEN AND INCLUSIVE INNOVATION ASTIKI MI KERDOSKOPIKI ETAIREIA	INCL	EL
19	ABB OY	ABB	FI
20	ENERVA OY	ENERV	FI
21	JARVI-SUOMEN ENERGIA OY	JSE	FI
22	DIMOSIA EPICHEIRISI ILEKTRISMOU ANONYMI ETAIREIA	PPC	EL
23	DIACHEIRISTIS ELLINIKOU DIKTYOU DIANOMIS ELEKTRIKIS ENERGEIAS AE	HEDNO	EL
24	INDEPENDENT POWER TRANSMISSION OPERATOR SA	IPTO	EL
25	ELLINIKO HRIMATISTIRIO ENERGEIAS	HENEX	EL
26	HARDWARE AND SOFTWARE ENGINEERING EPE	HSE	EL
27	QUE TECHNOLOGIES KEFALAIOUCHIKI ETAIREIA	QUE	EL
28	ARETI S.P.A.	ARETI	IT
29	APIO S.R.L.	APIO	IT
30	ACEA ENERGIA SPA	AE	IT
~~31~~	~~VOLKERWESSELS ICITY B.V.~~	~~VWICI~~	~~NL~~
32	ARNHEMS BUITEN BV	AB	NL
33	STICHTING VU	VU	NL
34	COOPERATIVE ELECTRICA DO VALE DESTE CRL	CEVE	PT
35	REN - REDE ELECTRICA NACIONAL SA	REN	PT
36	MC SHARED SERVICES SA	SONAE	PT
37	ELES DOO SISTEMSKI OPERATER PRENOSNEGA ELEKTROENERGETSKEGA OMREZJA	ELES	SI
38	ELEKTRO GORENJSKA PODJETJE ZA DISTRIBUCIJO ELEKTRICNE ENERGIJE DD	EG	SI
39	OPERATO DOO	OPR	SI

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

40	SVEUCILISTE U ZAGREBU FAKULTET ELEKTROTEHNIKE I RACUNARSTVA	UNIZG	HR
41	INSTITUT JOZEF STEFAN	JSI	SI
42	KONCAR - DIGITAL DOO ZA DIGITALNE USLUGE	KONC	HR
43	DS TECH SRL	DST	IT
44	CYBERSOCIAL LAB SRL IMPRESA SOCIALE	CSL	IT

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

DISCLAIMER

This deliverable is subject to final acceptance by the European Commission. The content and results of the publication herein are the sole responsibility of the publishers, it reflects only the contributors' view, and it does not necessarily represent the views expressed by the European Commission or its services, neither is the European Commission responsible for any use that may be made of the information it contains.

While the information contained in the documents is believed to be accurate, the contributor(s) or any other participant in the HEDGE-IoT Consortium make no warranty of any kind regarding this material including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Neither the HEDGE-IoT Consortium nor any of its members, their officers, employees, or agents, shall be responsible or liable in negligence or otherwise howsoever in respect of any inaccuracy or omission herein. Without derogating from the generality of the foregoing, neither the HEDGE-IoT Consortium nor any of its members, their officers, employees, or agents shall be liable for any direct or indirect or consequential loss or damage caused by or arising from any information advice or inaccuracy or omission herein.

COPYRIGHT NOTICE

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

EXECUTIVE SUMMARY

HEDGE-IoT (Holistic Energy Decentralized Grid for Enhanced IoT) is a project, funded by the European Union's Horizon Europe research and innovation program, that proposes a novel digital framework to explore the intrinsic data value of the energy grid, by deploying IoT assets at different levels of the energy system, from behind-the-meter, up to the Distribution System Operator (DSO) and Transmission System Operator (TSO) level. By deploying these assets, this novel framework positions itself to add intelligence to the edge and cloud layers through advanced AI/ML tools and ensures a cloud/edge continuum by introducing federated and decentralized applications governed by an advanced swarm-based computational orchestration framework.

The project will be demonstrated in 6 different countries through pilot sites. To address the multiple use cases in these pilots, the partners from the HEDGE-IoT consortium will develop services that will be integrated with the project's digital framework in different ways. Most services plan to integrate the project's interoperability framework and ML (Machine Learning) services that focus on the edge layer can easily integrate the computational orchestration framework.

The focus of this deliverable is on detailing the first specification of all services that will be deployed in the context of the project and on providing a first look into how these services will integrate the project's novel digital framework. Such connections are shown by detailing how the service will be implemented in the energy system and how its relationship with interoperability is envisioned.

As Table 1 shows, almost all services have a plan to integrate with the project's interoperability framework. The most common integration method envisioned will be via Data Spaces, using the Eclipse Data Connector (EDC), which is the most entry-level option available to perform interoperable data exchanges. A few other services opted by a more advanced form of interoperability, namely semantic interoperability, by applying ontologies to their data and using a semantic data broker.

The Technology Readiness Level (TRL) of the services specified is fairly high for the first stage of specification, as it is described in Table 1. Most services sit at a TRL of 4 or higher, which shows a high level of reusability. The expected TRL by the end of the project is also quite high, most services planning to reach TRL 6 or 7, showing their ambitious nature.

A total of 14 service specifications are detailed:

- 6 energy sector services for residential users, 3 of which have non-energy functionalities.
- 7 services in which the **main** stakeholders are grid operators (TSOs and DSOs).
- 1 computational orchestration framework to ensure edge/cloud continuum.

These services provide a rich landscape of applied intelligence to the edge and cloud layers. From the 14 described, 2 are federated learning approaches, 6 are edge-to-cloud AI/ML services with some degree of computation on the edge and the remaining 6 are cloud-based AI and non-AI services that support the pilots.

In conclusion, this document exposes the efforts of Work Package (WP) 3 and its tasks to provide high quality specifications of the services that help the HEDGE-IoT project achieve its goals. It highlights their innovations, their contributions to decentralize the grid and add intelligence to it and their plans to participate in an interoperable ecosystem with meaningful data exchanges. These contributions will be leveraged by the project and will help to enable the efficient integration of services and IoT assets into an edge/cloud continuum ecosystem where they can be solutions towards a more decentralized and intelligent energy system.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

TABLE 1 – SUMMARY OF THE SERVICES IN THE DELIVERABLE

Service Name	Type	Integration with Project's Interoperability Framework	Current TRL	Expected TRL (by the end of the project)
Federated Learning for Energy Forecasting & Disaggregation (ICCS)	FL	**Yes:** Open Data Connector (ODC) and semantic data standards	2	7
Vector Autoregressive Model for Energy Time Series Forecasting (INESC)	FL	**Yes:** ODC	4	5
Enhanced Network Management and Planning (UNIZG)	Edge-to-Cloud	**Yes:** ODC and semantic interoperability via the Common Information Model (CIM) standard	4	7
DTR-DLR on the Edge (JSI)	Edge-to-Cloud	**Industry standards:** SUMO bus and IEC 61850	5	7
Anomaly Detection and Predictive Maintenance on the Grid (VU)	Edge-to-Cloud	**Yes:** Semantic interoperability via the SAREF standard	2	4
APIO IoT Platform (APIO)	Edge-to-Cloud	**Yes:** ODC and semantic interoperability via the SAREF standard	4	6
Anomaly Detection and Fault Forecasting to Increase Distribution Network Resilience (VTT)	Edge-to-Cloud	**Yes:** Still studying how	4	7
Real-Time Congestion Management (TAU)	Edge-to-Cloud	**Yes:** ODC	6	8
EdgeConnect (INESC)	Cloud	**Yes:** ODC and semantic interoperability	6	7
Flexibility Optimization Service (ICCS)	Cloud	**No:** not being contemplated at the moment	3	5

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Real-Time Reserve Market Simulator (NESTER)	Cloud	**Yes:** ODC	5	7
Predictive Congestion Management (TAU)	Cloud	**Yes:** ODC	5	7
Energy Community Management Service for Frequency Restoration Reserve (INESC)	Cloud	**Yes:** ODC	6	7
Computational Orchestration Framework (TUC)	Edge/cloud continuum	**Yes:** ODC and semantic interoperability	5	6

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

TABLE OF CONTENTS

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

LIST OF TABLES

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

LIST OF FIGURES

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

ABBREVIATIONS

Abbreviation	Explanation
ABAC	Attribute-Based Access Control
aFRR	Automatic Frequency Restoration Reserve
AI	Artificial Intelligence
API	Application Programming Interface
BMS	Building Management System
BSM	Balancing Service Market
BSP	Balancing Service Provider
BTM	Behind the Meter
BiLSTM	Bidirectional Long Short-Term Memory
BUC	Business Use Case
CNN	Convolutional Neural Network
CM	Congestion Management
DER	Distributed Energy Resource
DL	Deep Learning
DLR	Dynamic Line Rating
DoA	Description of Action
DoEAP	Digitalization of Energy Action Plan
DRL	Deep Reinforcement Learning
DSO	Distribution System Operator
DTR	Dynamic Transformer Rating
EC	Energy Community

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

EDC	Eclipse Data Connector
EMS	Energy Management System
EV	Electric Vehicle
FL	Federated Learning
GA	Grant Agreement
GDPR	General Data Protection Regulation
IED	Intelligent Electronic Device
IoT	Internet of Things
KER	Key Expected Result
KPI	Key Performance Indicator
LASSO	Least Absolute Shrinkage and Selection Operator
LFM	Local Flexibility Market
LV	Low Voltage
LSTM	Long Short-Term Memory
mFRR	Manual Frequency Restoration Reserve
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
MV	Medium Voltage
MVP	Minimum Viable Product
NEMO	Nominated Energy Market Operator
NILM	Non-Intrusive Load Monitoring
PGUI	Power Grid User Interface
PV	Photovoltaic
QoS	Quality of Service

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

RBAC	Role Based Access Control
REC	Renewable Energy Community
RES	Renewable Energy Source
SCADA	Supervisory Control and Data Acquisition
SAREF	Smart Applications REFerence
SUC	System Use Case
TLS	Transport Layer Security
TSO	Transmission System Operator
TRL	Technology Readiness Level
UC	Use Case
UI	User Interface
VPN	Virtual Private Network
VAR	Vector Autoregression
WP	Work Package

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1 INTRODUCTION

The HEDGE-IoT project aims to explore the intrinsic data value by implementing intelligent data services that apply artificial intelligence across relevant sectorial use cases in the pilot demos. This implies the adoption of interoperable solutions to facilitate efficient data movement and content reasoning among pilot services. Moreover, the firm adoption of the data spaces concept in this project, namely by equipping digital services with data space connectors for inter-organizational exchanges, promotes data sharing while ensuring the concept of data sovereignty, in which organizations owning data remain fully in control of what, by who and for what purpose data is used by other organizations. Altogether, HEDGE-IoT establishes an interoperable, AI-focused development kit that explores data, computational infrastructures with commonly used services for data analysis and value sharing.

Work Package 3 establishes the technical specifications, design and development stages for the digital enablers in HEDGE-IoT, namely by : (a) addressing an extension of the digital interfaces and covering proprietary and open data interfaces of the tools considered in the pilot activities; (b) covering the design and specification of user-centric services that use artificial intelligence; (c) specifying federated learning tools, services and platforms that **operate at the edge** while ensuring (d) governance and integration orchestration operation of these services in the **edge-cloud continuum**.

The cornerstone of the technology and services in HEDGE-IoT lies precisely in exploring the continuum between the edge and the cloud, where data and services adopt distinct requirements in terms of **data availability**, **privacy**, **latency**, **scalability** and **energy efficiency** (from the computing perspective). Thus, it's relevant to introduce the key differences along the edge-cloud continuum which services explore, by covering the underlying meaning of key terms as **edge**, **fog**, **cloud**, and the continuum and what they represent in this context.

Edge and Fog computing favor a closer placement of computation and storage nodes to data generating devices (e.g., sensors, actuators, grid assets), providing several advantages: shorter distance to the origin of requests; lower latency responses and possibly increased privacy. The terms edge and fog are often mentioned in parallel with the cloud. The distinction between edge and fog computing varies. Some consider edge computing to place systems and move all tasks to edge nodes, effectively replacing the centralized resources with the edge. Others consider fog computing as mixed concepts, combining the cloud with edge nodes to extract advantages from both planes. Finally, some definitions assume no difference among them and use these terms interchangeability, representing only an option to conduct computation in a distributed fashion other than doing it at a central location.

The **cloud computing** paradigm makes access to all kinds of media (e.g., documents, presentations, photographs, videos or other raw formats) widely available due to the immediate, continuous and configurable service level agreement. It implies the transparent scalability and elasticity of the underlying supporting infrastructure, allowing services to seamlessly increase or reduce resources according to demand in computing and data storage planes.

The **edge-cloud continuum** integrates both approaches, allowing for processing and data storage to move between the edge and cloud (as swiftly as possible), depending on the computing requirements and resource availability.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

In this work package, we consider the **edge** to enclose the fog concept, as a location in the vicinity of the data producing assets, where computations occur in devices with lower capacity than the cloud, thus with limited scalability options, being characterized by reduced data storage, processing and often limited connectivity capabilities. The cloud is considered as previously described, as a compute paradigm with (virtually) unlimited capabilities that may be scaled up or down according to demand.

These definitions cascade down to how services adopt and use the edge and the cloud. We consider the **edge-cloud continuum** as it holds a unique feature reflected in dynamic allocation of computational tasks to the cloud or edge resources based on the expected workload and latency requirements. The cloud computing infrastructure takes on computationally intensive and data hungry tasks like AI model training and validation. On the other hand, the edge computing is responsible for performing AI inference closer to the data sources and real-time decision making required by latency sensitive tasks. This strategy implies resource utilization optimization in the continuum. Data replication is minimized and transferring unnecessary data to remote locations can also be avoided while maintaining the same level of contextual insight into the problem space at hand. This saves both storage and communication bandwidth. Finally, the cloud continuum enhances security and privacy by allowing sensitive data to remain at the edge devices while still enabling cloud-based insights. The continuum systems are better aligned with the strict privacy regulations like the GDPR since they can ensure that data stays at the source and maximally avoid sending sensitive data to remote locations.

The digital AI services considered in this work package, which will later be considered in the project pilots, are specialized to be deployed in the edge, in the cloud, or to operate along the edge-cloud continuum. They are organized in several classes, alluring the key techniques they consider in terms of artificial intelligence, namely: Federated Learning, Auto-Regression, Deep Learning or Anomaly Detection, being then applied for use cases such as forecasting, network management and planning, predictive maintenance and flexibility optimization.

1.1 ABOUT THIS DOCUMENT

This deliverable reports a first version of all services considered in HEDGE-IoT, which are detailed and developed in scope of WP 3 –Technological Enablers Specification, Design and Deployment.

The main objective of this document is to identify which services will be designed, developed or evolved to be later deployed in the project pilots and provide their first specifications. Namely, this document is organized to describe the AI services within the project (T3.2 & T3.3), describing their functionalities, the innovative aspects, the input and output data, the implementation details and how the service plans to integrate with the HEDGE-IoT's interoperability framework (WP4), concluding with the next steps for the service development. To support the deployment of AI services, the document also covers the main digital platforms and cloud services (T3.4) that will act as a ground base, supporting and/or linking the pilots' infrastructures for data collection, communication and mediation required for the AI services to operate. Finally, it presents the orchestration options considered to ensure that services are deployed along the continuum between the edge and the cloud (T3.5).

1.2 INTENDED AUDIENCE

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This deliverable is written considering the following audience groups:

- Service providers, developing or deploying services which are interested or require information on surveyed federated learning methods, digital platforms and services that support these techniques and how services are expected to be deployed in edge-cloud continuum.
- Potential end-users of services (i.e., consumers, other service providers, industrial partners and technology up takers) willing to understand how they may benefit from using these services.
- Researchers who are interested in understanding the potential applicability of the innovative aspects each service encompasses and how they are planned to be materialized in the context of the project pilots.
- Service providers and data space operators interested in considering the services introduced in this project to create/explore value out of the data available and to establish cross-domain use-cases.

1.3 READING RECOMMENDATIONS

This document is divided into 8 chapters. Chapter 1 provides the introduction to the document. Chapter 2 presents the federated learning mechanisms overall considered for the services. Chapter 3 approaches the data driven edge-to-cloud services that operate in the continuum, particularly applied in energy use cases. Chapter 4 covers the cloud services and digital platforms considered. Chapter 5 goes through the AI/ML tools required to govern the use of digital services and required coordination at the cloud level, while Chapter 6 approaches the deployment and orchestration across the cloud-continuum. Finally, chapter 7 concludes the document.

1.4 RELATIONSHIP WITH OTHER WORK PACKAGES

Work package 3 is defined as a pivot work package devoted to the identification, evolution and or development of digital services that provide AI capabilities to pilot activities, equipping or migrating their data interfaces to interoperable capabilities at a syntactic or semantic level.

As depicted in Figure 1, work package 3 has a direct relationship with the concept from the services to be developed, which were designed in scope of work package 2 and the baseline for the activities and existing digital platforms and services brought by project partners. Departing from these resources, work package 3 splits into several tasks devoted to the identification and extension of demo specific IoT proprietary digital interfaces, platforms and tools (T3.1). These services are then split into services that focus in consumers (T3.2) and services that focus on using Artificial Intelligence, namely Federated Learning Algorithms, Tools and Services on the edge (T3.3).

The services delivered by WP3, namely T3.3, will be integrated in the project-level digital interoperable platforms within the scope of WP4. There, services will be equipped with the Open Data Connector for interoperable inter-organization data exchanges that comply with the new data space protocol; thus, ensuring data sovereignty in all data exchanges. To support this interoperable vision, services from WP3 will be mapped into Data Apps (which under the new data space protocol may be a collection of specific control-plane policies and specific/optimized data planes) that are deployed in tandem with the Open Data Connect (T4.2) (HEDGE-IoT uptake of the Eclipse Data Connector (EDC)), discoverable, deployed and hosted (opt-in) at the App Store (T4.1). Semantic

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

interoperability features such as reasoning will be considered by service providers and developed in scope of T4.3.

Finally, the cumulative contributions of technical work packages 3 and 4 are delivered for integration in the starting digital platforms and services of work package 5 for piloting.

D3.3 focuses on the work carried out in scope of work package 3, namely in T3.3 which delivers AI services, specifically considering Federated Learning strategies at the edge.

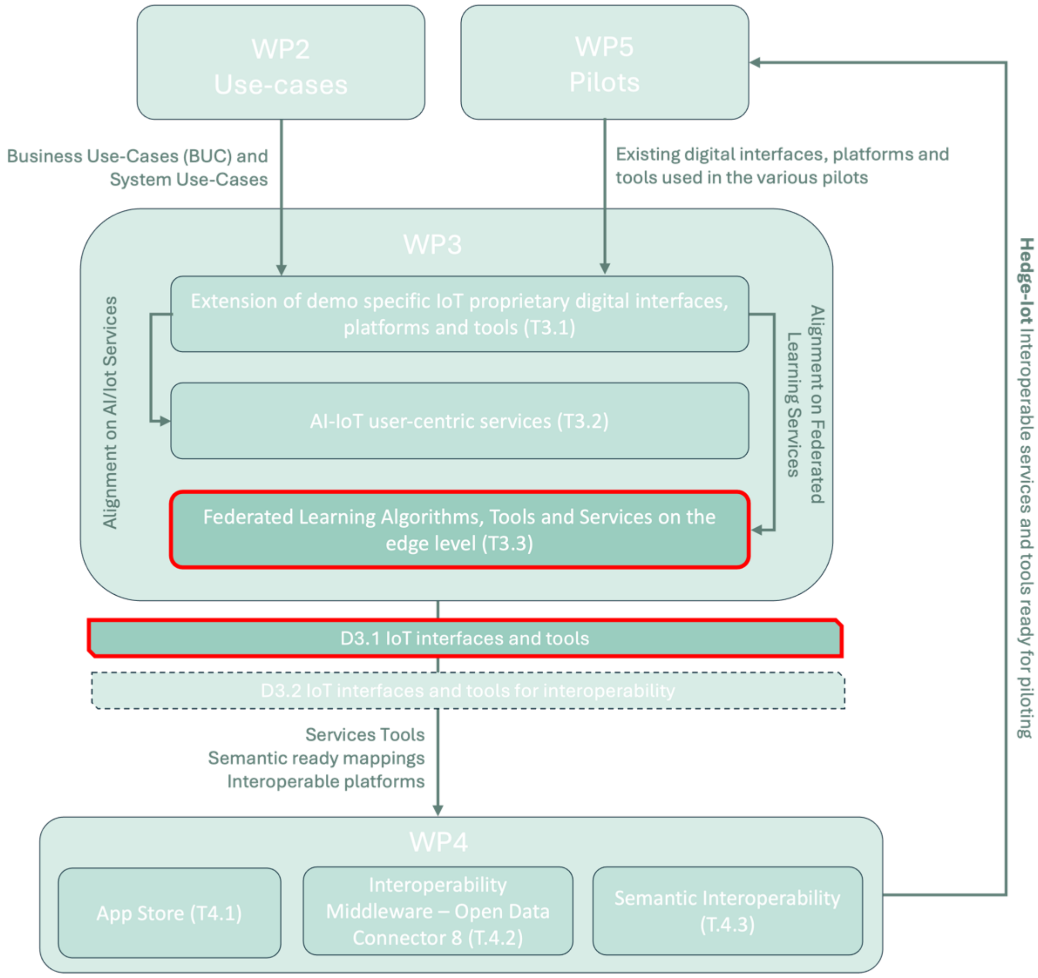

FIGURE 1 – RELATIONSHIP WITH OTHER WORKPACKAGES

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

2 FEDERATED LEARNING METHODS

Federated Learning (FL) is a decentralized approach to machine learning where multiples nodes, usually edge devices with some computational power, collaboratively train a shared model without sharing raw data, keeping said data in the local node.

FL is useful for ensuring data privacy, decentralization and scalability in low latency environments. These aspects are fundamental to achieve the goals of the project, namely the decentralization of the grid and adding intelligence to the edge layer.

Decentralized models can be useful for demand forecasting and renewable energy optimization, as they enable nodes to learn from different variable behaviors from each other, such as local weather forecasts.

