



HEDGE-IoT

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

D3.4 HEDGE-IoT Technological Enablers (Intermediate Release)

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.	D3.4
Deliverable Title	HEDGE-IoT Technological Enablers (Intermediate 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	INESC ED ICCS UNIZG JSI VUA APIO VTT TAU NESTER

	HENEX ABB ELERG SONAE KONC TUC RWTH TNO CEL ETB HOPS ENERV PPC HEDNO HSE QUE ARETI AB AE DST							
Due date	31/07/2025							
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							

DOCUMENT REVISION HISTORY

Version	Date	Description of change	List of contributor(s)
V0.1	11/03/2025	First version – template	INESC
V0.2	04/04/2025	Executive summary table	All authors
V0.3	30/05/2025	Technology enablers template	INESC
V0.4	20/06/2025	First round of technology enablers inputs	All authors
V0.5	16/07/2025	INESC’s review and second round of inputs	All authors
V0.6	21/07/2025	Consolidation of draft	INESC
V0.7	23/07/2025	Consolidation and final touches	INESC
V0.8	29/07/2025	Quality review	ICCS, ED
V0.9	30/07/2025	Draft for submission	INESC
V1.0	31/07/2025	Ready for submission	ED

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 ENGENHARIA DE 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

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.	VWIC	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

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 SRLS	CSL	IT

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

© HEDGE-IoT, 2025

This deliverable and its content are the property of the HEDGE-IoT Consortium. This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation, or both. Reproduction is authorised provided the source is acknowledged. The content of all or parts of these documents can be used and distributed provided that the HEDGE-IoT project and the document are properly referenced.

EXECUTIVE SUMMARY

HEDGE-IoT (Holistic Approach towards Empowerment of the Digitalization of the Energy Ecosystem through adoption of IoT solutions) 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.

This deliverable presents the intermediate release of the Technological Enablers (TEs) of HEDGE-IoT, which are designed based on requirements collected in WP2, integrated with the interoperability tools from WP4 and deployed in the pilots of WP5. These TEs establish the project’s digital multi-level framework, which creates an ecosystem that increases the resilience of the grid, creates new market opportunities and promotes advances in IoT standardization.

In scope of WP3, TEs address two objectives of the project, namely: to **leverage IoT solutions to add intelligence to the edge/cloud layers establishing an edge/cloud continuum** and to **design AI/ML tools for increased flexibility, resilience and observability**. These objectives are addressed by the designs, specifications and implementations of TEs in this deliverable, which are detailed in Table 1 - HEDGE-IoT’s Federated Learning Technology Enablers, Table 2 - HEDGE-IoT’s Data-Driven Edge-to-Cloud Technology Enablers, Table 3 - HEDGE-IoT’s Cloud Technology Enablers and Table 4 - HEDGE-IoT’s Computational Orchestration Framework. A summary and maturity overview follows.

Summary: 2 federated learning technology enablers targeting residential end-users

TABLE 1 – HEDGE-IOT’S FEDERATED LEARNING TECHNOLOGY ENABLERS

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Federated Learning for Energy Forecasting (ICCS)	Decentralized federated learning architecture using LSTM/BiLSTM models for forecasting energy demand and production. Initial training conducted using open datasets such as StoreNet. Model updates exchanged via MQTT secured by TLS.	BiLSTM models, trained with pilot data from residential apartments. Models converted to TensorFlow Lite for deployment on Shelly 3EM meters. MinIO used for model storage and secure model distribution. Integrated secure MQTT-based communication.	Pilot deployment and real-world testing in 100 Greek apartments. Validation of real-time performance on edge devices. Iterative model refinement based on pilot results. Energy footprint monitoring for edge model inference.

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Vector Autoregressive Model for Energy Time Series Forecasting (INESC)	<p>Proof of concept algorithm observed with synthetic data.</p> <p>Data processing pipeline developed and connected to real pilot data, for a select group of users.</p>	<p>First lab concept solution in a fully functioning demo environment with 4 different simulated edge agents.</p> <p>Design an initial development of integration of pilot users' data pipeline for each agent.</p>	<p>Integrate the lab prototype into a larger scale environment according to the number of users participating in the pilot at the given date.</p> <p>Integrate the edge nodes with the computational orchestrator tool being developed in WP3.</p>

Summary: 7 data-driven edge-to-cloud technology enablers addressing energy and non-energy use cases like predictive maintenance, congestion management, home energy management and building occupant's comfort. Furthermore, exploiting 5 different types of IoT/edge devices to leverage them for edge intelligence.

TABLE 2 – HEDGE-IOT'S DATA-DRIVEN EDGE-TO-CLOUD TECHNOLOGY ENABLERS

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Enhanced Network Management and Planning (UNIZG)	Investigation of techniques used in anomaly detection, demand forecast and initial development and testing of algorithms.	<p>Testing anomaly detection and demand forecast algorithms on synthetic data.</p> <p>Research on techniques to detect PVs in secondary distribution networks from limited measurements.</p>	Testing anomaly detection and demand forecast algorithms on realistic data and PV detection algorithm on synthetic data.
DTR-DLR on the Edge (JSI)	Initial overview of the DLR/DTR and weather forecast algorithms.	DLR/DTR algorithms adapted to IoT use case, improved interoperability plan.	Improved overall robustness of DLR/DTR algorithms, weather forecast algorithms deployed and updated, improved interoperability.
Anomaly Detection and Predictive Maintenance on the Grid (VU)	Learn patterns of nominal device behaviour from SAREFised data streams.	Recognize and report unnatural deviations from learned patterns (via SIF).	Improve overall robustness and minimize false error rate; add admin UI and export function.
Anomaly Detection and Fault Forecasting to Increase Distribution Network	Deep Learning Model to Analyze data patterns, correlations, and identify anomalies.	HLSTM model to analyze data patterns, correlations, identify anomalies, and forecast future anomalies.	Improved identification of data patterns and anomaly detection with high accuracy and forecasting of future anomalies.

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Resilience (VTT)			
Real-Time Congestion Management (TAU)	Breaking down the service into micro services to enhance modularity of algorithms and to facilitate coordination and cooperation of developers. Specifying the steps in which the implementations could be realized. Those steps are data preparation and harmonization, 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.	Micro services and the hardware where they are going to be executed have been decided. Initial agreement on the timeline of microservice developments done by different organizations aimed at piloting the whole solution. Harmonizing Input data has been completed. (validating the input data is expected to be completed by the end of April). It is expected to have the first version of state estimation micro service necessary for real-time congestion management ready for this deliverable. The first version of the test environment where the developed micro services will be tested is also expected to be ready.	Real-time CM is also developed. State estimation and real-time CM micro services are integrated for the full-service testing in Lab with real-world input data (grid and load data). The tested setup is piloted in the pilot site and some results are already achievable.
HEMS (INESC)	TE not described in D3.3	Integrations with PT pilot DERs: Heat pump, EV charger and inverter. Core functionalities developed and initiated deployment.	Improvement of the integration between DER and Cloud service. New functionalities developed: flexibility actions, user notifications and incentives.
Digital platform capabilities for distribution automation (ABB)	TE not described in D3.3	Implemented virtualized protection solution and pre-processing modules to calculate Fast Fourier Transforms.	Modules to provide communication capabilities to the solution, namely to SCADA, to the congestion management and anomaly detection solution. Data storage for the results of the application.

Summary: 9 cloud technology enablers addressing use cases like flexibility optimization, demand response, energy poverty alleviation, congestion management and manual frequency restoration reserve.

TABLE 3 – HEDGE-IOT'S CLOUD TECHNOLOGY ENABLERS

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
EdgeConnect (INESC)	Design presentation of the existing platform to Hedge.IoT configuration and System Operator types.	This platform established an ecosystem for stakeholders across the flexibility value chain, enabling integration, qualification and market participation, to unlock flexibility potential from LV and MV grids. The platform unlocks a multi-stakeholder environment where all relevant roles coexist.	Bilateral agreement functionalities and full communication over standard protocols such as IEC 62325. Deployment complete and operational with integration with partner's platforms and services completed or tested.
Flexibility Optimization Service (ICCS)	Initial Overview of the algorithms of bidding strategies and flexibility dispatch.	IoT data cleaning and processing to fit into the decision-making algorithms and an initial implementation of specific algorithm to one defined scenario.	Define more scenarios and identify more decision-making algorithms.
Real-Time Reserve Market Simulator (NESTER)	Initial internal release allowing basic single bid submission, clearing and response.	Integration of IEC 62325 standard for energy market communication and file exchange. Multiple bid submission and analysis implementation. Release for pilot's initial integration testing. Improvements in bid validation, clearing and activation signal.	Full deployment of both services mFRR and aFRR with complete data stream tested and implemented among the pilot's members.
Predictive Congestion Management (TAU)	Breaking down the service into micro services to enhance modularity of algorithms and to facilitate coordination and cooperation of developers. However, this service has been in the design stage and implementation work has not yet started.	Define the data exchange's payload between cloud and edge to be able to develop adaptors (cloud-edge data exchange adaptor). The progress in real-time congestion management on edge will have an over-lapping benefit for this service for example in the areas of input data because the same grid and load data could be used in the cloud during the simulations. For the state forecast algorithm, the state estimation algorithm on the edge could be already a good starting point. The first version of the test environment (virtual machine running on the cloud) is ready.	The simulated predictive CM is tested and then executed on the cloud and the results are already available for analysis.
Energy Community Management Service for Frequency Restoration Reserve (INESC)	Platform's main functionalities developed and being tested in a relevant environment (another European project pilot).	Incorporate new functionality for flexibility provision to the reserve market and test it with mock data provided by the TSO.	Integrate the functionality in the platform and with the market platform using data spaces. Test the prototype with the new functionalities and integrations in the relevant environment (PT pilot).
Local Flexibility Market Platform (HENEX)	TE not described in D3.3	Developed an algorithm for daily generation and configuration of MTUs and for the automated operation of trading gates according to market schedule.	Implement asset registration and pre-qualification and portfolio management. Deployment of settlement and clearing mechanisms.

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
		Implemented algorithm for the execution of the market clearing.	
Energy Community Platform (APIO)	TE not described in D3.3	Development and implementation of core functionalities, like baseline computation, flexibility offer optimization and energy community power management.	Connection to local flexibility market platform. Integration with the interoperability framework of the project.
TurnGreen – OptiFlex (ELERG)	TE not described in D3.3	Integrated new EnergyBox solutions into supermarkets and into the OptiFlex platform.	Asset integration. Develop optimization algorithm for each asset. Create automated responses for flexibility requests.
PowerCIM tool (KONČAR)	TE not described in D3.3	Imported preliminary data and created appropriate measurements for the time series for DTR calculations.	Updates to the CIM models to reflect measurement points of the substation and for DTR integration.
<p>Summary: 1 computational orchestration framework, tailored for multiple different use cases, to facilitate computational sharing between the edge and cloud levels, contributing to establish edge/cloud continuum.</p>			

TABLE 4 – HEDGE-IOT’S COMPUTATIONAL ORCHESTRATION FRAMEWORK

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Computational Orchestration Framework (TUC)	High level design of the orchestration framework architecture and first specifications for two use cases: edge offloading for low-latency data processing and federated learning orchestration.	Detailed design of orchestration framework along with current implementation for each of the two use cases, and for an additional use-case that was identified: application/federated learning models rolling out at edge. Orchestrator integration with the HEDGE IoT framework focusing on available services, and data space connector.	Integration of orchestration framework with services available in pilots. Complete the integration with blockchain platform. Dashboard for real-time metric monitoring and performance insights.

D3.4 showcases the intermediate release of designs and specifications for the technology enablers of the HEDGE-IoT project. Furthermore, it shows the work package is progressing towards addressing its proposed objectives and that the technology enablers ecosystem is diverse, increasing in maturity and ready to progress to the final stage of getting ready for integration in the pilots under HEDGE-IoT’s multi-dimensional framework.

With the submission of this deliverable, 3 out of 5 Key Exploitable Results (KERs) directly linked with work package 3 are considered to be fully achieved, while the remaining 2 are considered to be

partially achieved. These have been considered partially/fully achieved based on the means of validation previously agreed upon on the project's grant agreement.

TABLE OF CONTENTS

TABLE OF CONTENTS

1	INTRODUCTION	27
1.1	About this document.....	28
1.2	Intended audience.....	28
1.3	Reading recommendations	28
1.4	Relationship with other work packages	29
2	FEDERATED LEARNING TECHNOLOGY ENABLERS	31
2.1	Federated Learning Service for Energy Forecasting.....	33
2.1.1	Description of technology enabler.....	33
2.1.2	Innovative aspects	34
2.1.3	Service data	34
2.1.4	Integration with HEDGE-IoT interoperability framework.....	35
2.1.5	Implementation details.....	35
2.2	Vector Autoregressive Model for Energy Time Series Forecasting	37
2.2.1	Description of technology enabler.....	37
2.2.2	Innovative aspects	38
2.2.3	Service data	38
2.2.4	Integration with HEDGE-IoT interoperability framework.....	38
2.2.5	Implementation details.....	39
3	DATA-DRIVEN EDGE-TO-CLOUD ENERGY TECHNOLOGY ENABLERS	42
3.1	Enhanced Network Management and Planning	45
3.1.1	Description of technology enabler.....	45
3.1.2	Innovative aspects	46
3.1.3	Service data	46
3.1.4	Integration with HEDGE-IoT interoperability framework.....	47
3.1.5	Implementation details.....	47
3.2	DTR-DLR on the Edge	50
3.2.1	Description of technology enabler.....	50
3.2.2	Innovative aspects	50
3.2.3	Service data	51
3.2.4	Integration with HEDGE-IoT interoperability framework.....	51

3.2.5	Implementation details.....	51
3.3	Anomaly Detection and Predictive Maintenance on the Grid.....	54
3.3.1	Description of technology enabler.....	54
3.3.2	Innovative aspects	55
3.3.3	Service data	55
3.3.4	Integration with HEDGE-IoT interoperability framework.....	56
3.3.5	Implementation details.....	57
3.4	Anomaly Detection and Fault Forecasting to Increase Distribution Network Resilience	60
3.4.1	Description of technology enabler.....	60
3.4.2	Innovative aspects	60
3.4.3	Service data	61
3.4.4	Integration with HEDGE-IoT interoperability framework.....	62
3.4.5	Implementation details.....	62
3.5	Real-Time Congestion Management	65
3.5.1	Description of technology enabler.....	65
3.5.2	Innovative aspects	66
3.5.3	Service data	66
3.5.4	Integration with HEDGE-IoT interoperability framework.....	67
3.5.5	Implementation details.....	67
3.6	HEMS - Home Energy Management System.....	71
3.6.1	Description of technology enabler.....	71
3.6.2	Innovative aspects	72
3.6.3	Service data	73
3.6.4	Integration with HEDGE-IoT interoperability framework.....	73
3.6.5	Implementation details.....	74
3.7	Digital platform capabilities for distribution grid automation.....	77
3.7.1	Description of technology enabler.....	77
3.7.2	Innovative aspects	78
3.7.3	Service data	78
3.7.4	Integration with HEDGE-IoT interoperability framework.....	78
3.7.5	Implementation details.....	79
4	CLOUD TECHNOLOGY ENABLERS.....	82

4.1	EdgeConnect.....	85
4.1.1	Description of technology enabler.....	85
4.1.2	Innovative aspects	86
4.1.3	Service data	86
4.1.4	Integration with HEDGE-IoT interoperability framework.....	87
4.1.5	Implementation details.....	88
4.2	Flexibility Optimization Service	91
4.2.1	Description of technology enabler.....	91
4.2.2	Innovative aspects	92
4.2.3	Service data	93
4.2.4	Integration with HEDGE-IoT interoperability framework.....	94
4.2.5	Implementation details.....	94
4.3	Real-Time Reserve Market Simulator	98
4.3.1	Description of technology enabler.....	98
4.3.2	Innovative aspects	98
4.3.3	Service data	99
4.3.4	Integration with HEDGE-IoT interoperability framework.....	100
4.3.5	Implementation details.....	100
4.4	Predictive Congestion Management.....	103
4.4.1	Description of technology enabler.....	103
4.4.2	Innovative aspects	103
4.4.3	Service data	104
4.4.4	Integration with HEDGE-IoT interoperability framework.....	105
4.4.5	Implementation details.....	105
4.5	Energy Community Management Service for Frequency Restoration Reserve.....	109
4.5.1	Description of technology enabler.....	109
4.5.2	Innovative aspects	109
4.5.3	Service data	110
4.5.4	Integration with HEDGE-IoT interoperability framework.....	111
4.5.5	Implementation details.....	111
4.6	Local Flexibility Market Platform.....	114
4.6.1	Description of technology enabler.....	114
4.6.2	Innovative aspects	116

4.6.3	Service data	117
4.6.4	Integration with HEDGE-IoT interoperability framework.....	117
4.6.5	Implementation details.....	118
4.7	Energy Community Platform	121
4.7.1	Description of technology enabler.....	121
4.7.2	Innovative aspects	122
4.7.3	Service Data.....	122
4.7.4	Integration with HEDGE-IoT interoperability framework.....	123
4.7.5	Implementation details.....	123
4.8	TurnGreen – Optiflex	126
4.8.1	Description of technology enabler.....	126
4.8.2	Innovative aspects	127
4.8.3	Service Data.....	128
4.8.4	Integration with HEDGE-IoT interoperability framework.....	128
4.8.5	Implementation details.....	129
4.9	PowerCIM tool.....	131
4.9.1	Description of technology enabler.....	131
4.9.2	Innovative aspects	131
4.9.3	Service data	132
4.9.4	Integration with HEDGE-IoT interoperability framework.....	132
4.9.5	Implementation details.....	132
5	COMPUTATIONAL ORCHESTRATION TO ENSURE THE CLOUD-EDGE CONTINUUM..	135
5.1	Description of the orchestration framework.....	135
5.2	Energy services offloading at edge.....	135
5.2.1	Description.....	135
5.2.2	Implementation details.....	137
5.3	Orchestration of Federated AI-driven Energy Services	142
5.3.1	Description.....	142
5.3.2	Implementation details.....	143
5.4	Energy services rolling out at edge.....	147
5.4.1	Description.....	148
5.4.2	Implementation details.....	148
5.5	Integration with Hedge-IoT framework	149

5.5.1	Italian pilot	149
5.5.2	Portuguese pilot	150
5.5.3	Finish pilot	150
5.6	Integration with Hedge-IoT Interoperability framework	151
5.7	Next steps	152
6	CONCLUSION	153
7	APPENDIX A	156
8	APPENDIX B	165
8.1	USER INTERFACES (CLOUD SERVICES)	165
8.1.1	Description of Service	165
8.1.2	Innovative Aspects	166
8.1.3	Service Data	167
8.1.4	Integration with Hedge-IoT Interoperability Framework	168
8.1.5	Implementation Details	169

LIST OF TABLES

TABLE 1 – HEDGE-IOT'S FEDERATED LEARNING TECHNOLOGY ENABLERS.....	8
TABLE 2 – HEDGE-IOT'S DATA-DRIVEN EDGE-TO-CLOUD TECHNOLOGY ENABLERS.....	9
TABLE 3 – HEDGE-IOT'S CLOUD TECHNOLOGY ENABLERS.....	11
TABLE 4 – HEDGE-IOT'S COMPUTATIONAL ORCHESTRATION FRAMEWORK	12
TABLE 9 – STATUS OF TECHNOLOGY ENABLERS	31
TABLE 10 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING FUNCTIONALITIES.....	35
TABLE 11 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERVES FORECASTING FUNCTIONALITIES	39
TABLE 12 – STATUS OF IMPLEMENTATION OF THE TE'S.....	42
TABLE 13 – ENHANCED NETWORK MANAGEMENT AND PLANNING FUNCTIONALITIES.....	48
TABLE 14 – DLR/DTR ON THE EDGE FUNCTIONALITIES.....	52
TABLE 15 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID FUNCTIONALITIES	57
TABLE 16 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE FUNCTIONALITIES.....	63
TABLE 17 – REAL-TIME CONGESTION MANAGEMENT FUNCTIONALITIES	69
TABLE 18 – HEMS (HOME ENERGY MANAGEMENT SYSTEM) FUNCTIONALITIES	74
TABLE 19 – DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION FUNCTIONALITIES 79	
TABLE 20 – STATUS OF IMPLEMENTATION OF TE'S.....	82
TABLE 21 – EDGECONNECT BASE DATA REPRESENTATION	86
TABLE 22 – EDGECONNECT FUNCTIONALITIES	89
TABLE 23 – FLEXIBILITY OPTIMIZATION SERVICE FUNCTIONALITIES	95
TABLE 24 – REAL-TIME RESERVE MARKET SIMULATOR FUNCTIONALITIES	100
TABLE 25 – PREDICTIVE CONGESTION MANAGEMENT FUNCTIONALITIES.....	106
TABLE 26 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE FUNCTIONALITIES	111
TABLE 27 – LOCAL FLEXIBILITY MARKET PLATFORM FUNCTIONALITIES.....	119
TABLE 28 – ENERGY COMMUNITY PLATFORM FUNCTIONALITIES.....	123
TABLE 29 – TURNGREEN – OPTIFLEX FUNCTIONALITIES.....	129
TABLE 30 – POWERCIM TOOL FUNCTIONALITIES	133
TABLE 31 – INPUT AND OUTPUT DATA (SERVICE 2.1)	156
TABLE 32 – INPUT AND OUTPUT DATA (SERVICE 2.2)	156
TABLE 33 – INPUT AND OUTPUT DATA (SERVICE 3.1).....	157

TABLE 34 – INPUT AND OUTPUT DATA (SERVICE 3.2)	157
TABLE 35 – INPUT AND OUTPUT DATA (SERVICE 4.1).....	158
TABLE 36 – INPUT AND OUTPUT DATA (SERVICE 4.2)	159
TABLE 37 – INPUT VARIABLES (FROM D3.3) (SERVICE 4.7)	161
TABLE 38 – OUTPUT VARIABLES (FROM D3.3) (SERVICE 4.7).....	161
TABLE 39 – CHARACTERIZATION OF VARIABLES (SERVICE 4.7)	161
TABLE 40 – INPUT AND OUTPUT DATA (SERVICE 4.9)	162
TABLE 41 – INPUT AND OUTPUT DATA (SERV. APPENDIX B)	163
TABLE 42 – USER INTERFACES FUNCTIONALITIES	170

LIST OF FIGURES

FIGURE 1 – RELATION BETWEEN WORK PACKAGES	30
FIGURE 2 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING ARCHITECTURE DIAGRAM – ICCS.....	33
FIGURE 3 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING TIMELINE OF DATASETS – ICCS	34
FIGURE 4 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING PILOT SGAM DIAGRAM – ICCS	36
FIGURE 5 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIVES FORECASTING ARCHITECTURE DIAGRAM – INESC.....	37
FIGURE 6 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIVES FORECASTING TIMELINE OF DATASETS - INESC	38
FIGURE 7 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIVES FORECASTING PILOT SGAM DIAGRAM – INESC	41
FIGURE 8 – ENHANCED NETWORK MANAGEMENT AND PLANNING ARCHITECTURE DIAGRAM – UNIZG	45
FIGURE 9 – ENHANCED NETWORK MANAGEMENT AND PLANNING TIMELINE OF DATASETS – UNIZG..	47
FIGURE 10 – ENHANCED NETWORK MANAGEMENT AND PLANNING PILOT SGAM DIAGRAM – INESC..	49
FIGURE 11 – DLR/DTR ON THE EDGE ARCHITECTURE DIAGRAM – JSI	50
FIGURE 12 – DLR/DTR ON THE EDGE PILOT SGAM DIAGRAM – JSI.....	53
FIGURE 13 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID ARCHITECTURE DIAGRAM – VU.....	54
FIGURE 14 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID PILOT SGAM DIAGRAM – VU.....	59
FIGURE 15 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE ARCHITECTURE DIAGRAM – VTT	60
FIGURE 16 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE TIMELINE OF DATASETS – VTT	62
FIGURE 17 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE: NEXT-PHASE OF IMPLEMENTATION – VTT	62
FIGURE 18 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE PILOT SGAM ARCHITECTURE	64
FIGURE 19 – REAL-TIME CONGESTION MANAGEMENT ARCHITECTURE DIAGRAM – TAU.....	65
FIGURE 20 – REAL-TIME CONGESTION MANAGEMENT PILOT SGAM DIAGRAM – TAU.....	70
FIGURE 21 – HEMS (HOME ENERGY MANAGEMENT SYSTEM) ARCHITECTURE DIAGRAM – INESC.....	72
FIGURE 22 – HEMS (HOME ENERGY MANAGEMENT SYSTEM) PILOT SGAM DIAGRAM – INESC.....	76
FIGURE 23 – DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION ARCHITECTURE DIAGRAM – ABB.....	77