This section details the first specifications of two FL methods that will respond to two use cases in different pilots:

- **Federated learning for energy forecasting (ICCS):** solution designed to enhance energy management by providing accurate predictions of energy consumption and production, coupled with detailed energy disaggregation capabilities. It uses a decentralized horizontal federated learning approach to ensure privacy by keeping raw data localized on local IoT devices, such as smart meters, while enabling centralized model training. The approach leverages advanced time-series models like LSTM and BiLSTM for forecasting and NILM (Non-Intrusive Load Monitoring) techniques for disaggregating energy consumption at the device level. It aligns seamlessly with the HEDGE-IoT interoperability framework, by adopting the project's open data connectors and semantic interoperability protocols. Currently, the service sits at technology readiness level (TRL) 2.
- **Vector Autoregressive model for energy time series forecasting (INESC):** employs a vertical federated learning approach to perform short-term energy predictions while preserving data privacy. By keeping data localized on devices and leveraging privacy-preserving encryption, the method allows data owners to collaboratively estimate VAR model coefficients without sharing sensitive information. The approach leverages VAR's multivariate component to extract relevant correlations with non-energy data, such as weather forecasts and humidity. LASSO regularization enables efficient variable selection and model optimization, addressing the high-dimensional nature of energy data. The integration with HEDGE-IoT's interoperability framework will be done via the adoption of the project's open data connector, ensuring secure and interoperable data exchanges. Currently, the approach sits at TRL 4.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

2.1 FEDERATED LEARNING FOR ENERGY FORECASTING & DISAGGREGATION

2.1.1 Description of FL method

Federated Learning for Energy Forecasting and Disaggregation is designed and developed by ICCS to provide accurate energy consumption and production forecasts, alongside detailed energy disaggregation, enabling improved energy management and planning. The service is being applied to the Greek Pilot with access to real data from approximately 100 residential apartments, collected via IoT devices and smart meters.

This federated learning model will adopt a decentralized horizontal federated learning approach, where the same type of data (energy usage metrics) is distributed across multiple nodes (smart meters in different houses). Central training will ensure consistency and accuracy while maintaining data privacy. All training will be conducted on the central server, which aggregates and optimizes the model based on data characteristics. The application of the trained model occurs locally on each participating node (main meters – e.g. Shelly 3Ems[1]), allowing them to utilize the centrally trained model without sharing raw data. This approach minimizes data transfer, protects privacy, and leverages central training for robust model performance across all nodes.

This setup ensures:

- **Data Privacy:** Raw data remains on local devices, reducing privacy risks.
- **Reduced Data Transfer:** Only model updates are exchanged, minimizing bandwidth requirements. Updates or gradients are exchanged using MQTT (Message Queuing Telemetry Transport), secured with TLS (Transport Layer Security) encryption. This ensures data in transit is fully protected against unauthorized access and eavesdropping, meeting high standards of data security and compliance.
- **Scalable Model Training:** Collaboration across nodes enhances model robustness and generalizability.
- **Energy Disaggregation:** Integration of NILM (Non-Intrusive Load Monitoring) techniques allows the system to disaggregate overall energy consumption into device-level insights.
- **Real-Time Operation**: Predictions and disaggregation occur directly on the nodes, enabling immediate insights without the need for constant cloud connectivity. "Real-time" refers to the ability to provide forecasts and disaggregation within a few seconds or minutes, suitable for actionable energy management.

The architecture diagram (Figure 2) illustrates the workflow and data flow of the Federated Learning system for energy forecasting and disaggregation. It utilises a decentralized approach to model training and deployment, ensuring data privacy and scalability.

- **Step 1 - Data Collection:** IoT devices (smart plugs) and smart meters (main electricity meter, e.g. Shelly 3EM) deployed in residential apartments collect energy consumption and production data. This data is stored locally, ensuring raw data remains private.

[1] https://www.shelly.com/blogs/documentation/shelly-pro-3em

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- **Step 2 - Central Server:** The central server aggregates updates received from local nodes (main meters), such as model parameters or gradients, rather than raw data. This ensures consistency and coordination across the distributed learning system.
- **Step 3 - Model Training and Optimization**: The central server performs the heavy lifting of model training and optimization using aggregated data characteristics. The models used for forecasting include LSTM (Long Short-Term Memory) and BiLSTM (Bidirectional LSTM) networks, which are well-suited for time-series forecasting tasks due to their ability to capture sequential dependencies. For disaggregation, the system employs Factorial Hidden Markov Models (FHMM) and Combinatorial Optimization techniques to identify device-specific energy consumption. The trained and optimised models are then distributed back to the local nodes for application.
- **Step 4 - Federated Model Deployment**: The trained model is deployed to the local nodes (main meters). Each node receives the updated global model for localized application without compromising data privacy.
- **Step 5 - Local Nodes**: The local nodes utilise the deployed model for energy forecasting and disaggregation.
- **Step 6 - Energy Forecasts and Disaggregation**: The system produces energy forecasts and disaggregates energy consumption into device-level insights. These outputs provide actionable insights for energy management, enabling smarter energy use and planning.

This architecture is highly scalable, allowing new nodes to join the system seamlessly without major increases in data transfer or computational burden. Additionally, it enables real-time insights by deploying models on edge devices, which perform energy forecasting and disaggregation locally, thereby reducing latency and reliance on cloud connectivity.

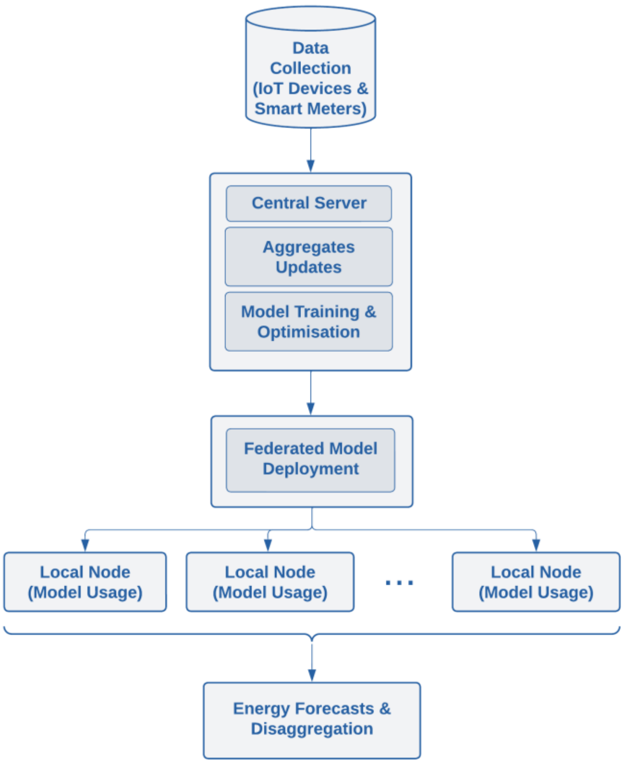

FIGURE 2 – FEDERATED LEARNING ARCHITECTURE DIAGRAM

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

2.1.2 Innovative Aspects

By the end of the project, the Federated Learning for Energy Forecasting and Disaggregation service aims to deliver an advanced, decentralized energy forecasting and disaggregation platform. This service will combine high accuracy in predicting both energy consumption and production with optimized computational efficiency, reducing latency and resource usage on edge devices. Robust security measures will be implemented to safeguard local data integrity and privacy, ensuring compliance with data protection standards. The integration of advanced NILM (Non-Intrusive Load Monitoring) techniques will enable precise device-level energy consumption disaggregation, providing granular insights into energy usage patterns. Additionally, adaptive learning models will fine-tune predictions based on localized data trends, ensuring responsiveness to varying conditions across deployment environments. Dynamic federated optimisation strategies will further enhance energy demand-supply balance, improving overall system resilience and performance. The deployment of these techniques on edge devices ensures localized data storage and efficient computation, minimizing data transfer while maintaining real-time operational capabilities. Together, these innovations will drive significant advancements in decentralized energy forecasting and disaggregation, supporting smarter, more resilient energy systems.

2.1.3 Input and Output Data

The table below categorizes the input and output data used in the Federated Learning for Energy Forecasting and Disaggregation service. The input data includes time-series energy consumption and production data, weather information, and contextual metadata, while the output data provides user-level energy forecasts and device-specific energy disaggregation insights.

TABLE 2 – INPUTS AND OUTPUTS OF FEDERATED LEARNING FOR ENERGY FORECASTING

Data Group	Variable	Variable Description	Units	Format Type	Type of Data
Input	Main residential energy consumption timeseries data	Time-series data representing overall residential energy consumption	W	Timeseries, JSON	Numerical
Input	Energy production timeseries data	Time-series data showing energy production patterns	W	Timeseries, JSON	Numerical
Input	Device-specific energy consumption timeseries data	Energy consumption data for specific devices or appliances	W	Timeseries, JSON	Numerical
Input	Weather data	Weather-related data, including temperature, humidity, etc.	°C, %, m/s	Structured, CSV	Numerical /Text
Input	Contextual metadata	Information about appliance types, usage patterns, etc.	N/A	Categorical, JSON	Textual
Output	Energy consumption forecast per user	Forecasted energy consumption at the user level	W	Timeseries, JSON	Numerical
Output	Energy production forecast per user	Forecasted energy production at the user level	W	Timeseries, JSON	Numerical
Output	Disaggregated device-specific energy consumption timeseries	Forecasted energy consumption disaggregated at the device level	W	Timeseries, JSON	Numerical

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

2.1.4 Implementation Details

At this stage in the project, the Federated Learning for Energy Forecasting and Disaggregation service is currently at TRL 2 (Technology Concept Formulated), where the foundational concepts of federated learning methodologies, decentralized model deployment, and data privacy mechanisms have been established. Initial efforts are focused on developing forecasting models using open datasets for training and evaluation, ensuring the models can accurately predict energy consumption and production trends. Alongside this, research into NILM (Non-Intrusive Load Monitoring) techniques is being conducted to enable granular energy disaggregation at the device level, enhancing the precision and usability of the forecasts.

As the project progresses, the service is expected to reach TRL 5 (Technology Validated in Relevant Environment). During this phase, pilot testing of forecasting techniques will be conducted using real datasets collected from approximately 100 residential apartments in the Greek pilot. The focus will be on refining the models for real-time short-term energy predictions across multiple horizons, including intervals of up to 6 hours. Concurrently, NILM techniques will undergo validation and refinement in real-world pilot environments, ensuring their effectiveness and reliability under operational conditions.

By the end of the project, the service aims to achieve TRL 7 (System Prototype Demonstrated in an Operational Environment). At this stage, federated training deployment will be fully implemented across the pilot, enabling decentralized models to operate independently on edge devices while benefiting from centralized optimization. The deployment will include iterative refinement cycles driven by performance metrics and feedback loops, ensuring continuous improvement in model accuracy and computational efficiency. Integration with the Hedge-IoT interoperability framework will be prioritized, using MQTT protocols and semantic data standards to ensure seamless data exchange and alignment with project-wide interoperability requirements. Additionally, real-time energy monitoring on edge devices will be conducted using Shelly 3EM meters, allowing precise measurement of incremental energy usage during forecasting tasks. This will provide critical insights into the computational energy footprint of the models, supporting further optimization. By advancing to TRL 7, the Federated Learning service will demonstrate its readiness for large-scale deployment, offering accurate energy forecasts, granular device-level disaggregation, and improved edge device performance.

2.1.5 Integration with the HEDGE-IoT Interoperability Framework

The Federated Learning for Energy Forecasting and Disaggregation service will fully adopt open data connectors and semantic interoperability protocols provided by the project. These connectors will enable seamless data exchange across different platforms, ensuring smooth communication and integration within the broader IoT ecosystem. The utilization of semantic data standards will guarantee compatibility and uniform interpretation of data across various systems and services. Throughout the project's duration, continuous validation and alignment with the Hedge-IoT data ecosystem will be carried out to enhance interoperability and foster cross-service synergy. These measures collectively ensure that the Federated Learning service operates cohesively within a smart energy management infrastructure, enabling robust, scalable, and interoperable energy forecasting and disaggregation capabilities.

2.1.6 Next Steps

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The service will progress through several key phases to ensure successful deployment and operational efficiency. First, NILM (Non-Intrusive Load Monitoring) techniques will be finalized and validated using real pilot datasets (M19) to guarantee precise energy disaggregation at the device level, enabling granular insights into energy consumption patterns. Simultaneously, forecasting techniques will be optimized and validated with real pilot data (M19) to enhance short-term energy consumption and production predictions across multiple time horizons, including intervals of up to 6 hours.

Once these validations are complete, the federated models will be deployed across the pilot, with performance closely monitored and assessed to ensure stability, accuracy, and reliability in real-world operational environments. This deployment phase will also include integration with IoT frameworks and interoperability protocols, ensuring seamless data exchange and alignment with broader system architectures.

Finally, an iterative optimization process will be implemented, enabling continuous refinement of the models based on real-world performance data. This phase will focus on improving prediction accuracy, computational efficiency, and data privacy, ensuring the service maintains robust performance and reliability throughout its deployment lifecycle.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

2.2 VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING

2.2.1 Description of FL method

The Vector AutoRegressive (VAR) model is generally used to capture linear dependencies between time series data from multiple parties. In the energy sector, VAR has been used for some time, usually associated with short-term forecasts (15 minutes to 6 hours ahead), as well as to forecast time series that originate from different data owners, which pair well with the context of forecasting residential households' energy consumption, namely when it is a part of a tertiary reserve market that allows intraday offers. To put it simply, a VAR model describes the dynamic between multiple time series by assuming that each time series in the system depends on its own past values and the past values of all the other time series in the system. As is common in this type of models, VAR captures coefficients that quantify the relationships between the different variables represented in the system across time lags. These coefficients are essential to understand the effects of the past values of a variable and the past values of other variables on itself.

Nowadays, and especially in the context of the HEDGE-IoT project, there is ample real-time access to a variety of IoT data sources that can output all sorts of data in very small intervals. Energy metering devices, such as Shelly meters, allow for the capture of energy, temperature and humidity data that, together with weather forecasts, can be leveraged to enhance energy forecasts.

In a scenario dealing with an ever-increasing number of data sources that contain personal information about households, data privacy is a major concern. For that reason, in HEDGE-IoT, we adopt a federated learning (FL) approach, by leveraging the computational power of (virtually simulated) edge devices. This approach keeps the data in the local devices, while locally estimating the coefficients of the VAR model. The main issue is how to estimate the coefficients without sharing sensitive information.

As the number of data owners increases, as well as the number of lags, regularization techniques must be used to perform automatic variable selection and remove redundant data. The Least Absolute Shrinkage and Selection Operator (LASSO) is convenient to use when handling high-dimensional data since it shrinks some of the coefficients to zero, performing variable selection. Instead of assuming that all lagged multivariate time series are contributing to the model, LASSO extracts, with a small computational effort, the predictors with the strongest contribution to forecast the target variable. In the LASSO-VAR approach, the coefficients are estimated by solving the minimization problem:

$$\widehat{\boldsymbol{B}}_{LASSO} = \arg min_{\boldsymbol{B}} \|\boldsymbol{Y} - \boldsymbol{ZB}\|_2^2 + \lambda \|\boldsymbol{B}\|_1$$

Where $\|.\|_r$ represents both vector and matrix L_r norms, and $\lambda > 0$ is a scalar penalty parameter to be tuned. A literature review [1] shows that to solve this optimization problem without compromising data privacy, existing privacy-preserving techniques are unsatisfactory. As a result, a novel solution has been put forth to encrypt the data is such a way that data owners can solve the LASSO-VAR model [2]. In summary:

1. Collaboration among data owners is established to encrypt their respective data using a collectively unknown key where only each data owner knows a portion of the key.

Additionally, covariates are locally transformed using a secret key, crucial for safeguarding the coefficient values.

2. Data owners estimate the coefficients by sharing the encrypted version of their local models. A main advantage of this algorithm is that global model's coefficients are not affected or deteriorated as is the case with differential privacy techniques.

2.2.2 Innovative Aspects

While VAR models can be very effective for very short-term forecasting and the proposed encryption technique is crucial to maintain a high level of privacy, they both come with a disadvantage: scalability. Since VAR models require the data to be in the form of a matrix, we have a much larger dimensional space than with regular time series forecasting techniques. Furthermore, the privacy-preserving encryption technique requires that multiple operations be performed on the matrices, such as inversion and multiplication, which grow exponentially, more costly, with the number of data owners (clients) and with the size of the historical horizon used in training.

During the project, the plan is to apply a multitude of techniques to optimize this approach, namely: apply parallelization methods, replace Python functions that perform mathematical calculations with Cython[2] functions and experiment with clustering FL techniques [3].

2.2.3 Input and Output Data

Table 3 categorizes the input and output data of the method, according to the different steps of the model:

TABLE 3 – INPUT AND OUTPUT DATA

Step	Data Group	Variable	Variable Description	Type of Data
Initialization	Input	Peer API URL	API URL of data owner node (peer)	String
		API version	Version of the API of the peer	String
	Output	Error message	Possible error message from the API	String
		Ledger log		List of Strings
Data Encryption	Input	Timestamps	Timestamps from the input time series	Integers
		Local data anonymized	Data on the peer without sensitive information	Float
		Random matrices	Randomly generated matrix for the encryption operation	Float
	Output	Timestamps	Timestamps from the input data	Integers
		Encrypted data	Data encrypted with the key from each peer	Float
Coefficients	Input	Encrypted data	Data encrypted with the key from each peer	Float

[2] https://github.com/cython/cython

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

		Rho and lambda	Hyperparameters to control the VAR model fit to the data (lag decay and regularization)	Float (positive value)
	Output	Encrypted coefficients	Coefficients encrypted with a secret key	Float
Forecast	Input	Decrypted coefficients and local features	Coefficients and local features after decrypting them	Float
	Output	Forecast	Final result – energy time series forecast for a given horizon	Float

2.2.4 Implementation Details

So far, an initial proof of concept with real and large-scale data has been done with a Python script that simulates collaboration between multiple agents using power data from wind farms. Therefore, it is assumed that the tool is at TRL 4, as it has been validated in a lab environment, with variables that are relevant to its use case.

By the end of D3.4 (M19), it is expected that the tool will be validated in a similar environment, as before, but using real data from the PT pilot. Therefore, the approach will be validated using variables that are relevant to the PT pilot energy community use case (BUC-PT-01), such as residential energy consumption, weather forecasts and humidity percentages.

Finally, by M30, a full implementation of the optimization techniques that make the tool scalable to the number of residential users in the pilot is expected. Furthermore, the plan is to have the tool fully operational in the pilot, working with real IoT data in a relevant environment (energy communities). The expected TRL is, therefore, 5.

2.2.5 Integration with Hedge-IoT Interoperability Framework

The FL approach is expected to fully adopt the Eclipse Data Space Connector[3] to exchange forecast data with the Energy Community Management Platform (described in section 4.5) in an interoperable manner. The idea is that Eclipse Data Space Connector will allow seamless integration of the service into the HEDGE-IoT ecosystem and ensure the possibility of seamless communication with other platforms that adopt the same standards.

2.2.6 Next steps

The next steps for this approach include the following actions:

- M15 → Finalize the data processing pipeline for the IoT data relevant to the FL approach.
- M19 → Test the solution with the relevant data and benchmark its execution times, to get a better insight its scalability.
- M20 → Implement optimization techniques to improve the scalability of the solution and test it with real IoT data:
    - Test them and benchmark the execution times to understand the evolution of the approach.

[3] https://github.com/eclipse-edc/Connector

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- M25 → Integrate the FL approach in the context of the pilot by communicating its forecasts to the Energy Community Management Platform.
- M30 → Integrate the solution with the HEDGE-IoT ecosystem:
    - Adopt the Eclipse data space connector for interoperable data exchanges.
    - Connect it to the computational orchestration module (section 6), if applicable.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3 DATA-DRIVEN EDGE-TO-CLOUD ENERGY SERVICES

This section outlines the first specifications of the data-driven edge-to-cloud energy services provided by HEDGE-IoT. These services focus on bringing intelligence to the edge and cloud layers and will be integrated into the digital framework of the project, both via interoperability (data spaces and semantic) and via the computational orchestration framework (see section 6), which will facilitate their integration by discovering available edge devices and offloading their workloads in an efficient manner.

Intelligence and decentralization

The importance of these services to the projects lies not only with their addition to the intelligence of the energy system, but also with its decentralization. Besides equipping the grid with IoT assets that cover a large variety of data sources (substations, behind-the-meter resources and feeders), these services also perform edge-level computation on local devices, while reserving the computationally intensive training of the models to their cloud layer.

Stakeholders across the energy system value chain

The services in this chapter span a large part of the energy system value chain and targets different stakeholders, namely:

- **End users:** platform that focuses on delivering aggregated IoT data to support AI algorithms with cross-sector aspects.
- **TSOs/DSOs:** Dynamic thermal rating of overhead lines based on ML-powered local weather forecasts and real-time congestion management of the distribution grid.

AI/ML solutions

A variety of AI/ML solutions are described in this chapter, demonstrating the ample aspect of the digital framework of the project and how these different services can be part of the same interoperable ecosystem, while being applied to different use cases across the energy system value chain. The algorithms described in the services include deep learning using convolutional neural networks for anomaly detection, time-series ML approaches for short-term weather forecasting and isolation forests for anomaly detection in the distribution grid.

This section details the first specifications of 6 data-drive edge-to-cloud energy services that will respond to multiple use cases across 4 different pilots:

- **ML Algorithm for Enhanced Network Management and Planning (UNIZG):** utilizes a combination of machine learning techniques to address anomaly detection, distributed energy resource (DER) capacity assessment, and forecasting of electrical quantities in secondary substations. Focused on improving planning and operational efficiency for DSOs, the service operates on edge devices to minimize reliance on centralized data storage and cloud computing. With all calculations performed close to the data source, this approach ensures low-latency decision-making and enhances data privacy. Interoperability is achieved through the PowerCIM tool for data exchange between energy stakeholders and by integrating with the project's interoperability framework. Currently, the service is at TRL 4.
- **DTR-DLR on Edge (JSI):** is an edge service that utilizes dynamic thermal rating algorithms to optimize the capacity and operational efficiency of overhead lines and transformers by

leveraging real-time weather and operational data. Using IoT devices like the maxx GW-4100 gateway, the service performs localized computations for ampacity, thermal states, and short-term forecasts, providing a contribution to decentralizing the grid. The service integrates the notion of interoperability through the SUMO bus, which uses standardized protocols (e.g., IEC 61850) to enable real-time data exchange with other grid components. This decentralized and interoperable approach reduces latency and ensures secure operation within DSO and TSO infrastructures. Currently, the service is at TRL 5.

- **Anomaly Detection and Predictive Maintenance on the grid (VU):** designed to identify anomalies and faults in the local grid through real-time analysis of streaming data from IoT devices. The AI algorithm of the service learns the nominal behavior of energy nodes through online learning, allowing it to adapt to new environments and detect anomalies in both structural and non-structural graph data. The service implements semantic interoperability by using the SAREF ontology combined with TNO's Knowledge Engine as a semantic data broker. By leveraging SAREF-compliant Smart Connectors, the service integrates with the HEDGE-IoT Interoperability Framework, facilitating transparent data exchange and ensuring compatibility with third-party dashboards. Currently, the TRL of the service is 3.
- **APIO IoT Platform (APIO):** cloud-native, multi-tenant platform designed for managing time-series data and supporting machine learning applications in the energy domain. It integrates edge devices like PGUIs (Power Grid User Interfaces) that aggregate, sign, and securely transmit data from the edge to the cloud, without running any local algorithms themselves. Moreover, the PGUIs power consumption will be estimated by analyzing the behavior of the SoC (System on Chip) under several conditions, ensuring that the energy consumption of edge devices stays low. Data privacy is ensured through encryption protocols (CHAIN2 and MQTTS), rotating credentials, and strict access controls, safeguarding the integrity of data throughout its journey. The integration with the HEDGE-IoT Interoperability Framework will be achieved by applying SAREF-based ontologies to the platform's data and by adopting the project's open data connectors. Currently, the TRL of the service is 4.
- **Anomaly detection and fault forecasting to increase distribution network resilience (VTT):** enhances grid resilience by analyzing high-resolution, real-time data streams from Intelligent Electronic Devices (IEDs) within substations. The service uses advanced deep learning to establish a baseline of the grid's normal status and uses a Convolutional Neural Network (CNN) as a primary anomaly detection model. A Deep Reinforcement Learning (DRL) model is used for fault forecasting, to identify deviations from the grid's normal state and predict the faults before they occur. By processing data locally on the edge devices, the service ensures data privacy and does not rely on central storage. To secure interoperability with grid management systems, the service uses IEC 61850 and open data standards and will also leverage the project's interoperability framework. Currently, the service is at TRL 4 having been validated using historical and synthetic data.
- **Real-time congestion management (TAU):** is a tool to manage grid congestion by gathering real-time data from primary substation and using edge nodes with significant computational power. Its modular architecture is designed to facilitate the development of algorithms. The system integrates microservices for load and generation estimation, state estimation, and congestion management. By combining active and passive grid management approaches, the service enables grid operators to use flexibility resources effectively while improving grid observability. In terms of flexibility, the service adopts the Eclipse data space connector to ensure secure and interoperable data exchanges between edge nodes and the cloud. Moreover, its data is based on the IEC61850 standard. Currently, the service is at TRL 4.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.1 ENHANCED NETWORK MANAGEMENT AND PLANNING

3.1.1 Description of the service

Machine learning algorithms

The service consists of three main parts. Each part is defined with the machine learning algorithm to solve a specific problem that DSOs face in the planning and operation of distribution networks. The service relies on the data provided by the DSO collected from IoT devices installed in secondary substations. It alleviates the planning and operational challenges for DSOs created by the increased share in renewable energy sources and other low-carbon technologies.

The first component of the service (first ML algorithm) is an anomaly detection algorithm in the initial data set. The functionality of the algorithm is to analyze collected historical measurements and to learn whether some of them do not fit well with the rest of the set, e.g., if measured voltage values are around 400 V, an occurrence of 4,000 V can be identified as an outlier. Once such anomalies are detected, they are removed and replaced by more suitable ones by applying statistical analysis and techniques. Anomaly detection can typically be approached through statistical analysis, machine learning methods, or signal processing techniques, depending on the problem and the nature of anomalies in the data. Common statistical methods include Z-score anomaly detection, moving average filters, Fourier filters, the interquartile range (IQR) technique, and the Chi-square test. Among machine learning methods, the Isolation Forest algorithm has been developed. Additional techniques can be implemented as needed to address specific data requirements. This functionality is applied each time measured data is requested and collected. Furthermore, data pre-processing is important before applying any network analysis since using invalid measurements can lead to wrong conclusions and endanger the safe and reliable operation of the distribution network.

The second ML tool is used in the DER detection and capacity assessment. Based on the measured electrical quantities and weather data collected at the substations or available at public weather services, the result of this part of the service is the detected capacity of DERs installed in secondary distribution networks. The applicability of the algorithm is not in the near real-time operation since the installed capacity of DERs does not change rapidly. The tool is more useful in the planning stage since the DSO can apply the product in assessing the contribution of DERs to the aggregated demand curve measured at the substation level. Based on already conducted research in the area, experience from multiple DSOs and publicly available measurements, PVs are the most likely technology to be detected by the algorithm since their dependency on the weather data is most significant due to indicators such as irradiance and air temperature. Therefore, the focus of this part of the service will be on identifying the installed PV capacity in a secondary distribution network.

The final ML algorithm that is part of the service is the one for forecasting values of measured electrical quantities. IoT devices installed in secondary substations measure voltage, current, and active and reactive power values. The algorithm will be trained on the historical data set for a certain quantity and provide forecasts to allow DSOs better planning and operation of a secondary distribution network. Adequate forecasting can help prevent certain unwanted events, such as transformer congestion. Accurate forecasts are the prerequisite for the safe operation of distribution networks, especially in ones with a high share of intermittent and hardly predictable DERs.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

All algorithms that are part of the service will be centrally trained in the cloud. Devices installed in secondary substation measure specific electrical quantities which are currently stored in a central database. However, that does not be the case in future since devices are installed at the "edge" of the network, meaning that simulations can be performed closer to the source of data, without the need to transmit data to the central database.

The architecture of the service is shown in Figure 3**Σφάλμα! Το αρχείο προέλευσης της αναφοράς δεν βρέθηκε.**.

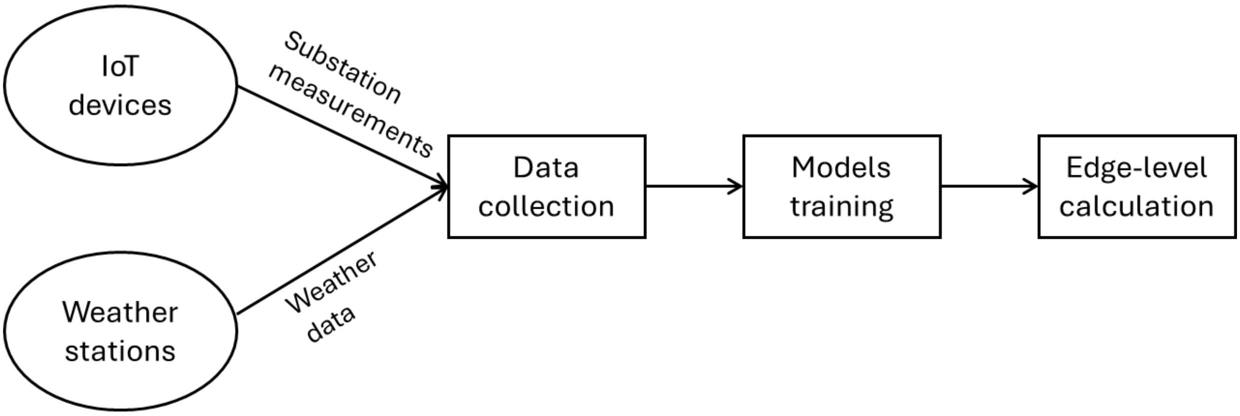

FIGURE 3 – ENHANCED NETWORK MANAGEMENT AND PLANNING SERVICE ARCHITECTURE

Functional requirements

- Historical measurements of electrical quantities and certain weather indicators are needed to train and validate the model
- The service requires access to needs, real-time measurements and weather data to calculate the solution for each of the three algorithms
- The service calculates output values on the edge-level, avoiding the transfer of data to a central cloud database.

Data privacy

Most shared data do not contain sensitive information without knowing the specific context that is not shared. However, information about the geographical location of substations will be shared to enable the identification of solar irradiation from public services on the specific location. However, no specific locations of the substations will be publicly shared. Specific identifiers of substations will be anonymized in such a way that their unique label is replaced by a randomly generated string that does not contain any sensitive information.

Low-maintenance aspect

Since the basis of the service is machine learning algorithms, using training data cannot be avoided. For each secondary substation, historical measurements needed for training need to be centrally stored and models will be separately trained and validated for each substation. However, once the training process is finalized, models will rely on data stored at edge-level and the need for a central database and exchanging data will no longer exist since all calculations will be performed close to the data source. In this specific case, the edge-level means that there is no need for devices that

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

are the source of measurements to transmit collected measurements to the central database. In the initial phase, there is still the need for a database since the ML algorithms will be centrally trained in the cloud. However, once the training process is finalized, data transmission can be stopped and all calculations can take place closer to the data source, i.e., the edge computing is enabled.

3.1.2 Innovative Aspects

Even though the application of ML algorithms in distribution networks planning and operation is rapidly increasing, in most cases such algorithms are tested using synthetic measurements that often do not represent real-world situations. Furthermore, the solutions are often missing details about accuracy and applicability aspects. This is not the case with this specific service, since for some algorithms that are part of it, e.g., anomaly detection, multiple solutions will be tested and compared. The results will be communicated with the DSO and the solution that best fits the needs of a DSO will be implemented on the demo-site.

Furthermore, there are many cases in which ML algorithms use all available data. On the other hand, there are solutions that are tested on a limited set of input data. In the first case, the algorithms become too complex and are almost impossible to train and validate within reasonable time while in the second case the information that is the result of the algorithm cannot be taken for granted. The service in the demo-site will be tested on multiple cases, relying on different input data. As a result, service users will know which measurements are needed in future since the digitalization of distribution systems raises the concern about the space needed to store all measurements and communication protocols that are needed to adequately connect to IoT devices.

Finally, similar solutions have not yet been fully tested on the edge level. Since the service relies on edge measurements, the service can be implemented close to the data source. That way, cloud calculations and data transfers to a central database can be avoided.

3.1.3 Input and Output Data

The input data consists of the measured weather data and operational data. Weather data is collected from local weather stations, secondary substations in which certain weather indicators are measured and public services from which additional data is available. Furthermore, measurements of electrical quantities from IoT devices installed in secondary substations are also input data. Based on the description of the service, IoT devices are the main source of data while weather data is used to additionally increase the accuracy of algorithms and to determine which weather data should be collected in future scenarios. Table 4 shows the input and output data needed for this service.

TABLE 4 – SERVICE INPUT AND OUTPUT DATA

Data Group	Variable	Variable description	Units	Format type	Type of data
Weather data	Wind speed	float	m/s		Input
	Wind direction	float	°		Input
	Ambient temperature	float	°C		Input
	Solar irradiance	float	W/m2		Input

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

	Pressure	float	Pa	Input
	Relative humidity	float	%	Input
	Rain Intensity	float	mm/h	Input
Substation measurements	Current	float	A	Input
	Voltage	float	V	Input
	Active power	float	W	Input
	Reactive power	float	Var	Input
Simulation results	DER production /consumption	float	kWh (kW)	Intermediate output
	Active power forecast	float	W	Intermediate output

3.1.4 Implementation Details

All algorithms that are part of the service are tested in the controlled laboratory environment, i.e., their TRL is currently 4. Each algorithm that is part of the service has been tested on synthetic data, resulting in verifying the algorithms' accuracy and applicability on a specific test case. However, the goal is to increase the TRL of the service to 7 by the end of the project. To reach TRL 7, the solution needs to be implemented by a DSO in a fully operational environment, in which all calculations are performed close to the data source. Also, the results of the calculation need to be used by a DSO in actions taken in the planning and operation of distribution networks. Each algorithm that is part of the service can be implemented on the edge level since it relies only on the specific data collected at substations. However, relatively complex and time-consuming training and validation of ML models will take place on the central cloud platform.

The solutions will be mostly used by the DSO since it is the provider of all input data and they are the entity that will benefit the most from implementing the solution in the operational environment.

3.1.5 Integration with the HEDGE-IoT Interoperability Framework

Interoperability framework is ensured through using two approaches: data space connectors and semantic interoperability. The service relies on the data exchange between multiple entities, i.e., from measurement providers, service providers, etc. Data space connectors are being considered as a solution for creating the platform that will enable data exchange and access to relevant measurements. The other measurements will be exchanged by following standard communication protocols, such as the MQTT protocol. Semantic interoperability will be ensured by relying on the specific CIM solution[4] developed for this service.

3.1.6 Next Steps

The following actions are being planned as next steps of the service:

[4] https://intellexi.hr/project/powercim-repository/

- Testing the service (ML algorithms) on data set representing the pilot
- Adaptation of the algorithms to specific characteristics of the demo site
- Implementation of the service in the fully operational environment and its application in the planning and operation of secondary distribution networks defined as a part of the demo site
- There is a possible increase in the number of output values since the applicability of some parts of the services might be extended.

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.2 DTR-DLR ON THE EDGE

3.2.1 Description of Service

The service computes the dynamic thermal rating of overhead lines (DLR) and transformers (DTR) based on the weather measurements or ML-based local weather forecast and operation data. We have two separate use cases. In the DLR case, we are interested in the thermal state, ampacity and time to overheat of the overhead power line, and in the DTR case, we are interested in the top oil temperature and ampacity of the transformer. In both cases, the simulations are running in the background, and users can access results upon request.

Architecture diagram

The architecture diagram that includes the DLR/DTR service is shown in Figure 4**Σφάλμα! Το αρχείο προέλευσης της αναφοράς δεν βρέθηκε.**. It involves an API interfacing with a DLR/DTR simulation engine, which performs computations by leveraging inputs from a local weather forecast engine (connected to a weather station API) and data from a local database. The outputs are integrated with the SUMO BUS for further processing. SUMO BUS provides standardized communication protocol to relay critical information, such as ampacity calculations and thermal capacity updates, to other grid components.

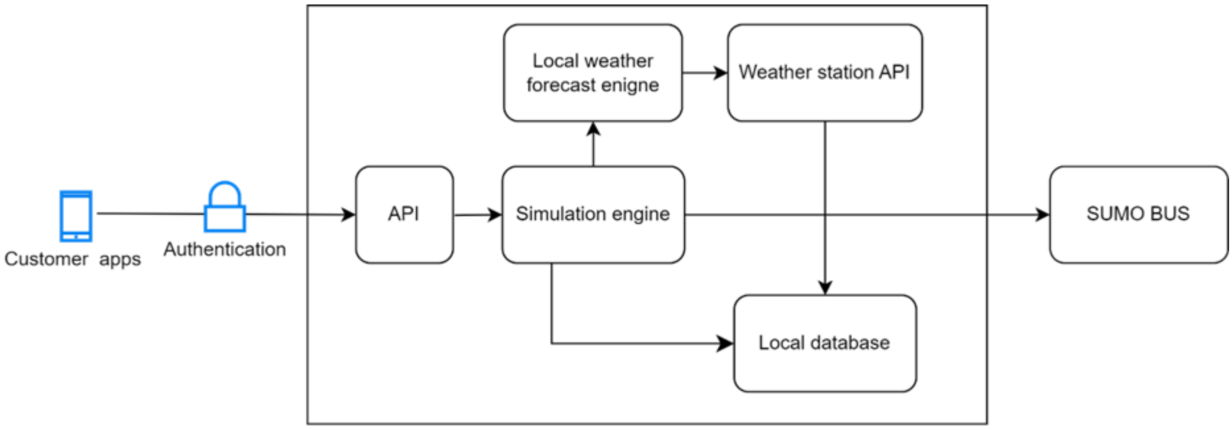

FIGURE 4 – DLR/DTR SERVICE ARCHITECTURE

Functional requirements

- The system should collect and process real-time data from local weather stations and power meters.
- The system should compute current and/or short-term production forecasts of the power line skin temperature, ampacity and time to overheat (in the DLR use case) and the top oil temperature and ampacity (in the DTR case).
- The system should be able to compute ampacity without external data sources.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Edge devices

We will use IoT maxx GW- 4100 gateway[5], as a local computational power on the edge, to run the algorithms.

Algorithms

Short-term local weather forecasts are calculated using ML algorithms, that use weather station measurements. DLR/DTR simulation computes the thermal state of the overhead power line /transformer, taking into account the weather forecast and operational data. The DLR computations are based on the standard CIGRE and IEEE thermal models, and DTR is based on a three-mass (core and winding, oil and transformer station) model mathematically described with three coupled partial differential equations and corresponding boundary conditions describing the heat exchange with surroundings. To achieve real-time response, the algorithms are developed for hierarchical execution using edge computing.

Data privacy

Although no personal or business-sensitive data will be processed, ensuring data privacy is a a priority during the development of new IoT services. We will ensure strict user-level access controls to prevent unauthorized access. All data transmissions will also be made over encrypted data channels(additional encryption to the existing encryption in the transport and lower network layers). These measures ensure a comprehensive and layered approach to data security.

Edge devices energy consumption

We aim to estimate the power consumption of algorithms by analyzing the time the System on Chip (SoC) spends in various power modes. Since neither the hardware nor the operating system directly measures the SoC's power consumption. Therefore, we can only provide an estimate of the device's power consumption by leveraging kernel-provided power state counters. Nevertheless, this method enables us to focus exclusively on the SoC's power usage, effectively filtering out the power consumption of external hardware components that fall outside our control (e.g., network interfaces).

Low-maintenance aspects

The nodes are not powerful enough to perform complete model learning, therefore the model training is centralized, while still maintaining data privacy. In this case, each node, i.e. a local weather station, will securely store its own data and periodically send relevant aggregated data to the central server. The central server will then use this aggregated information to train and improve the global model.

3.2.2 Innovative Aspects

The proposed solution will be piloted on the energy network and its assets in an advanced beyond state-of-the-art asset management concept, utilizing new forecasting algorithms for DLR/DTR. By

[5] https://www.iotmaxx.com/en/gateways-router-iotmaxx-produkte/gateways-iotmaxx-gmbh/maxx-gw4100-basic-gateway-lte-4g

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

enabling real-time, edge-based computations on IoT devices (IoT maxx GW- 4100 gateway) within substations, these services move away from static and conservative rating methodologies. Instead, they leverage localized weather data, operational parameters, and advanced algorithms to dynamically adjust transformer and line capacities. In the context of these services, "real-time" refers to the immediate processing and utilization of local data as it is received, enabling dynamic adjustments to transformer and line capacities based on current conditions. This enhances grid efficiency by accommodating fluctuating network conditions, improving the utilization of existing assets while mitigating the risk of overload or underutilization. Additionally, the integration of short-term forecasts and compatibility with standardized communication protocols further strengthens their role in modern, data-driven grid management.

The added value to the pilots, in which the service is going to be deployed, lies in demonstrating the viability of edge computing for grid applications. For the DTR use case, the edge-based DTR solution optimizes transformer operations, enabling distribution system operators to delay costly infrastructure upgrades while maintaining reliability. Similarly, the DLR case with DLR, used by transmission system operators, ensures that line capacities are dynamically adapted in real-time, reducing congestion and improving response times during critical conditions. The service exemplifies the practical benefits of decentralized computation, such as reduced latency, enhanced security, and greater resilience in network operations. The use of real-world data additionally ensures that the solution is rigorously tested under realistic conditions.

3.2.3 Input and Output Data

The input data consists of the measured weather data and operational data. Weather data is collected from local weather stations and provides the input for the DLR/DTR model and the ML local weather forecast algorithm. The weather data or forecast is input into DTR and DLR simulation to produce the main outputs. The validation data is used during the test phase to assess the operation. Table 5 shows the input and output data.

TABLE 5 – INPUT AND OUTPUT DATA OF DTR-DLR SERVICE

Data Group	Variable	Variable description	Units	Format type	Type of data
Weather data	Wind speed	float	m/s		Input
	Wind direction	float	°		Input
	Ambient temperature	float	˚C		Input
	Solar irradiance	float	W/m2		Input
	Pressure	float	Pa		Input
	Relative humidity	float	%		Input
	Rain Intensity	float	mm/h		Input
Operational data	Current	float	A		Input
Simulation results	Wind speed	float	m/s		Intermediate output

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

	Wind direction	float	°	Intermediate output
	Ambient temperature	float	˚C	Intermediate output
	Solar irradiance	float	W/m2	Intermediate output
	Pressure	float	Pa	Intermediate output
	Relative humidity	float	%	Intermediate output
	Rain intensity	float	mm/h	Intermediate output
	Ampacity	float	A	Output
	Skin temperature	float	˚C	Output
	Top oil temperature	float	˚C	Output
	Time to overheat	float	s	Output
Validation data	Skin temperature	float	˚C	Input
	Top oil temperature	float	˚C	Input

3.2.4 Implementation Details

At this stage in the project, the DLR/DTR on edge service has reached TRL 5, with key components deployed and tested in controlled environments. As the project progresses, we plan to advance the TRL of the service to level 7 by the end of the project. This will be achieved by optimizing DLR and DTR algorithms for deployment on edge devices, enabling fully independent operation without reliance on external data sources, and testing this approach within a significant real-world environment. Machine-learning-based local weather forecast models will handle predictions at the edge, while computationally intensive training of global models will be conducted on central servers. These central servers will aggregate data to train and refine global models, which will then be deployed across local stations to ensure consistent performance and improvement.

The service will be demonstrated through large-scale pilot deployments, showcasing the ability to deliver operational efficiency, flexibility, and resilience across DSO and TSO infrastructures. In particular, the DLR/DTR on edge service will be deployed in BUC-SI-0—Maximizing asset capacity for increased lifetime of DSO and TSO equipment—and their associated SUCs: SUC-SI-01.1 (Dynamic Thermal Rating edge calculation) and SUC-SI-01.2 (Dynamic Line Rating calculation).

3.2.5 Integration with the HEDGE-IoT Interoperability Framework

At this stage, it remains unclear whether the service will integrate the HEDGE-IoT Interoperability Framework, since it uses a dedicated SUMO bus developed by ELES/Operator, which is not yet contemplated by the project's framework. The SUMO bus serves as a middleware layer, ensuring seamless interoperability between edge devices, local algorithms, and other DSO systems.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The DLR/DTR on edge service leverages this framework for data exchange and synchronization between edge-based computations and cloud-based applications. The SUMO bus provides a standardized communication protocol (e.g., IEC 61850) to relay critical information, such as ampacity calculations or thermal capacity updates, to other grid components. This interoperability framework enhances the scalability and efficiency of the solution, enabling both the DSO and TSO use cases to operate cohesively within their respective infrastructures while supporting advanced functionalities like predictive maintenance and system optimization.