FIGURE 24 – DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION PILOT SGAM DIAGRAM – ABB.....	81
FIGURE 25 – EDGECONNECT PILOT SGAM DIAGRAM – INESC.....	85
FIGURE 26 – EDGECONNECT: COMPONENT DIAGRAM – INESC	88
FIGURE 27 – FLEXIBILITY OPTIMIZATION SERVICE ARCHITECTURE DIAGRAM – ICCS.....	91
FIGURE 28 – FLEXIBILITY OPTIMIZATION PILOT SGAM DIAGRAM – ICCS	97
FIGURE 29 – REAL-TIME RESERVE MARKET SIMULATOR: TIMINGS FOR DOCUMENT EXCHANGE IN THE MFRR SERVICE – NESTER.....	99
FIGURE 30 – REAL-TIME RESERVE MARKET SIMULATOR PILOT SGAM DIAGRAM – NESTER.....	102
FIGURE 31 – PREDICTIVE CONGESTION MANAGEMENT ARCHITECTURE DIAGRAM – TAU.....	103
FIGURE 32 – PREDICTIVE CONGESTION MANAGEMENT PILOT SGAM DIAGRAM – TAU	108
FIGURE 33 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE ARCHITECTURE DIAGRAM – INESC.....	109
FIGURE 34 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE TIMELINE OF DATASETS – INESC	110
FIGURE 35 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE PILOT SGAM DIAGRAM – INESC.....	113
FIGURE 36 – LOCAL FLEXIBILITY MARKET PLATFORM ARCHITECTURE DIAGRAM – HENEX.....	115
FIGURE 37 – LOCAL FLEXIBILITY MARKET PLATFORM TIMELINE OF DATASETS – HENEX.....	117
FIGURE 38 – LOCAL FLEXIBILITY MARKET PLATFORM PILOT SGAM DIAGRAM – HENEX	120
FIGURE 39 – ENERGY COMMUNITY PLATFORM ARCHITECTURE DIAGRAM – APIO.....	121
FIGURE 40 – ENERGY COMMUNITY PLATFORM TIMELINE OF DATASETS – APIO.....	123
FIGURE 41 – ENERGY COMMUNITY PLATFORM PILOT SGAM DIAGRAM – APIO.....	125
FIGURE 42 – TURNGREEN – OPTIFLEX ARCHITECTURE DIAGRAM – SONAE/ELERG.....	127
FIGURE 43 – TURNGREEN – OPTIFLEX PILOT SGAM ARCHITECTURE – SONAE/ELERG.....	130
FIGURE 44 – POWERCIM TOOL PILOT SGAM DIAGRAM – KONC	131
FIGURE 45 –POWERCIM TOOL TIMELINE OF DATASETS – KONC.....	132
FIGURE 46 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR ENERGY SERVICES OFFLOADING AT EDGE – TUC.....	136
FIGURE 47 – MONITORING STACK ARCHITECTURE	137
FIGURE 48 – GRAPH MODEL FOR UNIFIED TASK AND INFRASTRUCTURE REPRESENTATION	139
FIGURE 49 – ORCHESTRATION OF ENERGY SERVICES AT EDGE WORKFLOW.....	140
FIGURE 50 – BLOCKCHAIN INTEGRATION WORKFLOW	141
FIGURE 51 – TOKEN ID STRUCTURE	142
FIGURE 52 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR FEDERATED AI-DRIVEN SERVICES.....	143

FIGURE 53 – ORCHESTRATION FRAMEWORK OVERVIEW FOR HIERARCHICAL FEDERATED LEARNING PROCESS 144

FIGURE 54 – ORCHESTRATION FRAMEWORK OVERVIEW FOR P2P FEDERATED LEARNING PROCESS 146

FIGURE 55 – ORCHESTRATION CLUSTERING FLOW 147

FIGURE 56 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR ENERGY SERVICES ROLLING OUT AT EDGE 148

FIGURE 57 – SERVICE ORCHESTRATION WORKFLOW 149

FIGURE 58 – USER INTERFACES PILOT SGAM DIAGRAM – ICCS173

ABBREVIATIONS

Abbreviation	Explanation
aFRR	Automatic Frequency Restoration Reserve
AI	Artificial Intelligence
AIOTI	Alliance for AI, IoT and Edge Continuum Innovation
API	Application Programming Interface
B2B	Business to Business
B2C	Business to Consumer
BEMS	Building Energy Management System
BMS	Building Management System
BTM	Behind the Meter
BUC	Business Use Case
CIM	Common Information Model
CM	Congestion Management
CSV	Comma Separated Values
DER	Distributed Energy Resource
DNP3	Distributed Network Protocol 3
DoA	Description of Action
DoEAP	Digitalization of Energy Action Plan
DSO	Distribution System Operator
EC	European Commission
EDC	Eclipse Dataspace Component
EMS	Energy Management System
ETSI	European Telecommunication Standardization Institute

EU	European Union
EV	Electric Vehicle
FSP	Flexibility Service Provider
GA	Grant Agreement
GDPR	General Data Protection Regulation
HEDGE-IoT	Holistic Energy Decentralized Grid for Enhanced IoT
HTTP	Hypertext Transfer Protocol
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IP	Intellectual Property
JSON	JavaScript Object Notation
KER	Key Expected Result
KPI	Key Performance Indicator
LFM	Local Flexibility Platform
LLM	Large Language Model
LV	Low Voltage
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
OCPP	Open Charge Point Protocol
OWL	Web Ontology Language

PV	Photovoltaic
QoS	Quality of Service
RDF	Resource Description Framework
RES	Renewable Energy Source
REST	Representational State Transfer
SAREF	Smart Applications REFerence Ontology
SHACL	Shapes Constraint Language
SIF	Semantic Interoperability Framework
SME	Small and Medium-sized Enterprise
SOSA	Sensor, Observation, Sample and Actuator
SUC	System Use Case
TRL	Technology Readiness Level
TSG	TNO Security Gateway
TSO	Transmission System Operator
UC	Use Case
UI	User Interface
VPN	Virtual Private Network
WoT	Web of Things
WP	Work Package
XML	eXtensible Markup Language
YAML	Yet Another Markup Language

1 INTRODUCTION

The HEDGE-IoT (Holistic Approach towards Empowerment of the Digitalization of the Energy Ecosystem through adoption of IoT solutions) project continues its mission to unlock the value of data and IoT assets across the energy ecosystem by advancing the development and deployment of intelligent data services. The HEDGE-IoT Framework will upgrade the RES-hosting capacity of the energy systems and will unleash a previously untapped flexibility potential. It will increase the resilience of the grid, create new market opportunities and promote advances in IoT standardization, by introducing and managing a plethora of diverse and interoperable energy services over an infrastructure of scalable and highly distributed data platforms. Altogether, HEDGE-IoT establishes an interoperable AI-focused development kit that explores data and computational infrastructures with innovative services for data analysis and value sharing.

The multi-dimensional framework of HEDGE-IoT comprises the following pillars: (a) the **Technology Facilitator Pillar**, exploiting the computational sharing by offloading applications to the grid edge, as to provide a set AI/ML federated learning and swarm computing applications; (b) the **Interoperability Pillar**, which leverages on leading-edge interoperable architectures, such as the Data Spaces; (c) the **Standardisation Pillar** which enables platforms, systems, tools and actors to seamlessly communicate and exchange data through widely adopted standardized data representations (e.g. SAREF); (d) the **Digital Energy Ecosystem Enabling Pillar** that establishes a resilient ecosystem facilitating the integration of RES.

The cornerstone of the technology and services in HEDGE-IoT lies precisely in exploring the edge-cloud continuum [1], where data and services adopt distinct requirements in terms of **data availability, privacy, latency, scalability** and **energy profile** (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** in this context.

Building on the foundations laid in D3.3, this deliverable further details the technological enablers that support the deployment of AI at the edge and throughout the edge-cloud continuum. In this deliverable, new services/technology enablers are described and their positioning regarding their pilot and other transversal services is laid. These specifications and designs play a big part in uncovering potential liaisons between services and, therefore, enabling the exploration of **cross-pilot** aspects (**Digital Energy Ecosystem Enabling Pillar**).

Furthermore, the **edge-cloud continuum** angle (**Technology Facilitator Pillar**) is explored via the computational orchestrator, which holds a unique feature in dynamic allocation of computational tasks based on cloud and edge resources available. This paradigm is designed and specified in three use cases for three different pilots.

The digital AI services considered in this deliverable, 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 AI, namely: Federated Learning, Auto-Regression, Deep Learning or Anomaly Detection, being then applied for use cases such as energy forecasting, network management and planning, predictive maintenance and flexibility optimization.

1.1 ABOUT THIS DOCUMENT

This deliverable reports an intermediate version of all technology enablers considered in HEDGE-IoT, which are detailed and developed in the scope of WP 3 –Technological Enablers Specification, Design and Deployment.

The main objective of this document is to **identify new technology enablers** that were not reported on the previous deliverable (D3.3) and **highlight the delta** (new developments) of previously specified technology enablers, which will later be deployed in the project pilots. Namely, this document is organized to describe the AI services in the project (T3.2 & T3.3), describing how they fit in the use case of their pilot (WP2), their goals, software architecture, innovative aspects, datasets, implementation details, functionalities, dependencies within the pilot and how it plans to integrate with the HEDGE-IoT's interoperability framework (WP4). Lastly, it presents a detailed description of the computational orchestration framework, as well as its implementation details. Furthermore, it details the steps and activities taken to ensure that services from three different pilots, prone to benefit from workload offloading, are deployed along the continuum between the edge and the cloud (T3.5).

1.2 INTENDED AUDIENCE

This deliverable is written considering the following audience:

- 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 the 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.
- General audience as soon as the document is available for dissemination.

1.3 READING RECOMMENDATIONS

This document is divided into 8 chapters and two appendices. Chapter 1 introduces the document. Chapter 2 presents the federated learning approaches considered as technology enablers. 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 approaches the deployment and orchestration across the edge-cloud continuum. Chapter

6 concludes the document. Finally, Appendix A contains details of input and output variables of relevant technology enablers, and Appendix B has a characterization of a service that has a different nature from the others in this delivery (User Interface).

1.4 RELATIONSHIP WITH OTHER WORK PACKAGES

Work package 3 is defined as a pivot work package devoted to the identification, evolution and development of digital technology enablers 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, as Business (T2.2) and System use cases (T2.6), and the baseline for the activities and existing digital platforms and technology enablers 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 technology enablers are then split into services that focus on consumers (T3.2), on using AI on the edge level, namely Federated Learning Algorithms, Tools and Services on the edge (T3.3) and AI/ML services on the cloud layer, that focus on operation and planning of infrastructure (T3.4). Lastly, to achieve cloud-edge continuum, a task is devoted designing, developing and implementing a computational orchestrator for workload offloading of real-time AI applications (T3.5).

The services delivered by WP3, namely in scope of T3.3, will be shipped and integrated in the project-level digital interoperable platforms in 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, 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 Connector (T4.2) (HEDGE-IoT uptake of the Eclipse Data Connector (EDC)), made discoverable, deployable and hostable (opt-in) in the App Store (T4.1). Semantic 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.

D3.4 focuses on the work carried out in scope of work package 3, namely in T3.2, T3.3, T3.4 and T3.5, from M13 to M19, delivering specifications, designs and implementation of technology enablers across the energy ecosystem and the edge/fog/cloud layers, empowered by a computational orchestrator to benefit from the edge-cloud continuum.

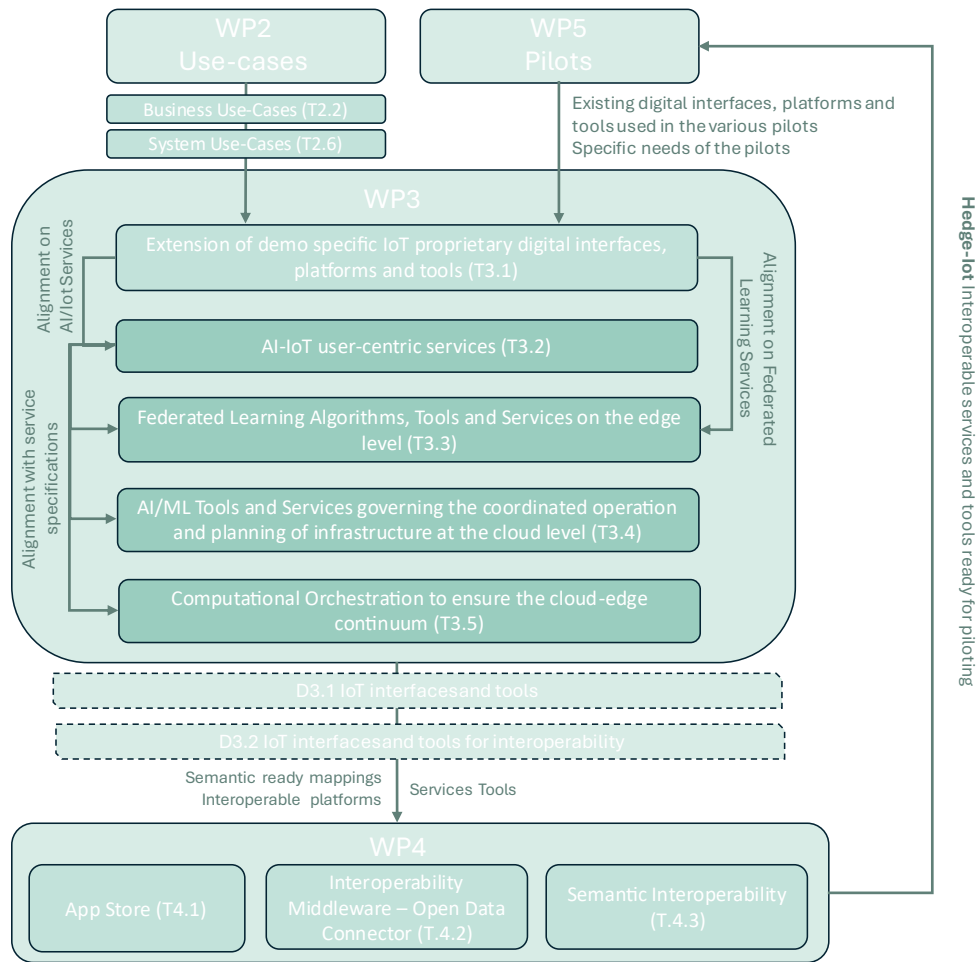


FIGURE 1 – RELATION BETWEEN WORK PACKAGES

2 FEDERATED LEARNING TECHNOLOGY ENABLERS

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. This technique mitigates the need for centralized data aggregation, which is often infeasible or undesirable due to privacy, legal, or bandwidth constraints [2].

FL is useful for a variety of reasons, most importantly data privacy, decentralization and scalability in low latency environments [3]. These aspects are fundamental to achieve the goals of the project, namely the decentralization of the grid and adding intelligence to the edge layer.

In this section, the HEDGE-IoT federated learning technology enablers (TEs) intermediate specifications and designs are outlined. These TEs are currently being developed and will be deployed and tested to be part of the HEDGE-IoT framework and its pilots.

Table 5, below, presents a summary of the current developments and describes the upcoming updates for each TE.

TABLE 5 – STATUS OF TECHNOLOGY ENABLERS

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Federated Learning for Energy Forecasting (ICCS)	Decentralized federated learning architecture using LSTM/BiLSTM models for forecasting energy demand and production. Initial training conducted using open datasets such as StoreNet. Model updates exchanged via MQTT secured by TLS.	BiLSTM models, trained with pilot data from residential apartments. Models converted to TensorFlow Lite for deployment on Shelly 3EM meters. MinIO used for model storage and secure model distribution. Integrated secure MQTT-based communication.	Pilot deployment and real-world testing in 100 Greek apartments. Validation of real-time performance on edge devices. Iterative model refinement based on pilot results. Energy footprint monitoring for edge model inference.
Vector Autoregressive Model for Energy Time Series Forecasting (INESC)	Proof of concept algorithm observed with synthetic data. Data processing pipeline developed and connected to real pilot data, for a select group of users	First lab concept solution in a fully functioning demo environment with 4 different simulated edge agents Design an initial development of integration of pilot users' data pipeline for each agent	Integrate the lab prototype in a larger scale environment according to the number of users participating in the pilot at the given date Integrate the edge nodes with the computational orchestrator tool being developed in WP3

The TEs described in the table above provide a means of **validation for the accomplishment of the project's objectives**, namely by:

- **Exploiting 2 different types of IoT/edge devices:**
 - Shelly 3EM and Raspberry Pis

- **Designing 2 applications for residential users for the energy sector:**
 - Household energy consumption forecasting applications

2.1 FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING

2.1.1 Description of technology enabler

On D3.3, the previous iteration of this deliverable, a detailed description of this approach was provided. Therefore, section 2.1 of D3.3 is the reference for the description of this service. The Federated Learning (FL) for Energy Forecasting service aims to enable privacy-preserving, decentralized energy forecasting. It was originally designed to train forecasting models (e.g., LSTM, BiLSTM) centrally and deploy them on low-resource residential edge devices such as Shelly 3EM.

As described in D3.3, the concept relied on transmitting encrypted model updates from edge nodes to a central server and pushing trained models back to the devices. However, during testing since D3.3, it was found that cloud-trained models were too computationally intensive to run on the edge devices due to their limited resources.

In response, the service is now undergoing some alterations. Instead of running full models on edge devices, the updated goal is to translate trained models into lightweight matrix operations that can be executed locally. This mathematical simplification would allow for local inference without the burden of model complexity, while still preserving user data privacy.

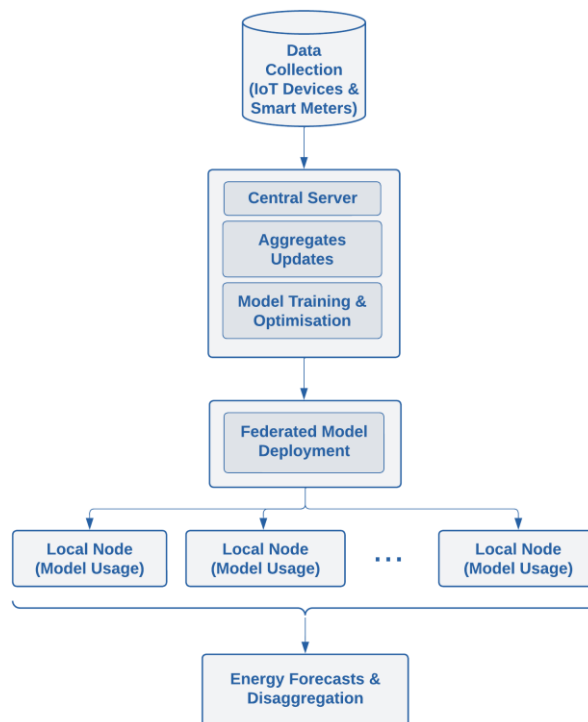


FIGURE 2 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING ARCHITECTURE DIAGRAM – ICCS

2.1.2 Innovative aspects

PRIVACY-BY-DESIGN

Raw time-series energy data remains on the device; only anonymized or encrypted parameters are exchanged. This ensures GDPR compliance and aligns with HEDGE-IoT’s distributed architecture principles.

PIVOT TO MATHEMATICAL LIGHTWEIGHT INFERENCE

A major shift has been initiated to translate complex forecasting models into simplified matrix operations, allowing constrained devices to perform inference using basic math instead of full neural networks.

ENERGY-EFFICIENT DESIGN

Edge devices such as Shelly 3EM are used to minimize computational overhead and energy consumption during model inference.

2.1.3 Service data

The FL service uses real-time and historical datasets from ~100 households in the Greek pilot. However, the system has not yet reached the stage of real-time inference or successful model execution at the edge.

Input Data:

- Time-series energy consumption (per household)
- PV production where applicable
- Environmental weather variables from external weather service (e.g, OpenMeteo¹)

Output Data (Planned):

- Local forecasts for energy consumption and production.



FIGURE 3 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING TIMELINE OF DATASETS – ICCS

¹ <https://open-meteo.com/>

2.1.4 Integration with HEDGE-IoT interoperability framework

As of D3.4, no functional integration has been completed yet.

2.1.5 Implementation details

As of deliverable D3.4, testing revealed that running cloud-trained forecasting models directly on Shelly 3EM edge devices is not feasible due to computational constraints. As a result, the architecture is being restructured to support lightweight inference via pure mathematical formulations (e.g., matrix multiplications) derived from trained models.

This change represents a significant shift in the implementation strategy. The new approach focuses on translating trained models into simplified matrix-based operations suitable for execution on constrained edge hardware.

2.1.5.1 Functionalities

Table 6 shows the list of functionalities comprehended by this service.

TABLE 6 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Privacy-preserving architecture	FL architecture designed to keep raw data local	100%	M13
Cloud model training	Central training using real data	70%	M26
Mathematical model translation	Planned simplification of models to matrix operations for edge inference	0%	M30
Integration with the App Store	Service registered and packaged in the HEDGE-IoT App Store	0%	M32

2.1.5.2 Integration and dependencies

Current integration remains theoretical. The architecture anticipates the following dependencies:

1. Shelly 3EM meters:

Collect and send real-time production and consumption data from residential households. These devices are also intended to apply centrally trained models (translated to matrix operations) to local data in order to perform inference directly on the device.

2. Timescale Database:

Used to store and access collected energy data for training and validation of forecasting models.

Figure 4 shows an architecture diagram that maps the integrations and dependencies of the service with the remaining Greek pilot services.

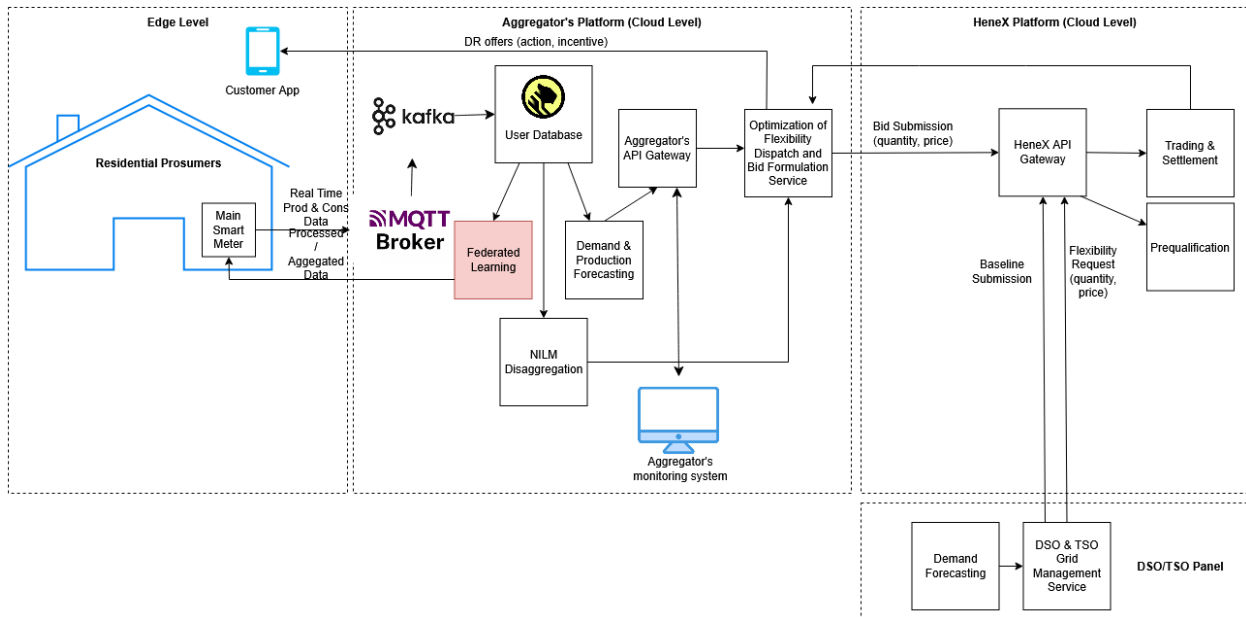


FIGURE 4 – FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING PILOT SGAM DIAGRAM – ICCS

2.2 VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING

2.2.1 Description of technology enabler

On D3.3, the previous iteration of this deliverable, a detailed description of this approach is provided. Therefore, section 2.2 of D3.3 is the reference for the description of this service.

This service has not changed in its core [4], but it does have a new functionality, which is already being developed, aimed at integrating with the HEDGE-IoT ecosystem. The functionality is the integration with the computational orchestrator (Section 5). The orchestrator will communicate with the “central” server of the FL application (see Figure 5) and organize the queue of P2P agents, as well as try to divide them into clusters.

Regarding updates to existing functionalities, a new approach for the integration of the pilot data in the P2P agents was designed (see Figure 5) and is now being developed. The approach will rely on data consumers streaming data directly, in real-time, to the agents, rather than the agents fetching the data from the data lake, in an offline manner.

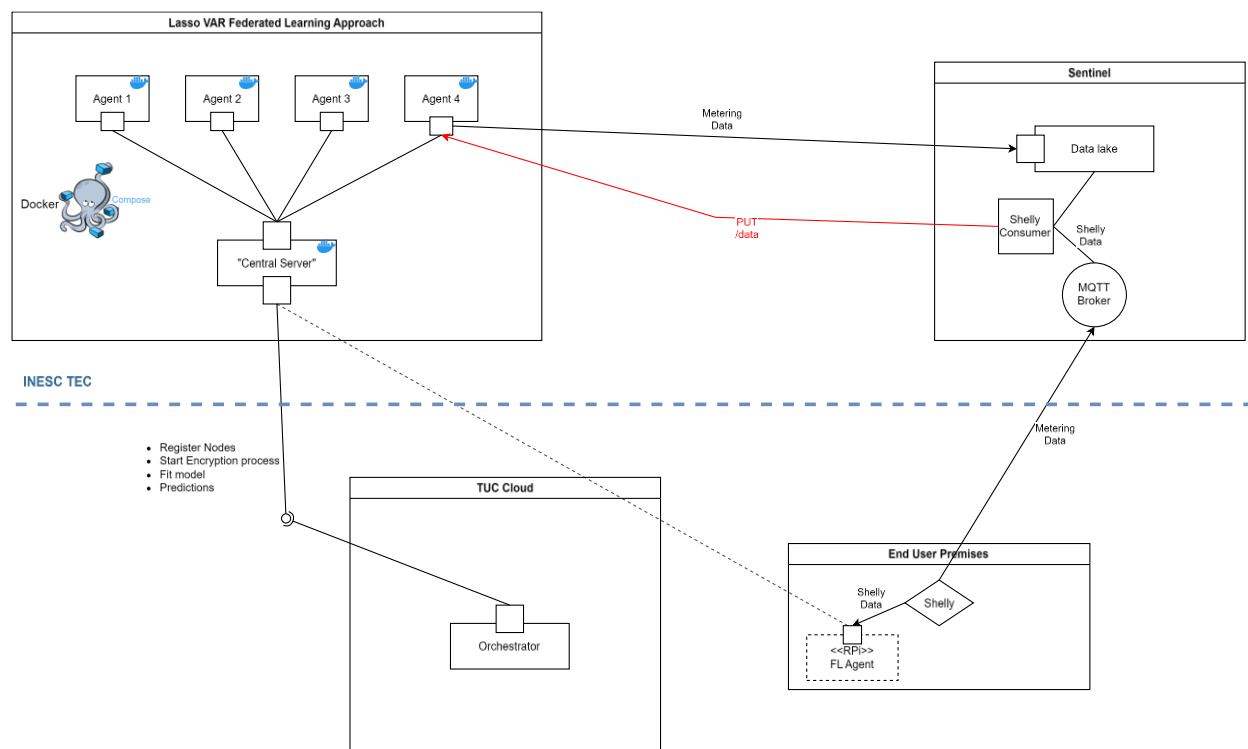


FIGURE 5 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING ARCHITECTURE DIAGRAM – INESC

Finally, for the next and final iteration of this deliverable (D3.5), the plan is to finalize the integration of both the functionalities described above: Integration with orchestrator and integration with data consumers. Furthermore, the FL model will continue to be optimized and trained to improve its results and its integration with more pilot end users will steadily continue.