3.2.6 Next Steps

The following actions are being planned as next steps of the project:

- Initial adaptation of DLR and DTR algorithms for edge-based devices.
- Establish access to high-resolution environmental data from weather stations for ML model development and initial testing.
- Development of prototype machine-learning-based local weather models.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.3 ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID

3.3.1 Description of Service

This service focuses on anomaly and fault detection on the local grid, by performing real-time analysis of current and past data streams that are broadcast by registered energy nodes and other IoT devices. TNO's Knowledge Engine[6] acts as a broker between the service and these energy nodes, facilitating seamless communication using the SAREF data standard[7]. The Figure 5, below, provides a high-level overview of this service. A dedicated AI algorithm and model will take these data as input (detailed in section 3.3.3) and will learn the nominal behaviour of an energy node over a sliding window of visited data points (online learning), allowing for adaptation to new environments and gradual natural shifts, and to provide timely warnings for possible future anomalies (Predictive Maintenance).

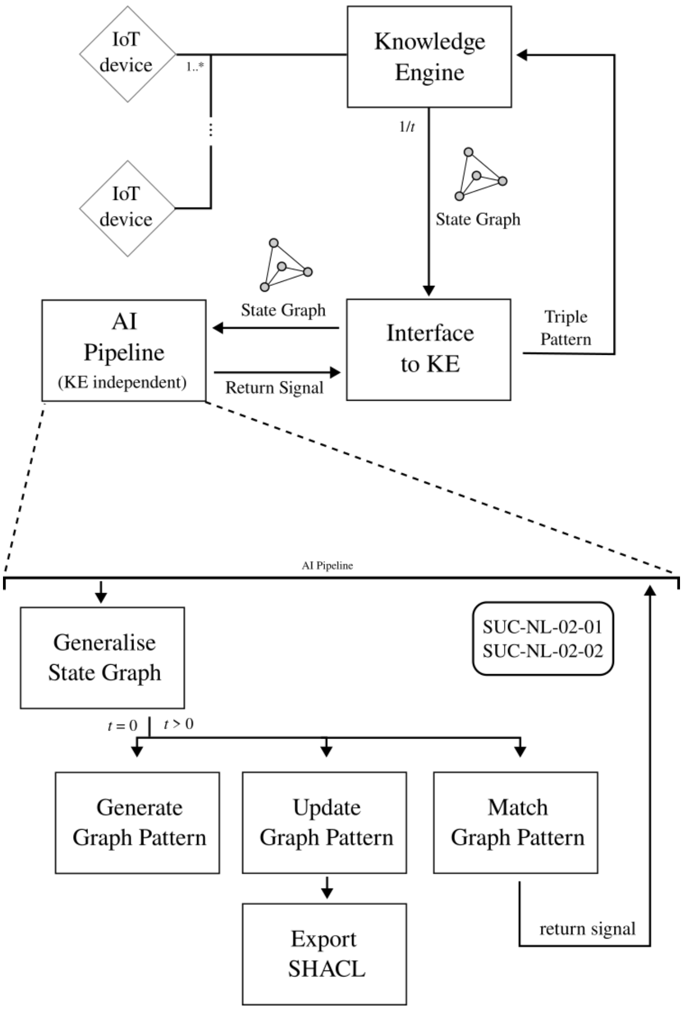

FIGURE 5 – A HIGH-LEVEL OVERVIEW OF THE ANOMALY AND FAULT DETECTION ARCHITECTURE.

---
[6] https://github.com/TNO/knowledge-engine
[7] https://saref.etsi.org/

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.3.2 Innovative Aspects

The main contributions of this service provide the ability to learn both structural and non-structural regularities from graph data. While methods exist that perform one or the other, the method provided by this service will integrate both capabilities, allowing for the detection of anomalies in the relational components of the graph (structure) as well as in its attribute values (non-structural). The ability to learn the nominal behaviour bottom-up, directly from the data, without the need for explicit input from stakeholders or other end-users. This reduces the overhead otherwise necessary to design and implement behaviour templates and allows the method to adapt to unforeseen forms of nominal behaviour.

Moreover, it considers the ability to perform online learning from streaming graph data. Instead, most other methods assume a static and complete dataset, which is an unrealistic assumption for learning with IoT devices. The method provided by this service is specifically designed to accommodate such data pattners

3.3.3 Input and Output Data

The input to this service is expected to be a stream of graphs in RDF format[8] that report on the state of one or more IoT devices to which the service is subscribed via TNO's Knowledge Engine[6]. The output likewise is a stream of graphs in RDF format, one per input graph, but which contains reports on the behaviour of the corresponding IoT devices, together with error codes and possible explanations. This return graph is published via a REST API and received by listening Knowledge Engine runtime(s). Any dashboard or interface subscribed to a listening Knowledge Engine runtime can then process and display the report.

3.3.4 Implementation Details

Development on the algorithm and model has recently begun, after having completed the design and conceptualization phases in Q4 of 2024. A prototype implementation is currently being built using the Python programming language, and will support the latest versions of RDF, SAREF, and REST. Upon completion, the service will provide a command-line interface for initiating, terminating, and configuring the service. A containerized version will be provided for easy deployment.

3.3.5 Integration with the HEDGE-IoT Interoperability Framework

The method which underlies this service is specifically designed to process streaming graph data that uses the SAREF standard for semantic interoperability, allowing it to work with all IoT devices in HEDGE that make use of Smart Connectors. The service will also incorporate a dedicated interface for seamless communication with TNO's Knowledge Engine, and will publish all results via that same interface and using open data standard, allowing third-party dashboards to receive and process reports on anomalies.

This service is conceptualized to run the cloud axis of the continuum, where service and the semantic data broker, TNO's Knowledge Engine has at least one instance. The sensing of the

[8] https://www.w3.org/RDF/

environment, in which sensor devices collect metrics on the physical medium, namely energy metrics are placed at the edge.

3.3.6 Next Steps

The next steps will have to be taken before our method reaches the desired TRL:

- Completement of the prototype implementation with the core functionality, facilitating the learning of structural and non-structural regularities from graph data (Anomaly Detection).
- Refinement of the prototype implementation to support streaming graph data and the challenges pertaining to such data (e.g. robustness to signal loss and overflows).
- Extending the method to provide timely warnings of possible future anomalies (Predictive Maintenance).
- Performing initial experiments on synthetic data to assess learning and adaptation performance.
- Performing experiments in a real-world setting at Arnhems Buiten.
- Finetuning of configurations setting to provide sane defaults to end users.
- Documenting service and containerizing method.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.4 APIO IOT PLATFORM

3.4.1 Description of Service

The Apio IoT Platform is a cloud-native, multi-tenant, IoT data platform focused on the energy domain. It spans through the whole edge-cloud continuum by integrating our edge suite and supporting edge and fog computing models. Within the platform time series data is treated as a first-class product targeted at Machine Learning based use cases.

In most common scenarios on the platform, edge devices, including PGUIs (Power Grid User Interface), will lack the computational power required to run algorithms locally. Instead, they will focus on collecting, signing, and aggregating data from the field before forwarding it to the cloud.

Architecture Diagram

In the simplified architecture diagram shown in Figure 6, the focus is on the centrality of the IoT Platform between the Edge Layer and the ML-ready Time Series Data Layer.

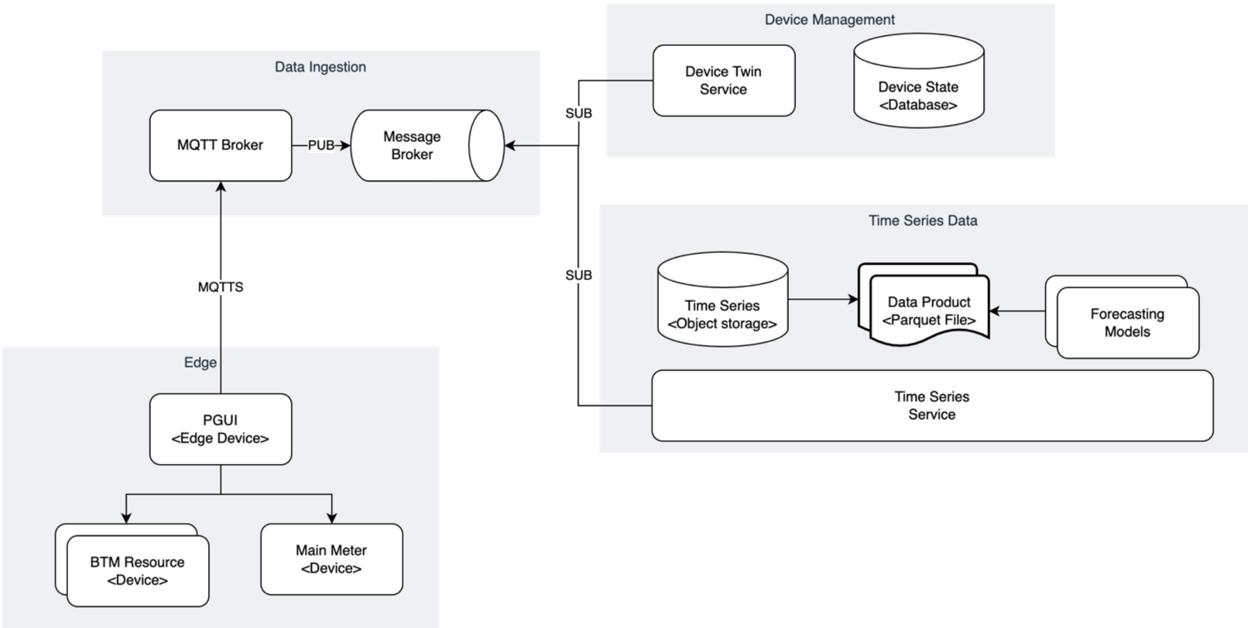

FIGURE 6 – APIO IOT PLATFORM SERVICE ARCHITECTURE

Functional Requirements

- The system should collect and process real-time data from the field devices (Meters and Behind-The-Meter resources, if available), ensuring the data is immediately (or as soon as connection is restored) transmitted to the cloud for being persisted by the time series domain.
- The system should update the device twin representation in a timely manner.
- The system should store field data in a format suitable for efficient consumption by ML Models.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- The system should be able to provide forecasts for Energy Load and Production with Low Latency.

Data Privacy

The PGUI (Power Grid User Interface) collects data from the main meter through the CHAIN2 protocol, a regulated and encrypted powerline-based communication standard, which is an encrypted channel. The collected meter data is then securely transmitted to the IoT Platform via the MQTTS protocol over an encrypted channel, utilizing rotating credentials to enhance security. Upon reaching the IoT Platform, the data is safeguarded by authentication and authorization mechanisms to prevent unauthorized access.
The data used in ML algorithms does not contain sensitive information, as the datasets are correlated with API objects using non-descriptive identifiers.

Edge devices energy consumption

The PGUI's power consumption will be estimated by analyzing the behavior of the SoC (system on chip) under several conditions, primarily during normal observability activities (sampling every 15 minutes) and during activations, where behind-the-meter resources may stress the SoC by frequent readings.

Low-maintenance aspects

PGUIs are designed to be low maintenance since they act as data aggregators both for uplink and for downlink data. Furthermore, since the devices are already used in a quite large deployment within project RomeFlex, no active maintenance is anticipated.

Cross sector aspects

The PGUI's capabilities are being extended by enabling a bidirectional Modbus communication profile, which broadens the range of use cases of the PGUIs, but it is still under active development.

PV Load / Production Forecasting

To produce meaningful forecasts for PV load and production, a dedicated data pipeline will be architected, developed, and deployed:

- The **Data Cleaning** stage will fetch device data available as Parquet Files in the Object Storage and remove outliers/invalid datapoints.
- The **Data Preparation** stage will turn raw time-series into features by enriching time-series data with other data sources such as device/plant details.
- The **Model Training** stage will train the models using prepared data. It will also store and manage model artifacts and versioning.
- **The Model Serving** stage will serve real-time (e.g. serving user queries) or batch predictions based on the deployed models.

While the generation of the data used by this system is distributed among the field devices and edge devices, the model lifecycle will be handled in cloud.
Regarding the algorithm choice, the performances and ergonomics of deep learning models, such as LSTM and transformer-based methods, will be compared:

- Prophet: fits well with seasonal data like PV production but should not perform very well.
- ARIMA family: classic approach to PV forecasting.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- LSTM: flexible enough to factor in several aspects from weather to plant layout and device characteristics.
- TimeGPT (or similar): It is interesting to understand the performances of this category of models in this scenario.

3.4.2 Innovative Aspects

While not directly edge-related, the chosen approach to time series data management delivers a performant and general-purpose data pipeline that bridges the gap between generic IoT workflows and Machine Learning algorithms.

Time-series data is ingested, processed, and stored efficiently. By utilizing object storage systems and storing data in the Parquet file format, a cost-effective, predictable, and scalable solution for time-series data storage is provided. This approach ensures the necessary performance for real-time queries by microservices, while also supporting analytical queries on large datasets, which is critical for ML workflows. As a result, it offers a future-proof solution for data-intensive applications.

Novel, transformer-based ML models for time-series forecasting will be explored, with their results compared to other approaches, such as LSTM-based techniques and more traditional, non-ML methods like ARIMA/ARIMAX.

3.4.3 Input and Output Data

TABLE 6 – DATA FOR APIO IOT PLATFORM

Group	Variable Name	UoM	Data Type
Main Meter	Absorbed Active Energy	Wh	Integer
	Injected Active Energy	Wh	Integer
	Absorbed Reactive Energy	Varh	Integer
	Injected Reactive Energy	Varh	Integer
	Mean Absorbed Active Power	W	Integer
	Mean Injected Active Power	W	Integer
	Mean Absorbed Reactive Power	Var	Integer
	Mean Injected Reactive Power	Var	Integer
Storage	State of Charge	%	Float
	Total Capacity	Wh	Integer
	SelfConsumptionEnergy	Wh	Integer
Sunmeter	Irradiation	KWh/m^2	Float
EMS / Inverter	Solar Power	W	Float
	Solar Energy	Wh	Float
	String currents	A	Float
	String voltages	V	Float
	Module Temperature	C	Float

Group	Variable Name	UoM	Data Type
Forecast	Total Produced Energy	Wh	Float
	Total Absorbed Energy	Wh	Float

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

All the relevant input data is stored for later consumption, so output data is the processed input data with the same shape but in a format more suitable for ML toolchain.

3.4.4 Implementation Details

The IoT Platform core is the foundation of Apio's ecosystem, and it is an established component. In the scope of Hedge-IoT, the focus will be on establishing the Time Series Service Domain, which will also support the Load/Production forecasting component. The architecture was validated in a research environment, with a controlled dataset and without a properly engineered data pipeline, so the current TRL is 4.

A proper Machine Learning data pipeline will be designed, implemented and deployed as mentioned in 4.7.1, allowing the usage of real-time data (Expected latency in the order of seconds) to produce real-time and batch forecasts.

Performance evaluations will include diverse algorithmic approaches, such as LSTM, transformer-based models, and ARIMA family methods, ensuring robustness across various conditions.
By testing the system under the, including device data, the algorithm's reliability, scalability, and practical usability will be demonstrated, achieving the necessary maturity for TRL 6.

3.4.5 Integration with the HEDGE-IoT Interoperability Framework

The PGUI will adopt a SAREF-based ontology. Furthermore, the use of an open data connector at the IoT Platform level is being considered.

3.4.6 Next Steps

- Increase the interoperability of the solution by adopting SAREF.
- Align time series data with the ontology.
- Implement PV Load/Production forecast with the described methods.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.5 ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE

This service offers new tools for collecting network data and processing it for further usage to support network operators in managing grid faults and improving system resilience. The opportunity revolves around modern Intelligent Electronic Devices (IEDs) such as protection relays, which can measure different parameters with extremely high resolutions that allow detailed analysis. The purpose is to apply the latest edge and AI capabilities within a substation, where the full resolution data stream is available, and perform online analysis that can use full details and thus forecast events that are slowly building up in the grid.

3.5.1 Description of Service

The service will improve grid operators' awareness and possibilities for preparing for and reacting to disturbances and faults in a timelier manner. The service focuses on possibilities of utilizing full-resolution data stream on substation bus level for improved analytics.
Currently, grid operators mostly monitor the grid and take actions based on data available through SCADA system. While SCADA provides real-time and accurate data, it lacks the high-resolution details available from modern IEDs such as protection relays. The SCADA level data is always aggregated and filtered, since transferring and storing full-detail data is practically impossible. In case of faults, more detailed data is saved through buffer system that is triggered by the fault and includes also moments before fault occurred. At the same time, IEDs have advanced significantly and are now capable of supporting highly sophisticated data analytics. Currently these possibilities are underused due to data collection, transfer and storing related challenges. Applying the latest developments in edge computing and AI based analysis offers exciting possibilities for taking the analysis on site on substations and being able to utilize all details available.
Different faults and incidents taking place in Medium Voltage (MV) distribution grid often develop over time. Typical examples are cable insulation faults that are evolving slowly or breaker malfunctions which normally first show slower opening or closing times before actual malfunctions. Also, transformer faults are often preceded by temperature rise. Many of these parameters are measured all the time; for instance, cable feeder harmonics or partial discharge measurements can indicate developing insulation fault. Likewise, breaker operation times are measured. Especially when combining different data sources, new ways for identifying phenomena in the grid could be found. However, monitoring of faults evolving on different time steps is still challenging.

The approach of this service is to provide an early warning system for the grid operator. In case there is a warning for a coming fault on certain grid part, the operator can take preparatory measures like performing grid topology changes to move important customers on other feeders or otherwise trying to limit the risk area. At the same time, manual inspections or closer look at measurements can be done to detect the problematic part. This could be implemented as “traffic light” model for the operator room: the system would show green-yellow-red light for specific grid parts and thus improve awareness of the operator. As the system is a warning system and does not have a direct protection function, false triggers can also be tolerated to a certain level. In practice the system would be self-learning so that reinforcement learning type solutions are used for operator feedback to improve the model.

With this approach, the main idea is to leverage machine learning for defining the baseline normal status of the grid, and deviations from this normal state can trigger the yellow light and warnings

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

(see Figure 7). In this sense the system does not need to exactly know what kind of fault is about to occur but instead be able to detect that the system is not in the normal state.

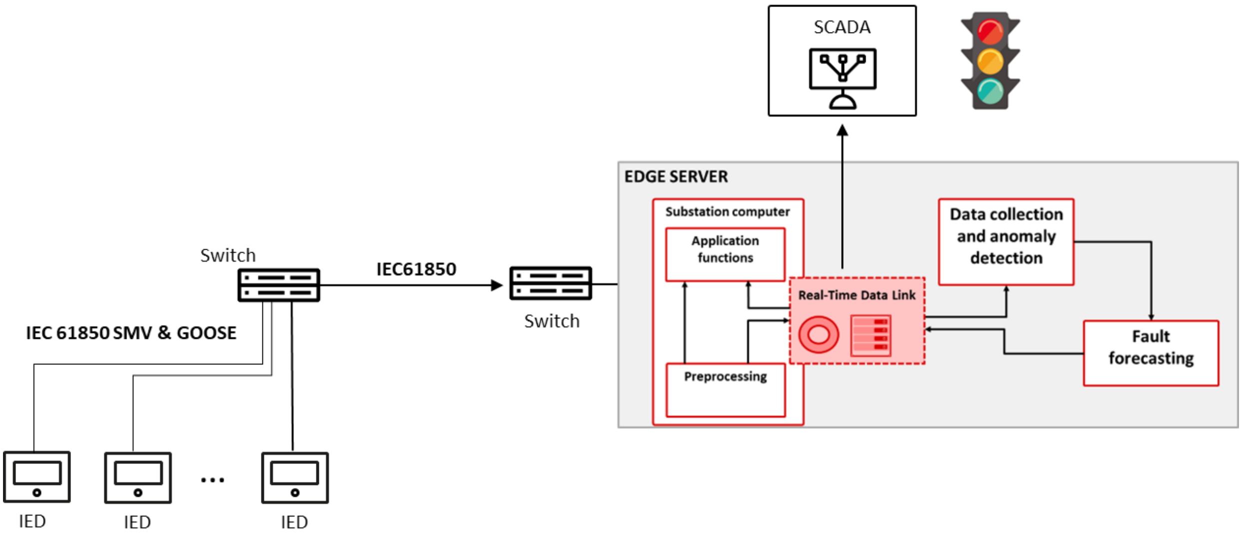

FIGURE 7 – ANOMALY DETECTION AND FAULT FORECASTING SERVICE ARCHITECTURE

Functional requirements

- The service needs to interface properly with the Real-Time Data Link and support the data format provided through the link.
- The service must be capable of handling high-resolution real-time data without delays in a reliable manner.
- The service must be able to manage loss of data, corrupted data and missing data stream situations properly, flagging them for the next steps.
- The service needs to analyze the processed data stream according to defined characteristics and calculate event indicators on a real time basis.
- The service must detect and flag all anomalies for which the indicators exceed the defined thresholds as well as report them.
- The service must support accurate time stamping for including timing information.
- The service must trigger reporting and disturbance recording where needed.

3.5.2 Innovative Aspects

The key innovation is around modern IEDs such as protection relays, which can measure different parameters with extremely high resolutions that would allow detailed analysis. However, this data cannot be stored or transferred to SCADA systems with full details due to big data masses. The data brought to operator room level through SCADA is currently filtered and less detailed, and only for occurred faults a more detailed fault information package (disturbance recording package) is saved. The purpose of this service is to apply the latest edge and AI capabilities within the substation where full resolution data stream is available and perform online analysis that can use full details and could thus forecast events that are slowly building up in the grid. The solution does not store full-detailed data but rather processes it online and keeps logs on any potentially deviating measurements or events. This analysis serves two purposes: (1) providing early warnings to grid operators for specific grid sections and (2) triggering detailed fault package recordings on correct timing.

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

For data collection and analysis, the approach involves deploying a data processing system capable of handling real-time, high-resolution IEC 61850 or similar format data streams via a data link interface. This system identifies anomalies and logs indicators in separate counters for future use in the subsequent SUC step. At this stage, the system delivers proper data streams format and identifies as the first anomaly indicators. The core concept is to employ Deep Learning (DL) to establish a baseline of the grid's normal status and flag abnormalities. Currently, the approach is not fully deterministic, as the final decisions on fault flagging will be addressed in the subsequent SUC steps. Upon completion of the project, the DL service in MV distribution grids delivers advanced anomaly detection and faults prediction solutions to enhance grid resilience and operational efficiency. This service will combine high accuracy in detecting anomalies and forecasting faults with optimized computational efficiency, minimizing latency and resource usage on edge devices. Robust security measures will be implemented to ensure the integrity and privacy of the power systems, maintaining compliance with the protection standards. The integration of multivariate time-series data from IEC 61850 and sensors will enable anomaly detection and precise fault detection, including gradual insulation faults and sudden breaker malfunctions, providing granular insights in grid performance.

The Convolutional Neural Network (CNN) serves as the primary anomaly detection model. CNN learns to recognize and classify a wide range of anomalies and faults based on the real-time data received from the distributed energy systems. Once trained, the CNN model optimized accuracy, precision, and recall metrics to ensure robust detection and classification performance in substations. The designed model on edge devices ensures localized data processing, reducing the need for large-scale data transfer while maintaining real-time operational capabilities. To further refine its self-learning-based decision-making capabilities, the system incorporates a feedback loop using DL. For fault forecasting, the DL provides an early warning system for the grid operator. In practice, the DRL would be self-learning which enables the model to continuously adapt to operator feedback, improving its ability to detect new anomaly and faults scenarios over time in the substations. Together, these methods will significantly enhance the resilience of distribution grids, enabling smarter, more efficient grid management and improving faults detection and forecasting accuracy.

Table 7 Table 7categorizes the input and output data used in the Deep Learning for anomalies detection service. The input data includes time-series data received from IEDs located at different positions in the distributed feeders, while the output data provides operator level device-specific anomalies detection insights.

3.5.3 Input and Output Data

TABLE 7 – ANOMALY DETECTION SERVICE FROM VTT INPUT AND OUTPUT DATA

Data Group	Variable	Variable description	Units	Format type	Type of data
Process bus data stream	Current	float	A	IEC61850-9-2 Sampled Value	Input
	Voltage	float	V	IEC61850-9-2 Sampled Value	Input

	Harmonics	float	Hz	IEC61850-9-2 Sampled Value	Input
Warning / traffic light signal Event information package Disturbance recording triggering	Green/ red/ yellow	Integer			Output
	Triggered values	Text			Output
	Trigger signal	Integer			Output

3.5.4 Implementation Details

The Deep Learning service is designed to provide anomaly detection in distribution grids to enhance grid resilience and operational efficiency. The service is being applied to the Finish Pilot with access to historical and real-time data from feeders and substations, collected from IEC 61850 and sensors. A centralized server model architecture integrates multivariate time-series data from IEC 61850 and sensors, such as voltage, current, and harmonics readings. Historical and live-streaming data from power system components, stored on the centralized data storage system, ensures reliable and efficient access for simulation and testing purposes. The data includes specific types such as missing, inappropriate, duplicate, fake, outliers, noise, and corrupted data. Both historical and real-time data undergo preprocessing steps such as normalization and segmentation into time-series snapshots. CNN serves as the primary anomaly detection model. Its training involves using historical labeled data stored in the centralized server. This dataset includes examples of both normal and abnormal grid conditions, such as missing, duplicate, fake, noisy, and corrupted data. The diversity of the training data ensures the CNN learns to recognize a wide range of anomalies. The model is optimized using categorical cross-entropy as the loss function and evaluated with accuracy, precision, and recall metrics to ensure reliable and robust detection performance.

Once trained, the CNN model is deployed on real-time data received from the distributed power system devices located at feeders and substations. This deployment enables the system to analyze high-frequency data streams in real-time, minimizing latency and ensuring prompt detection of anomalies. The use of edge computing allows for localized processing, reducing the dependency on centralized resources and enhancing system responsiveness. To further refine its detection capabilities, the system incorporates a feedback loop powered by DL. The DL model continuously improves by interacting with the grid environment and incorporating feedback from operators. For instance, if an operator confirms or refutes an anomaly and faults, the DL model adjusts its predictions and strategies accordingly. This self-learning capability allows the model to adapt to new grid conditions, detect previously unseen anomalies, and refine its predictions over time.

The system also includes a feature extraction process that identifies critical attributes from the multivariate time-series data. These extracted features are processed in distributed blocks to perform tasks such as anomaly classification and clustering in anomaly counter logger. Anomalies detected are subjected to root cause analysis, allowing the system to identify underlying issues and suggest actionable insights. Based on these insights, a set of corrective actions (e.g., Actions 1, 2, …, n) is generated to maintain grid stability and operational efficiency. To simplify decision-making, the system provides a user-friendly interface that visualizes the health of the power system components using a traffic light system. Green indicates normal operations, yellow signals potential

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

warnings, and red highlights faults requiring immediate attention. This intuitive visualization enables operators to quickly assess the grid's status and take appropriate actions.

This deep learning service delivers a comprehensive, efficient, and adaptive solution for secure and more resilient grid operations (see **Σφάλμα! Το αρχείο προέλευσης της αναφοράς δεν βρέθηκε.Σφάλμα! Το αρχείο προέλευσης της αναφοράς δεν βρέθηκε.**).

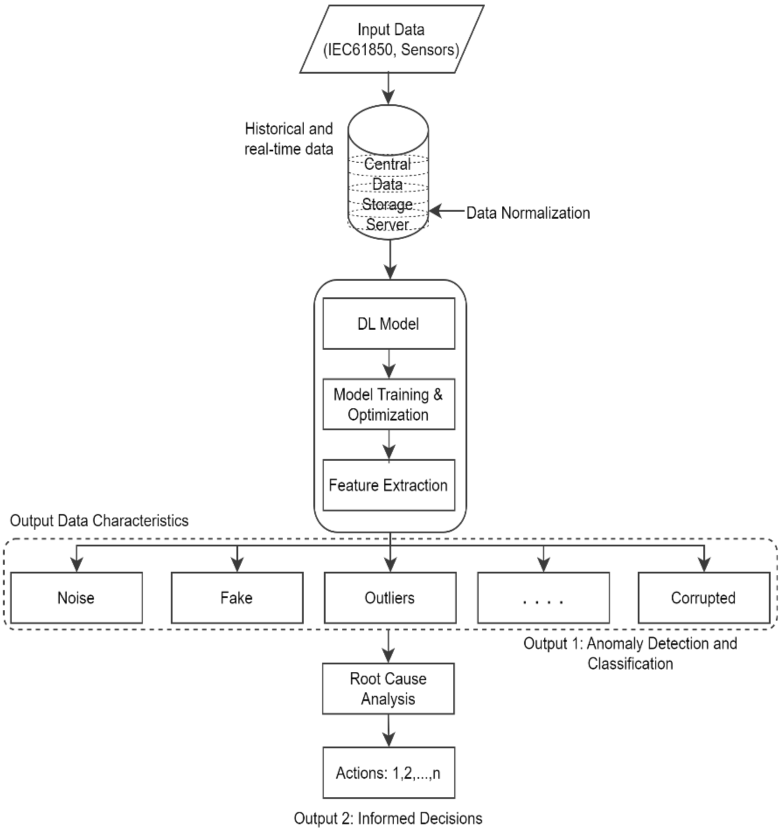

FIGURE 8 – DEEP LEARNING ARCHITECTURE DIAGRAM FOR ANOMALIES DETECTIONS AND FAULTS FORECASTING

3.5.5 Integration with the HEDGE-IoT Interoperability Framework

The designed method specifically processes high-resolution, real-time data streams from IEDs in distribution grids, leveraging the HEDGE-IoT Interoperability Framework for seamless integration. Details on how this framework will be integrated with the service need to be further studied, since its details have only been clarified recently. By utilizing standardized communication protocols, including IEC 61850, the service ensures compatibility with a wide range of IoT devices in distribution grids. The service incorporates a dedicated interface for smooth data exchange, enabling real-time anomaly detection and fault forecasting while supporting communication with other grid management systems. All results are published via this interface using open data

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

standards, allowing grid operators to receive, process, and visualize anomaly and faults reports for improved grid resilience and operational efficiency.

3.5.6 Next Steps

By the next deliverable, the focus will be on developing an initial model that will lead to early-stage implementation and validation guidelines. Key activities will include:

- Assessing the historical and real-time data streaming characteristics.
- Developing early-stage model for implementation and validation.
- Testing the service on historical energy system datasets.
- Adapting algorithm to align with the specific characteristics of the demonstration site for real-time data streaming.
- Deploying the service in a fully operational environment at the demonstration site.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

3.6 REAL-TIME CONGESTION MANAGEMENT

The primary objective of this service is to provide real-time congestion management (CM) for distribution grid operators (DSOs), allowing them to utilize grid flexibility in real-time to mitigate congestion issues and improve overall grid operations, and enhance grid hosting capacity. In the context of this service, "edge" refers to the distribution grid's control unit, located at either at the primary substation or a remote location to which substation measurements (IEC 61850-9-2 sampled values) and substation status information and control signals (IED 61850-8-1 GOOSE) can be transmitted from the substation.

3.6.1 Description of Service

Figure 9 illustrates the architecture of the entire CM ecosystem, emphasizing the role of real-time CM at the edge. The real-time in this context means from seconds up to a minute. The real-time CM service has been divided into several micro-services—such as load and generation estimation, state estimation, and real-time CM—following a service-oriented design. This modular and maintainable approach facilitates the efficient development and management of algorithms. Consequently, the edge node requires significant computational power to execute these micro-services effectively. Neither computational power nor the energy consumption of edge is a critical concern because the system employs only one edge server per each primary substation. Despite its limited quantity, this server has a considerable impact on real-time grid operations.

The real-time CM service on the edge relies on three key data sources: **(1)** real-time data from the grid, transmitted through a dedicated real-time data link, and **(2)** harmonized data from the cloud, accessed via the Dataspace connector adopted in HEDGE-IoT - Eclipse Data Space Connector - and processed through an edge-cloud data exchange adaptor, and **(3)** internal data storage which contains data concerning the grid, and profiles of loads and generations.

The real-time CM service offers significant benefits for the DSO by enabling two key functionalities: **(1)** real-time observation of the grid state and **(2)** the utilization of flexibility for CM. This approach enhances the grid's hosting capacity, allowing new end-customers to connect to the grid more quickly than traditional CM methods, such as grid reinforcement, would permit.

Furthermore, the real-time CM service enhances stakeholder data exchange by incorporating multiple data sources. For instance, smart meter data from end-customers is utilized to create more accurate load and generation estimations necessary for state estimation, improving the DSO's grid observability. Additionally, data from flexibility service providers (FSPs) regarding the characteristics of flexible resources is integrated into the process. Ultimately, achieving the goal of increased grid hosting capacity relies on effective data sharing among various stakeholders.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

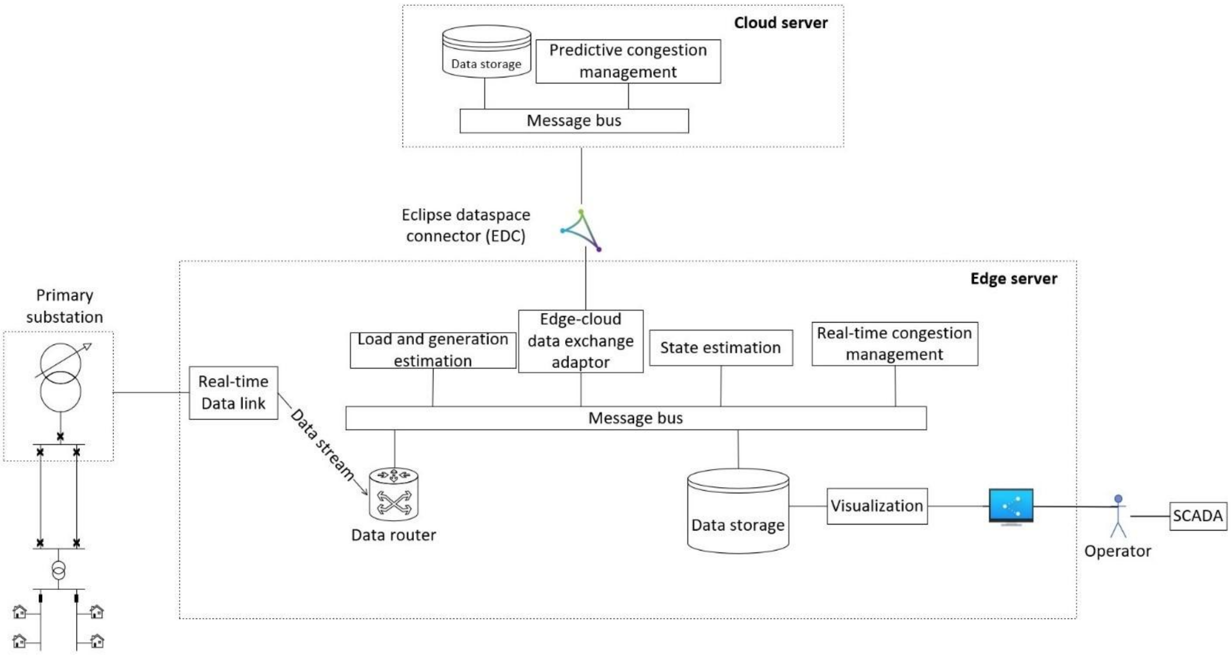

FIGURE 9 - REAL-TIME CONGESTION MANAGEMENT SERVICE ARCHITECTURE

3.6.2 Innovative Aspects

Currently, most DSOs have yet to fully/partially adopt active grid management solutions necessary to address the growing levels of congestion in their grids. Several factors contribute to this, including longstanding national regulations in EU member states that have traditionally favored passive solutions, such as grid reinforcement, over active approaches. However, this has rapidly evolved with the introduction of the Clean Energy Package's guidelines and regulations in 2019, which prioritizes the use of flexibility and incentivize DSOs to adopt active solutions wherever technically and economically feasible.

The slow pace of adoption can be attributed to weak motivation among DSOs, which is closely tied to regulatory frameworks, as mentioned, as well as the inherent complexities of implementing active grid management systems. These systems require seamless access to harmonized data from diverse sources, advanced algorithms for decision-making, and the integration of controllable devices to execute control actions. Accomplishing those steps is challenging and slow because each step has several sub-steps and some involve the contribution of more than the DSO itself, which requires time and effort to have all required stakeholders on board. Meanwhile, the rising penetration of renewables is placing increasing stress on distribution grids, prompting DSOs to explore a hybrid approach that combines passive with active solutions. This dual strategy is a key innovation of the proposed service from an industrial perspective, offering a path to more efficient and adaptable grid management so that DSOs would leverage both active and passive grid solutions for their grid's CM.

From a research perspective, active grid management has been a widely discussed topic by academics for some time. However, practical solutions that can be implemented and widely accepted by industry stakeholders remain limited. The proposed service aims to bridge this gap by jointly developing the real-time CM service by partners within the Finnish pilot, with diverse expertise and knowhow including Tampere University (TAU) as academic partner, ABB as

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

technology provider and industry pioneer, JSE as DSO and Enerva as JSE's close partner that operates JSE's network. By developing the service collectively, it is ensured that the gap between academia and industry is reduced leading to a CM service that is viable in practice.

The service also builds upon the extensive insights gained from previous EU projects, such as INTERRFACE and IDE4L, in which TAU played an active role. INTERRFACE focused on exploring market-based CM, while IDE4L concentrated on utilizing grid-side flexibility for CM. By integrating the learnings from both projects, the proposed service combines market-based and grid-side flexibility for CM, representing a novel and comprehensive approach to active grid management

3.6.3 Input and Output Data

The service will use the real-time data stream as the key input. The data is based on IEC61850 that is provided through the real-time data link. The data includes sampled values of voltages and currents measured by merging units at the beginning of each feeder and on some locations along the feeders. Also, the service requires historical load and generation profiles that could be static in the data storage. The edge requires data coming from the cloud as well, such as electricity market price data, default settings of controllable resources calculated by predictive CM. The service's output is also logged into the data storage to be used for visualization. The visualized data is then used by the grid operator to make decisions according to the real-time CM recommendations. This means that the real-time CM service will not act automatically in the pilot when congestion arises, because that would require rigorous validation and testing before being put into use in the real grid operation.

3.6.4 Implementation details

The implementation of the service will be break down into steps to facilitate coordination and cooperation between pilot partners. The steps are input data preparation and harmonization, including grid data and load and generation data, work related to edge server (engineering the hardware, establishing virtual machines, networking), component development and testing, integration testing of developed components, and finally, piloting the solutions. Each step has several tasks involved, and those are assigned during the development to the pilot partners depending on their interests and expertise.

3.6.5 Integration with the HEDGE-IoT Interoperability Framework

The service will implement Eclipse Dataspace Connectors as its data connectors for edge-to-cloud and cloud-to-cloud data transfer. The Eclipse components are a comprehensive framework (concept, architecture, code, samples) providing a basic set of features (functional and non-functional) that dataspace implementations can re-use and customize by leveraging the framework's defined APIs and ensure interoperability by design. The HEDGE-IOT framework is not yet known at this point however, it is expected that utilizing Eclipse data space components could facilitate integration with HEDGE-IOT framework.

3.6.6 Next Steps

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This service is currently at the design stage and implementation work has just started. The service realizes SUC-FI-02.03 and SUC-FI-02.03, which is the second business use case of the Finnish pilot (BUC-FI-02). The implementation would need the following steps:

- Input data preparation (M12-M21)
    - Grid data (grid topology, characteristics)
    - Historical load and generation data
    - Validation of grid data
    - Validation of historical load and generation data
    - Associating load and generation data to the grid topology
- Edge-cloud connection through Eclipse (M18-M30)
    - Developing the necessary adaptor for edge-cloud exchange
- Architectural design (M12-M21)
    - Location of micro-services (cloud/edge)
    - Data flow between micro-services
- Component development and testing (M18-M30)
    - Load and generation estimation
    - State estimation
    - Real-time CM
    - Visualization
- Integration testing and piloting (M30-M42)

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

4 CLOUD SERVICES

This section introduces the cloud services provided by HEDGE-IoT. These services play pivotal roles in the pilots they are associated with, since they usually work to support or in tandem with other services closer to the edge. Moreover, they will show how the project's framework can support even services that are fully cloud based, adding intelligence to this layer. Ultimately, cloud services are important to heavily lift the computation of models and provide larger software infrastructure to support large scale and high maturity pilots, therefore it is important that the HEDGE-IoT framework supports them and enables their interoperability with other services.

This section details the first specifications of 5 cloud services that will respond to multiple use cases across 3 different pilots:

- **EdgeConnect (INESC):** provides an ecosystem for stakeholders across the flexibility value chain, enabling integration, qualification and market participation, to unlock flexibility potential. The service facilitates onboarding and certification of users, registration and pre-qualification of flexible assets, sharing of flexibility needs, baselines and bids and activation and settlement of flexibility services, allowing consumers to actively participate in energy markets. EdgeConnect ensures data privacy with role-based data access, while having critical information anonymized. Service to service data exchange interoperability is guaranteed via the integration of the project's data space connector. Furthermore, semantic interoperability will be integrated using the approach defined in the project. Currently, the service is at TRL 6, having already been tested in a controlled environment in a different European project.
- **Flexibility optimization service (ICCS / HENEX):** is comprised of four modules that enable consumers to participate and place bids in local flexibility markets: 1) short and long-term forecast for energy demand and production; 2) calculation of incentives optimization and formulation of optimal bid; 3) communication of flexibility requests to consumers and 4) submission of bids in the local flexibility market. The service ensures data privacy by enforcing encryption and access control mechanisms like Attribute-Based and Role-Based Access Control (ABAC and RBAC). Currently, the service is at TRL 3 and the focus is on establishing the foundational architecture.
- **Real-Time reserve market simulator (NESTER):** is designed to emulate manual Frequency Restoration Reserve (mFRR) and automatic Frequency Restoration Reserve (aFRR) market operations, providing TSOs and Balancing Service Providers with a platform to test and optimize bidding strategies. The simulator validates bids, performs market simulations, and delivers outputs such as settlement curves and activation setpoints. The main advantage of the tool is enabling real-time market simulations with real consumer and TSO data, but without needing to comply with the full restrictions of the actual market. The tool will integrate HEDGE-IoT's interoperability framework by adopting the project's data space connector and performing data exchanges with the previously mentioned EdgeConnect service. Currently, the application is containerized using Docker and deployed in AWS.
- **Predictive congestion management service (TAU):** is composed of several micro-services for load, generation and grid state forecasting with the aim of enabling grid operators to procure flexibility. It uses 3 data sources: weather, market, and historical grid data to make predictive analyses and foster market participation. The service will implement the project's data space connector, therefore integrating its interoperability framework. Currently, the service is still at the design stage, as no implementation work has started yet.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- **Energy community management service for frequency restoration reserve (INESC):** enables Renewable Energy Communities (RECs) to participate in Balancing Service Markets (BSMs) by provisioning mFRR and aFRR, using a set of modules to manage an energy community, including sizing, energy management and settlement. Since they act as natural aggregators, this service can coordinate a REC's members' flexibility while adhering to strict TSO requirements for frequency restoration reserve markets. The service will integrate HEDGE-IoT interoperability framework by adopting the project's data space connector, which will be used for interoperable data exchanges with other frequency restoration reserve market stakeholders. Currently, the frequency restoration reserve service is at TRL 2, while the underlying energy community management platform is at TRL 4, having been tested in a lab environment in the scope of another project.

4.1 EDGECONNECT

4.1.1 Description of the Service

The EdgeConnect platform provides stakeholders (i.e., consumers, flexibility service providers, aggregators, DSOs and TSOs) along the value chain of flexibility provision with an integrated ecosystem to support all main activities in this value chain, to help identify, unlock and make use of all available flexibility potential. As a multi-stakeholder platform, it comprehends several views, providing distinct value propositions for each stakeholder. Below, we provide a brief description of each functionality and how it can create value for the stakeholders.