2.2.2 Innovative aspects

In this section, the innovative aspects of the service will be detailed. In the previous iteration of this deliverable, two innovative aspects are already mentioned: privacy-preserving encryption and optimization of the matrix approach. For this iteration, one more innovative aspect has been observed and is detailed below.

INTEGRATION WITH COMPUTATIONAL ORCHESTRATOR

This P2P FL approach focuses on using edge nodes (real and/or virtual), storing data locally and performing computation on them. These edge nodes have varying computational resources and, while it is not possible to leave any of them out of the fitting process, it is possible to organize the queue to maximize the efficiency of the resources, which is where the computational orchestrator enters.

Most approaches ignore the real complexity and variations of edge computing with resources that have different computational capabilities. This integration is a step in the direction of making edge-level FL solutions more computationally aware.

2.2.3 Service data

The FL service is a P2P approach that makes predictions about the energy consumption of its users. The data is stored locally on each agent and only an encrypted portion is shared with the other agents.

For that purpose, two datasets are stored on each agent, as Figure 6 shows:

- **Residential households' energy consumption:** a time series of the aggregated energy consumption of the user in kWh.
- **Individual assets energy consumption:** time series of the energy consumption of the relevant and controllable flexible assets of each agent, for example, heat pumps.




FIGURE 6 - VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING TIMELINE OF DATASETS - INESC

2.2.4 Integration with HEDGE-IoT interoperability framework

In the previous iteration of this deliverable (D3.3), it was described that this service would integrate with the HEDGE-IoT interoperability framework by connecting to its data space via the Eclipse data

space connector. While this is still the plan, a new integration is being planned and it is detailed below.

Under the HEDGE-IoT interoperability framework, this service will directly integrate with a transversal component of the project called the app store. The exact architecture for this integration is still unclear, as the app store is under development, but it is important to mention that the service will be described in this store and, if possible, packaged under it as well.

2.2.5 Implementation details

The implementation of the integration with the computational orchestrator has also started and is in the early stages, which include the preparation of the FL service to send/receive inputs from the orchestrator.

Lastly, the expectation for M30 has not changed much. The plan is to have the service fully implemented and scaled to a larger number of pilot users (TRL 5), ready for the beginning of the large-scale pilot on M33. Additionally, the plan is to have, at least, a rough integration with the interoperability components of HEDGE-IoT, namely the data space and the App Store.

2.2.5.1 Functionalities

Table 7 shows the list of functionalities comprehended by this service.

TABLE 7 – VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Data Storage	Each agent in the FL approach stores their data locally in an encrypted way.	100%	M18
Data encryption	P2P procedure where data is shared and partially encrypted by all peers.	100%	M13
Privacy preserving model fitting	P2P procedure where each agent trains its model locally and then shares encrypted coefficients with the remaining peers	70%	M26
Forecast energy consumption	Each agent computes its decrypted forecast of the energy consumption data stored in it	70%	M30
Compute node resources and energy usage	Agents compute their resource usage percentage and energy consumption of the algorithm	80%	M22

Interoperable integration with computational orchestrator	Computational orchestrator defines the order and cluster of the agents used in the P2P approach. Communication is done using data spaces and the Eclipse Data Connector	10%	M32
Integration with the App Store	Service registered and packaged in the HEDGE-IoT App Store	0%	M32

2.2.5.2 Integration and dependencies

This section describes how the FL service integrates with other services in the Portuguese Pilot.

3. [Service 4.5] Energy Community Management Platform (RECreation):

- a. Sends consumers energy consumption forecasts

4. [Service 5] Computational Orchestrator:

- a. Sends resource analytics of each node
- b. Receives orchestration of nodes in the form of clusters and queue orders
- c. Communicates using the Eclipse data space connector

5. App Store/Data Space Integration:

- a. Deployed to the app store
- b. Integrated into the project's data space

Figure 7 **Error! Reference source not found.** shows a SGAM diagram that maps the integrations and dependencies of the service with the remaining PT pilot services, as well as showcases the status of the implementation of each integration.

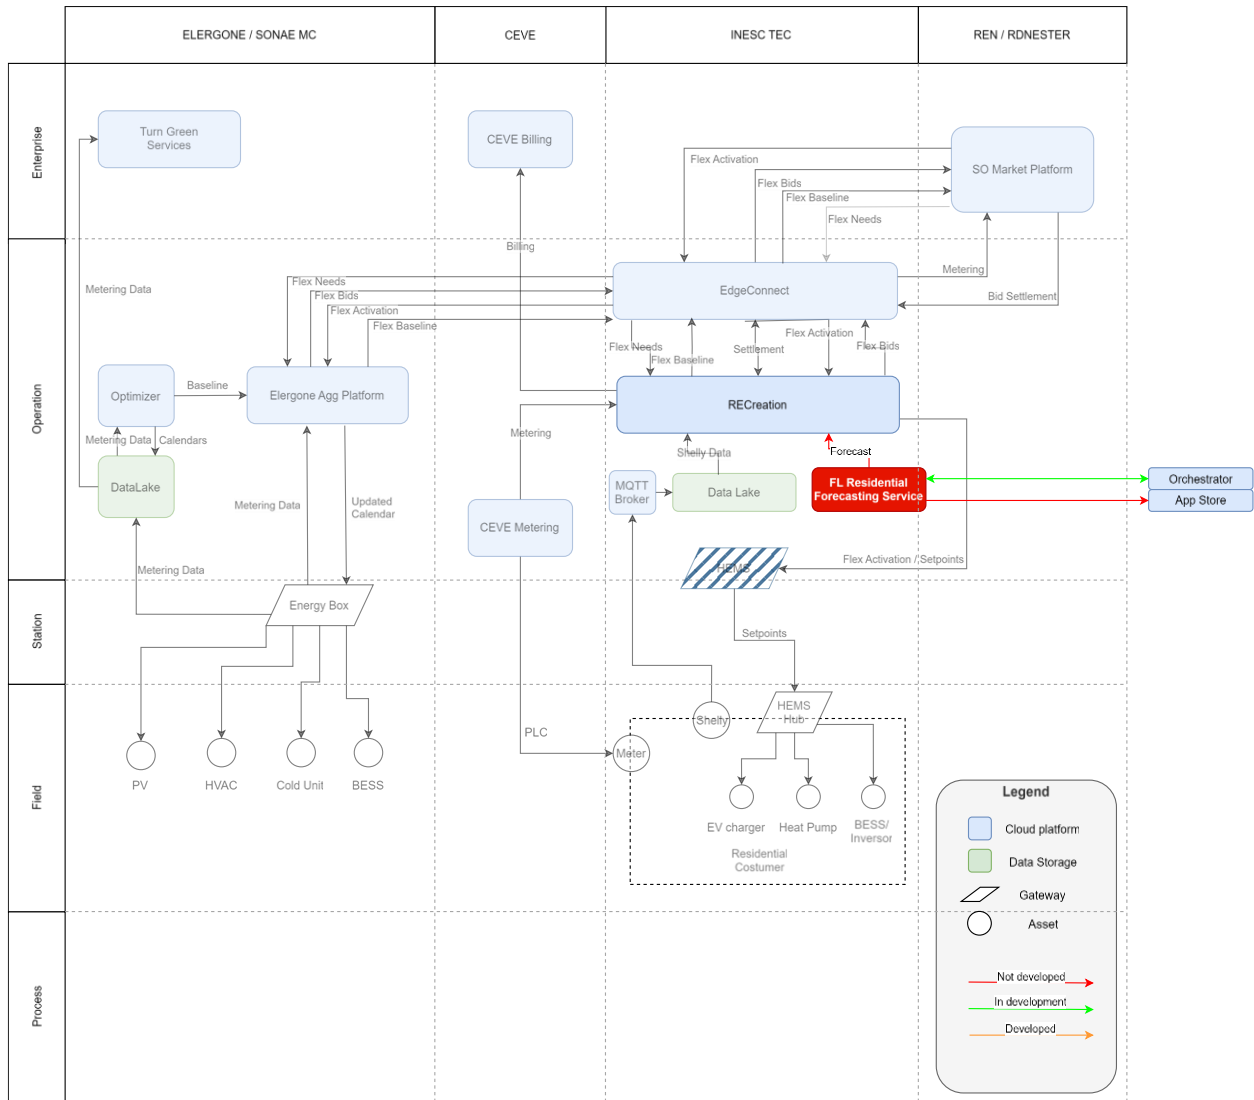


FIGURE 7 - VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERVICES FORECASTING PILOT
 SGAM DIAGRAM - INESC

3 DATA-DRIVEN EDGE-TO-CLOUD ENERGY TECHNOLOGY ENABLERS

This section outlines the intermediate specifications of the data-driven edge-to-cloud energy technology enablers (TEs) 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 5), which will facilitate their integration by discovering available edge devices and offloading their workloads in an efficient manner.

Table 8, showcased below, shows a summary of the status of implementation of the TEs, while mentioning what has been improved since the previous deliverable and what is expected to be complete in the final one.

TABLE 8 – STATUS OF IMPLEMENTATION OF THE TE'S

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Enhanced Network Management and Planning (UNIZG)	Investigation of techniques used in anomaly detection, demand forecast and initial development and testing of algorithms.	Testing anomaly detection and demand forecast algorithms on synthetic data. Research on techniques to detect PVs in secondary distribution networks from limited measurements.	Testing anomaly detection and demand forecast algorithms on realistic data and PV detection algorithm on synthetic data.
DTR-DLR on the Edge (JSI)	Initial overview of the DLR/DTR and weather forecast algorithms.	DLR/DTR algorithms adapted to IoT use case, improved interoperability plan.	Improved overall robustness of DLR/DTR algorithms, weather forecast algorithms deployed and updated, improved interoperability.
Anomaly Detection and Predictive Maintenance on the Grid (VU)	Learn patterns of nominal device behaviour from SAREFised data streams.	Recognise and report unnatural deviations from learned patterns (via SIF).	Improve overall robustness and minimize false error rate; add admin UI and export function.
Anomaly Detection and Fault Forecasting to Increase Distribution Network Resilience (VTT)	Deep Learning Model to Analyze data patterns, correlations, and identify anomalies.	HLSTM model to analyze data patterns, correlations, identify anomalies, and forecast future anomalies.	Improved identification of data patterns and anomaly detection with high accuracy and forecasting of future anomalies.

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
Real-Time Congestion Management (TAU)	Breaking down the service into micro services to enhance modularity of algorithms and to facilitate coordination and cooperation of developers. Specifying the steps in which the implementations could be realized. These steps are data preparation and harmonization, work related to edge servers (engineering the hardware, establishing virtual machines, networking), component development and testing, integration testing of developed components, and finally, piloting the solutions.	Micro services and the hardware where they are going to be executed have been decided. Initial agreement on the timeline of microservice developments done by different organizations aimed at piloting the whole solution. Harmonizing Input data has been completed. (validating the input data is expected to be completed by the end of April). It is expected to have the first version of state estimation micro service necessary for real-time congestion management ready for this deliverable. The first version of the test environment where the developed micro services will be tested is also expected to be ready.	Real-time CM is also developed. State estimation and real-time CM micro services are integrated for the full-service testing in Lab with real-world input data (grid and load data). The tested setup is piloted in the pilot site and some results are already achievable.
HEMS (INESC)	TE not described in D3.3	Integrations with PT pilot DERs: Heat pump, EV charger and inverter Core functionalities developed and initiated deployment	Improvement of the integration between DER and Cloud service New functionalities developed: flexibility actions, user notifications and incentives
Digital platform capabilities for distribution automation (ABB)	TE not described in D3.3	Implemented virtualized protection solution and pre-processing modules to calculate Fast Fourier Transforms	Modules to provide communication capabilities to the solution, namely to SCADA, to the congestion management and anomaly detection solution Data storage for the results of the application

These technology enablers are pivotal in achieving the project's goals, namely by adding **AI/ML intelligence to the edge**, enabling **decentralized approaches** and providing **energy and non-energy applications to stakeholders across the energy system value chain**.

Furthermore, the TEs described in the table above provide a means of validation for the accomplishment of the project's objectives, namely by:

- **Exploiting 5 different types of IoT/edge devices:**
 - IoT maxx GW- 4100 gateways, Raspberry Pis, smart devices, smart meters and intelligent electronic devices (IEDs)

- **Designing 1 applications for residential users for the energy sector:**
 - HEMS(Energy Home Management).
- **Designing 6 applications for grid operators:**
 - Network management and planning, DTR-DLR, Anomaly detection and fault forecasting, congestion management and grid automation.
- **Designing 2 non-energy applications:**
 - Predictive maintenance, DTR-DLR (advanced weather forecasting), HEMS (comfort of building occupants)

3.1 ENHANCED NETWORK MANAGEMENT AND PLANNING

3.1.1 Description of technology enabler

On D3.3, the previous iteration of this deliverable, a detailed description of this approach is provided. Therefore, section 1.1 of D3.3 is the reference for the description of this service. The core of the service remains unchanged; however, its current focus has shifted from conceptual development (as presented in D3.3) to verification and validation. All algorithms comprising the service have now been tested on historical datasets formatted to match those used in the pilot.

No new functionalities have been introduced. The aim of the current phase is to assess the accuracy of the service's algorithms in preparation for potential deployment at the edge level.

During the initial phase, measurements from IoT devices and weather stations are stored in a central database (see Figure 8), and all algorithmic computations are performed in the cloud. This serves to establish a validated benchmark for accuracy.

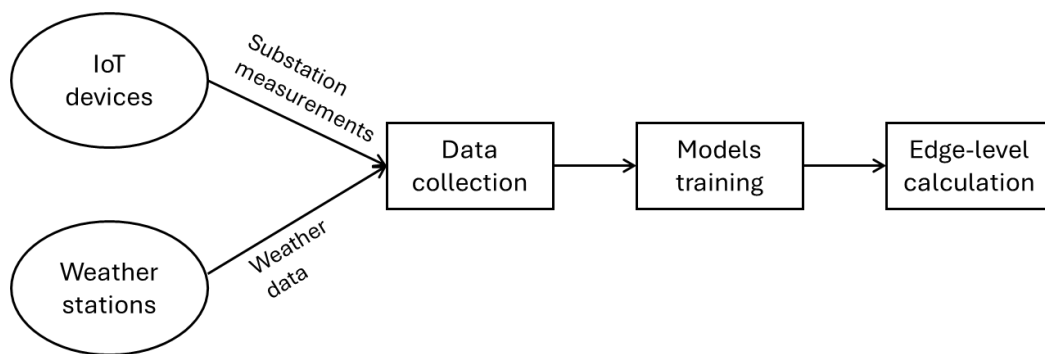


FIGURE 8 – ENHANCED NETWORK MANAGEMENT AND PLANNING ARCHITECTURE DIAGRAM – UNIZG

In the final phase that will be presented in D3.5, the concept will be evaluated at the edge level to explore the feasibility of offloading computations from the cloud to the edge infrastructure. While data sources are consistently located at the edge level in all phases, it remains an open question whether computational tasks can also be effectively decentralized.

Validation of algorithm accuracy on the cloud level is a critical prerequisite before attempting any edge-level deployment.

There are three algorithms that are part of the service, each presenting a specific functionality:

- Firstly, PV consumption is estimated from cumulative substation measurements using a combination of the Savitzky-Golay filter for signal smoothing and XGBoost for refining results based on supervised learning. The approach has been tested using PV production data generated for real-world locations and is designed to preserve key signal features while accurately identifying embedded PV generation.
- Anomaly detection is performed on three-phase voltage measurements from the DTR substation using the Isolation Forest algorithm, which identifies data points that are rare and different from the rest of the dataset. By recursively isolating observations through random features and split selection, the algorithm calculates an anomaly score for each data

point, enabling the detection and removal of point anomalies based on a defined contamination rate.

Finally, a deep learning artificial neural network (ANN) is used to predict active and reactive power consumption by learning complex patterns in historical measurement data. The model is trained through a structured process including data preprocessing, architecture design, and hyperparameter tuning, and is optimized using the Mean Squared Error loss function with the Adam optimizer, incorporating techniques such as L2 regularization to improve generalization and ensure accurate consumption forecasts.

3.1.2 Innovative aspects

DEVELOPMENT AND VERIFICATION OF PLANNING AND OPERATIONAL ML ALGORITHMS

This service comprises three machine learning algorithms: Anomaly Detection, Electricity Consumption Forecasting, and Photovoltaic (PV) Detection, i.e., an algorithm that detects PV generation in a secondary distribution network, from substation-level measurements. The innovation and contribution of these algorithms to the planning and operation of distribution networks were described in detail in D3.3. In this deliverable, the focus is on the validation and refinement of these algorithms using real-world data, which eliminates issues commonly associated with synthetic datasets, such as unrealistic results or potential manipulation. Multiple machine learning techniques were evaluated to determine the most effective approach for each task, further strengthening the robustness and applicability of the service.

3.1.3 Service data

All data (see Figure 9) used in this service is measured at the edge level, specifically by IoT devices installed in secondary substations. These measurements are then stored centrally, allowing for the verification and validation of the service's algorithms at the cloud level, which represents the first step in the overall implementation strategy. Currently, only historical measurements are available, which were used to test and validate the developed service. In the upcoming months, IoT devices that measure needed values will be installed at new locations and an additional set of measurements will be available for additional improvements of the service. The development and testing of data-driven models were conducted using two real-world datasets:

- **DTR substation measurements and weather data**
 - This dataset contains calculated thermal limit values for transformers and different weather data, including ambient temperature outside the substation.
- **Power quality measurements dataset**
 - This dataset contains electrical measurements, i.e., values of active and reactive power, voltage magnitude and current measured at secondary substations.

These datasets (Figure 9) served as the foundation for training, evaluating, and optimizing the machine learning algorithms included in the service.

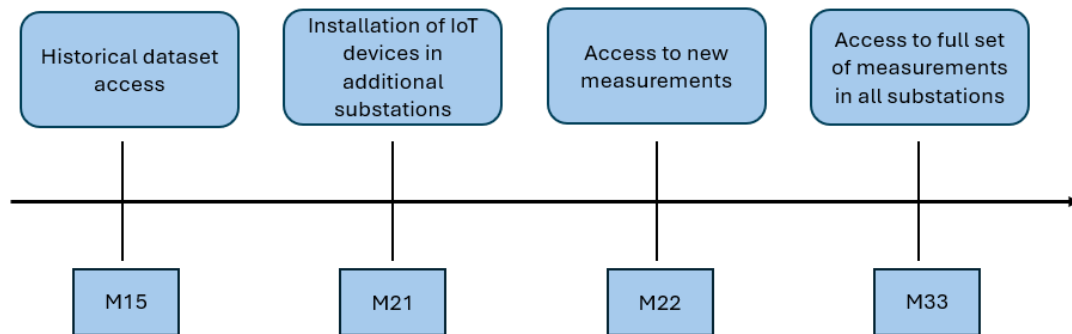


FIGURE 9 – ENHANCED NETWORK MANAGEMENT AND PLANNING TIMELINE OF DATASETS – UNIZG

3.1.4 Integration with HEDGE-IoT interoperability framework

The service ensures interoperability through two main approaches: data space connectors and semantic interoperability. **Integration with the HEDGE-IoT interoperability framework has already been defined and fully aligns with the project’s architecture, making further changes unnecessary, as they would not bring additional value.** Data space connectors are being considered to enable secure and standardized data exchange between entities such as measurement and service providers. For other data transfers, standard communication protocols like MQTT are used. Semantic interoperability is achieved through a custom CIM-based solution developed specifically for this service.

3.1.5 Implementation details

After initial testing on synthetic data, which corresponded to TRL 4, additional testing has been conducted using real world measurements from a secondary substation. These measurements match the type and format of the final dataset expected to be available in M30, enabling a more realistic evaluation of the service’s algorithms. As a result, the current TRL has increased to 5. The goal is to continue the development and validation of the service and reach TRL 7 by M30. Achieving this level means the service will be ready for deployment by a Distribution System Operator (DSO) in a fully operational environment. This includes performing all necessary calculations close to the data source and enabling DSOs to use the results to support decisions in the planning and operation of distribution networks, thereby offering a significant enhancement to current practices. While each algorithm can be deployed at the edge level as it relies solely on data collected at substations, the training and validation of the machine learning models will remain on the central cloud platform due to their complexity and computational requirements. The primary end user of the service remains to be the DSO, as they provide the input data and stand to gain the most from its application in real world scenarios.

According to the initial plan, there was not supposed to be any established connection with the interoperability components of HEDGE-IoT, namely the data space and the App Store, and for this reason it was not reported in D3.3 or in this deliverable. However, some changes are anticipated in the upcoming period, and the plan for D3.5 is to achieve at least a rough integration of the service with the dataspace connectors.

3.1.5.1 Functionalities

Table 9 describes the functionalities of this technology enabler.

TABLE 9 – ENHANCED NETWORK MANAGEMENT AND PLANNING FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Anomaly detection	ML technique is applied to identify and remove measurements that do not fit with the given dataset. The functionality has been tested and verified using synthetic and real-world data, but additional real-world datasets are required to further test potential application in a real-world environment.	85%	M30
Electricity demand forecast	ANN algorithm is implemented for purposes of active and reactive power consumption prediction. Same as for the previous functionality, the algorithm has been tested on both synthetic and real-world data with the request for additional testing on more real-world datasets.	85%	M30
PV detection in secondary distribution networks	Installed capacity and generation of PVs that are located in secondary distribution networks are estimated based on cumulative measurements at the substation level and relevant weather data. Testing of the functionality and plans for further development remain unchanged compared to other functionalities.	70%	M32

3.1.5.2 Integration and dependencies

This section describes how this cloud-edge service integrates with other services in the Slovenian Pilot.

1. **[Service 4.9] PowerCIM tool:**

- a. Sends semantic substation model of substations, enabling the access to power quality measurements

2. **[Service 3.2] SUMO Dynamic Rating System:**

- a. Sends DTR measurements dataset, containing information relevant for verifying the service

3. **App Store/Data Space Integration:**

- a. Deployed to the app store
- b. Integrated into the project's data space

Figure 10 shows a SGAM diagram that maps the integrations and dependencies of the service with the remaining SI pilot services, as well as showcases the status of the implementation of each integration.

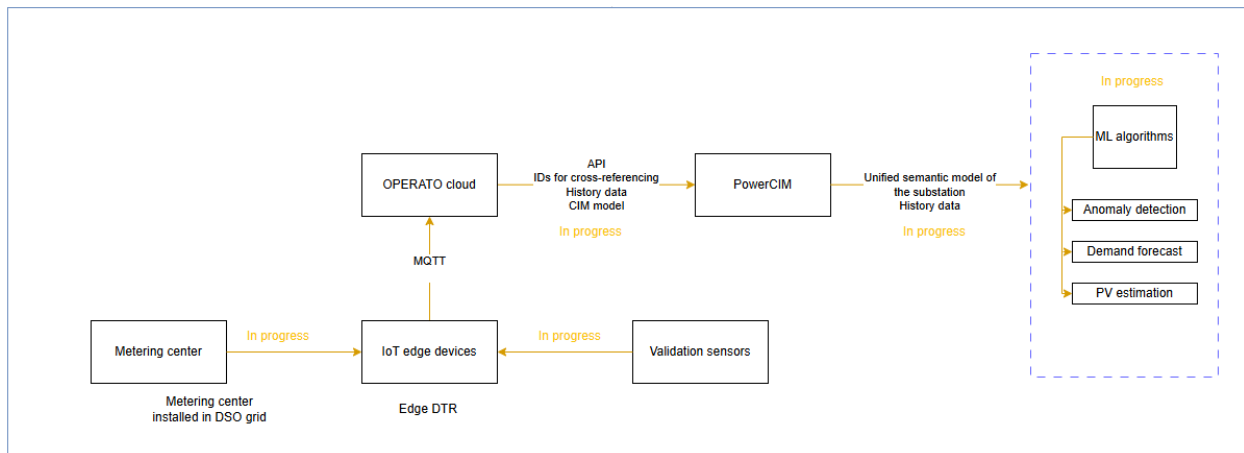


FIGURE 10 – ENHANCED NETWORK MANAGEMENT AND PLANNING PILOT SGAM DIAGRAM – INESC

3.2 DTR-DLR ON THE EDGE

3.2.1 Description of technology enabler

This service computes the dynamic thermal rating of overhead lines (DLR) and transformers (DTR) based on weather measurements or ML-based local weather forecast and operation data. It was described in detail in D3.3, the previous iteration of this deliverable.

The service in its core has not changed from D3.3, the main distinction is that in D3.3, the focus was on the DLR/DTR and weather forecast algorithms, and in D3.4, SUMO, a modular and productized Dynamic Line Rating solution, is also described in a bit more detail. The development of the service has been ongoing, and the progress will be described in the following sections.

Architecture diagram

The architecture diagram of the service is the same as in D3.3 and is shown in Figure 11. It involves an API interfacing with a DLR/DTR simulation engine, which performs computations by leveraging inputs from a local weather forecast engine and data from a local database. The outputs are integrated with the SUMO BUS for further processing.

SUMO BUS is a part of the SUMO system, which will be responsible for connecting the DLR/DTR and weather forecast algorithms and its outputs to other systems outside of the service.

SUMO is a modular and productized Dynamic Line Rating solution developed by Operato, a subsidiary of ELES. It's designed to increase the capacity of existing transmission lines, power transformers, and phase-shift transformers by providing important information for power system operation and planning based on real-time and forecasted weather conditions. It significantly improves grid resiliency, safety, and reliability, mitigating hundreds of network events annually, and identifies higher transmission capacity 92-96% of the time, often increasing nominal capacity by 15-20%.

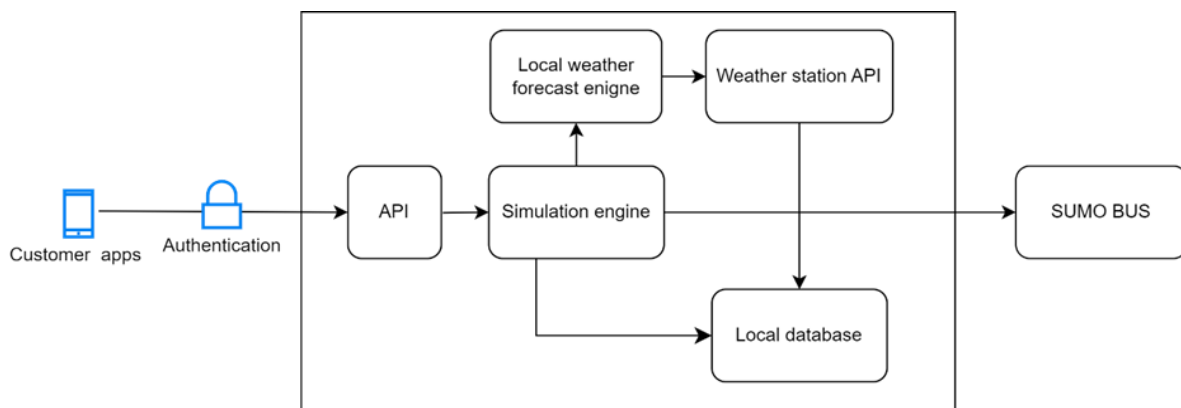


FIGURE 11 – DLR/DTR ON THE EDGE ARCHITECTURE DIAGRAM – JSI

3.2.2 Innovative aspects

Innovative aspects of this service have already been described in D3.3.

The primary innovative aspect is that this service 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.

Additionally, the deployment within this pilot will demonstrate the viability of edge computing for grid applications.

3.2.3 Service data

The service relies on two types of data for the input:

- **Weather data** consisting of historic data for a specific site, and current weather measurements.

They serve as the input and training data for the ML weather forecast algorithms and as the possible input for DLR/DTR algorithms. There are 2 subsets of this data, one for the DTR and one for the DLR case, and they may vary in the number of variables they include. The whole list of variables is presented in Appendix (3.2).

- Currently, for the DTR case, historic data is being used for developing algorithms, and full datasets are expected to be available when the algorithms are deployed by M26.
- **Operational data** from the observed line (DLR case) or transformer (DTR case), which serves as an input for the DLR and DTR algorithms. They consist of the real-time electrical current.

Currently, with historic measurements being used and live data is expected to be available when the distribution and transmission system operators deploy the IoT devices in their grid.

3.2.4 Integration with HEDGE-IoT interoperability framework

It was noted in D3.3 that integration with the HEDGE-IoT Interoperability Framework was under consideration. At this stage, it remains unclear whether such integration will be feasible, as the service currently relies on a dedicated SUMO bus developed by ELES/Operato for communication with the outside world.

The SUMO bus acts as a middleware layer that enables seamless communication between edge devices, local algorithms, and other distribution and transmission operators' systems. The DLR/DTR on-edge service uses this bus for data exchange and synchronization between edge-based processes and cloud-based applications. By implementing standardized protocols such as IEC 61850 or MQTT, the SUMO bus supports the transmission of critical grid information, including ampacity calculations and thermal capacity updates. This setup enhances the scalability and operational efficiency of the solution, supporting both DLR and DTR use cases and enabling functionalities such as predictive maintenance and system optimization.

3.2.5 Implementation details

Since D3.3, the DLR and DTR algorithms were adapted for the use on the edge device and successfully deployed on the chosen edge device IoT maxx GW- 4100 gateway in the lab

environment. Most of the data specifications and test sensors that will be used in the field in the later project stages have been obtained, and we are currently working on integrating the DLR/DTR algorithms on the IoT device with the sensors and the SUMO bus.

In parallel, we have also started developing the ML-based local weather forecast algorithms. We are currently testing different configurations. We have also obtained some historic weather data for relevant locations which is used to train and validate the algorithms. Moreover, we have explored the options for porting the algorithms to the IoT device. Due to the inherent limitations in computing power and storage capacity on such devices, the final implementation may rely on algorithms that, while not the top ones in terms of performance, still offer satisfactory results and are feasible to deploy on resource-constrained hardware. The development is ongoing.

3.2.5.1 Functionalities

Table 10 describes the functionalities of this technology enabler.

TABLE 10 – DLR/DTR ON THE EDGE FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
DLR/DTR algorithms	The DLR and DTR algorithms were adapted for running on the IoT device and deployed on the IoT device	100%	M20
Integration with the edge sensors and SUMO bus	The data from the edge sensors will serve as an input for the algorithms. The SUMO BUS will connect this service to the outside world.	80%	M20
ML-based local weather forecast algorithms	The ML-based local weather forecast algorithms will be developed for this service and trained on the historical data.	80%	M21
Deployment of ML-based local weather forecast algorithms	The ML-algorithms will run on the edge, while computationally intensive training of global models will be conducted on central servers.	30%	M26

3.2.5.2 Integration and dependencies

- **External dependencies**

- o Weather station service that provides ambient temperature, air pressure, wind speed and direction, solar irradiance, and precipitation.

- o Device temperature service that provides temperature readings from the sensors installed at various locations on the transformer.
- **Integration with the data space**
 - o To be determined.

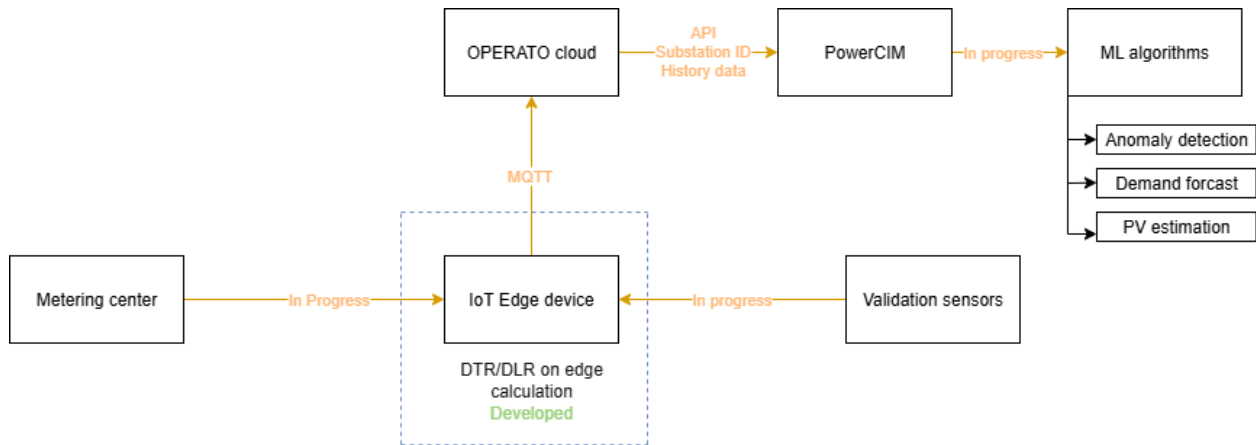


FIGURE 12 – DLR/DTR ON THE EDGE PILOT SGAM DIAGRAM – JSI

3.3 ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID

3.3.1 Description of technology enabler

The *Anomaly Detection* and *Predictive Maintenance* services provide real-time monitoring of graph data streams between smart devices, aiming to detect technical failures and other anomalies within the system on time, providing real-time alerts and explanations. Both services are part of *AI for Local Grid Resilience*.

Since the Predictive Maintenance service is an extension to the Anomaly Detection service and cannot function on its own, both services are treated as a single service for the purpose of this document. Figure 13 depicts how the two services interact, as well as depict a full overview of their architecture and the integration with the Semantic Interoperability Framework (SIF). A full description can be found in D3.3, under service 3.3.

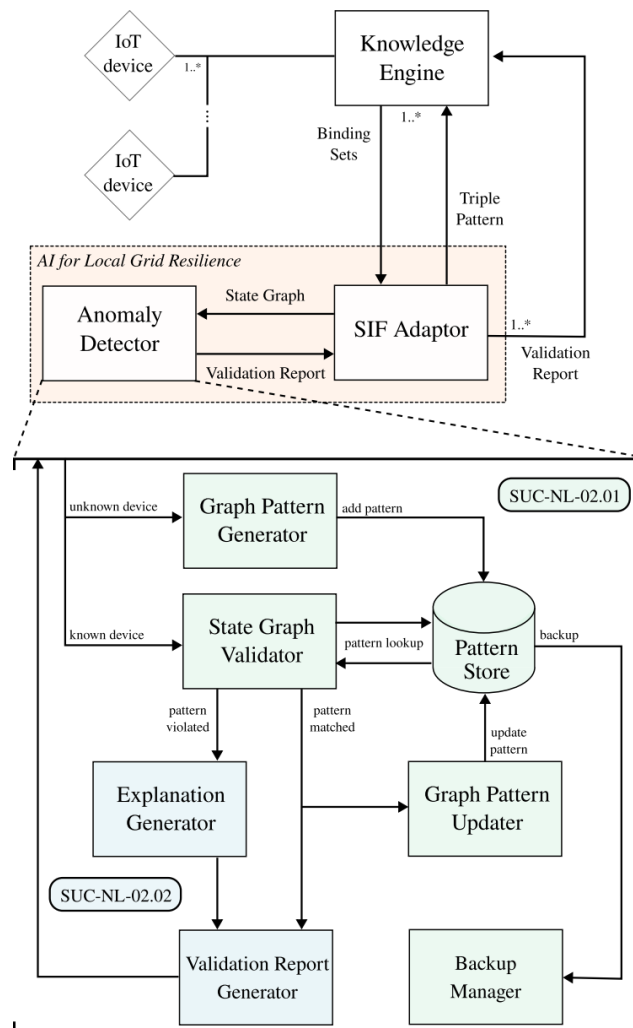


FIGURE 13 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID ARCHITECTURE DIAGRAM – VU

3.3.2 Innovative aspects

CONTINUOUS ONLINE LEARNING FROM STREAMING GRAPH DATA

Learning methods for graph data generally assume the availability of a static and complete dataset, which is loaded into memory and learned over. This is an unrealistic assumption for learning with smart devices in a distributed IoT environment, where data is dynamic, potentially incomplete, and is observed via discontinuous streams. The method provided by this service is specifically designed to accommodate this situation.

COMBINED SYMBOLIC AND SUB SYMBOLIC LEARNING

Learning methods for graph data generally employ symbolic or sub symbolic learning techniques, yet only few support both types of techniques in the same model. The method developed for this service is capable of integrating both techniques, allowing for the detection of anomalies in the relational components of the graph (symbolic) as well as in its attribute values (sub symbolic).

BOTTOM-UP AUTONOMOUS PATTERN RECOGNITION

The method provided by this service learns the nominal behaviour bottom-up, directly from the data streams, without the need for explicit training datasets or input from stakeholders and 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.

3.3.3 Service data

This service employs online continuous learning by listening for incoming graph data via SIF. There is therefore no training in the traditional sense of the word, but rather continuous updating of model parameters as new data is observed. Likewise, there is no dataset for training nor for testing. For this reason, the next part will provide descriptions and examples of the expected input and output data.

Both the input and output are streams of graph data in the *Resource Description Framework* (RDF) format². The RDF format is an open W3C³ standard in which knowledge, information, and data are encoded using binary statements, often called triples. Triples relate an object o to its subject s via a relationship p , and are represented via the infix notation: s, p, o . Here, both subject s and relationship p are denoted via *Universal Resource Identifiers* (URI) which are a generalisation of the URL, whereas object o can be either a URI or a raw value, called a literal.

To facilitate reuse and interoperability, the RDF best practices encourage the separation of schema-level and instance-level data. The schema-level data can then be published online and reused by others who are in need for knowledge frameworks to model their instance-level data with. More often, these shared frameworks are called ontologies or vocabularies. An example of a

² <https://www.w3.org/TR/rdf12-concepts/>

³ <https://www.w3.org/>

vocabulary for modelling energy use in IoT environments is SAREF⁴; a vocabulary for modelling anomalies is SHACL⁵. The following examples demonstrate their use.

The following graph (Text 1) is an example of the data that is expected as input, using the Turtle serialization format⁶:

```
@PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@PREFIX saref: <http://saref.etsi.org/core/> .
@PREFIX xsd: <http://www.w3.org/2001/XMLSchema#> .
@PREFIX ex: <http://example.org/demo/> .

ex:Uwmi0dxkgf8qjrd9bmsjc rdf:type saref:Observation .
ex:Uwmi0dxkgf8qjrd9bmsjc saref:madeBy ex:Uqwodyj51zp144lnzuqtw> .
ex:Uwmi0dxkgf8qjrd9bmsjc saref:hasTimestamp "2024-12-30T16:55:33.257238"^^xsd:dateTime .
ex:Uwmi0dxkgf8qjrd9bmsjc saref:hasResult ex:Uwm2wgp669iloqrcjz8j7> .

ex:Uwm2wgp669iloqrcjz8j7 rdf:type saref:PropertyValue .
ex:Uwm2wgp669iloqrcjz8j7 saref:isValueOfProperty ex:Energy .
ex:Uwm2wgp669iloqrcjz8j7 saref:hasValue "-1.93"^^xsd:float .
ex:Uwm2wgp669iloqrcjz8j7 saref:isMeasuredIn "kWh"^^xsd:string .
```

Text 1: Example of input data - a single state graph

The following graph (Text 2) is an example of the data that is expected as output, using the Turtle serialization format:

```
@PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@PREFIX sh: <http://www.w3.org/ns/shacl#> .
@PREFIX ex: <http://example.org/demo/> .

ex:report rdf:type sh:ValidationReport ;
sh:conforms false ;
sh:result
[ rdf:type sh:ValidationResult ;
sh:resultSeverity sh:Violation ;
sh:sourceConstraintComponent sh:InConstraintComponent ;
sh:focusNode ex:wmi0dxkgf8qjrd9bmsjc ;
sh:value :Unknown ;
sh:resultMessage "Observed value deviates too much from expected behaviour:
Expected - 26.5
observed - 32.6 (27.3% > 10%, p = 0.013)"]
```

Text 2: Example of output data - a single validation report

3.3.4 Integration with HEDGE-IoT interoperability framework

The Anomaly Detection and Predictive Maintenance services fully support the SIF (Semantic Interoperability Framework) and depend on its functionality for routing and transcoding the data published by the available smart devices. Validation reports made up about these data are likewise published to and routed by the SIF.

⁴ <https://ontology.tno.nl/saref/>
⁵ <https://www.w3.org/TR/shacl/>
⁶ <https://www.w3.org/TR/turtle/>

3.3.5 Implementation details

All core functionalities of the service have been developed and are now being tested and fine-tuned, and preparations are underway for validation in the controlled environment. This validation is scheduled to be completed in M22. The remaining months will focus on the validation with real-world data at the Dutch pilot site in Arnheems Buiten.

3.3.5.1 Functionalities

Table 11 lists the core functionalities from the combined service together with their implementation status.

TABLE 11 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Learn nominal patterns	Learn patterns of nominal behaviour for data from unknown devices.	100%	M12
Update learned patterns	Update learned patterns of nominal behaviour based on newly observed data from known devices.	100%	M14
Validate observed state	Validate incoming graph data from known devices against learned patterns.	100%	M17
Integration with SIF	Adaptor to SIF which facilitates data receiving, translating, and sending.	90%	M14
Generate validation report	Generation of validation report in SHACL format given new observations from known devices	80%	M18
Generate explanation	Generation of interpretable explanations for the detected anomaly in natural language.	50%	M23
Backup management	Periodically writing the learned patterns to disk and restoring them after shutdown.	10%	M24

3.3.5.2 Integration and dependencies

Integration

- *Semantic Interoperability Framework*

The service subscribes to various smart devices via SIF and continuously listens to new observations from these devices (routed by the Knowledge Engine runtime(s)⁷). Validation reports are likewise published via SIF, which are then forwarded to subscribed consumers. The SIF is available on premises and transcends various organizational layers due to the reach of the local area network (Figure 14).

- *Smart Dashboards*

The service publishes validation reports via SIF which can be received and displayed on the various smart dashboards and related interfaces, enabling stakeholders to take note of detected anomalies and take action. Both dashboards and related interfaces are part of the on-premises local grid Energy Management System (EMS), as depicted in Figure 14.

- *Smart Devices*

The services continuously listen to messages published by the smart devices it is subscribed to via SIF. Each known device is registered by their identifier in the service. A selection of currently available devices is shown in Figure 14, as part of the Building Management System (BMS) at the customer's premises.

Dependencies

- *Semantic Interoperability Framework*

The service is dependent on SIF to route and convert the messages sent by the subscribed smart devices. The service also requires the Knowledge Engine runtime(s) to keep track of the devices that are available on the network.

Pilot Architecture

The following SGAM diagram (Figure 14) depicts the pilot architecture.

⁷ <https://knowledge-engine.eu/>

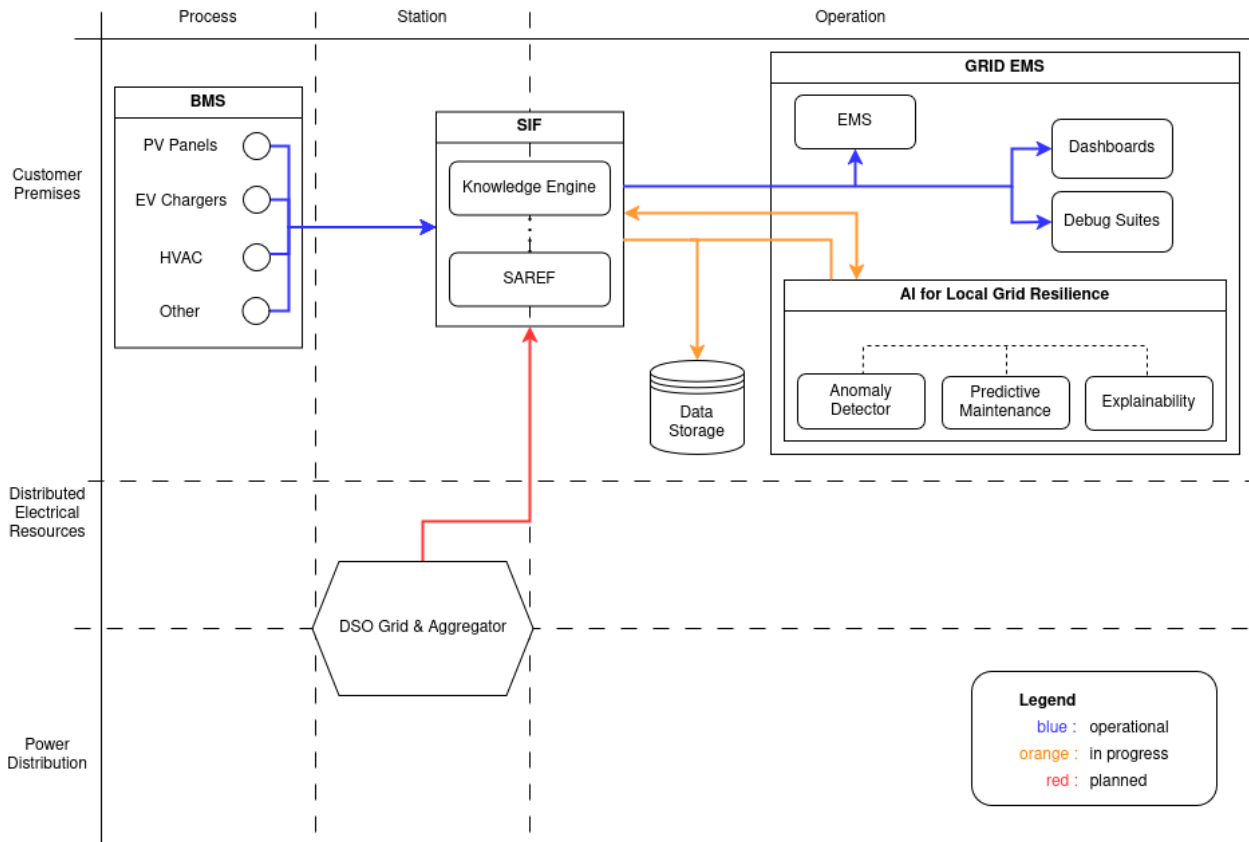


FIGURE 14 – ANOMALY DETECTION AND PREDICTIVE MAINTENANCE ON THE GRID PILOT SGAM DIAGRAM – VU

3.4 ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE

3.4.1 Description of technology enabler

A detailed description of this method is provided in previous iteration of this deliverable (D3.3). Therefore, Section 3.5 of D3.3 is referred to as a comprehensive overview of the service.

The core functionality of this service remains unchanged in deliverable (D3.4). The service aims to enhance grid operators' ability to anticipate and respond to disturbances more effectively by leveraging high-resolution data streams at the substation bus level—data typically underutilized due to transfer and storage limitations. The system acts as an early warning tool, using a traffic-light model (green-yellow-red) to alert operators to abnormal conditions, allowing preemptive actions such as reconfiguring grid topology or conducting inspections. Incorporating this hybrid long short-term memory (HLSTM) mode, the system adapts over time based on operator feedback, aiming to detect deviations from normal grid behavior without needing to identify exact fault types in advance (see Figure 15).

The HLSTM model will be continuously trained, refined, and optimized using large datasets and streaming data with a focus on its initial integration with the pilot end user.

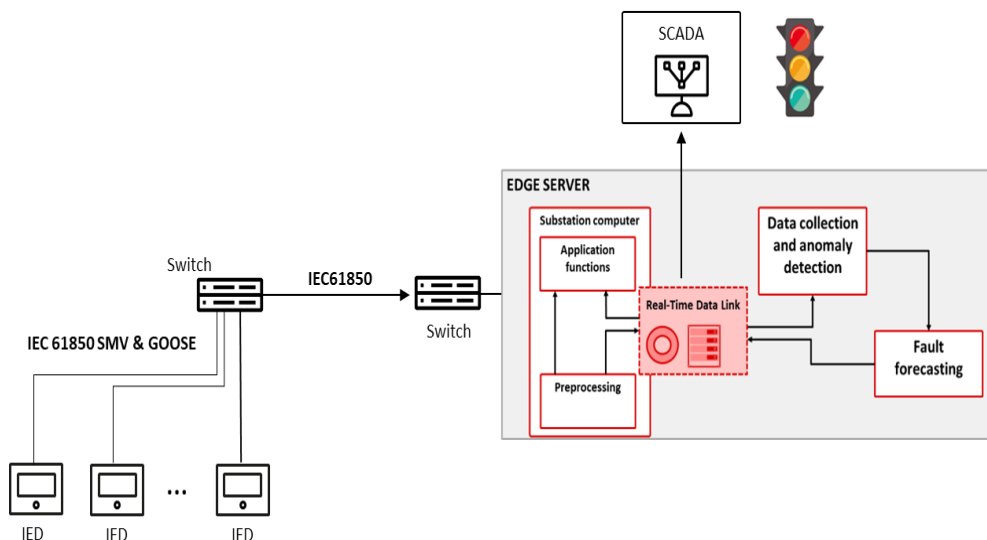


FIGURE 15 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE ARCHITECTURE DIAGRAM – VTT

3.4.2 Innovative aspects

In the previous iteration of this deliverable (D3.3), the Hybrid Long Short-term Memory (HLSTM) algorithm implements an initial pipeline for spotting outliers in time-series data, visualizing them, persisting the results, and then training the model to both classify and forecast anomalous events. In the current iteration (D3.4), various following functions have been modified for the HLSTM algorithm.

The HLSTM model has been refined with a layered, hierarchical architecture, enabling it to effectively capture both short-term fluctuations and long-term dependencies in multivariate time-series data—particularly electrical signals such as current and voltage measurements observed in substations. This structural improvement enhances the model’s ability to accurately distinguish between normal operational variations and real anomalies with high accuracy at low latency. The updated HLSTM model constructs temporal sequences directly from raw electrical parameters such as IL1-IL3 (current) and U1-U3 (voltage). By learning from these parameters without manual feature engineering, the model exhibits strong adaptability to handle various datasets and configurations. This self-learning capability reduces development effort while improving flexibility and scalability across diverse environments.

In addition, the HLSTM model has been enhanced to not only detect current anomalies but also forecast future deviations. In the updated HLSTM, predictive functionality allows early detection of faults, which reduces the system downtime. The model’s design remains lightweight and computationally efficient, making it suitable for deployment in various big data scenarios. The combination of hierarchical learning, direct feature extraction, continuous online adaptability, and predictive intelligence makes the HLSTM model a comprehensive and innovative solution for anomaly detection and fault forecasting in energy systems.

3.4.3 Service data

In the previous iteration of this deliverable (D3.3), the datasets

- (DS1) provided in M12 originates from a feeder or distribution center system and captures fault-related measurements, including phase currents and voltages. In particular, the datasets record phase currents for the A-phase, B-phase, and C-phase, denoted as IL1, IL2, and IL3, respectively. It also includes measurements for neutral current (I_0) and voltages U_0 , U_1 , U_2 , and U_3 , all expressed in kilovolts (kV). These values represent various electrical parameters such as zero, positive, and negative sequence currents and voltages, as well as neutral, ground, and residual currents and voltages. The datasets are structured such that the first line contains metadata including a device ID, an IP address, and a timestamp indicating the year of recording. The data types found in the dataset include floating-point numbers, fixed precision numbers, binary values, and textual information. These may correspond to real-time measurements, event logs, configuration parameters, or raw binary data from monitoring equipment. Using the designed algorithm, the datasets (Comtrade format) are converted into .CSV files and normalized, so that algorithm can read the datasets files.

In the current iteration (D3.4), new datasets

- (DS2) have been provided for the purpose of further algorithm training, analysis, and optimization in M16.

Live data streaming from the simulation setup on Real-time Discrete Simulator (RTDS) will be configured and integrated with the HLSTM model, probably in M22 (see Figure 16).




FIGURE 16 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE TIMELINE OF DATASETS – VTT

3.4.4 Integration with HEDGE-IoT interoperability framework

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

3.4.5 Implementation details

The core functionalities of the tool and all required modules have been defined, and implementation has been started on small scale datasets. In the next phase, the HLSTM algorithm will be deployed on the server and will undergo training on large-scale datasets. Next, a simulation use case(s) will be configured on the RTDS and input will be directly fed up to the developed algorithm.