Flexibility Enablement

Consumers with flexible assets are incentivized to participate in the flexibility value chain, leading to the onboarding of their assets.
Consumers / prosumers (asset owners) register their flexible assets to participate in the value chain to take their flexibility potential to market, in exchange for benefits/incentives provided by the service providers that want to capture their flexibility potential.
The expectation is that asset owners want to collect a return on investment of the installed assets, and the service providers want to utilize the offered flexibility potential to solve grid problems.

The currently implemented system architecture is shown on Figure 10.

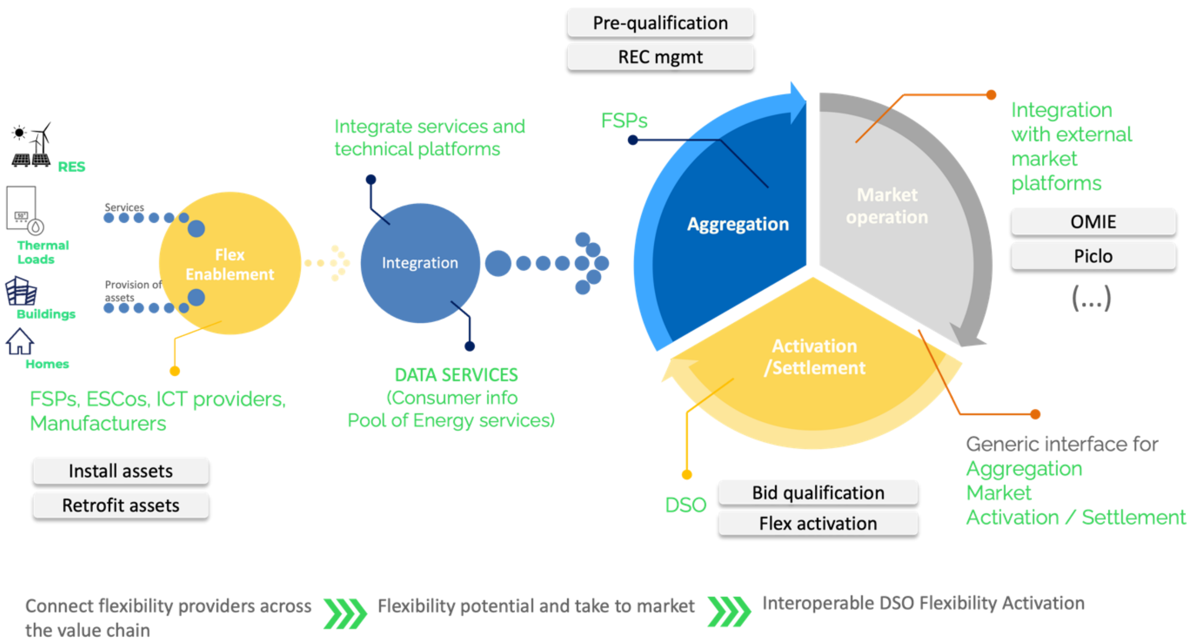

FIGURE 10 – EDGECONNECT SYSTEM CONCEPT

As the figure above shows, the main functionalities of EdgeConnect consist of:

Integration of services and assets

Flexible assets and services are integrated into the value chain via technical assessments. Consumers / prosumers (asset owners) allow technical platforms available to communicate information to be used to compute their flexibility potential and to send commands to the flexible asset that realize said flexibility potential. The expectation is that participation in the flexibility market can be fully conducted through a service provider by sending commands to flexible assets to change their consumption profiles.

Pre-qualification

Certification of aggregators and flexible assets that ensures financial and technical capabilities to participate in the flexibility market. This functionality ensures that service providers adhere to the financial requirements of market participants and that flexible assets have the technical capacity to provide the flexibility that the consumer offers.

These pre-qualification processes are assumed to be performed outside of the EdgeConnect ecosystem. Our service has the role of facilitating and increasing the interoperability of these processes.

Flexibility needs and bids

During the negotiation phase, EdgeConnect allows the System Operator (SO) to publish their flexibility needs to the market, where the process of acquiring flexibility using the local market starts and the aggregators can respond to it by submitting their own flexibility bids, together with a baseline.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Select winning bids

During the market operation phase, the market clearing process is performed inside the market platforms (outside EdgeConnect). Afterwards, the SO selects the winning bids and communicates the results to EdgeConnect, which will relay it to the aggregators.

Activation

In the activation phase, EdgeConnect will inform of the activated flexibility, in the form of setpoints to the aggregators. It is assumed that the disaggregation of the flexibility for each asset is the responsibility of the aggregators, while the activation of the asset will be done by the consumer/prosumer (for example, with an Energy Management System).

Settlement

The settlement functionality uses the consumer metering data provided by the SO, and the baseline sent by the aggregator, to calculate the aggregator's delivered flexibility and calculate its remuneration and penalties.

This functionality can be done using the built-in methods from EdgeConnect or the SO can choose to settle directly with the aggregators and simply relay the final financial settlement considerations via EdgeConnect.

The service focuses heavily on **data privacy** ensuring that critical information is either prevented from being stored persistently or that anonymization is implemented as much as possible. Other information, which is stored, but not anonymized, is subjected to the terms and conditions applicable. Due to the multi-stakeholder environment, a role-based controlled access approach is adopted where each stakeholder can only access the data that is relevant to them. When using the interoperability enablers in place, the same restrictions apply, and a cascade of the permissions are pushed as usage filters before the hand-over to the interoperability enablers.

Reusability is also a strong focus of this platform. Edge Connect's main development departs from the Grid and Market Hub platform developed in context of HEU BeFlexible. New bilateral contract services for the provision of flexibility services will be included, together with enabling the platform with the new generation of data space connectors. Moreover, the platform provides the possibility for services to be included, allowing each stakeholder type to host them in EdgeConnect.

4.1.1.1 Functional Requirements

- The system must be able to onboard all its stakeholders
- The system must allow consumers to register their flexible assets
- The system must match consumers with the relevant energy community service provider
- The system should allow technical platforms to test activation commands on the consumers flexible assets, to verify their capability to provide flexibility
- The system must support the certification of Aggregators to ensure they are financially and technically eligible to participate in the flexibility market

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- The system must recognize that market and pre-qualification processes can occur outside its ecosystem and provide functionalities to interface with these external processes
- The system must support crucial flexibility value chain operations, namely share flexibility needs, share selected winning bids, distribute activation

4.1.1.2 Non-Functional Requirements

- The system must store data securely in an encrypted manner
- The system must exchange data securely
- The system should have an uptime of 99%
- The system must use the OAuth 2.0 protocol for user authorization and role management
- Requests should not take longer than 10 seconds to produce a response

4.1.2 Innovative Aspects

The main innovation of this service is the concept itself. EdgeConnect allows the integration of energy market stakeholders into a single ecosystem that promotes transparency and interoperability and allows them to perform common market operations. The integration of aggregators, instead of simply considering each asset as a standalone, opens the door for a local energy community to participate in the flexibility market, which is a differentiating factor.

As a multi-tenant, cloud-based application, EdgeConnect can operate several energy markets in Europe, providing the interlink for day-ahead and long-term market bids. Thus, Edge Connect can provide services for stakeholders (e.g., Aggregators) that hold operations in European regions where different markets are in place. This is as regulations governing energy markets may differ from region to region, or the operator holding a digital platform to receive and short-list flexibility bids is in principle different. The services provided then offer seamless control and integration.

Another innovative aspect is the empowerment of the consumer, enabling its participation in the market and the choice to monetize it. The consumer is also free to choose the service provider with the most attractive conditions.

4.1.3 Input and Output Data

The inputs and outputs of this service are all related to the flexibility value chain, its stakeholders and operations. In terms of stakeholders, EdgeConnect allows for the creation of users, services and assets and interlinking data repositories. In terms of operations, EdgeConnect allows for the qualification of an object, creation flexibility needs, submission flexibility bids, submission baselines, selection and activation bids. Table 8 shows the input and output data for some transactions.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

TABLE 8 – SAMPLE OF EDGECONNECT INPUTS AND OUTPUTS

Object	Variable	Variable type	Units	Format type	Type of data
Asset	Id	UUID	-	JSON	Input
	Owner id	UUID	-	JSON	Input
	Active power import capacity	Integer	W	JSON	Input
	Active power export capacity	Integer	W	JSON	Input
	Power	Integer	W	JSON	Input
	Connected	Boolean	-	JSON	Input
	Divisible	Boolean	-	JSON	Input
Bids	Business partner id	UUID	-	JSON	Service parameter
	Date interval	Date-time	-	JSON	Service parameter
	Id	UUID	-	JSON	Output
	Minimum quantity	Integer	kW	JSON	Output
	Maximum quantity	Integer	kW	JSON	Output
	Start datetime	Date-time	-	JSON	Output
	End datetime	Date-time	-	JSON	Output
	Activation Price	Float	€	JSON	Output
	Reservation price	Float	€	JSON	Output
	Divisible	Boolean	-	JSON	Output
	Status	String	-	JSON	Output
	Recovery time	Date-time	-	JSON	Output
	Flexibility needs id	UUID	-	JSON	Output
Activation	Flexibility bid id	UUID	-	JSON	Service parameter

The full description of the service interface can be found on Appendix A and the publicly available swagger documentation for EdgeConnect can be accessed through this link.

4.1.4 Implementation Details

At this stage in the project, most of the features of this service have already been deployed and tested on a controlled environment, putting the TRL of the tool at level 6. Moreover, this tool is set to collect a TRL 7 with 2 ongoing field piloting activities. This comprises functionalities across the whole flexibility value chain from all stakeholders. Registering and onboarding all stakeholders, registering and pre-qualifying assets and services, registering needs, baselines and bids and relaying activation and settlement messages.

EdgeConnect will be implemented and tested on the Portuguese pilot of HEDGE-IoT and it embodies BUC 1 and their corresponding SUCs.

As Figure 10**Σφάλμα! Το αρχείο προέλευσης της αναφοράς δεν βρέθηκε.** shows, most of the c omponents are already implemented and ready to be deployed on HEDGE-IoT. We plan to advance the service TRL to level 7 by the end of the project by completing three main tasks:

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Develop a bilateral agreement functionality.
- Integrate most functionalities with the interoperability framework of the project.
- Demonstrate the service in a large full-scale pilot.

4.1.5 Integration with the HEDGE-IoT Interoperability Framework

Services as fully qualified applications (i.e., with their own private data sources, workflows, decision processes and interfaces) relate to a data space like data producers, data consumers, or both. Services adopted a data space compliant connector that will serve as an interface. With each service equipped with a data space connector, service to service interoperable data exchange is enabled. Other than becoming a data space compliant service, via the integration of the HEDGE-IoT project's open data connector, there are other optional components, available on the project's interoperability framework, that enhance the interoperability of services, such as the App Store and the Knowledge Engine.

EdgeConnect will be integrated with the Semantic capabilities enabled by WP4, most importantly to sponsor the discovery of services based on the individual metadata descriptions deriving from. the assets, location of those assets and the purpose of the services available.

4.1.6 Next Steps

The following actions are planned:

- Integration of the project's open data space connector.
- Development and tests of a bilateral agreement functionality.
- Integration with the semantic interoperability framework of the project for certain interactions.
- Field test of an Alpha version of EdgeConnect on the pre-demo phase of the project.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

4.2 FLEXIBILITY OPTIMIZATION SERVICE

4.2.1 Description of Service

The flexibility optimization service includes the following four key components:

1. **Demand and Production Forecasting:** Accurate short-term and long-term forecasts for energy demand and production using historical data, weather patterns, and real-time IoT sensor data from buildings or renewable sources to assess the flexibility potential of customers within a specified timeframe.
2. **Incentives Optimization and Bid Formulation**: Calculating optimal incentives to encourage consumer participation in flexibility actions and formulates an optimal bid (quantity, price) for the Local Flexibility Market.
3. **Consumer Interaction for Flexibility Requests**: Communicate flexibility requests to consumers and capture responses.
4. **Bid Submission and Overview**: Managing the submission of flexibility bids in the Local Flexibility Market (LFM).

4.2.1.1 Functional Requirements

1. The system should be able to use a variety of different data sources (consumption data, production data, weather and climate data).
2. The system should be able to provide day-ahead forecasts per hour and periodic forecasts (15 minutes, 30 minutes etc.).
3. The system should support multiple forecasting models for different short-term time horizons (e.g., weekly, monthly).
4. The system should be able to retrieve demand forecast for a specified time range.
5. The system should be able to retrieve production forecasts for a specified time range.
6. The system should be able to provide demand and production forecasts for specific customers as well as aggregate forecasts.
7. The system should determine the optimal incentives for consumer participation.
8. The system should be able to provide specific incentives in euros for all the different customer groups.
9. The system should be able to provide such incentives that it makes sense for the customers to participate.
10. The system should be able to provide such incentives that the aggregator has profit.
11. The system should consider the behavior of the customers and their response probability.
12. The system should calculate an aggregated bid in terms of quantity and price.
13. The system should handle real-time data from multiple consumers efficiently to quickly formulate incentives and bids.
14. The system should be able to send flexibility requests to consumers.
15. The system should capture consumer responses in real time.
16. The system should include customers' responses in the decision-making process.
17. The system should allow the aggregator to submit a bid with specific details including bid ID, quantity, price for a specific MTU.
18. The system should be able to retrieve the details of the bid using a unique bid ID.
19. The system should provide real-time updates on bid status, allowing users to check whether a bid is "pending," "approved," or "rejected".

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Connection with the WP3 goals:

Data Privacy Concerns
The service ensures data privacy through robust identity and access management mechanisms, including Attribute-Based Access Control (ABAC) and Role-Based Access Control (RBAC). These methodologies regulate access to sensitive data based on predefined attributes and roles, minimizing unauthorized access risks.

Additionally, data encryption techniques are applied both at rest and in transit to safeguard the integrity and confidentiality of information exchanged between the edge and cloud systems. This aligns with the project's commitment to secure and private AI/IoT interactions while maintaining compliance with relevant data protection standards.

Low-Maintenance Aspects of Methodology
The service is designed for minimal maintenance, a key concern outlined in Task 3.3. Features such as pre-trained models that require infrequent updates, hybrid AI approaches blending machine learning with domain knowledge, and the capability for on-device fine-tuning ensure sustained functionality with minimal intervention. The low-maintenance design aligns with the project's goal to create energy-efficient, scalable, and human-readable intelligent systems.

Cross-Sector Aspects of the Service
This service transcends energy applications by addressing multiple societal and environmental objectives, including:

- *User Comfort:* By optimizing energy usage without compromising user convenience, the service supports personalized energy management solutions, enhancing living and working environments.

- *Sustainability*: By enabling efficient use of Distributed Energy Resources (DERs) and enhancing renewable energy integration, the service contributes to reducing carbon footprints and supporting climate goals.

4.2.2 Innovative Aspects

Contributions to the State of the Art
This service pushes the boundaries of energy flexibility management by combining AI and IoT technologies with a modular edge-cloud architecture. Unlike the traditional centralized approaches, it harnesses distributed intelligence and ensures data security with advanced methods like Attribute-Based Access Control (ABAC), Role-Based Access Control (RBAC), and encryption. By blending domain expertise with machine learning, the system offers a hybrid optimization framework that delivers fast and accurate energy predictions while optimizing operations.

What truly sets this solution apart is its focus on the user. With features like tailored flexibility offers and real-time interactions through pop-up notifications, it adapts seamlessly to the unique needs of different users, creating a more personalized and efficient approach to energy management.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Added Value to the Pilot/ Impact of the results

Enhanced Decision-Making: By providing accurate and privacy-preserving energy forecasting and flexibility optimization, the service enables pilots to make data-driven, real-time decisions, improving operational efficiency and user satisfaction.

Energy Efficiency and Sustainability

The service significantly enhances the pilot's ability to manage Distributed Energy Resources (DERs) and integrate renewable energy, contributing to the overall energy efficiency and sustainability goals.

Social welfare

Enables the customers to participate in the LFM with flexibility actions leading to reduced energy bills.

4.2.3 Input and Output Data

In the table below the main input and output data for the service are presented:

TABLE 9 – INPUT AND OUTPUT DATA FOR THE FLEXIBILITY OPTIMIZATION SERVICE

Category of data	Variable name	Description	Data Format	Real or Synthetic
Input Data	Customer_id	The identifier of each customer	Numeric	Real
	Main_energy_consumption	The current residential energy consumption in Kwh	Numeric (Kwh)	Real
	mtu	Market Time Unit (15 mins)	Categorical	Real
	PV Production	The current PV production in kwh	Numeric (Kwh)	Real
	Battery_capacities	The total capacity of the batteries of the customers	Numeric (Kwh)	Real
	Battery_control_levels	The control level of the batteries from the aggregator	Numeric (%)	Real
	Weather_data	Weather conditions affecting energy demand and production	Numeric/ Categorical	Real
	Electricity_price	The cost of electricity in the day-ahead and intraday market	Numeric (€/kWh)	Real
	Incentive_Limits	Maximum and minimum incentives allowed per customer	Numeric (€/kWh)	Real
	Historical_energy_consumption	Historical energy consumption data for the forecasting task	Numeric (kWh)	Real
	Historical_energy_production	Historical energy production data from the PVs for the forecasting task	Numeric (kWh)	Real

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

	Acceptance_probabilities	The probability of each customer to accept a flexibility request	Numeric (%)	Synthetic
Output data	Forecasted_consumption	Timeseries of forecasted data for the main energy consumption and the energy consumption per device	Numeric (kWh)	Synthetic
	Forecasted_production	Timeseries of forecasted data for the energy production and the energy consumption per device	Numeric (kWh)	Synthetic
	Flexibility_offers	The flexibility offers to the customers (load shifting, incentive)	Numeric (kWh, €)	Synthetic
	Bid price	The optimal bid price for the aggregator to submit to the LFM	Numeric (€/kWh)	Synthetic
	Bid_quantity	The optimal bid quantity for the aggregator to submit to the LFM	Numeric (kWh)	Synthetic

4.2.4 Implementation Details

Current Status of Implementation: The service is currently in the design phase. Activities are focused on establishing the foundational architecture, refining the data from the IoT devices, initial testing and preparing for implementation.

4.2.5 Integration with the HEDGE-IoT Interoperability Framework

Currently N/A

4.2.6 Next Steps

By the next deliverable (M19), the focus will be on transitioning from design to early-stage implementation and validation, including:

- Prototypes: Development of initial prototypes focusing on the backend functionalities to implement the initial logic for generating flexibility offers based on customer acceptance probabilities and predefined incentive limits. Also, forecasting techniques will be optimized and validated with real pilot data to improve short-term energy consumption and production predictions across multiple time horizons.
- Testing Frameworks: Conducting basic tests to evaluate the initial prototype's functionality and performance using synthetic data to simulate real-world conditions.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

4.3 REAL-TIME RESERVE MARKET SIMULATOR

4.3.1 Description of Service

The Real-Time Reserve Market Simulator is an advanced cloud-based solution designed to emulate the processes of the manual Frequency Restoration Reserve (mFRR) and the automatic Frequency Restoration Reserve (aFRR) markets. The simulator provides stakeholders, such as Transmission System Operators (TSOs) and market participants, with the capability to test and analyse bid submissions for specific timeframes. By leveraging historical bid data retrieved from the ENTSO-E Transparency Platform, it enables performance evaluation in a realistic and reliable market environment. The service is built to facilitate better understanding and optimization of balancing market operations, ensuring alignment with European balancing market frameworks such as MARIi(Manually Activated Reserves Initiative) and PICASSOii (Platform for the International Coordination of Automated Frequency Restoration and Stable System Operation).

The simulator receives bid submissions in XML format from Balancing Service Provider (BSPs), validates them against the relevant market requirements, and simulates market operations using historical bids from the ENTSO-E Transparency Platform. If the BSP's bids are cleared during the simulation, the application generates and sends scheduling and activation signals back to the BSP. By enabling these functionalities, the simulator allows BSPs to evaluate their compliance with bid requirements (e.g., Full Activation Time (FAT)), the potential revenues and imbalances resulting from their market participation as well as the potential impacts of their bids on the market performance.

4.3.2 Innovative Aspects

The service provides residential and industrial energy communities with a user-friendly tool to explore and test bidding strategies, fostering greater awareness of their potential to actively participate in ancillary services markets. By simulating realistic market scenarios, the simulator enables users to optimize their energy consumption profiles, evaluate the economic value of their flexibility, and make informed decisions to maximize revenue opportunities. Additionally, the tool bridges the gap between large-scale market participants and smaller energy actors, promoting inclusivity in ancillary service markets and encouraging diverse stakeholder participation.

4.3.3 Input and Output Data

The Real-Time Reserve Market Simulator operates by processing a range of data inputs that enable comprehensive bid validation, market simulation, and performance evaluation. By integrating bid submissions with historical and system data, the simulator creates a realistic market environment where Balancing Service Providers (BSPs) can test their strategies and refine their participation approach:

- A bid document submitted by a Balancing Service Provider (BSP) intending to participate in a specific ancillary services market.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Historical market data from the ENTSO-E Transparency Platform, used for benchmarking and scenario simulation.
- System and market parameters, including reserve requirements, activation requests, market rules, and compliance criteria (e.g., Full Activation Time (FAT)).

A collection of JSON messages is delivered according to the internal tasks being completed. These outputs include:

- Validation results indicating whether the BSP's bid complies with relevant market requirements.
- Scheduling and activation signals sent to the BSP if their bid is cleared during the simulation.
- Simulation outcomes, including revenue estimates, imbalance metrics, and bid performance compared to historical data.

4.3.4 Implementation Details

The Real-Time Reserve Market Simulator is currently at TRL 5, which means that it has been validated in a relevant environment. By the end of the project, we aim to elevate it to TRL 7, by demonstrating its operational capability in a real environment. It is implemented as a modular and scalable cloud service, leveraging modern technologies to ensure efficiency and reliability. The application has been containerized with Docker software and has been deployed in AWS using Amazon Elastic Container Service (ECS). A collection of end-point URLs is available for authorized users to interact with the service.