Consequently, the algorithm will then be trained and optimized using real-time data streams generated from the use case(s) on RTDS, which will replicate the data produced at the actual pilot site as shown in Figure 17. The final functionalities will be tested in lab environment by M30 and will be taken into use at the piloting areas by M32.

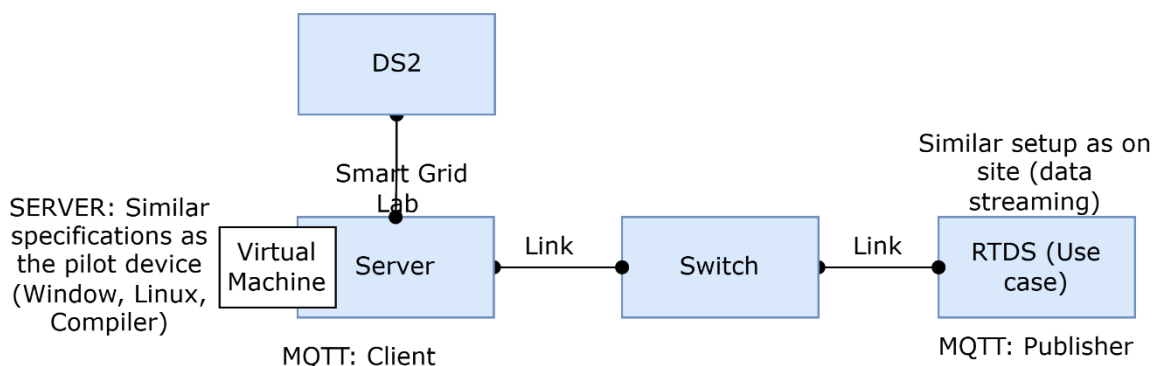


FIGURE 17 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE: NEXT-PHASE OF IMPLEMENTATION – VTT

3.4.5.1 Functionalities

Table 12 shows the list of functionalities comprehended by this service.

TABLE 12 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Data Analysis and Validation	Small-scale data analysis and validation.	100%	M13
LSTM Model	LSTM model was designed and running on the small-scale datasets for anomaly detection and fault forecasting.	100%	M14
HLSTM Model	LSTM algorithm was enhanced by integrating sequential model features and running on the small-scale datasets for accurate anomaly detection and fault forecasting.	100%	M16
Big datasets Integration	HLSTM algorithm for running on new large-scale datasets for more accurate anomaly detection and fault forecasting.	60%	M22
Live Data Streaming	HLSTM algorithm for running on live data streaming for anomaly detection and fault forecasting.	40%	M30
Service Integration	By utilizing standardized communication protocols, including IEC 61850, the service registered and packaged in the HEDGE-IoT.	0%	M32

3.4.5.2 Integration and dependencies

At the current stage, the service operates independently of other components in the Finish pilot architecture. However, in the next phases, it may require access information from connected substations or feeders that generate data in real time in the network. Figure 18 shows an architecture diagram that maps the integrations and dependencies of the service with the remaining Finish pilot services, as well as showcases the status of the implementation of each integration.

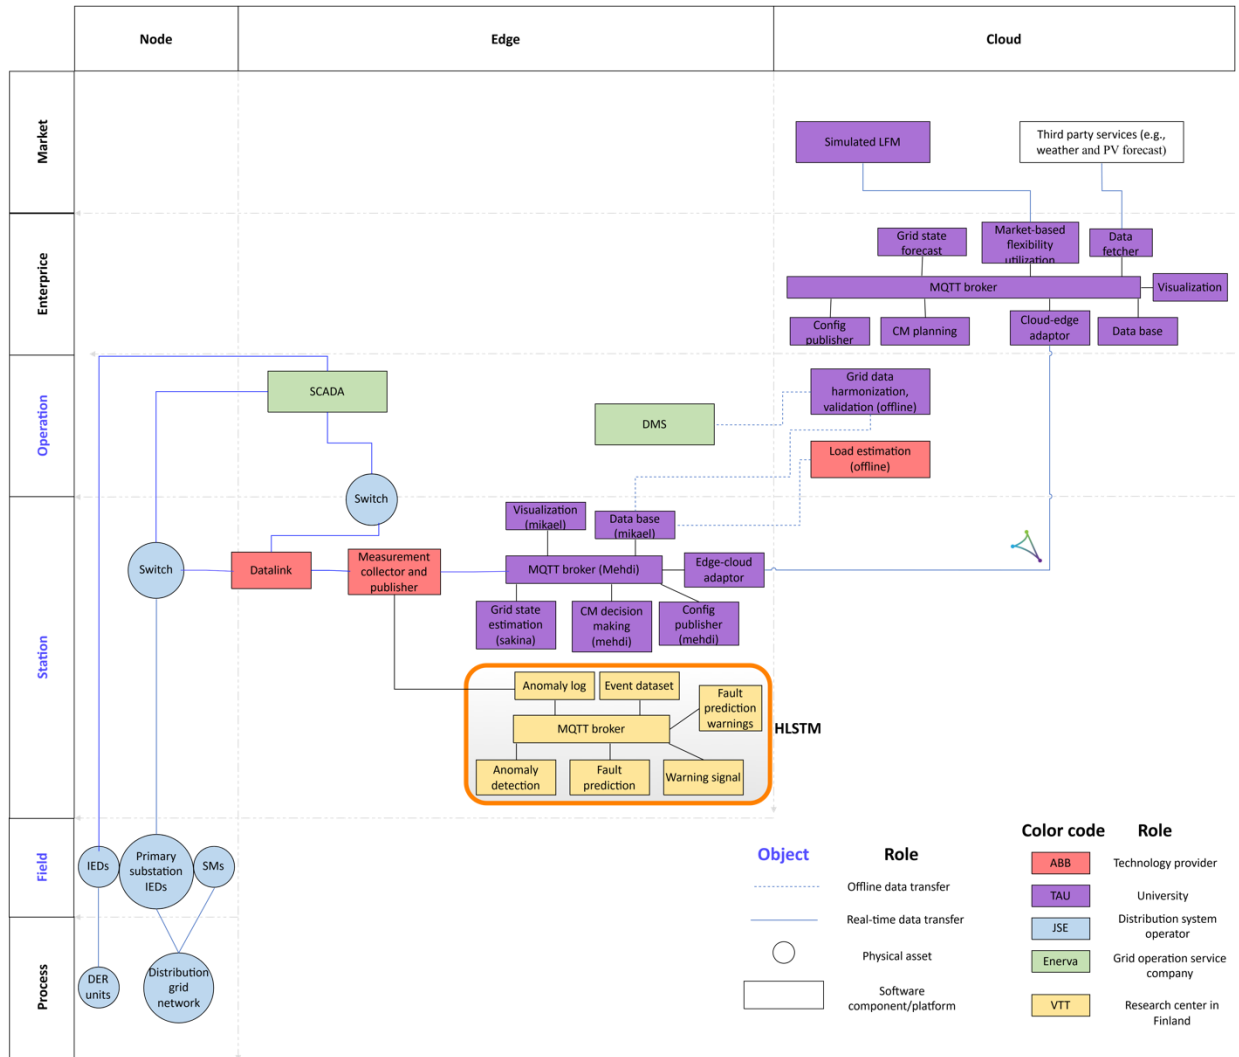


FIGURE 18 – ANOMALY DETECTION AND FAULT FORECASTING TO INCREASE DISTRIBUTION NETWORK RESILIENCE PILOT SGAM ARCHITECTURE

3.5 REAL-TIME CONGESTION MANAGEMENT

3.5.1 Description of technology enabler

Figure 19 illustrates the architecture of real-time congestion management (CM) service and its micro services. The real-time in this context is in minutes resolution. The real-time CM service of SUC-FI-02.03 and SUC-FI-02.04 have been divided into several micro-services – such as state estimation, and CM decision making – following a service-oriented design. For data exchange between edge and cloud, the edge-cloud adaptor is used.

Compared to the previous version detailed in D3.3 the following progress/changes have been made:

- A config publisher has been added to the architecture, which helps to parameterize and synchronize all the subservices involved in CM service.
- The MQTT protocol has been adopted for the message bus.
- The real-time measurements are passed to the CM micro services through a new component “measurement collector and publisher”. The component accesses the real-time measurements coming from primary substation and SCADA and forwards their RMS values to the CM micro services over MQTT protocol.

With the advances on the architecture diagram as mentioned above, the development is now moving to implementation of components and interfaces. In nutshell, among the micro services, state estimation and measurement collector and publisher are currently in progress with the plan to start working on congestion management decision making from August 2025. The work on config publisher, edge-cloud data exchange adaptor and database will be started before January 2026.

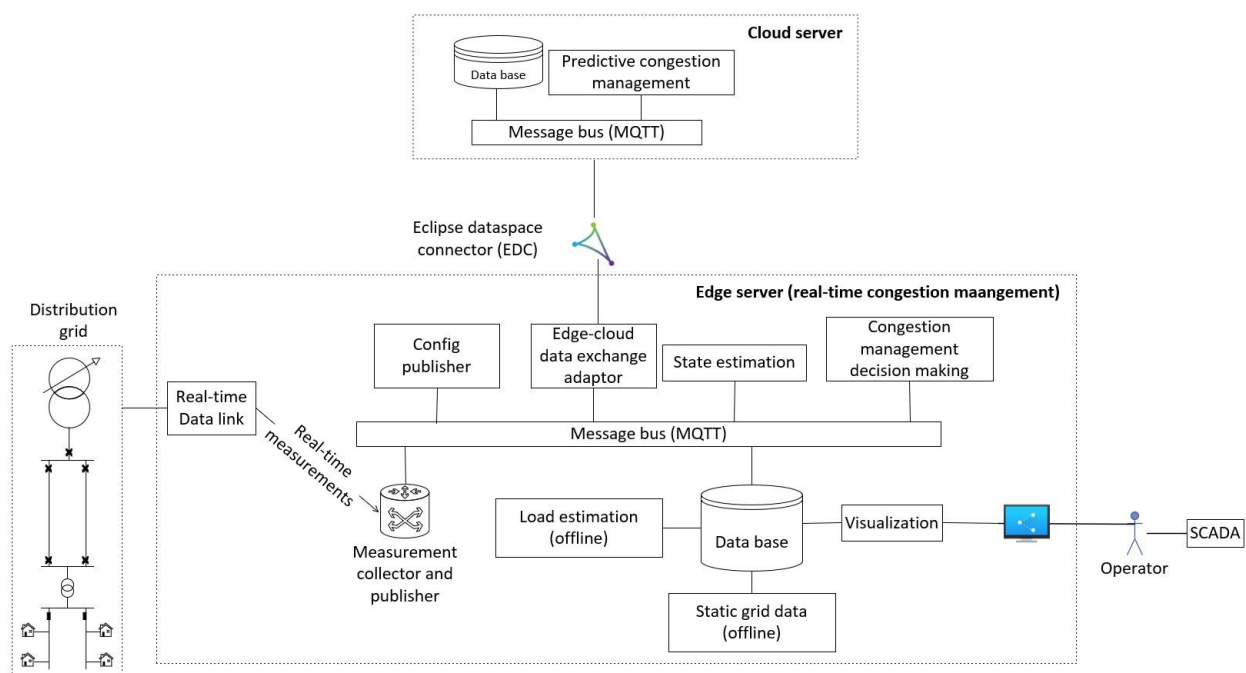


FIGURE 19 – REAL-TIME CONGESTION MANAGEMENT ARCHITECTURE DIAGRAM – TAU

3.5.2 Innovative aspects

ENHANCING PRACTICALITY OF REAL-TIME CM FOR REAL-WORLD APPLICATIONS

Most DSOs have not yet adopted active grid management to its full potential, particularly in the area of real-time CM, due to historic regulatory preferences for passive solutions, but EU policies since 2019 have begun to shift incentives toward active approaches. Adoption remains slow because of low motivation and complex implementation requirements, though increasing grid stress from renewables is encouraging DSOs to pursue a hybrid model combining active and passive solutions. While real-time CM is well-studied in academia, practical, industry-ready implementations remain scarce, and this pilot bridges that gap through collaboration between academic, technological, and operational partners in Finland.

MICROSERVICE INTEROPERABILITY

The proposed service draws on experience from prior EU projects like INTERFACE and IDE4L, integrating both market-based and grid-side flexibility into a unified, innovative active grid management approach. These were the innovative aspects mentioned earlier in D3.3. In addition to them, under the transversal use cases defined in T2.6, the interoperability of some of the micro services developed in the pilot will be examined in T3.5, which is expected to generate new insights in this area. Improved interoperability would enable subservices to be exchanged between DSOs based on their specific needs and capabilities—ideally through a friction-free process.

3.5.3 Service data

The service will use real-time data as a 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 as well as setting of components like on load tap changer at the primary substation. For the measurement data, root means square (RMS) is calculated first, and then published to MQTT broker through measurement collector and publisher component. The CM micro services, in particular state estimation, will subscribe to those data. For components' settings such as on load tap changers, those data are published to the message broker to be used by CM micro services, in particular CM decision making. Progress in this area has been moderate, meaning that the server has already been configured and sent to the pilot site, but still the interfaces must be defined and implemented. The plan is to define and implement the interfaces from August 2025 onwards, so that by the end of 2025, the initial interfaces could be tested.

Harmonization, validation and analysis of static data sets have been almost completed. Those data are as follows:

- **Load data for the last three years (2022,2023, and 2024)**
 - The load data has been obtained from smart meter measurement service provider, aggregated to secondary substation level (i.e., anonymized), and validated.

- **Grid data in QGIS format**

- Grid data has been converted to power technologies international (PTI) format useful for load flow calculations. The converted grid data has been harmonized and validated (there are still minor modifications needed).

- **Generation data**

- Generation data of tens of rooftop PVs, one solar park, and one hydro power plant has been collected and associated with the grid data.

The above-mentioned data sets have gone through analysis to assure that the previous steps for load, generation and grid data have been done correctly. The analysis has helped to find minor errors (some have been already addressed, and a few needs will be addressed soon).

For the dynamic data originating from third parties, for example ambient temperature which is necessary for state estimation, implementation has not yet started. On that regard, the discussion on whether the data is obtained from the cloud through eclipse data space connector is still in progress. As the state estimation subservice requires weather data, it is planned that weather data, its origin, how to access it and its data model be specified in August 2025, so the implementation of it can follow afterwards.

For the data that is the outcome of the real-time CM, the initial idea is to visualize the state of the grid to the operator and recommend CM solutions. It is necessary to define how that could be implemented, something that should be decided once the micro services produce some meaningful results and therefore, this part is expected to be done in later stages of the pilot, possibly early 2026.

3.5.4 Integration with HEDGE-IoT interoperability framework

The service will utilize Eclipse Dataspace Connectors (EDC) for edge-to-cloud and cloud-to-cloud data transfer, providing a reusable and interoperable framework via standardized APIs. The initial plan is to experiment with EDC by exchanging arbitrary data between edge and cloud environments. Once the edge and cloud services reach a near-ready stage, they will be connected through EDC. This experimentation phase is scheduled to begin in August 2025.

In parallel, the pilot will explore service interoperability by testing how a subservice—such as grid state estimation—can be pushed, pulled, and orchestrated via an app store environment.

3.5.5 Implementation details

Compared to the previous evaluation cycle, the pilot has progressed on different fronts including testing method, testing infrastructure and location, implementation of the testing environment, dataset harmonization validation and analysis, and algorithm implementation. The following explains the progress of each part in more detail:

- **Testing method**

- The partners have agreed on the method of testing. Testing will have three phases of unit testing, integration testing and system testing. Unit testing and integration testing will be done in isolation by ABB and TAU, and after completion, system

testing could be done so that ABB and TAU's systems are connected. The fully tested systems then will be finally migrated to the Järvi-Suomen Energia's pilot site.

- **Testing infrastructure and location**

- Unit and integration testing will happen at ABB and TAU separately, but system testing will be conducted in ABB's premises.

- **Algorithm development**

- Grid state estimation algorithm development has started. The CM decision making algorithm development will start in August 2025. Those two algorithms perform the main functionalities on the edge.

Specific developments of the current cycle (common with SUC-FI-02.01 & SUC-FI-02.02) concerning the data sets

- Demonstration grid's data has been migrated from the QGIS software (used by Järvi-Suomen Energia) to MATLAB to model it based on state estimation needs (PTI format).
- The missing data of grid components in QGIS (e.g., transformer impedances, cable impedances) has been found from component's documentations and added to the grid data model.
- The grid data has been validated meaning that grid components are modelled correctly.
- Historical load data (hourly resolution) has been collected and aggregated (secondary substation level) from smart meters in the demonstration area from 2022 to 2024.
- Historical generation data (hourly resolution) has been collected from their measurement devices from 2022 to 2024.
- Individual load profiles, which can be used for load estimation and forecasting, were calculated for secondary substations, considering the time of day, calendar data, and outdoor temperature. These profiles model hourly substation loads and their variance over the coming year.
- Load and grid data have been associated with each other, and a grid analysis has been done to evaluate the state of the grid in previous years of grid operation (e.g., 2022, 2023, 2024).

Based on the progress made so far and the remaining tasks, it is expected that by month 30, unit testing and integration testing of critical grid algorithms including grid state estimation and CM decision making is completed. The system testing of them will follow right after integration testing, in the pilot site. Also, by M30, the use of EDC will be first experimented to connect cloud and edge systems using arbitrary data, and then useful data will be transferred between edge and cloud using the connector. Also from an interoperability perspective, some initial testing with respect to sharing and reusing the subservices (e.g., state estimation) in the app store is expected.

3.5.5.1 Functionalities

Table 13 indicates the functionalities and their timeline to be completed for the real-time CM service.

TABLE 13 – REAL-TIME CONGESTION MANAGEMENT FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Database	It stores static and dynamic data	0%	M30
State estimation	It estimates the state variables of the grid	50%	M24
CM decision making	It makes CM decision depending on the grid state and aims to remove congestion	20%	M30
Config publisher	It publishes parameters necessary for all components on the edge.	0%	M30
Measurement collector and publisher	It received the measurement data from the data link calculates the RMS values and publishes them to the MQTT broker	20%	M30
Cloud-edge adaptor for using EDC	It sends and receives data to/from cloud	0%	M30
Integration with app store and orchestrator	Sending and receiving micro services through app store	0%	M30

3.5.5.2 Integration and dependencies

As shown in Figure 20, the real-time CM that is located on the edge, at the station level, has to connect to the cloud server where predictive CM is located through EDC. Therefore, the edge and cloud servers are dependent on each other, although the dependency has to be mitigated as much as possible in case of unavailability of EDC (N-1 principle), so those two servers can ideally work independently if situations demand that. In addition, the real-time CM service is dependent on real-time measurement data coming from ABB’s technology represented by components including Data link and measurement collector and publisher. The following section describes how the real-time CM service integrates with other services in the Finnish pilot:

1. [Service 4.4] Predictive CM on the cloud

- a. Receives ambient temperature data.
- b. Receives, component settings

2. **Real-time measurements (through measurement collector and publisher)**
 - a. Receives RMS value of real-time measurement of grid's state variables
 - b. Receives settings of controllable components (DERs, transformers, etc)
3. **EDC Connector**
 - a. Mainly receives but possibly publishes data from/to EDC

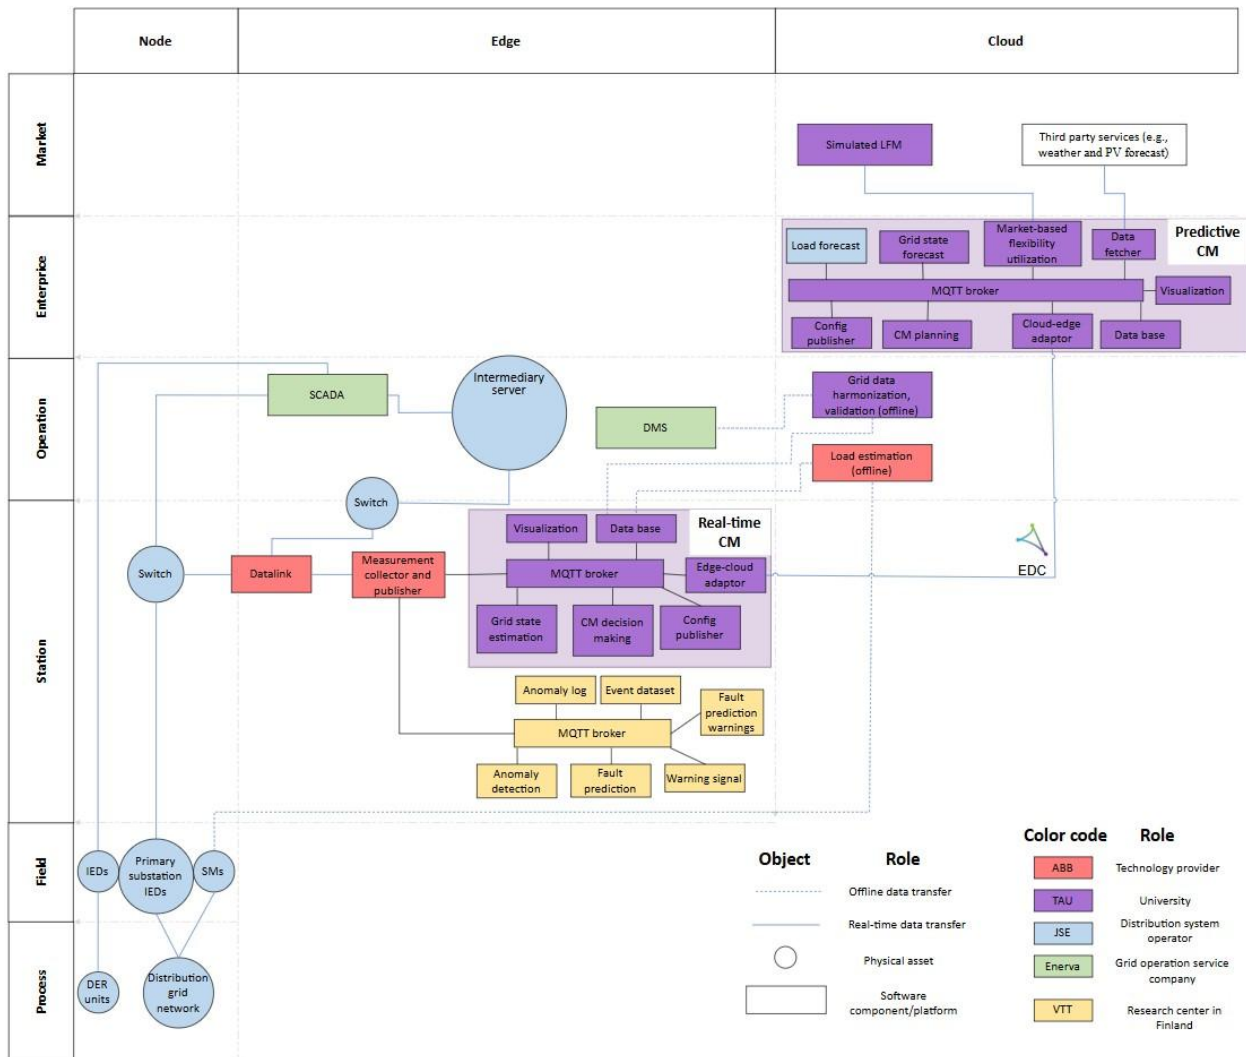


FIGURE 20 – REAL-TIME CONGESTION MANAGEMENT PILOT SGAM DIAGRAM – TAU

3.6 HEMS – HOME ENERGY MANAGEMENT SYSTEM

3.6.1 Description of technology enabler

EMSs (Energy Management Systems) play a key role in the flexibility enablement of consumers, residential and tertiary, which is paramount to accessing the previously untapped flexibility potential of residential DERs (Distributed Energy Resources) [5]. These resources, under the form of energy assets, are usually household appliances like heat pumps, EV chargers, dishwashers, PV inverters, batteries, etc. This is where the HEMS (Home Energy Management System) comes in. The goal of this technology enabler is to facilitate the user’s participation in the flexibility value chain, while providing them with incentives in a clear, explainable way.

To fulfill this goal in an effective and scalable way, the HEMS is designed with a micro-services architecture (Figure 21), orchestrated in a Kubernetes environment, where each micro-service is modular and can be replaced or expanded, without breaking the remaining logic. The entire suite of the HEMS is comprised of three applications:

- **HEMS Cloud:** Comprised by multiple cloud-based micro-services (account management, device management, flexibility optimization, etc.) that communicate with each other to provide all the backend logic to the HEMS suite. This is also where the “heavy lifting” is done for the flexibility optimization operations.
- **HEMS Hub:** An edge component (Raspberry Pi) that connects directly to the energy assets on the end user’s premises via low-level interoperable communication protocols (Modbus/OCPP).
- **HEMS Connect:** A mobile application, available for Android and iOS, that will provide a user interface for the PT pilot end users. This application will allow the user to manage their assets and to have a clear view of their flexibility incentives.

On the cloud front, the HEMS allows the user to discover their assets and integrate them, while, on the other end, integrating with RECreation, the Energy Community Management Platform (section 4.5) of the Portuguese pilot. This integration allows HEMS to receive activation signals for each asset, then factor in user comfort preferences and then relay the signals to the HEMS Hub for activation on the asset. Finally, during the settlement process, the HEMS will be responsible for sending the metering data to RECreation, to complete the flexibility process with the System Operator.

On the edge front, the HEMS Hub will receive an activation signal from the HEMS Cloud and relay it to the corresponding asset. At this moment, 3 different assets and manufacturers have been tested:

- **Daikin Heat Pumps:** Integration using a proprietary Modbus solution (DCOM⁸). The registries allow the HEMS Hub to, for example, change the temperature of the water.

⁸ https://www.daikin.eu/en_us/products/product.html/DCOM-LT/MB.html

- **Victron Inverter with Weko battery pack:** Integration using a Victron proprietary controller, connected via CAN and accessible using Modbus. The registries allow the HEMS Hub to, for example, change the power setpoint of the battery pack.
- **Wallbox EV charger:** Integration using the OCPP protocol, a very standard industry protocol for EVSE-EMS communication. This allows the HEMS Hub to, for example, set the maximum power flowing from the charger to the electric vehicle.