In relation to the diagram below, the service also includes a scheduler to delay the actual submission of bids to the market to the corresponding Market Time Unit (MTU). Also, a database for the management of submissions and storage for the historical bids and activated power were implemented. An internal back-office allows the developers to maintain the service.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

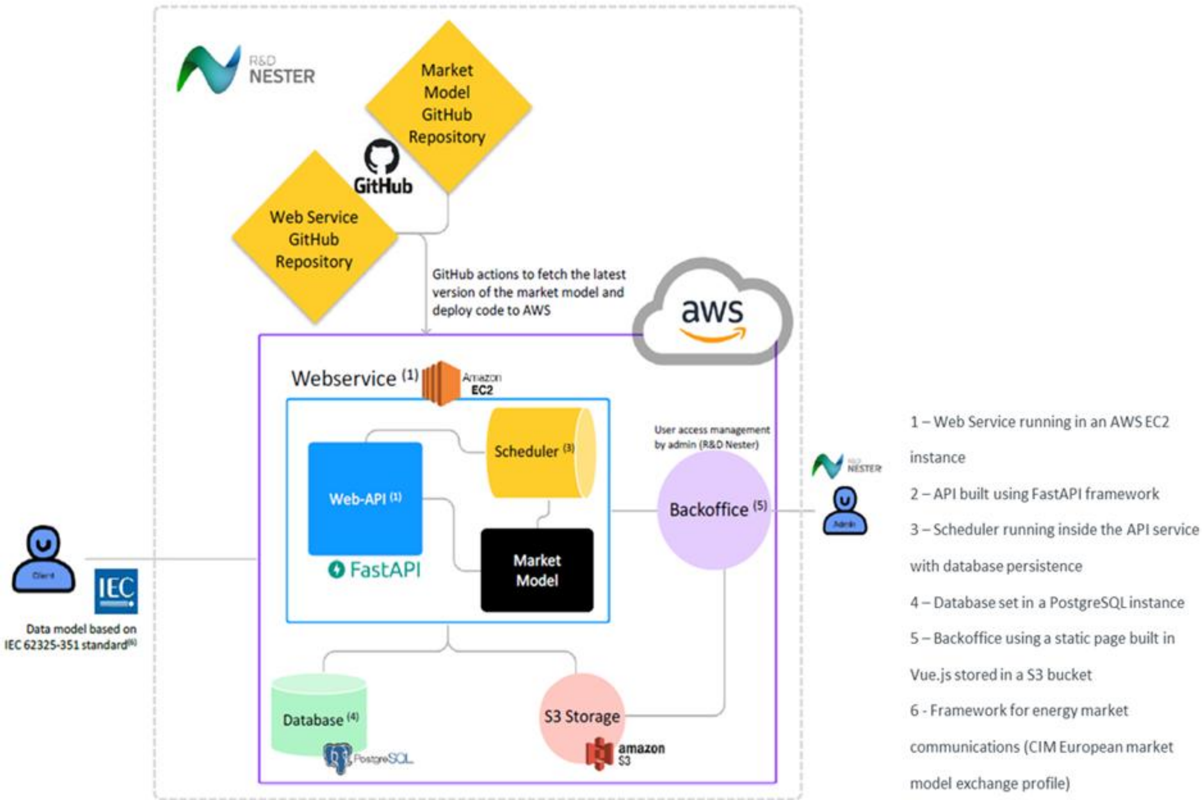

FIGURE 11 – REAL-TIME MARKET SIMULATOR ARCHITECTURE

4.3.5 Integration with the HEDGE-IoT Interoperability Framework

The service will be deployed on a server accessible through endpoints URLs using a Restful API. Those will be connected to a middleware called EdgeConnect (see section 4.1) which is a platform developed by another partner of the Portuguese Demo.

4.3.6 Next Steps

Continuous maintenance and developments on the service are being performed, namely, aFRR implementation, adoption of new available API connection with ENTSO-E transparency platform, and settlement process development. In addition, coordinated work inside the Portuguese Demo is carried on for a seamless integration of the various services and tools being developed.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

4.4 PREDICTIVE CONGESTION MANAGEMENT

The purpose of this service is to provide predictive Congestion Management (CM) for the distribution grid operator (DSO), enabling the proactive procurement of flexibility to mitigate potential congestion issues. This service complements the reactive real-time CM, creating a comprehensive CM approach that integrates both proactive and reactive capabilities. By leveraging these services together, DSOs can achieve a more robust and adaptive congestion management solution.

4.4.1 Description of Service

Figure 12 illustrates the architecture of the CM ecosystem, with a specific focus on predictive CM on the cloud. The predictive CM service is composed of several micro-services, including load and generation forecasting, state forecasting, and predictive CM. These micro-services follow a service-oriented design, enhancing the modularity and maintainability of the algorithms. Consequently, the cloud node requires significant computational power to execute these tasks efficiently. The predictive CM service on the cloud relies on three data sources: **(1)** weather data, retrieved via API requests, **(2)** static data stored in the system, such as grid data and historical load data, and **(3)** market data. Together, these data sources enable accurate forecasting and predictive analysis as well as market participation.

Overall, the predictive CM service allows the DSO to anticipate and prepare for potential grid congestion by utilizing either grid-side or market-based flexibility. This proactive approach enhances the resilience of the grid in addressing congestion challenges. Furthermore, it aligns with the objectives of improving grid flexibility, resilience, and observability.

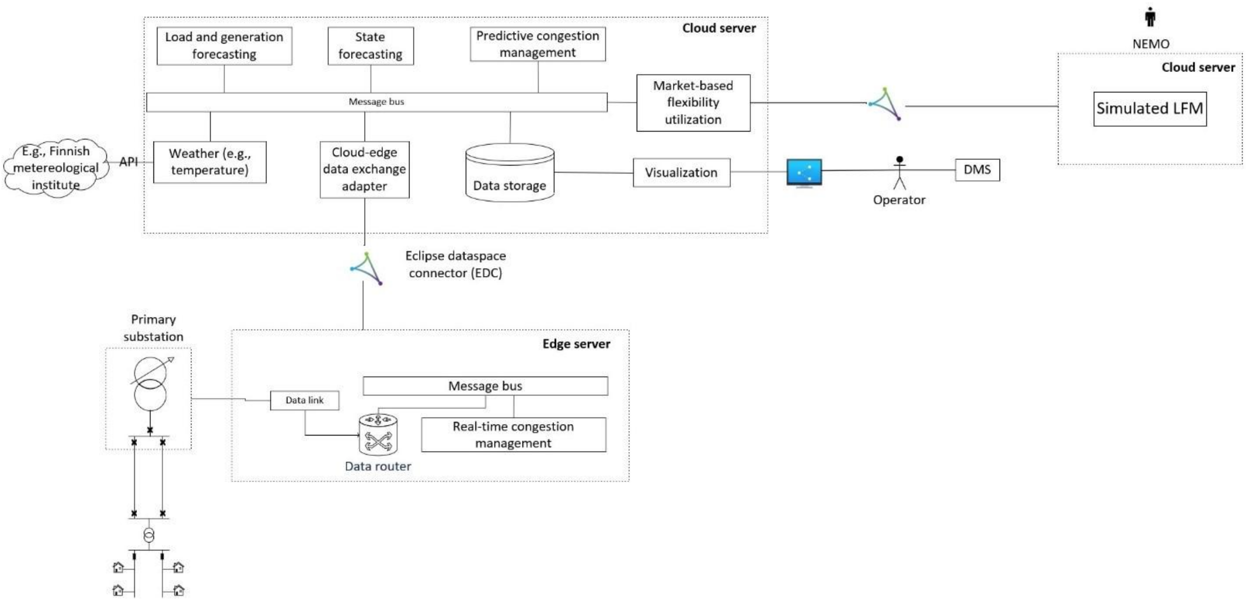

FIGURE 12 – PREDICTIVE CONGESTION MANAGEMENT SERVICE

4.4.2 Innovative Aspects

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Predictive CM is done by a limited number of DSOs in Europe, and they might utilize flexibility from local flexibility market, or plan switch statuses as a grid reconfiguration measure. Since predictive CM is at an early stage of adoption, this service in the pilot allows for further investigation of its potential and challenges.

The novelty of the service lies in its integration with real-time CM service in this pilot. This allows for harmonizing the predictive and real-time CM so that both market-based and grid-based CM solutions are utilized.

By designing the CM architecture to incorporate both predictive and real-time elements together, rather than separately, the pilot enhances their complementary strengths. Additionally, the distribution of predictive and real-time CM services across cloud and edge servers adds significant value. Traditionally, CM services are hosted either in the cloud or at the edge, but the cloud-edge continuum allows for optimal placement of algorithms, maximizing implementation efficiency and benefits. By connecting cloud systems with edge infrastructure, this approach unlocks new opportunities, such as enabling edge services to access data applications hosted in the cloud for interoperable services.

Extending this approach within the project and beyond could enable CM algorithms deployed on edge servers to be updated with verified and trustworthy applications from the cloud, ensuring continuous improvement and reliability of the service.

4.4.3 Input and Output Data

The service requires historical load and generation profiles that could be static in the data storage. Also, the service requires grid data (e.g., grid topology) stored in the data storage. Since the weather conditions affect the grid's consumption and production, therefore the service fetches weather data from a weather forecast service through an API request.

The service's output is also logged into data storage to be used for visualization. The visualized data is then used by the grid operator to make decisions according to the predictive CM recommendations. This means that the predictive CM service will not act automatically in the pilot when congestion is predicted, because that would require rigorous validation and testing before putting into use in the real grid operation. Another output of the service is a flexibility request, which is sent to the nominal electricity market operator (NEMO). That allows FSPs to access the grid needs and react to it accordingly.

4.4.4 Implementation Details

This service is currently in the design stage and implementation work has not yet started. Although the real-time and predictive CM services are designed together, the implementation of real-time CM precedes the predictive one. Nevertheless, the predictive CM algorithms will be simulated (not piloted) in the Finnish pilot. The service realizes SUC-FI-02.01 and SUC-FI-02.02, which is the second business use case of the Finnish pilot (BUC-FI-02). The work accomplished in the real-time CM could benefit the predictive service because there are commonalities between these two services.

4.4.5 Integration with the HEDGE-IoT Interoperability Framework

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The service will implement Eclipse Dataspace Connectors as its data connectors for edge-to-cloud and cloud-to-cloud data transfer. The Eclipse components are a comprehensive framework (concept, architecture, code, samples) providing a basic set of features (functional and non-functional) that dataspace implementations can re-use and customize by leveraging the framework's defined APIs and ensure interoperability by design. The HEDGE-IOT framework is not yet known at this point however, it is expected that utilizing Eclipse data space components could facilitate integration with HEDGE-IOT framework.

4.4.6 Next Steps

The next steps, which benefit both the real-time and predictive services, include:

- Input data preparation (M12-M21)
- Edge-cloud connection through Eclipse (M18-M30)
- Architectural design (M12-M21)
    - Location of micro-services (edge/cloud)
    - Data flow between micro-services

The next steps, which work towards this service specifically, include:

- Component development and testing (M18-M30)
    - Load and generation forecast
    - State forecast
    - Predictive CM
    - Market utilization
    - Visualization
- Integration testing and simulating (M30-M42)

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

4.5 ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE

4.5.1 Description of Service

The importance of Energy Communities (ECs) towards a clean energy transition and the decentralization of the grid is well known. Renewable Energy Communities (RECs), in particular, have a key role because they enable the sharing of resources between the community using local markets and allow support to decarbonization efforts. Another relevant aspect of ECs is their ability to contribute to external markets, like the flexibility and reserve ones, by pooling their resources together and leveraging their demand response capacity. The flexibility of the users in an EC allows them to dynamically respond to grid demands, even in near real-time as it is required to participate in frequency reserve markets (aFRR and mFRR).

However, coordinating all the different goals, members and assets of a community poses several challenges. Namely in the optimization of the available assets according to a specific objective, the participation in markets with strict rules, like the frequency reserve markets and the actual activation of residential household assets without inconveniences to the end user.

The service described in this section will focus on enabling an Energy Community to participate in a frequency reserve market. The service itself is integrated into a larger platform called **RECreation**. RECreation is a digital platform for the integration of the different energy services needed to manage a REC, including:

- **User Interface:** Front end with several REC configuration options and a dashboard to inform the REC manager and the REC members on the operation of the REC.
- **Data base:** with data regarding the REC structure definition, and economic and energy data such as the opportunity costs of the REC members, consumption and generation data, and settlement results.
- **Transactions:** to compute the local energy transaction and prices according to the selected model
- **Settlement:** computes the internal compensations derived from the energy transactions, and the energy allocation performed by the System Operator (SO) and the grid access tariffs charged to the REC for verification. From the settlement results, RECreation prepares a billing guide in Excel format with the information needed by an invoicing system to invoice the internal compensations among members for each settlement period (typically monthly).
- **Energy management:** computes the setpoints of the flexible resources (pre-delivery optimization) and can also compute the optimal transactions according to predefined criteria (post-delivery optimization).
- **Sizing:** for sizing new resources for the energy community.

The Figure 13 provides an overview of how all these energy services interact in the context of RECreation.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

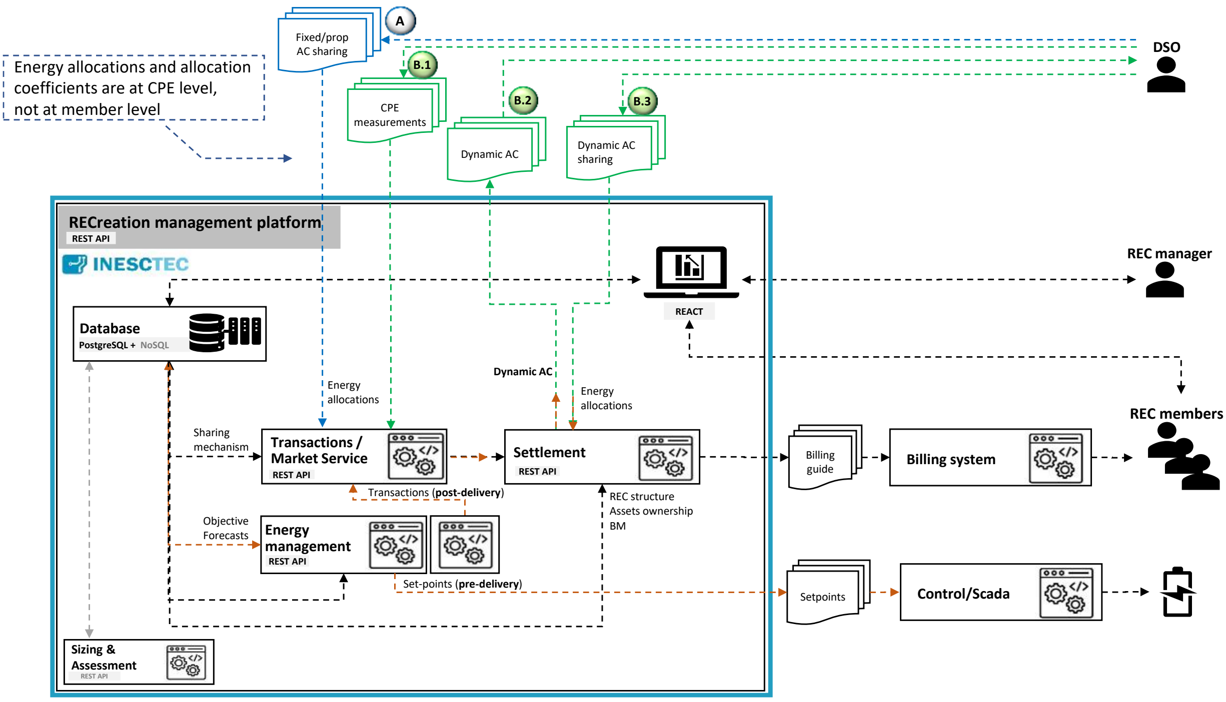

FIGURE 13 – RECREATION ARCHITECTURE

Dynamic energy sharing

This running mode considers the possibility of dynamic energy sharing and the use of the energy management module to dispatch flexible resources.

In the pre-delivery phase, the energy management module can compute the optimal setpoints of the flexible assets, for example, for the day-ahead.

In the post-delivery phase, the DSO shares with the REC manager the energy measurements of the members and a default energy allocation based on the proportional-to-consumption approach.

The transactions module computes the optimal transactions and prices according to some specific criteria, for example using a post-delivery pool market. With these transactions and prices, the settlement module computes the dynamic allocation coefficients corresponding to these transactions, that are shared with the System Operator (SO). The SO verifies the allocation coefficients and if no error is detected, it applies them to perform the dynamic energy allocation. Otherwise, the proportional-to-consumption allocation is used.

Dynamic energy sharing with flexibility provision

Energy communities can act as natural aggregators of their members, with the potential of providing explicit flexibility to third parties, such as grid operators. This running mode complements the dynamic energy sharing with the provision of flexibility to third parties. Providing flexibility implies additional steps in the pre-delivery and post-delivery optimization processes and settlement.

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

RECreation provides a baseline for flexibility verification, which reflects the REC's optimal operation (pre-delivery) if no flexibility is requested. When the hours and values of the flexibility needs are known, an iterative optimization procedure (pre-delivery) determines the optimal flexibility to be offered based for a set of expected flexibility prices, resulting in the construction of a flexibility bidding curve to address flexibility prices uncertainty.

Once the final flexibility requested is activated, the aggregator determines the final optimal setpoints (pre-delivery) for the flexibility assets to ensure the optimal delivery of the requested flexibility.

Following flexibility delivery (post-delivery), the local energy transactions, reflecting the energy shared among REC members while accounting for the flexibility provided, must be calculated. This enables the determination of allocation coefficients sent to the DSO to allocate energy and inform energy suppliers of the final energy supplied, reduced by the energy self-consumed.

Dynamic energy sharing with flexibility provision for frequency restoration reserve

Still acting as natural aggregators, energy communities can provide flexibility to participate in Balancing Service Markets (BSMs). This running mode follows a similar logic of the previous one but adding the restrictions of participating in BSMs. The EC can participate by provisioning manual Frequency Restoration Reserve (mFRR) or automatic Frequency Restoration Reserve (aFRR).

Since provisioning aFRR and mFRR has stricter rules than provisioning flexibility, namely in terms of activation timing, when participating in them RECreation focuses solely on them, rather than trying to coordinate different running modes at the same time.

For mFRR the same principle of pre and post-delivery will be applied, simply in shorter periods of time. The needs are received from the TSO directly and the baseline and optimal flexibility are sent to the TSO. The activations will be received by the TSO and RECreation will send them to an external platform that will enforce the setpoints directly in the assets. The settlement will be received and distributed along the members of the EC.

For aFRR the rules for provisioning are even stricter. RECreation will allow the TSO to directly have control over specific types of assets that work well with such strict rules, such as batteries and EV chargers.

Relation with EdgeConnect

EdgeConnect (described in section 4.1) is a service that allows stakeholders all across the flexibility value chain to have access to an integrated ecosystem that enables all the steps across the value chain, in order to unlock its full flexibility potential.

This platform will have a strong impact on RECreation, since it will allow other stakeholders on the flexibility value chain, such as the TSO, to easily discover and interact with it. This integration will also enable the integration with HEDGE-IoT's interoperability framework, namely by being equipped with data space connectors.

4.5.2 Innovative Aspects

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The main innovative aspect of this service lies in enabling an energy community to provision frequency restoration reserve and participate in the balancing reserve market. The service will be a part of a larger platform (RECreation), so its inclusion in that ecosystem will also be an improvement over previous versions of it, namely from the previous version which contains an initial version of the DSO flexibility provision mode and is currently being tested in the H2020 BeFlexible project.

Provisioning both mFRR and aFRR is a clear improvement over the latest versions of the platform and will require work in terms of scalability and process optimization, to comply with all the rules established by the TSO. Therefore, optimizing the forecasting and energy management processes will provide significant contributions.

4.5.3 Input and Output Data

Table 10 describes the types of data shared by the energy community management platform (RECreation) that are relevant to the Frequency Restoration Reserve service and its provisioning.

TABLE 10 – INPUT AND OUTPUT DATA FOR THE ENERGY COMMUNITY MANAGEMENT PLATFORM

Type	Variable	Description	Data Format
Input	Energy curve forecast	Energy curve forecast of EC members	JSON
Input	Assets	Pre-qualified assets registered for each member	JSON
Input	Members	Members registered by EC manager	JSON
Input	Activation	Activation setpoints received by the SO	JSON
Input	Settlement	Settlement values received by the SO	JSON
Output	Baseline	Flexibility baseline for the EC	JSON
Output	Flexibility bid	Flexibility bid of the EC offered to market	JSON
Output	Metering for settlement	Metering data from the DSO used for settlement purposes	JSON

4.5.4 Implementation Details

Currently, the platform is entering TRL 6, since it is being tested at the same time in a pilot from another European project. However, the service dedicated to the frequency restoration reserve is at TRL 2, since its conceptualization and development started in the HEDGE-IoT project and will be tested alongside the RECreation platform.

By M19, we plan to have the service up to TRL 4 by validating the technology in a lab environment with real IoT data and synthetic market data.

By M30, the service will be at TRL 7 and ready to be demonstrated in the full-scale pilot. Besides using real IoT data from EC members in Portugal, it will also be connected to the TSO via the real time market simulator (section 4.3). The TRL of the energy community management platform itself will accompany that of the frequency restoration reserve service.

This service is integrated into the Portuguese pilot directly in BUC 01 and it enables SUC 01, 02 and 03.

4.5.5 Integration with the HEDGE-IoT Interoperability Framework

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The RECreation platform will use the Eclipse data space connector, as per the guidelines of the HEDGE-IoT project. Mainly, the data exchange will be done with EdgeConnect (described in section 4.1) which works a flexibility value chain enabler, rather than communicating directly with the connectors of other stakeholders.

The adoption of this data space connector will ensure seamless integration with the HEDGE-IoT ecosystem and enforce interoperability of the data being exchanged. Furthermore, the service will be catalogued and specified in the project's App Store, to further increase its interoperability.

4.5.6 Next Steps

The next steps are expected to follow the timeline below:

- M15 → Integrate RECreation with the data sources from the HEDGE-IoT pilot.
- M19 → Develop and integrate the mFRR service into RECreation
    - M22 → Test the service with real pilot user's data.
- M25 → Integrate the platform with the real time market simulator from the TSO
    - M30 → Test it with real data from a select group of pilot users.
- M27 → Integrate the Eclipse data connector.
- M32 → Scale service and platform to accommodate the needs of the full-scale pilot (30 residential users in an energy community).

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

5 AI/ML TOOLS AND SERVICES GOVERNING THE COORDINATED OPERATION AND PLANNING OF INFRASTRUCTURE AT THE CLOUD LEVEL

5.1 INTRODUCTION

The design of AI/ML algorithms, tools and services are based on stakeholder requirements and system specifications defined in WP2, service specifications of Task 3.2 and specific needs of demonstrations in WP5. Cloud level algorithms, tools and services focus on stakeholder coordination, information exchange, analysis capabilities, flexibility trading, etc. to **1)** enable efficient integration of RES, EVs, heat pumps, etc. as a part of electricity system and different markets, **2)** utilize services like forecasting, information exchange and analytics to optimize their operation and to facilitate technical management of resources and services, and **3)** coordinate decision-making between stakeholders to realize synergy benefits and to avoid conflicts. The design includes specification of core micro-services, common for several functionalities, and how those are utilized to build complete service to register, discover, orchestrate and authorize services. The second phase of the task is to develop the AI/ML algorithms, tools and services needed in the demonstrations.

5.2 INTEGRATION WITH HEDGE-IOT FRAMEWORK

The design of cloud services and afterwards their development, requires constant compatibility check with the HEDGE-IOT framework so that integration task can be done smoothly making the task important from the project point of view. The services designed by each partner of this task would be realized through several microservices and therefore, during microservice design stage, interoperability and scalability of the solutions should be emphasized. From the infrastructure design perspective, the integration with the HEDGE-IOT framework would be reinforced by:

- Microservice based integration for loose coupling and late binding of system components, enabling evolving the system and efficiently orchestrating configurations in different settings.
- Standardized data connectors (e.g., Eclipse data connector) and agreed data architecture practices for improved interoperability and scalability of the solutions, including eventually cloud to edge expansion using the same practices and control plane solutions.

5.3 NEXT STEPS

Action plan until the next iteration of this deliverable for the T3.4 is as follows:

- Specification and design (M12-M18)
    - BUCs and SUCs become implementation use cases (specifically what information is exchanged, and how, following the smart grid architecture model (SGAM) architecture).
    - Design the integration model, select service framework tools and components for microservice based integration (e.g. Eclipse Arrowhead framework and similar).

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Select, design and evaluate suitable data connector implementations (e.g. Eclipse Dataspace Connector).
- Implementation of specified services and tools (M14-M18)
    - We select some of the tools and services to be first implemented and check it with demo milestones.
- Infrastructure testing (M16-M18)
    - Having the server available and set up the virtual machines, docker and the data space connector.
- Deliverable writing (M18-M19)
- Draft timeline preparation for the next deliverable (M19)
    - Implementation and testing oriented.

6 COMPUTATIONAL ORCHESTRATION TO ENSURE THE CLOUD-EDGE CONTINUUM

In this section, a first overview of the computational orchestration framework to ensure the cloud/edge continuum is provided, by detailing its architecture and how it will integrate with HEDGE-IoT's digital and interoperability frameworks. A short summary is provided below:

Cloud/Edge continuum computational orchestration (TUC): To ensure a streamline, homogenous and efficient cloud-to-edge computational effort, a swarm-based computation orchestration framework is developed in the project and its first specification is provided in this document. This framework has two main goals: 1) edge offloading for low-latency data processing for energy cloud services and 2) to orchestrate federated learning and distributed computing processes across the edge-fog-cloud continuum. Built on KubeEdge, the framework extends Kubernetes capabilities to edge environments, incorporating swarm-based heuristics to optimize resource allocation. Regarding data privacy, it integrates blockchain for secure and transparent service management, taking advantage of smart contracts and tokens for traceability. A monitoring system using Kube Prometheus will be established, supporting real-time infrastructure insights. It integrates HEDGE-IoT's interoperability framework by 1) considering the set of data-driven cloud services available in the project's App Store, 2) by using semantic annotation to expose edge devices computation capabilities and 3) by adopting the project's data space connector to perform interoperable data exchanges with the edge devices that also adopt it.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

6.1 DESCRIPTION OF THE ORCHESTRATION FRAMEWORK

The computational orchestration platform for smart grids provides two main roles, namely:

- Enabling the edge offloading of energy cloud services to enable improved data processing and reducing latency at the network's edge.
- Orchestration of federated learning processes across the edge-fog-cloud continuum facilitating the efficient delivery of AI-driven applications by incorporating hyper-parameter optimization and tailored model aggregation to enhance the performance and reduce computational and data network overhead.

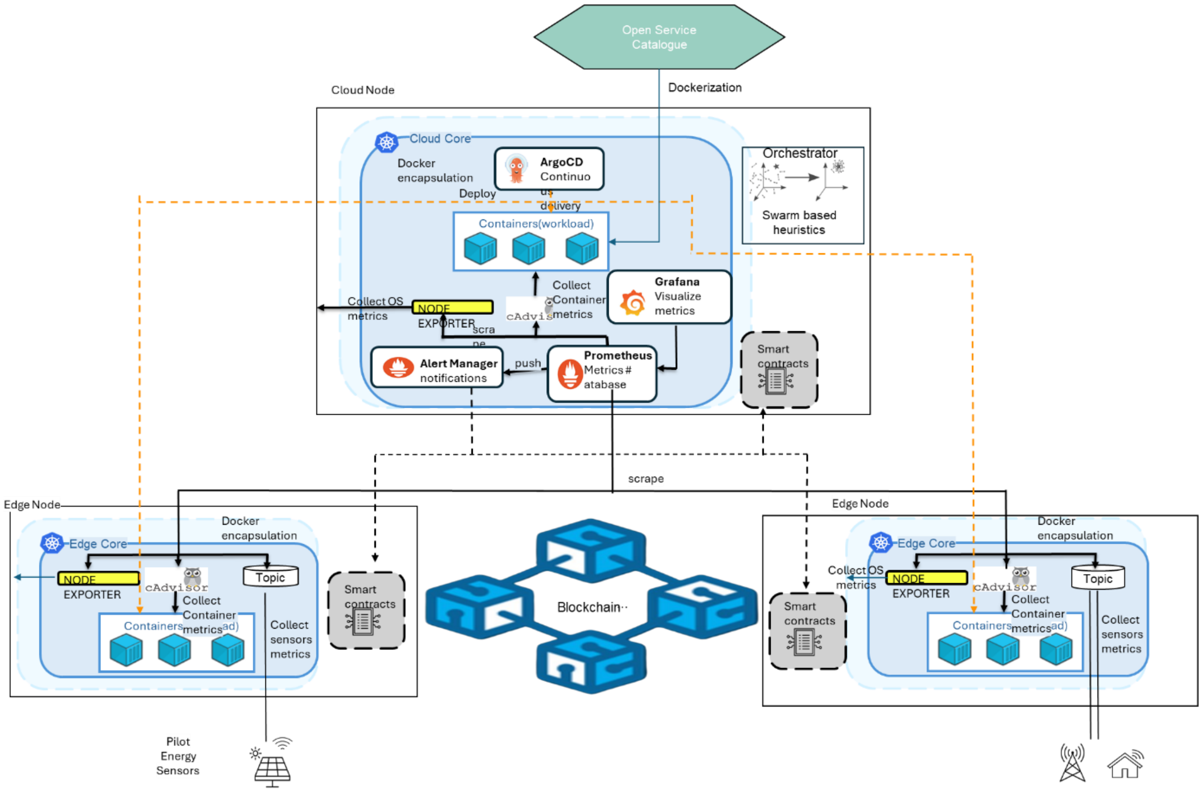

FIGURE 14 – ORCHESTRATOR ARCHITECTURE FOR EDGE OFFLOADING

The Figure 14 above shows the overall architecture of the **computational orchestration platform** that leverages on KubeEdge as cluster technology. KubeEdge is an open-source edge computing platform that extends the classic Kubernetes to the edge. To better support the unique requirements of Hedge-IoT services, the default Kubernetes Scheduler will be modified. These modifications will consider and trade-off service-specific requirements, such as network capabilities and computational demands, as defined by the Open Service Catalog with node available resources, latency and data availability. Additionally, a tailored swarm-based heuristic will be integrated to address optimization challenges related to offloading services to the edge.

The integration of blockchain technology into the platform offers a transparent and secure mechanism for managing resource allocation and services execution. Each dockerized service is uniquely identified and traceable within the blockchain network. Smart contracts are used to register service migrations or deployments, while a specific token acts as a globally unique identifier for the service, providing an immutable reference within the blockchain.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The orchestrator's computational monitoring system is built on Kube Prometheus, that offers a collection of tools designed to provide insights into the performance and health of the infrastructure's components.

- Node Exporter: deployed on each node, collects and exposes data about node resources, including metrics related to CPU, memory, disk and network usage.
- cAdvisor: deployed on each node, collects and exposes data about containers (workload) resources, such as CPU, memory and network usage. Ensures that workload receives adequate resources during execution, helping in maintaining the SLA compliance and triggering alerts when resource usage deviates from expectations.
- Prometheus: is the central monitoring component, deployed on the cloud side. It serves as both metric aggregator and a database to store time-series data. It scrapes metrics exposed by both types of nodes.
- Alerting Service: configured to define thresholds and conditions for system alerts. When certain metrics/queries, defined, exceed, fall or meet predefined thresholds, the alerting system is triggered to send notifications.
- Grafana: acts as the visualization layer. It connects to Prometheus and pulls metrics to generate visual representation of the data, making easier to monitor the system evolution over time.

The **orchestration of federated learning processes** deals with the optimization of the models' aggregation and hyper-parameters optimization (see Figure 15). Therefore, the Orchestrator will consider some federated learning processes and optimize them by leveraging on heuristics.

It employs heuristic methods to explore hyperparameter configurations, identifying the most promising candidates and sending them to validation nodes for evaluation. This approach ensures a streamlined process that maximizes performance while minimizing computational overhead. The orchestrator handles the scalability challenges of hyperparameter tuning in environments with massive datasets and diverse federated deployments. It effectively addresses the limitations of edge devices by considering their constrained computational resources during the tuning process.

To further improve efficiency, the Orchestrator incorporates a hierarchical optimization strategy, which aggregates updates at the fog and cloud layers. This method prioritizes and personalizes the updates from the best-performing models, improving overall training efficiency and accelerating convergence by focusing on the most impactful early updates.

In this way it will reduce the computational effort on edge nodes by processing only one hyperparameter configuration at a time for training and validation. Additionally, it may improve the federated learning processes by optimizing communication between edge devices and the fog or cloud layers.

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

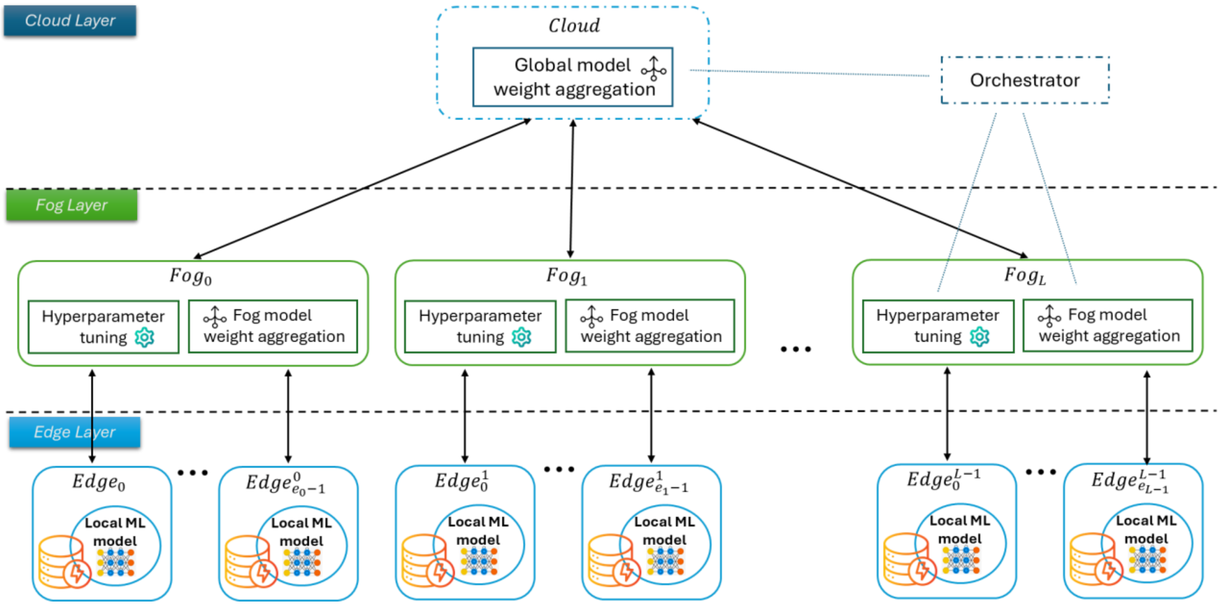

FIGURE 15 – ORCHESTRATION OF FEDERATED LEARNING METHODS

6.2 INTEGRATION WITH HEDGE-IOT FRAMEWORK

The edge offloading orchestration considers a set of the data driven cloud services (i.e. Available in Open Service Catalogue) to be offloaded in edge related environments considering pilot computational requirements, while the orchestration of federated learning processes will consider some of the federated methods specified in section 2. The current TRL level of the component is 5 with a planned TRL jump to 6 by the end of the project.

6.3 INTEGRATION WITH THE HEDGE-IOT INTEROPERABILITY FRAMEWORK

The orchestrator will consider the open data connectors for edge devices as well as the semantic interoperability annotation exposing edge devices computational capabilities and data availability, information that will facilitate the optimization process.

6.4 NEXT STEPS

Action plan until the next iteration of this deliverable (D3.4 à M19)

- Integration of model aggregation heuristics with federated machine learning methods defined in task 3.3 considering the specific local models available in pilots.
- Integration with Open Service Catalogue for gathering the specifications of the services to be offloaded.
- Implementation of swam heuristics for edge offloading and integration with the computational platform.
- Adaptation of smart contracts and tokens to workload tasks and blockchain integration with the computational platform

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Further development of the computational platform towards envisioned TRL scaling and allow for new components integration

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

7 CONCLUSIONS

Two of the main objectives of the HEDGE-IoT project are to enhance the intelligence across the energy system (at the edge and cloud levels) and to bridge the cloud/edge continuum in a secure, distributed and interoperable way. The first specifications of the services provided in this deliverable are a crucial step to achieve these goals. The current services ecosystem addresses critical challenges of the project, such as flexibility optimization, grid resilience, federated learning in the energy sector, energy communities' operations, RES optimization and TSO/DSO functions, while addressing the specific business and system use cases defined in WP2. However, it also focuses on some cross-sector challenges, such as user comfort and sustainability. The cloud/edge continuum challenge is addressed by providing the first specification to a swarm-based orchestrator that will enable edge nodes offloading and the orchestration of federated learning processes.

Interoperability is one of the pillars that will enable HEDGE-IoT's multi-dimensional framework, by leveraging cutting-edge interoperable architectures like data spaces and semantic interoperability. In this deliverable, services specify how they plan to be integrated with the project's interoperability framework defined by WP4, namely by adopting a data space connector, but also by adding semantic meaning to their data.

The following KPIs for the services were defined in the project's objectives for WP3:

- 6 open applications for residential users for energy and non-energy sector. **In D3.3: 4 energy sector services for residential users were described and 2 of them have non-energy components.**
- 4 open applications for grid operators. **In D3.3: 8 services whose main stakeholders are grid operators (TSOs and DSOs) are described.**
- A common App Repository for the HEDGE-IoT ecosystem. **In D3.3: The common app repository is not mentioned, since its specifications are still being discussed. However, the services will be containerized and fully prepared to be part of a registry.**

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

[1] C. Gonçalves, R. Bessa e P. Pinson, "A critical overview of privacy-preserving approaches for collaborative forecasting," *International Journal of Forecasting,* pp. 322-342, 2021.

[2] C. Gonçalves, R. Bessa e P. Pinson, "Privacy-preserving distributed learning for renewable energy forecasting," *IEEE Transactions on Sustainable Energy,* pp. 1777-1787, 2021.

[3] F. Sattler, K.-R. Muller e W. Samek, "Clustered Federated Learning: Model-Agnostic Distributed Multi-Task Optimization under Privacy Constraints," *CoRR,* 2019.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

APPENDIX A

Endpoint 1: Create a qualification request

URL: /api/v1/qualification

Method: POST

Description: Creates a qualification for a given requester.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Request Body:

```
{
        "requesterId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
        "objectToQualifyId": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
    }
```

Response Examples:

201 Created:

{ "message": "qualification created" }

400 Bad Request:

{ "message": "invalid data format" }

401 Unauthorized:

{ "message": "Unauthorized access" }

Endpoint 2: Retrieve a qualification

URL: /api/v1/qualification/{qualificationId}

Method: GET

Description: Retrieves a qualification by its id.

Headers:

- Content-Type: application/json

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Authorization: Bearer <token>

Path:

- qualificationId: string UUID

Response Examples:

200 OK:

{
"requesterId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"objectToQualifyId": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
“qualified”: “True”
}

401 Unauthorized:

{ "message": "Unauthorized access" }

404 Not Found:

{ "message": "Not Found" }

Endpoint 3: Create flexibility needs

URL: /api/v1/flexibility/needs

Method: POST

Description: Stores information about the flexibility needs of a DSO/TSO

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Request Body:

{
"contractDuration": 300,
"contractId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"type": "longFlexibilityNeeds",

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

"startDatetime": "2024-11-08T17:48:07.220Z",
"endDatetime": "2024-11-08T17:48:07.220Z",
"gateOpenTimestamp": "2024-11-08T17:48:07.220Z",
"gateCloseTimestamp": "2024-11-08T17:48:07.220Z",
"productId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"requesterId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"flexibilityZoneId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"volume": 500,
"maxReservationPrice": 300.5,
"maxActivationPrice": 1000.65
}

Response Examples:

201 Created:

{ "message": "flexibility need created" }

400 Bad Request:

{ "message": "invalid data format" }

401 Unauthorized:

{ "message": "Unauthorized access" }

Endpoint 4: Retrieve a flexibility need

URL: /api/v1/flexibility/needs/{flexibilityNeedsId}

Method: GET

Description: Retrieves a flexibility needs request by its id.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Path:

- flexibilityNeedsId: string UUID

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Response Examples:

200 OK:

{
"id": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"contractDuration": 300,
"status": "AGGREGATION",
"type": "longFlexibilityResponse",
"startDatetime": "2022-12-30T20:45:00Z",
"endDatetime": "2022-12-30T20:45:00Z",
"gateOpenTimestamp": "2024-11-21T12:39:14.162Z",
"gateCloseTimestamp": "2024-11-21T12:39:14.162Z",
"productId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"requesterId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"dsoName": "INESCTEC",
"flexibilityZoneId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"flexibilityZoneName": "Barcelos",
"flexibilityZoneState": "Braga"
}

401 Unauthorized:

{ "message": "Unauthorized access" }

404 Not Found:

{ "message": "Not Found" }

Endpoint 5: Create flexibility bid

URL: /api/v1/flexibility/bids

Method: POST

Description: Creates a new flexibility bid.

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Request Body:

{
"minQuantity": 2,
"maxQuantity": 4,
"startDatetime": "2024-11-08T17:50:18.482Z",
"endDatetime": "2024-11-08T17:50:18.482Z",
"maxNumberActivations": 4,
"activationPrice": 55.5,
"reservationPrice": 65.5,
"divisible": true,
"recoveryTime": 60,
"businessPartnerId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"flexibilityZoneId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"serviceId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
"flexibilityNeedsId": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
}

Response Examples:

201 Created:

{ "message": "flexibility bid created" }

400 Bad Request:

{ "message": "invalid data format" }

401 Unauthorized:

{ "message": "Unauthorized access" }

Endpoint 6: Retrieve a flexibility bid

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

URL: /api/v1/flexibility/bids/{flexibilityBidId}

Method: GET

Description: Retrieves a flexibility needs request by its id.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Path:

- flexibilityBidId: string UUID

Response Examples:

200 OK:


```
{
    "id": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
    "minQuantity": 2,
    "maxQuantity": 4,
    "startDatetime": "2024-11-21T12:41:44.358Z",
    "endDatetime": "2024-11-21T12:41:44.358Z",
    "maxNumberActivations": 4,
    "activationPrice": 55.5,
    "reservationPrice": 65.5,
    "divisible": true,
    "recoveryTime": 60,
    "status": "SELECTED",
    "createdDate": "2024-11-21T12:41:44.358Z",
    "businessPartnerId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
    "flexibilityZoneId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
    "serviceId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
    "flexibilityNeedsId": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
}
```


Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

401 Unauthorized:

{ "message": "Unauthorized access" }

404 Not Found:

{ "message": "Not Found" }

Endpoint 7: Create baseline

URL: /api/v1/flexibility/baselines

Method: POST

Description: Creates a baseline data entry.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Request Body:

```
{
        "startDatetime": "2022-10-30T12:45:00Z",
        "endDatetime": "2022-10-30T20:45:00Z",
        "type": "baseline",
        "periodNumber": 60,
        "serviceId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
        "assetIds": [
                "3fa85f64-5717-4562-b3fc-2c963f66afa6"
        ],
        "podId": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
        "intervalDataEntries": [
                {
                        "datetime": "2022-12-30T20:45:00Z",
                        "units": "w",
```


This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

"dataValue": "600",

"variableName": "power"

}

]

}

Response Examples:

201 Created:

{ "message": "flexibility baseline created" }

400 Bad Request:

{ "message": "invalid data format" }

401 Unauthorized:

{ "message": "Unauthorized access" }

Endpoint 8: Select a flexibility bid

URL: /api/v1/flexibility/selection

Method: PATCH

Description: Select a list of flexibility bids.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Request Body:

[

{

"flexibilityBidId": "3fa85f64-5717-4562-b3fc-2c963f66afa6"

}

]

Response Examples:

204 No Content:

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

401 Unauthorized:

{ "message": "Unauthorized access" }

404 Not Found:

{ "message": "Not Found" }

Endpoint 9: Flexibility bid activation

URL: /api/v1/flexibility/activation/flexibilityBidId}

Method: PATCH

Description: Activate a flexibility bid by its id.

Headers:

- Content-Type: application/json
- Authorization: Bearer <token>

Path:

- flexibilityBidId: string UUID

Response Examples:

204 No Content

400 Bad Request:

{ "message": "invalid data format" }

401 Unauthorized:

{ "message": "Unauthorized access" }

Co-funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

HEDGE-IoT