Lastly, the HEMS contains a service that collects residential household energy consumption data via a MQTT broker and a Python consumer, being directly fed into INESC's data lake.

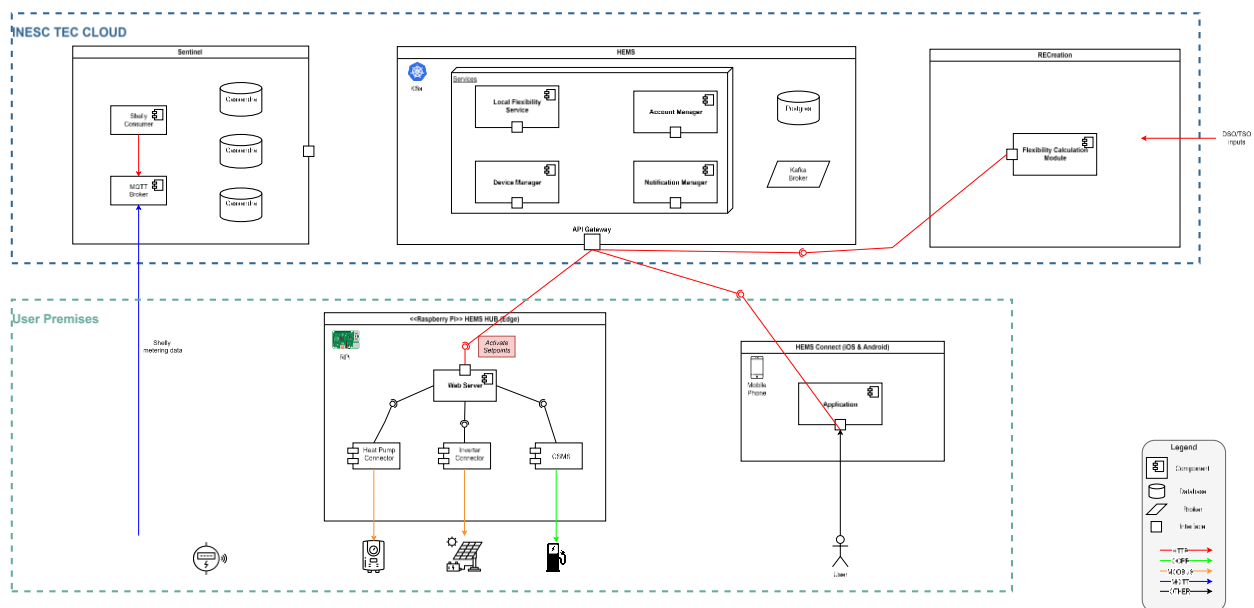


FIGURE 21 – HEMS (HOME ENERGY MANAGEMENT SYSTEM) ARCHITECTURE DIAGRAM – INESC

3.6.2 Innovative aspects

In this section, the innovative aspects of the service will be detailed.

FLEXIBILITY OPTIMIZATION

This is the core use case of HEMS. The innovation lies with the optimization algorithm that can combine user comfort constraints for the assets with the activation signals received from the System Operator, providing a smart schedule for any kind of appliance the user has associated with the HEMS (heat pumps, batteries and EV chargers).

INTEGRATION WITH ENERGY COMMUNITY AGGREGATOR

The HEMS provides an interface for seamless integration with the energy community aggregator of the PT pilot. This interface includes the programmatic registration and pre-qualification of the user's DERs (Distributed Energy Resources), the implementation of the activation signals and the metering information to complete the settlement process.

INTEROPERABILITY

The concept of interoperability is core for the HEMS application [6]. The idea is that this energy manager can include as many DERs as possible, from as many manufacturers as possible, as seamless as possible. For that reason, it implements well known communication standards, like OCPP for e-mobility and Modbus. Furthermore, HEMS Hub is deployed as an edge device to avoid the need to use wireless connections and third-party APIs, improving the interoperability and reliability of the system.

EXPLAINABILITY

One well-documented issue of Energy Management Systems that tackles the challenge of flexibility is explainability [7]. Namely, explaining to the end user the results of the model and why certain changes were made to the usage of their energy assets. For that reason, HEMS Connect was created and contains simplified information to facilitate the comprehension of the user regarding its flexibility changes. For example, an overview of financial incentives and disaggregated flexibility details for each asset.

3.6.3 Service data

Since this service does not contain any algorithm that requires training data, it does not store any datasets. However, it does store some data related to the users, devices and flexibility actions, namely:

- **User information:** Static information regarding the user's profile, such as email, energy tariff, city, etc.
- **Device information:** Manufacturer, model, type of device, comfort settings and online status.
- **Flexibility actions:** Financial incentives, activation signal impact on each asset and schedule of the activations.

There is no timeline for the data provided below, as these are not datasets, but rather data that will be acquired only once the pilot starts and the application is provided to the end users.

3.6.4 Integration with HEDGE-IoT interoperability framework

As shown on Figure 21, the HEMS application only integrates with RECreation (Energy community aggregator) and does not have any direct integration with the project's data space. Since the main goal of the HEMS is to integrate directly with the DERs of the user and they do not provide any connectors for the project's data space, the adopted interoperability approach focuses on technical interoperability using industry communication standards (OCPP and Modbus).

However, the idea is to integrate this technology enabler with the App Store of the project. At this stage, there are no guidelines for this integration, so it is still unclear how these proceedings will go.

3.6.5 Implementation details

At this stage, the tool stands at TRL 7 since a system prototype has been developed and demonstrated in an operational large-scale pilot in Portugal for the H2020 InterConnect project⁹. Since the beginning of this project, however, the application has suffered some changes that were planned from lessons learning from the InterConnect project’s pilot. Namely, the explainability functionalities and the interoperability without requiring wireless connections. These functionalities have been under development and integration until the current month (M19).

For M30, the expectation is that all the new functionalities will be implemented and tested, ready to be qualified in the large-scale pilot of the HEDGE-IoT project. After the large-scale pilot (M38), the expectation is that the KPIs collected allow the application to reach TRL 8.

3.6.5.1 Functionalities

Table 14 shows the list of functionalities comprehended by this service.

TABLE 14 – HEMS (HOME ENERGY MANAGEMENT SYSTEM) FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
User account management	Micro-service with functions related to account management, such as account creation, login, account deletion, password recovery, etc.	100%	M13
Flexibility optimization	Flexibility optimization algorithm that considers flexibility needs from the system operator and the comfort settings of the user to provide bids for the market.	60%	M25
Device management	Micro-service for the user to manage their DERs. Includes seamless device association using industry communication standards and user comfort settings	50%	M24
Flexibility actions explainability	Micro-service that takes the results of the flexibility optimization algorithm and transforms that output into explainable flexibility actions.	0%	M30
Mobile App (HEMS Connect)	Mobile app for Android and iOS that allows users to manage their	40%	M32

⁹ <https://interconnectproject.eu/news/portuguese-pilot-engages-with-customers-to-showcase-hems-app-capabilities/>

	DERs and their flexibility actions with ease		
Communication of activation signals to energy assets (HEMS Hub)	Integration of DERs into the HEMS Hub software and interface to communicate activation signals to them	50%	M28
Integration with the App Store	Service registered and packaged in the HEDGE-IoT App Store	0%	M32

3.6.5.2 Integration and dependencies

This section describes how the FL service integrates with other services in the Portuguese Pilot.

1. [Service 4.5] Energy Community Management Platform (RECreation):

- a. Receives activation signals, previously relayed by the TSO (Transmission System Operator)

2. App Store Integration:

- a. Deployed to the app store

Figure 22 shows a SGAM diagram that maps the integrations and dependencies of the service with the remaining PT pilot services, as well as showcases the status of the implementation of each integration.

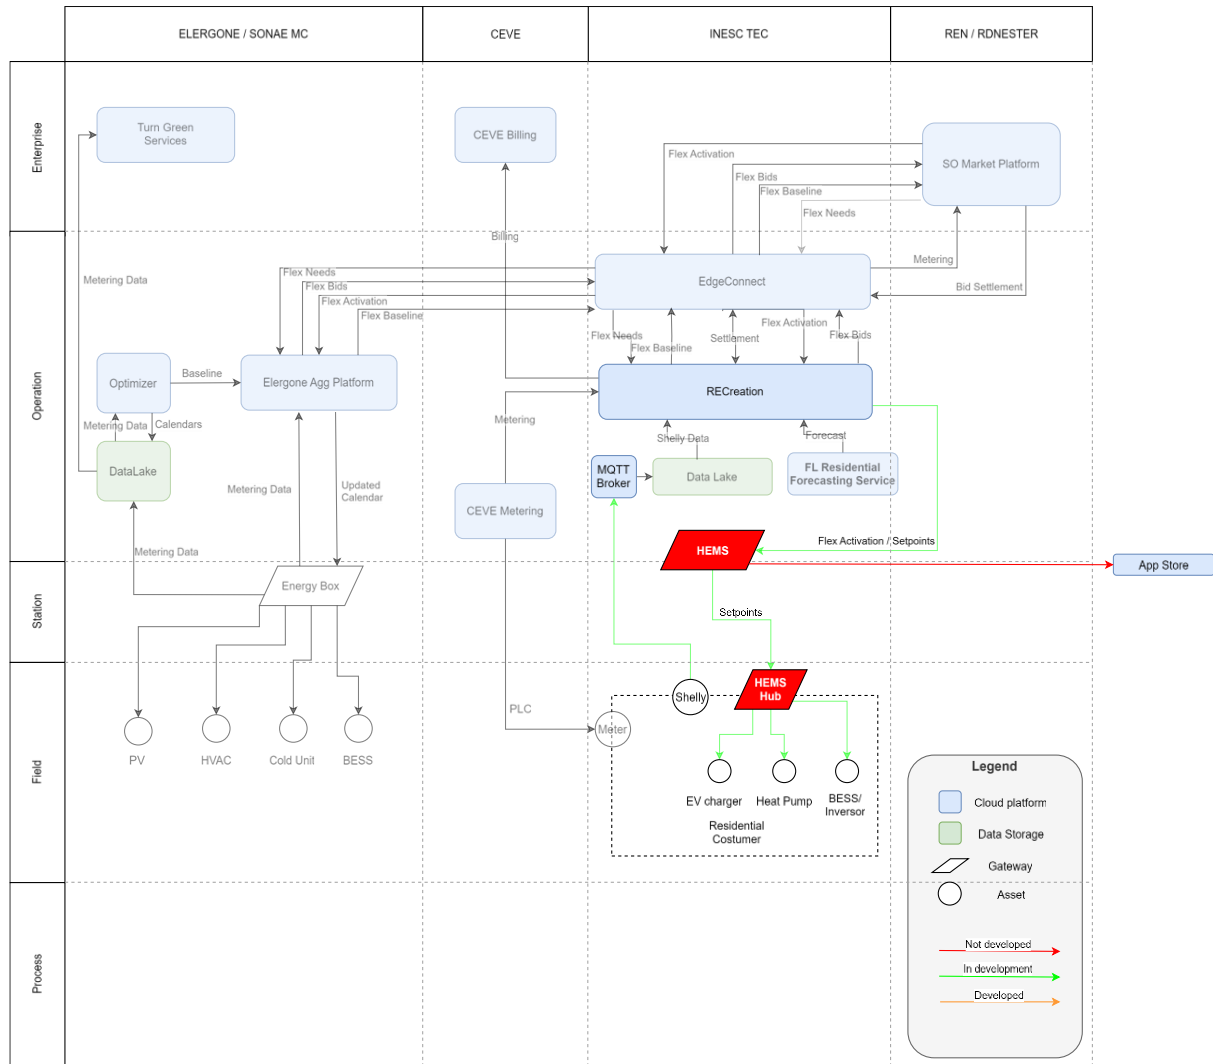


FIGURE 22 - HEMS (HOME ENERGY MANAGEMENT SYSTEM) PILOT SGAM DIAGRAM - INESC

3.7 DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION

3.7.1 Description of technology enabler

The edge platform for distribution grid automation enables sharing high volume data with strict real-time requirements between different applications that are executed on an edge server in a virtualization environment. The platform will have the capabilities to receive input data as IEC 61850-9-2 sampled values (SV), IEC 61850-8-1 Generic Object-Oriented Substation Event (GOOSE) and IEC 60870-5-104.

The data can then be shared onwards with different applications within the same virtual machine using real-time data link capabilities. Data streaming can be used for continuous delivery of data and is applicable especially for SV data that can be shared to different receivers either as raw or pre-processed data. Signal store can be used to share arbitrary signal data (supported formats: integer and floating-point numbers, Booleans, and limited-length strings) between applications.

The technology used to implement real-time data link capabilities is shared memory which enables fulfilling the requirements for strict real-time operation and large data volumes. Data sharing between different virtual machines or containers on the edge server is implemented using MQTT. It is possible to share both stream and signal data through MQTT but guaranteeing real-time characteristics will not be possible. For many applications, also this type of data exchange is adequate because not all applications have strict real-time requirements.

The edge platform is an enabler for power system use cases. Services 3.4 Anomaly detection and fault forecasting to increase distribution network resilience and 3.5 Real-time congestion management utilize the edge platform capabilities. Figure 23 represents the Finnish pilot architecture and the position of the edge platform as a part of it.

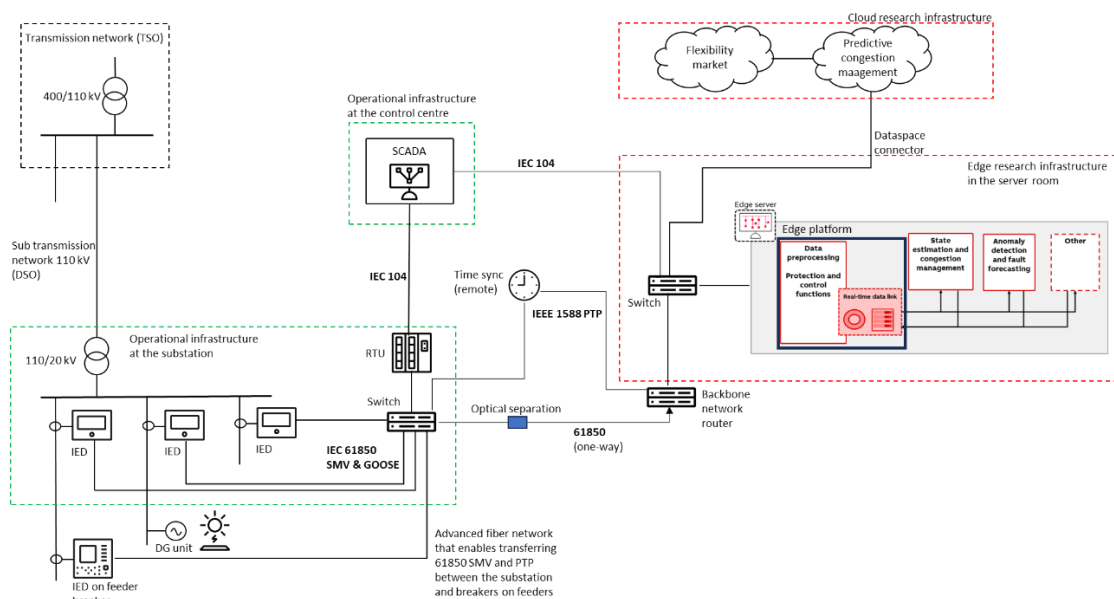


FIGURE 23 – DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION ARCHITECTURE DIAGRAM – ABB

3.7.2 Innovative aspects

INTEGRATING DIFFERENT FUNCTIONALITIES ON THE SAME EDGE HARDWARE

The changes in the energy system create a need to revise the existing protection, control and operation principles and to also add totally new functionalities that will, for instance, enable utilising flexibility from small-scale DERs both for distribution and transmission system needs. Currently each functionality is deployed as a separate device or cubicle at the substation. If also the new smart grid functionalities are deployed with this approach, the number of devices at the substation increases significantly. Another alternative is to utilise virtualisation technology and capabilities of digital substations with IEC 61850 process bus to enable combining several functionalities into one system. The edge platform capabilities enhance the options to integrate various functionalities on the same edge server hardware and to optimise usage of the computational resources and physical space.

EFFICIENT DATA SHARING

The platform capabilities enable efficient data sharing between different functionalities on the same hardware. Receiving and pre-processing of SV data does not need to be conducted separately by all functions needing the data, but the necessary parts of data can be shared between modules. Different modules can have their own cycle times and can be executed in their own running environments.

3.7.3 Service data

The edge platform itself does not use any datasets. Its role is to provide the necessary data to the actual power system services that are deployed on edge. The input data is received from intelligent electronic devices (IEDs) both at the substation and at the feeder breakers through IEC 61850 process bus. Measurement and status information from DER sites that do not have IEC 61850 process bus capable IEDs is obtained from the SCADA system using IEC 60870-5-104. The input data includes the following:

- Current and voltage measurements from each bay are obtained as IEC 61850-9-2 sampled values.
- Bay level breaker status information and control signals are obtained as IEC 61850-8-1 GOOSE.
- Measurement and status information from DER sites that do not have IEC 61850 process bus capable IEDs and switch statuses of the whole distribution grid (open/closed) are obtained from the SCADA system using IEC 60870-5-104.

3.7.4 Integration with HEDGE-IoT interoperability framework

The edge platform capabilities described here provide an efficient means for internal data exchanges and deployment of various functionalities. Standard interfaces are utilized for communication between the edge server and external devices (e.g. IEDs). Communication between

the real-time congestion management (service 3.5) and predictive congestion management (service 4.4) utilizes the Eclipse data connector as described in the respective sections.

3.7.5 Implementation details

The edge platform is still under development and, hence, it is not operational on the actual piloting sites yet. The data sources already exist and, hence, data can start flowing through the edge platform real-time data link as soon as it is deployed on field. At the moment, the functionalities are being developed and tested in development and lab environments.

At the moment, proof-of-concept of the functionalities exists and all required modules for the Finnish pilot have been defined, and implementation has been started in the development environment. Building the lab setup is ongoing. The functionality should be ready and tested in the lab environment by M30 (D3.5) and be taken into use at the piloting areas by M32.

3.7.5.1 Functionalities

TABLE 15 shows the list of functionalities comprehended by this service.

TABLE 15 – DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
SSC600 SW	Virtualised centralised protection solution	100 %	M12
Pre-processing module	Pre-processing of SV measurement data to calculate e.g. FFTs (fast Fourier transforms)	100%	M12
Data processing module for slower algorithms	Calculates RMS values for a longer time frame to have input data in the suitable format for slower algorithms such as state estimation	50 %	M22
IEC 60870-5-104 master	Module that is able to receive measurement and status data from SCADA	70%	M22
MMS client	Module that is used to communicate with the SSC600 SW using IEC 61850 MMS	50 %	M22
MQTT signal publisher to congestion management VM	Provides input data to the congestion management service (3.5)	20 %	M24
MQTT signal publisher to anomaly detection VM	Provides input data to the anomaly detection and fault forecasting service (3.5)	20 %	M26
Orchestrator for edge platform functionalities	Initialisation of shared memory and modules	40%	M28
Data storage	Database to which results are stored (needed for research purposes, would not be a	0 %	M28

mandatory part of a commercial
version of the platform)

3.7.5.2 Integration and dependencies

This section describes how the edge platform integrates with other services in the Finnish Pilot.

1. **[Service 3.4] Anomaly detection and fault forecasting to increase distribution network resilience:**

- a. The edge platform provides high resolution measurement data streams as inputs for the anomaly detection and fault forecasting service. This data originates from the 61850 SV and is shared between the different VMs or containers through MQTT. This is acceptable because the anomaly detection and fault forecasting service does not have as strict real-time requirements, as for instance, protection solutions.

2. **[Service 3.5] Real-time congestion management:**

- a. The edge platform provides processed data to the real-time congestion management service. This data is obtained from the IEDs using 61850 SV and GOOSE and from SCADA using 60870-5-104 and further processed to a form that is most suitable for the state estimation and congestion management functionalities. Data is shared between VMs using MQTT.

The Finnish pilot architecture is depicted in FIGURE 24.

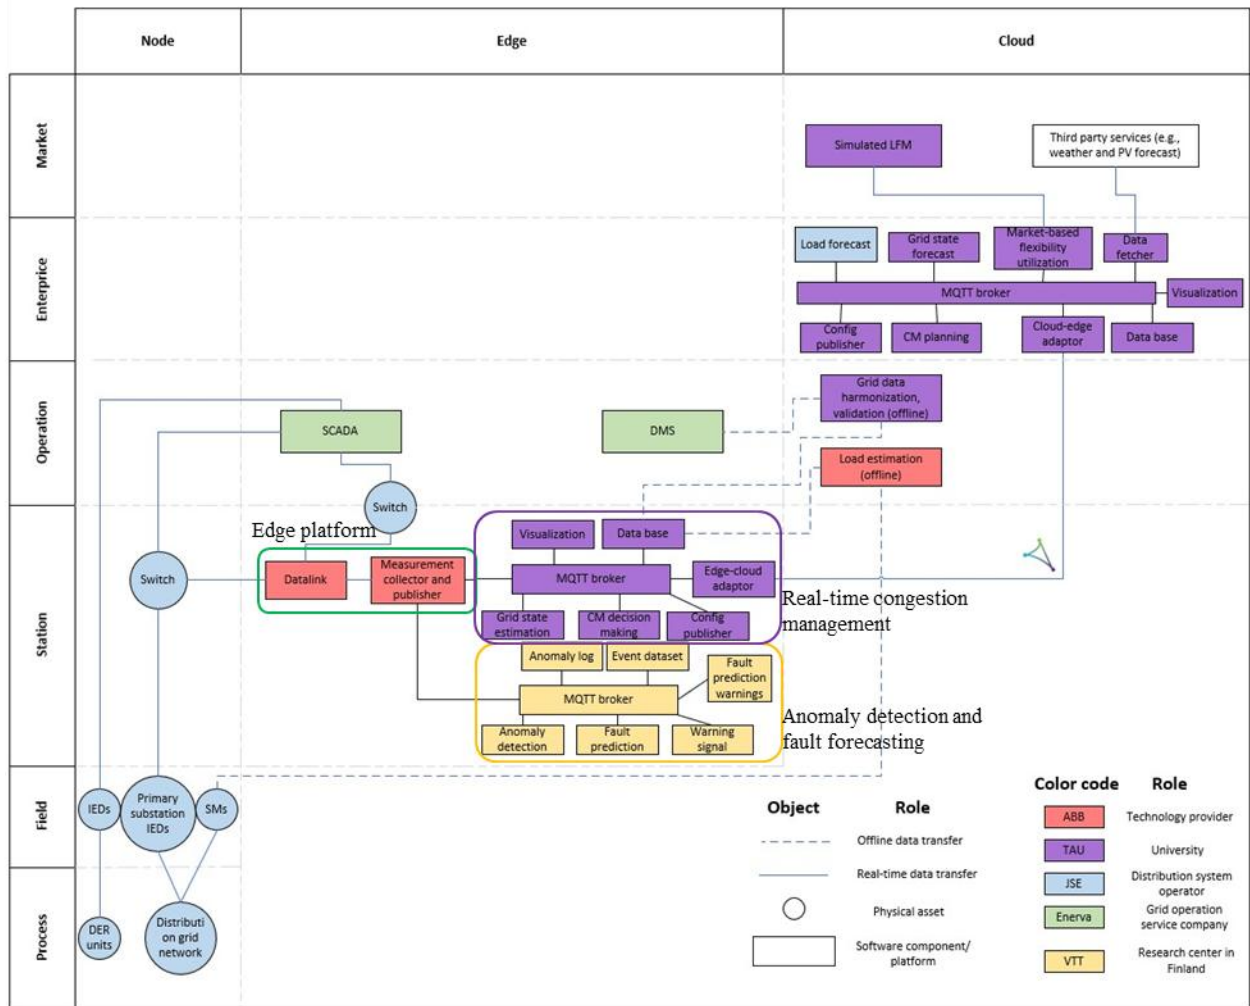


FIGURE 24 - DIGITAL PLATFORM CAPABILITIES FOR DISTRIBUTION GRID AUTOMATION PILOT SGAM DIAGRAM - ABB

4 CLOUD TECHNOLOGY ENABLERS

This section outlines the first specification of 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.

Table 16, shown below, shows a summary of the status of implementation of the TEs, while mentioning what has been improved since the previous deliverable and what is expected to be complete in the final one.

TABLE 16 – STATUS OF IMPLEMENTATION OF TE'S

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
EdgeConnect (INESC)	Design presentation and overall of the existing platform to Hedge.IoT configuration and System Operator types.	This platform established an ecosystem for stakeholders across the flexibility value chain, enabling integration, qualification and market participation, to unlock flexibility potential from LV and MV grids. The platform unlocks a multi-stakeholder environment where all relevant roles coexist;	Bilateral agreement functionalities and full communication over standard protocols such as IEC 62325. Deployment complete and operational with integration with partner's platforms and services completed or tested.
Flexibility Optimization Service (ICCS)	Initial Overview of the algorithms of bidding strategies and flexibility dispatch	IoT data cleaning and processing to fit into the decision-making algorithms and an initial implementation of specific algorithm to one defined scenario	Define more scenarios and identify more decision-making algorithms
Real-Time Reserve Market Simulator (NESTER)	Initial internal release allowing basic single bid submission, clearing and response.	Integration of IEC 62325 standard for energy market communication and file exchange. Multiple bid submission and analysis implementation. Release for pilot's initial integration testing. Improvements in bid validation, clearing and activation signal.	Full deployment of both services mFRR and aFRR with complete data stream tested and implemented among the pilot's members.
Predictive Congestion Management (TAU)	Breaking down the service into micro services to enhance modularity of algorithms and to facilitate coordination and cooperation of developers. However, this service has been in the design stage and	Define the data exchange's payload between cloud and edge to be able to develop adaptors (cloud-edge data exchange adaptor). The progress in real-time congestion management on edge will have an over-lapping benefit for this service for example in the areas of input data because the same grid and load data could be used in the cloud during the simulations. For the state forecast algorithm, the state estimation algorithm on the edge could	The simulated predictive CM is tested and then executed on the cloud and the results are already available for analysis.

Technology Enabler Name	Features in D3.3	New Features in D3.4	Future Features in D3.5
	implementation work has not yet started.	be already a good starting point. The first version of the test environment (virtual machine running on the cloud) is ready.	
Energy Community Management Service for Frequency Restoration Reserve (INESC)	Platform's main functionalities developed and being tested in a relevant environment (another European project pilot)	Incorporate new functionality for flexibility provision to the reserve market and test it with mock data provided by the TSO	Integrate the functionality in the platform and with the market platform using data spaces. Test the prototype with the new functionalities and integrations in the relevant environment (PT pilot)
Local Flexibility Market Platform (HENEX)	TE not described in D3.3	Developed an algorithm for daily generation and configuration of MTUs and for the automated operation of trading gates according to market schedule. Implemented algorithm for the execution of the market clearing.	Implement asset registration and pre-qualification and portfolio management. Deployment of settlement and clearing mechanisms.
Energy Community Platform (APIO)	TE not described in D3.3	Development and implementation of core functionalities, like baseline computation, flexibility offer optimization and energy community power management.	Connection to local flexibility market platform. Integration with the interoperability framework of the project.
TurnGreen – OptiFlex (ELERG)	TE not described in D3.3	Integrated new EnergyBox solutions into supermarkets and into the OptiFlex platform.	Asset integration. Develop optimization algorithm for each asset. Create automated responses for flexibility requests.
PowerCIM tool (KONČAR)	TE not described in D3.3	Imported preliminary data and created appropriate measurements for the time series for DTR calculations.	Updates to the CIM models to reflect measurement points of the substation and for DTR integration.

The TEs described in the table above provide a means of validation for the accomplishment of the project's objectives, namely by:

- **Designing 3 applications for residential users for the energy sector:**
 - Flexibility Optimization Service (flexibility provision), Energy Community Management Service (flexibility provision) and Energy Community Platform (demand response).
- **Designing 5 applications for grid operators:**

- EdgeConnect (flexibility provision), Real-Time Reserve Market Simulator (aFRR, mFRR), Predictive Congestion Management, Local Flexibility Market Platform and PowerCIM tool.
- **Designing 3 non-energy applications:**
 - Flexibility Optimization Service (comfort of building occupants), Energy Community Management Service (energy poverty alleviation) and Energy Community Platform (comfort of building occupants).

4.1 EDGECONNECT

4.1.1 Description of technology enabler

This platform established an ecosystem for stakeholders across the flexibility value chain, enabling integration, qualification and market participation, to unlock flexibility potential from LV and MV grids. The platform unlocks a multi-stakeholder environment where all relevant roles coexist; and provides services for each role. EdgeConnect facilitates onboarding and certification of users, registration and pre-qualification of flexible assets from aggregators, sharing of flexibility needs, baselines, bids and activation and settlement of flexibility services. Thus, it allows consumers to actively participate in energy markets. EdgeConnect ensures data privacy with role-based data access, while having critical information anonymized. Service to service interoperable data exchange is guaranteed via the integration of the project's EDC data space connector. Furthermore, semantic interoperability will be integrated using the approach defined in the project. Currently, the service is at TRL 6, with an active deployment for test purposes in the pilot.

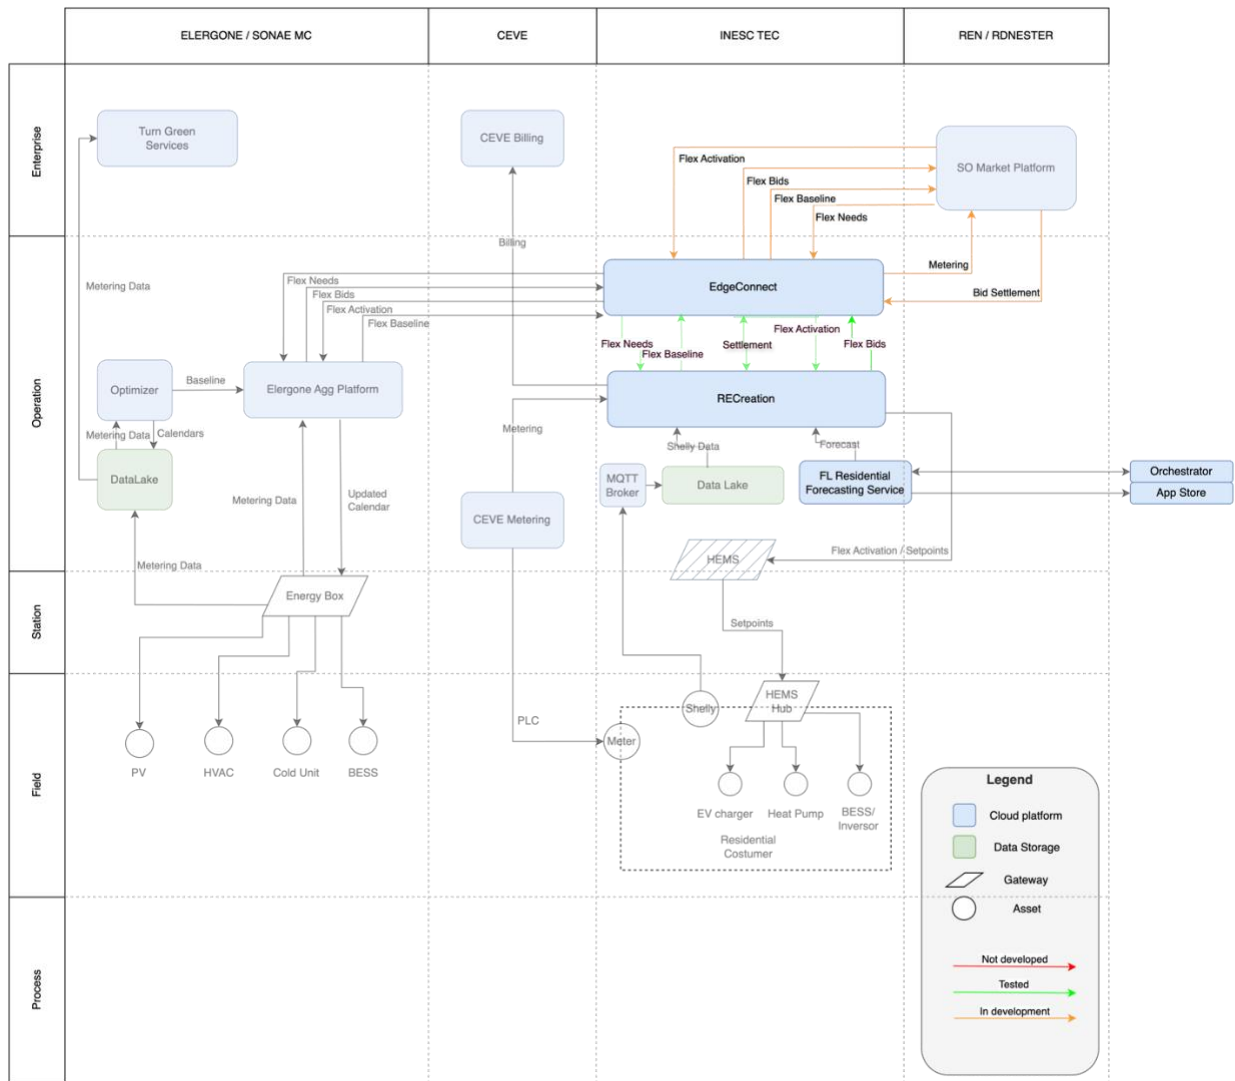


FIGURE 25 – EDGECONNECT PILOT SGAM DIAGRAM – INESC

4.1.2 Innovative aspects

MULTI-STAKEHOLDER ENVIRONMENT

Since the delivery of D3.3, the EdgeConnect platform consolidated its technical design to accommodate the interconnection with the digital platforms of all significant roles. That is, aggregators and system operators can now make use of EdgeConnect’s interoperable interfaces and its capability to connect to existing APIs on the stakeholder’s operational platforms. Figure 25 depicts the technical architecture of the Portuguese pilot considering a derivate SGAM overlay, where EdgeConnect is currently highlighted to depict its role as the value-chain enabler. From this perspective, interlinks are established towards the system operator market platform and the aggregators for residential consumers through an energy community management platform, or the commercial consumers through a building energy management service for tertiary buildings. Together, these key enablers establish an ecosystem for service co-creation and growth.

INTEROPERABILITY AND STANDARDS

Altogether, EdgeConnect provides a way for partners and stakeholders to become more flexible and allow their operation and its expansion to several regions, possibly abiding by several markets i.e., different geographies or member states. The same rational also brings the possibility for services that can interface with distributions system operators and transmission system operations under the same technological platform.

4.1.3 Service data

As a multistakeholder platform which congregates all actors in the flexibility value-chain, the platform holds data for the interacting roles of consumer, aggregator, system operator and the link with market platforms. Moreover, the transversal dimension of SO flexibility data for flexibility needs, Aggregator flexibility bids and operation baselines, market clearing processes and winning bids and flexibility activation requests (in location, aggregator, region, volume, time, penalties).

TABLE 17 – EDGECONNECT BASE DATA REPRESENTATION

Type of Data	Description
Consumer profile	HTTP – Basic information to register a consumer, including personal identification: name, surname, location, email account
Service provider profile	HTTP – Basic information to register a service provider, including identification: name, location, email (as username), contact, type of services
Service available by service provider	HTTP – Information that details a service, namely: name, API URL (if applicable), description, targets (client types, geographic locations), contact person, support centre contacts, conditions.
Metering Data	HTTP – Metering data account for active power consumption in periods of 15 minutes for spanning several hours (ideally 24)
Semantic representation of consumer	HTTP – Consumer representation in RDF format accounting to its characteristics namely location
Service List	HTTP – A list composed of service providers.

Flexible asset	HTTP – Description of a flexible asset
List of Flexible assets	HTTP – A list composed of flexible asset descriptions.
Consent	HTTP – One stakeholder provides consent to other stakeholder in scope of a service.
Flexibility capacity	HTTP – The measurable capacity to be flexible (i.e., change consumption profile) of an asset.
Aggregated flexibility capacity	HTTP – Grouping of flexibility capacity by aggregating several loads represented by Inf.10 messages. Can represent one household or a group of households.
Tips and advice for service	HTTP – Tips and direct advice towards flexibility and sustainable energy consumption.
Service subscription	HTTP – Service onboarding details that pair consumer with a service provider.
Payment invoice	HTTP – Cash flow output generating an invoice
Historical data	HTTP – Consumer’s historical metering data in 15 minutes periods.
Service terms and conditions	HTTP – Services specific terms and conditions
Form acknowledgement	HTTP – Stakeholder accepts/reject action. This is multi scope. Can be used for several purposes. Context is enclosed in request.
Acknowledgment notification	HTTP – Acknowledges operation success. This is multi scope. Can be used for several purposes. Context is enclosed in request.
Service Contract	HTTP – Service provider’s contract to onboard another stakeholder
Aggregator Profile	HTTP – Basic information to register an aggregator, including identification: name, location, email (as username), contact, type of services, geographical interest area.
Flexibility Model	HTTP – Flexibility model details included in a service.
Flexibility offering	HTTP – Flexibility offering resulting from the pre-qualification process.
Technical platform details	HTTP – The identification of a technical platform, including capabilities and description and capabilities
Technical platform capability	HTTP – Technical details on how to activate one given capability, namely type and endpoint details of the destination system or authentication details.
TSO profile	HTTP – the profile of a transmission system operator.
Service available by TSO	HTTP – service availability from the side of one TSO.
Consumer/ prosumer subscribes services	HTTP – Service onboarding details that pair a consumer with a service provider

4.1.4 Integration with HEDGE-IoT interoperability framework

Interoperable interfaces are established in three key directions:

- Adoption of established standards over common communication protocols. This is particularly the case for the interlink with the Market Simulator platform from system operators by allowing a connection via the IEC 62325 standard, with a connection established over web sockets.
- The interlink with approaches that abide by a data sovereign rationale such as the Data Spaces. This is particularly relevant for the non-regulated side of this platform, i.e., market results, consumer data, forecasts; with other organisations where there are established

incentives to consume and distribute data. Thus, EdgeConnect’s API will be connected to the EDC data space connector to provision this capability.

- EdgeConnect’s API departs from a conducted study and is derived from USEF’s modelling for energy flexibility.

On the IT side, this platform establishes also a large data hub, that congregates data from several stakeholder types and from the several stages of flexibility value chain, allowing for services that assist new players to assess their participation potential (i.e., assessing its impact by studying the available flexible resources or the lack thereof vs the economic opportunities available).

4.1.5 Implementation details

EdgeConnect core functionalities are 80% developed. The software is composed of a series of internal modules as depicted in Figure 26.

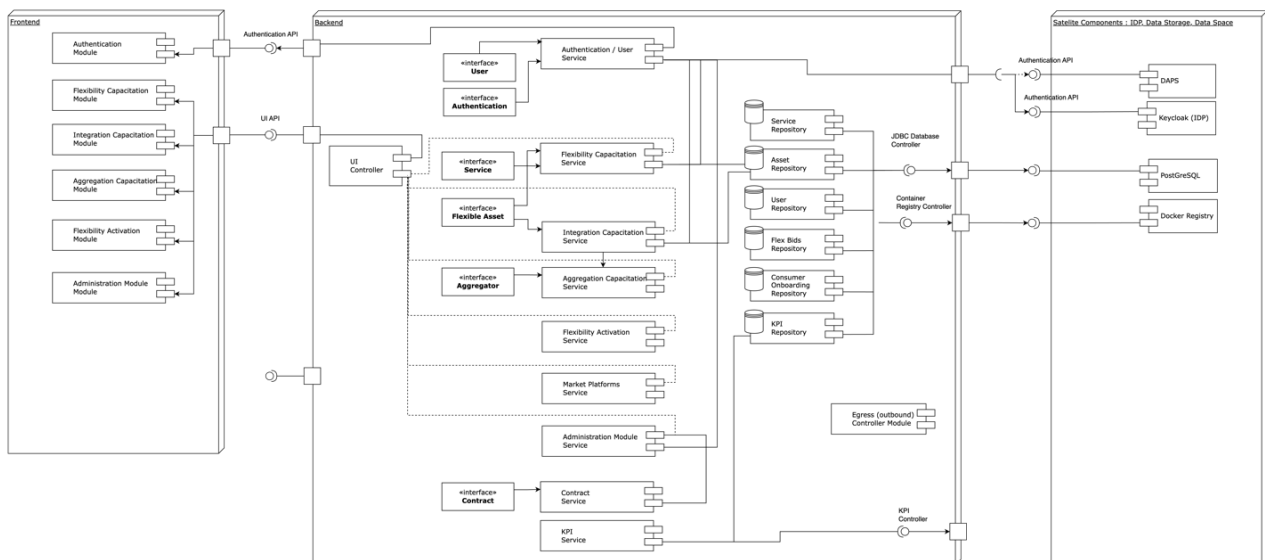


FIGURE 26 – EDGECONNECT: COMPONENT DIAGRAM - INESC

The next steps in this deployment will be focused on the complete validation of all technical System Use Cases (SUC) from Business Use Case (BUC #1), namely through:

- The expansion of the Markets Platform module to be able to interface with a market operator platform from a TSO (currently implementation is available for a DSO focused market, with an integration with OMIE Iberian market operator). This implementation will respect the IEC 62325 standard.
- Test and validation of technical integration with the RECreation community manager platform (also depicted as a service in this deliverable)
- Provision of support for integration with extra aggregators such as the Elergone Aggregation Platform for tertiary buildings.
- Integration with the EDC data space connector for interoperable and sovereign data exchange.
- Development of the bilateral agreement module to allow bilateral agreements among aggregators.

4.1.5.1 Functionalities

Table 18 represents the functionalities of EdgeConnect and its implementation status.

TABLE 18 – EDGECONNECT FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Create flexibility service	Procedure to create a new service in scope a given stakeholder	100%	M13
Publish flexibility needs	Procedure for System operator to submit new flexibility needs in each flexibility zone	100%	M13
Submit flexibility bids	Procedure for service aggregators to submit a flexibility bid	100%	M13
Submit aggregated flexibility bids	Deliver flexibility bids to the market operator	20%*	M24
Receive and forward	Collect flexibility activation requests	20%*	M24
Export value-chain data to the data space	Accept requests to grant access to data via request arriving through the data space	0%	M24

4.1.5.2 Integration and dependencies

This section describes how the Edge Connect integrates with other services in the Portuguese Pilot.

1. [Service 4.5] Energy Community Management Platform (RECreation):

- Sends flexibility request needs to the community platform.
- Receives bids
- Directs activation requests if sent bids are commanded for activation.
- Takes part in settlement processes.

2. [Service 4.3] Market simulator platform:

- Sends aggregated flexibility bids on behalf of the aggregators

* This service is already 100% completed for a market operated by a different system operator. This percentage refers to accommodating the required changes for the same behavior to occur in this specific market platform.

- b. Collect flexibility activation requests from the selected bids.

3. **[Service 2.2] FL service:**

- a. Provides an ecosystem and data hub that the service can consider for data acquisition and export through data spaces.
- b. Communicates using the Eclipse data space connector

4. **App Store/Data Space Integration:**

Interlink with services which are available as Apps the App Store of applications for the Data space.

4.2 FLEXIBILITY OPTIMIZATION SERVICE

4.2.1 Description of technology enabler

On D3.3, the previous iteration of this deliverable, a detailed description of this approach is provided. Therefore, section 4.2 of D3.3 is the reference for the description of this service. The fundamental architecture and logic of the service have remained consistent; however, notable updates and extensions have been made in D3.4, regarding the development of the subservices and their interoperability and interconnection among the Greek pilot components (Figure 27).

Among the most important additions is a new capability currently under development that enables the optimization of the aggregator’s bidding behavior in the Local Flexibility Market (LFM), using as input the optimized dispatch plans of the flexibility assets from a subset of prosumers. This process, nearing completion, has already yielded encouraging experimental results based on data from ten participating households.

In parallel, substantial progress has been achieved across several core subservices. In the area of demand forecasting, the service now leverages real-time and historical consumption data and IoT sensor readings to deliver accurate short and mid-term load forecasts. This functionality is implemented and specified in detail in SUC-GR-01.02, which outlines both the forecasting logic and system interfaces. Likewise, production forecasting has matured into a reliable service for estimating PV output based on dynamic environmental inputs, as detailed in SUC-GR-01.03, and will be fully integrated into the bidding and optimization pipeline.

Another important advancement in D3.4 is the incorporation of Non-Intrusive Load Monitoring (NILM), which allows device-level disaggregation of energy consumption from aggregate signals, using machine learning models. This functionality has evolved from early-stage research using open benchmark datasets to more advanced approaches based on denoising autoencoders and cluster-based optimization techniques applied to real Greek households. These improvements enhance the system’s ability to estimate and allocate flexibility potential with greater granularity. Moreover, early-stage implementations of the incentive’s optimization, consumer interaction, and bid formulation modules have progressed, enabling the dynamic tailoring of offers to prosumers and the construction of optimized bid strategies that incorporate customer acceptance probabilities and flexible incentive levels. Collectively, these enhancements bring the Flexibility Optimization Service closer to full operational deployment within the HEDGE-IoT platform, improving its responsiveness, accuracy, and value in local energy markets.

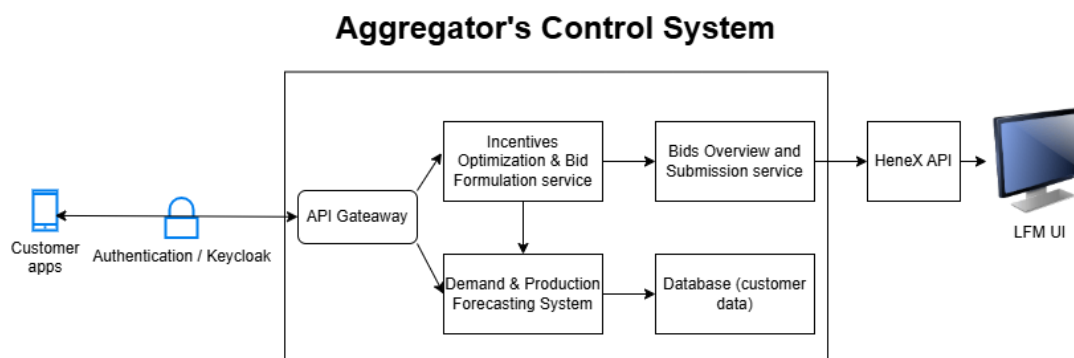


FIGURE 27 – FLEXIBILITY OPTIMIZATION SERVICE ARCHITECTURE DIAGRAM – ICCS

4.2.2 Innovative aspects

MODULAR EDGE-CLOUD ARCHITECTURE

The Flexibility Optimization Service adopts a modular edge-cloud design that distributes computational responsibilities across the HEDGE-IoT platform. Core forecasting tasks – such as demand and production forecasting – are executed in the cloud, ensuring scalability and model retraining with large datasets, while data ingestion and NILM disaggregation are increasingly handled at the edge to support real-time responsiveness and reduce network congestion. This separation of concerns improves system resilience and enables flexible deployment scenarios across diverse infrastructure settings.

AI + IOT INTEGRATION

At the heart of the service lies the integration of IoT-based sensing with artificial intelligence. Real-time load, PV production, and weather data, collected through edge IoT meters, are streamed into machine learning-based modules for demand forecasting, production forecasting, and NILM. These subservices continuously analyze usage patterns at both the household and device level, forming the predictive layer upon which incentive optimization and bid formulation operate. This data-driven loop enables dynamic, context-aware decision-making that adapts to evolving user behavior and environmental conditions.

HYBRID OPTIMIZATION FRAMEWORK

The service's optimization core combines rule-based logic (e.g., regulatory limits, operational constraints) with adaptive AI components for flexibility scheduling and bidding. Demand and production forecasts inform a real-time bidding mechanism that blends deterministic rules with probabilistic models, such as customer acceptance prediction and flexible incentive ranges. This hybrid approach ensures the system can react safely and quickly under known constraints while still learning and improving over time from historical data and user interaction outcomes.

ADVANCED SECURITY & PRIVACY

Each subservice operates within a secure data environment governed by Attribute-Based Access Control (ABAC) and Role-Based Access Control (RBAC). Access to raw and processed data—such as personalized forecasts or flexibility offers—is tightly regulated, and all exchanges between edge devices, cloud services, and user interfaces are encrypted. Especially for NILM, which derives fine-grained device usage from aggregate signals, privacy-preserving computation techniques are applied to ensure sensitive appliance-level data remains confidential and compliant with GDPR and platform-wide policies.

USER-CENTRIC PERSONALIZATION

The service embeds personalization through every stage of its workflow. Forecasting models are trained per user, NILM enables device-specific flexibility detection, and incentive optimization adjusts offers based on engagement likelihood. Flexibility requests and notifications are delivered

in real time via mobile interfaces, tailored to reflect user behavior, consumption habits, and preferences. This level of personalization ensures that flexibility offers are not only technically feasible but also behaviorally aligned, improving participation rates and end-user satisfaction.

4.2.3 Service data

The Flexibility Optimization Service utilizes a diverse set of real-time and historical input data streams from customers, IoT edge devices, and external platforms to deliver its core functionalities. Each data element supports one or more of the subservices—namely, demand and production forecasting, incentive optimization and bid formulation, NILM, and customer interaction:

- **Customer_id:** A unique identifier for each participating household or user, used across all subservices to personalize forecasting models, link incentives to user profiles, and associate NILM outputs with specific prosumers in the platform.
- **Main_energy_consumption:** Real-time total household energy demand (in kWh), directly ingested by the demand forecasting module for model inference and also used by NILM to disaggregate appliance-level loads for finer control of flexibility assets.
- **MTU (Market Time Unit):** A categorical timestamp (typically in 15-minute intervals) used to align real-time data across all subservices. It serves as the temporal backbone for forecasting horizons, bid construction, and incentive scheduling.
- **PV_Production:** Current solar generation data (in kWh), collected via IoT from PV inverters, and used by the production forecasting module to calibrate and validate short-term generation predictions for flexibility assessment.
- **Battery_capacities:** Maximum storage potential (in kWh) for each household's battery system. This input is used by the dispatch optimization logic to define feasible flexibility ranges and supports the bidding optimization module in estimating the available energy for shifting.
- **Battery_control_levels:** The real-time state-of-charge or control commands issued to batteries (in %), used in dispatch and incentive optimization to determine how much capacity can be mobilized during each bidding cycle.
- **Weather_data:** External meteorological parameters (e.g., temperature, irradiance, wind speed) retrieved from weather forecasting APIs¹⁰. This data is essential for production forecasting, as environmental conditions significantly influence load patterns and PV output.
- **Electricity_price:** Prices from the day-ahead and intraday wholesale markets (€/kWh), used by the bidding optimization engine to set economically efficient bid prices. These also serve as a benchmark to evaluate the cost-effectiveness of activating local flexibility.
- **Incentive_Limits:** Configurable thresholds defining the minimum and maximum incentive that can be offered to customers (€/kWh). These bounds are used by the incentive optimization module to ensure that flexibility offers are both compliant and competitive.
- **Forecasted_energy_consumption:** Timeseries of predicted consumption values (in kWh), generated by the demand forecasting module and used by the optimization engine to identify periods of high load where flexibility activation would yield the most benefit.
- **Forecasted_energy_production:** Timeseries of predicted PV generation (in kWh), output by the production forecasting module. These forecasts are crucial inputs for planning dispatch and forming accurate, risk-aware bids in the local market.

¹⁰ <https://open-meteo.com/>

- **Historical_energy_consumption:** Timeseries of historical household consumption values (in kWh), used by the demand forecasting module for model training.
- **Historical_energy_production:** Timeseries of historical PV generation (in kWh), used as input in the production forecasting module for model training.

Historical_device_consumption: Timeseries of historical appliance consumption (in kWh), used as input in the NILM for model training.

4.2.4 Integration with HEDGE-IoT interoperability framework

Based on the ongoing discussions within WP4 and the initial set of requirements, this service is expected to be integrated with the Hedge-IoT Interoperability Framework. The details around the architecture and integration approach are still evolving. We anticipate having a clearer picture and more concrete information in the next deliverable.

4.2.5 Implementation details

As of this deliverable (D3.4), the Flexibility Optimization Service's core components have been developed and validated in a controlled environment. Early-stage prototypes have been successfully executed using synthetic datasets that simulate real-world flexibility scenarios. These prototypes support core backend functionalities, including energy dispatch optimization, incentive calculation, and bid formulation, forming the minimal viable logic for the aggregator's participation in the Local Flexibility Market.

Since D3.3, on M15 a key milestone has been reached with the integration of real pilot data from Greek households into the optimization pipeline. This live data now feeds directly into several key subservices: demand and production forecasting, which provide predictive inputs for the optimization engine, NILM, which disaggregates appliance-level loads to improve the granularity of flexibility detection, and incentive optimization and bid formulation, which dynamically adjust offers and market bids in response to available flexibility and market conditions. The system has begun its first experimental runs using real IoT data, with 15-minute load and PV readings now partially ingested in near real-time, enabling validation of the complete optimization workflow under actual pilot conditions.

Another important implementation activity concerns the differentiation of customer scenarios. The optimization logic is being extended to handle various household configurations, including: (a) prosumers with both PV and batteries, (b) prosumers with only PV, and (c) pure consumers.

Each case introduces distinct flexibility profiles, incentive sensitivity, and dispatch constraints, which are now being modelled and tested in preparation for broader pilot deployment. In parallel, further enhancements are planned for the forecasting and NILM subservices. For forecasting, the upcoming work will concentrate on improving model accuracy and expanding backend capabilities to support integration with all real-time pilot data streams. In the case of NILM, future steps involve completing the backend implementation to ensure seamless integration with the aggregator platform and other services. The focus will be on finalizing the disaggregation pipeline, improving classification accuracy, and enabling consistent output formatting for downstream modules. Both subservices are on track to reach full functional maturity within the project timeline, contributing to the overall reliability and intelligence of the Flexibility Optimization system. In addition, the electricity price forecasting component will be developed to enable alignment with market signals

and increase the responsiveness of the optimization engine to pricing fluctuations in both day-ahead and intraday markets.

Looking ahead to M30, the plan remains aligned with earlier expectations. The service is progressing toward full algorithmic coverage for all flexibility scenarios and aims to be integrated with the aggregator’s backend by that milestone.

4.2.5.1 Functionalities

Table 19 shows the list of functionalities comprehended by this service.

TABLE 19 – FLEXIBILITY OPTIMIZATION SERVICE FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Dispatch	Optimize PV–battery charge/discharge	100%	M16
Electricity Price Forecasts Integration	Integrate the electricity price of day-ahead market as a benchmark for the optimization algorithms	90%	M26
Bidding Optimization	Real-time continuous quantity & price bids for the aggregator to participate in the local flexibility market	70%	M26
IoT Data Ingestion	15-min load & PV readings streamed from prosumers’ IoT meters for forecasting modules, 1-min load for NILM.	50%	M26
Integrate customer acceptance probabilities	Consider the probabilities of the customers to respond positively in flexibility offers	0%	M32
Integrate different levels of incentives	Consider different levels of incentives provided to the customers	0%	M32

4.2.5.2 Integration and dependencies

The Flexibility Optimization Service integrates with multiple components of the Greek Pilot, supporting real-time data processing, forecasting, user interaction, and market participation. Each integration point enables a specific subservice or function essential to delivering accurate, responsive, and market-aligned flexibility strategies.

1. Demand Forecasting

The demand forecasting service receives real-time and historical energy consumption data streamed from smart meters via MQTT and Kafka. The resulting consumption forecasts are used by the optimization engine to estimate available flexibility potential and to inform the timing and magnitude of load shifting and market participation.

2. Production Forecasting

Production forecasting leverages PV generation data and external weather inputs to predict solar output. These forecasts are essential for identifying surplus energy available for dispatch, especially for prosumers with PV systems. Forecasts are fed directly into the dispatch and bidding logic of the Flexibility Optimization Service.

3. NILM (Non-Intrusive Load Monitoring)

The NILM subservice is integrated within the aggregator's backend platform, operating on aggregate consumption data. It disaggregates household energy signals to identify appliance-level usage patterns using denoising autoencoder models. This granular insight improves the flexibility estimation, allowing the system to detect which specific loads are shiftable and when. The NILM outputs are also used to refine customer-specific flexibility offers and improve incentive targeting by understanding which devices respond to DR events.

4. User Interfaces

The aggregator interfaces with the optimization service through an API Gateway, where system operators can trigger optimization routines and monitor results. Flexibility offers generated by the system are communicated to end users through a mobile Customer App, allowing real-time feedback and engagement. This interaction loop ensures customer preferences and responses are captured and reintegrated into future optimization cycles.

5. Edge Processing

Edge-level IoT devices, including smart meters and PV inverters, continuously stream consumption and generation data to the aggregator platform. This data forms the basis for forecasting, NILM disaggregation, and real-time flexibility assessments. The modular edge-cloud architecture supports latency-sensitive processing while ensuring scalability and data consistency.

6. Trading Service (HenEx Platform)

Once the optimal bid is computed, including flexibility quantity and price, it is submitted via the HenEx API Gateway to the Local Flexibility Market. Integration ensures real-time alignment with market requirements for prequalification, baseline submission, bid entry, and settlement. The optimization service acts as the orchestrator for this pipeline, producing compliant bids based on live grid and asset conditions.

Figure 28 shows an architecture diagram that maps the integrations and dependencies of the service with the remaining Greek pilot services, as well as showcases the status of the implementation of each integration.

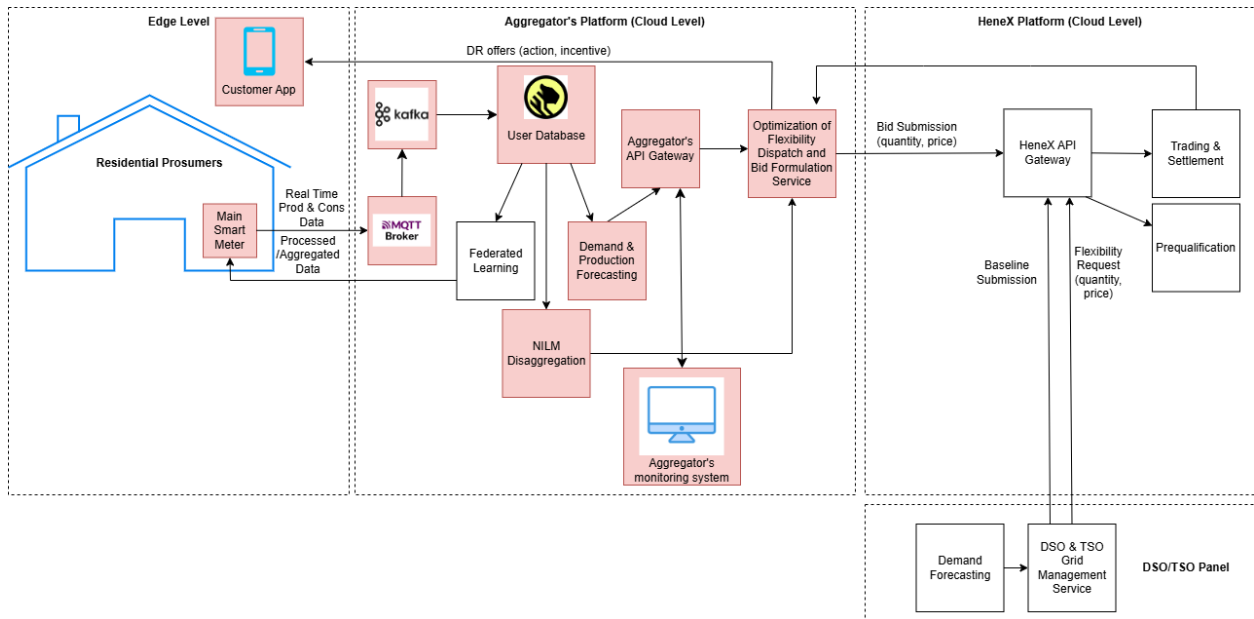


FIGURE 28 – FLEXIBILITY OPTIMIZATION PILOT SGAM DIAGRAM – ICCS

4.3 REAL-TIME RESERVE MARKET SIMULATOR

4.3.1 Description of technology enabler

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 tool 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. In addition, the Real-Time Reserve Market Simulator facilitates better understanding and the optimization of balancing market operations, ensuring alignment with European balancing market frameworks such as MARI¹¹ (Manually Activated Reserves Initiative) and PICASSO¹² (Platform for the International Coordination of Automated Frequency Restoration and Stable System Operation).

A detailed description of the service was presented in section 4.3 of Deliverable 3.3. The expectations for the final iteration of this deliverable (D3.5) is to completely integrate the real-time reserve market simulator into the whole Portuguese pilot ecosystem.

4.3.2 Innovative aspects

Following the goal of the tool to replicate the reserves energy market, there are some innovative aspects that the tool contributes to.

CLOSING THE GAP

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 inclusion in ancillary service markets and encouraging diverse stakeholder participation.

STANDARD-BASED

The service was improved by the adoption of the standard IEC 62325-315 CIM European market model exchange profile. It describes documents and their corresponding structure for a homogenous file creation and sharing inside the electricity market being simulated in this service. This allows the Real-Time Reserve Market Simulator completely mimics the process that will take place during an actual interaction with the electricity market. The standard is extensively detailed in the ENTSO-E platform¹³.

¹¹ https://www.entsoe.eu/network_codes/eb/mari/

¹² https://www.entsoe.eu/network_codes/eb/picasso/

¹³ <https://www.entsoe.eu/publications/electronic-data-interchange-edi-library/>

4.3.3 Service data

The real-time reserve market simulator leverages data available on the ENTSO-E transparency platform¹⁴. This dataset contains historical records of submitted bids and activated power for a specific market time unit in Portugal, as part of the pilot project. That data follows the structure described in standard IEC 62325 referred to previously.

In the perspective of the user, it needs to provide as many files as bids it wants to offer. Each file must follow a Reserve Bid Market Document¹⁵ structure which includes the information of the flexibility it wants to offer to the market for a specific market time unit. In the same perspective, the output received by the user is an Activation Market Document¹⁶ per bid provided in the beginning. Each “activation market document” includes the acceptance or not of the proposed flexibility.

In terms of timings for the submission and reception of the mentioned documents, Figure 29 depicts the process. It starts when the user submits the bids in the “Reserve Bid Market Document format” until 25 minutes before the analyzed Market Time Unit (MTU). After the BSP Gate Closure Time (GCT) and the TSO GCT, the actual clearing of the market and the Activation Optimization Function (AOF) is executed. Results are communicated in “Activation Market Document” format 7.5 minutes before the MTU.




FIGURE 29 – REAL-TIME RESERVE MARKET SIMULATOR: TIMINGS FOR DOCUMENT EXCHANGE IN THE MFRR SERVICE – NESTER

¹⁴ <https://transparency.entsoe.eu/dashboard/show>

¹⁵ https://eepublicdownloads.entsoe.eu/clean-documents/EDI/Library/cim_based/schema/Reserve_bid_document_UML_model_and_schema_v1.2.pdf

¹⁶ https://eepublicdownloads.entsoe.eu/clean-documents/EDI/Library/cim_based/schema/Activation_document_UML_model_and_schema_v1.2.pdf

Additionally, the service requires pre-submitted baseline profiles for the assets intended to provide flexibility. Upon bid acceptance, smart meter measurements are retrieved and compared against the initial baseline. This comparison enables the market operator to validate the actual delivery of the committed flexibility. If the validation is successful, the user receives remuneration; otherwise, discrepancies result in penalties.

4.3.4 Integration with HEDGE-IoT interoperability framework

The real-time reserve market simulator will be deployed on a server accessible through endpoints URLs using a Restful API. Those will be connected to a middleware called EdgeConnect (service 4.1) which is a platform developed by INESC TEC, another partner of the Portuguese Demo. This platform will enable the market simulator to integrate with the interoperability framework of the project, since it has an integration with the EDC connector.

4.3.5 Implementation details

D3.3 presented details concerning a deployment of the service on a TRL 5 stage. As for the current implementation, the TRL stage is the same as the developments are under Nester’s responsibility and integration with other partners for an increased TRL level is expected in the near future.

4.3.5.1 Functionalities

Table 20 shows the list of functionalities and tasks comprehended by this service.

TABLE 20 – REAL-TIME RESERVE MARKET SIMULATOR FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Document structure definition	Following the IEC 62325, example files were created and defined	100%	M18
Market rules definition	Study and definition of the applicable market rules to be implemented in the service	100%	M16
Bid submission and validation	A user is able to submit a bid with the correct format file in a dedicated endpoint	100%	M18
Market clearing	The service clears the market with the submitted bids and historical ones.	80%	M18
Activation message issue	The result of the simulation of the market is send to the user with the corresponding reply	50%	M19
Settlement calculation	After receiving the measurements, the settlement	20%	M20

calculation is performed, and fees
or payments are determined

4.3.5.2 Integration and dependencies

This section describes how the market simulator integrates with other services in the Portuguese pilot:

The main dependence of the service here described is with the EdgeConnect platform which is the middleware inside the Portuguese pilot allowing the connection among the industrial and residential energy communities and the flexibility buyer (e.g. DSO and/or TSO) acting as aggregator of flexibility for the participation in flexibility markets, such as the reserve's markets being simulated.

1. **[Service 4.1] EdgeConnect:**

- a. Sends market clearing results
- b. Sends activation signals
- c. Receives the bids from the aggregator of the residential and industrial energy players

Figure 30 shows a SGAM diagram that maps the integration of the Real-Time Reserve Market Simulator (red rectangle) inside the Portuguese Pilot.

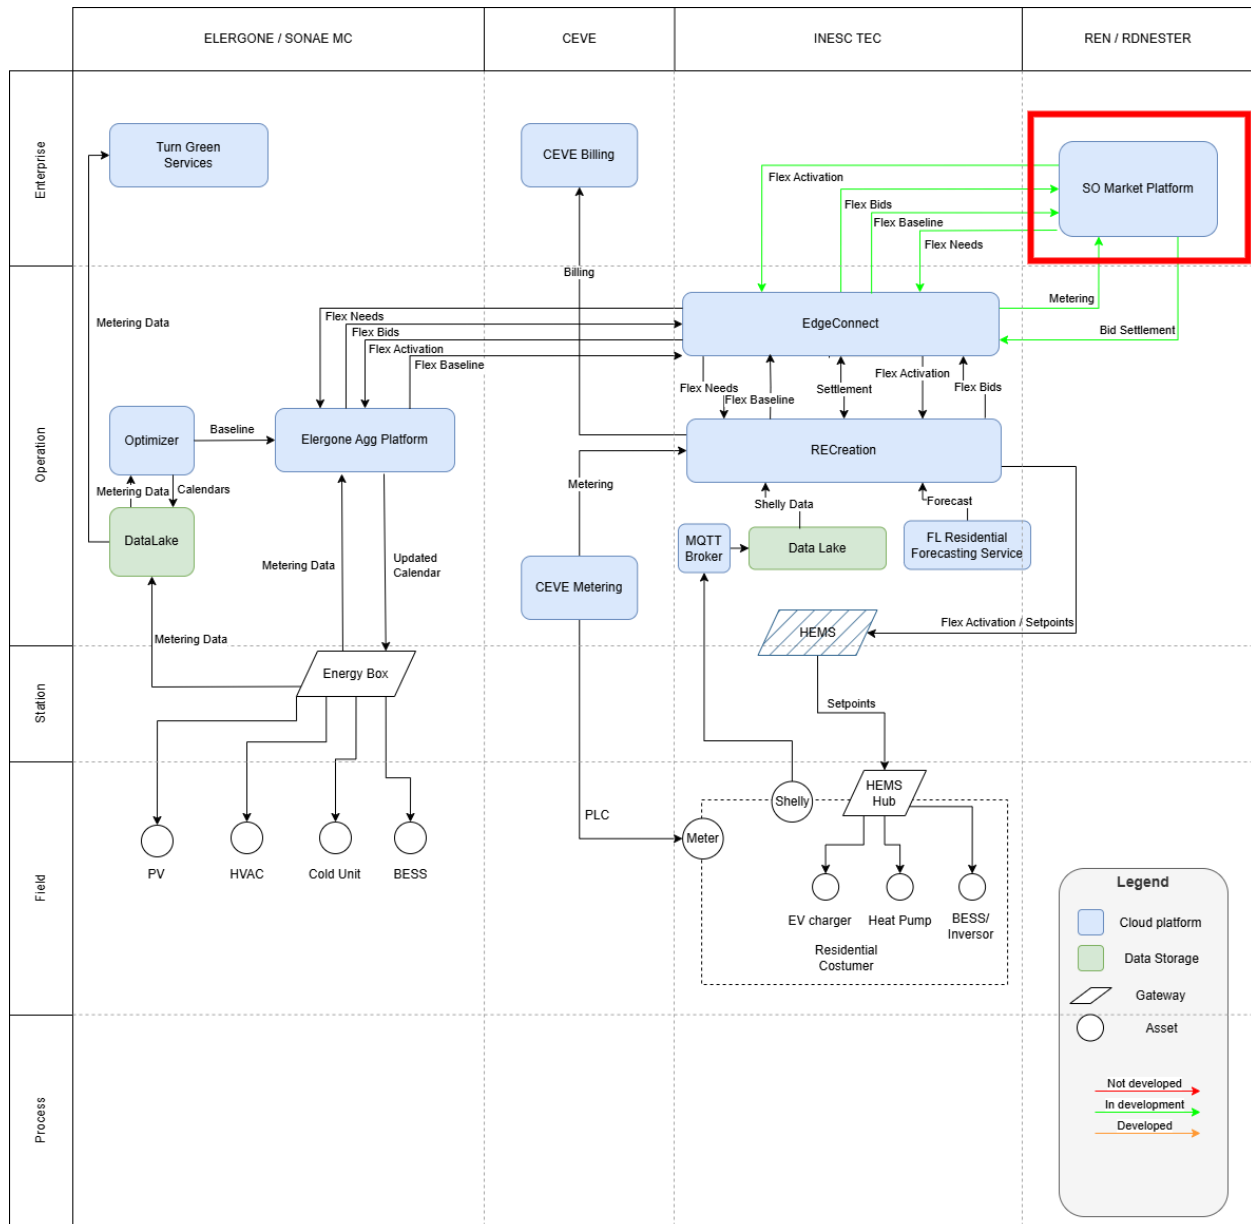


FIGURE 30 - REAL-TIME RESERVE MARKET SIMULATOR PILOT SGAM DIAGRAM - NESTER

4.4 PREDICTIVE CONGESTION MANAGEMENT

4.4.1 Description of technology enabler

Figure 31 illustrates the architecture of predictive CM service on cloud. The predictive in this context is in hourly resolution for the next 36 hours. The predictive CM service of SUC-FI-02.01 and SUC-FI-02.02 have been divided into several micro-services—such as grid state forecast, market-based flexibility utilization, CM planning, config publisher and data fetcher—following a service-oriented design.

For data exchange between edge and cloud, the edge-cloud adaptor is used. Compared to the previous version detailed in D3.3 the following progress/changes have been made:

- A config publisher has been added to the architecture, which helps to parameterize and synchronize all the subservices involved in CM.
- The MQTT protocol has been adopted for the message bus.

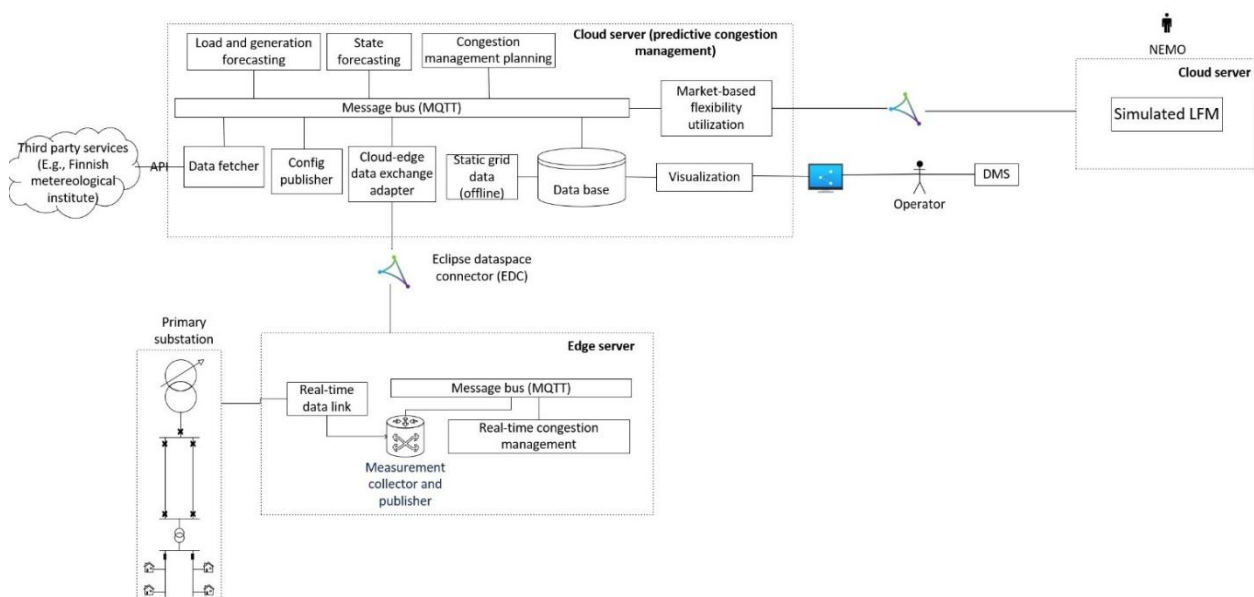


FIGURE 31 – PREDICTIVE CONGESTION MANAGEMENT ARCHITECTURE DIAGRAM – TAU

4.4.2 Innovative aspects

ENHANCING PRACTICALITY OF PREDICTIVE CM IN REAL-LIFE APPLICATIONS

Predictive Congestion Management (CM) is still in its early adoption stage among European Distribution System Operators (DSOs). It may involve using flexibility from local markets or grid reconfiguration strategies. This pilot project explores combining predictive and real-time CM services to maximize efficiency and leverage both market- and grid-based solutions.

INTEGRATING CLOUD AND EDGE COMPUTING FOR CM

A key innovation is integrating cloud and edge computing (the cloud-edge continuum), enabling optimized algorithm deployment and improved interoperability. This approach supports continuous improvement of CM services through reliable updates from cloud to edge devices.

4.4.3 Service data

The developments done for real-time CM (service 3.5) with respect to static data sets, are also benefiting predictive CM service, and therefore, they are repeated in this text. In that regard, harmonization, validation and analysis of static data sets have been almost completed. Those data are as follows:

- **Load data for the last three years (2022,2023, and 2024)**
 - The load data has been obtained from smart meter measurement service provider, aggregated to secondary substation level (i.e., anonymized), and validated.
- **Grid data in QGIS format**
 - Grid data has been converted to power technologies international (PTI) format useful for load flow calculations. The converted grid data has been harmonized and validated (there are still minor modifications needed).
- **Generation data**
 - Generation data of tens of rooftop PVs, one solar park, and one hydro power plant has been collected and associated to the grid data.

The above-mentioned data sets have gone through analysis to assure that the previous steps for load, generation and grid data have been done correctly. The analysis has helped to find minor errors (some have been already addressed, and a few needs will be addressed soon).

The predictive CM service requires ambient temperature forecasts as well. On that regard, the weather forecast data has to be fetched from a third-party forecaster server by data fetcher component. As the state forecast subservice requires weather data, it is planned that weather data forecast, its origin, how to access it and its data model be specified from August 2025 onwards, so the implementation of it can follow afterwards.

For the CM outputs, the initial idea is to visualize the state of the grid to the operator and recommend CM solutions. It is necessary to define how that could be implemented, something that should be decided once the subservices produce some meaningful results and therefore, this part is expected to be done in later stages of the pilot, possibly mid-2026.

Concerning Local Flexibility Market (LFM) and its integration into the predictive CM, there is still uncertainty whether data from a real LFM or a simulated LFM will be available. That depends whether LFM operator, as a third party, could participate in the demonstration through open call. That would affect the data necessary to be exchanged between market utilization component and LFM.

4.4.4 Integration with HEDGE-IoT interoperability framework

The service will utilize Eclipse Dataspace Connectors (EDC) for edge-to-cloud and cloud-to-cloud data transfer, providing a reusable and interoperable framework via standardized APIs. The initial plan is to experiment with EDC by exchanging arbitrary data between edge and cloud environments. Once the edge and cloud services reach a near-ready stage, they will be connected through EDC. This experimentation phase is scheduled to begin in August 2025.

In parallel, the pilot will explore service interoperability by testing how a subservice—such as grid state estimation—can be pushed, pulled, and orchestrated via an app store environment.

4.4.5 Implementation details

Compared to the previous evaluation cycle, the pilot has progressed in testing methods and infrastructure, and dataset harmonization validation and analysis. The following explains the progress of each part in more detail:

- **Testing method and infrastructure**

- Testing would have three phases of unit testing, integration testing and system testing. Unit testing and integration testing will be done using TAU's cloud infrastructure, and after completion, system testing could be done so that for example visualization of state forecast is realized at Järvi-Suomen Energia's pilot site meaning that the DSO at the pilot site would have access to TAU's cloud system and therefore can have access to predictive CM results.

- **Algorithm development**

- As the primary focus of the pilot is to demonstrate the CM on the edge, the algorithm development for cloud sub services has not yet started.

Specific developments of the current cycle (common with SUC-FI-02.03 & SUC-FI-02.04) concerning the data sets

- Demonstration grid's data has been migrated from the QGIS software (used by Järvi-Suomen Energia) to MATLAB to model it based on state estimation needs (PTI format).
- The missing data of grid components in QGIS (e.g., transformer impedances, cable impedances) has been found from component's documentations and added to the grid data model.
- The grid data has been validated meaning that grid components are modelled correctly.
- Historical load data (hourly resolution) has been collected and aggregated (secondary substation level) from smart meters in the demonstration area from 2022 to 2024.
- Historical generation data (hourly resolution) has been collected from their measurement devices from 2022 to 2024.

- Individual load profiles, which can be used for load estimation and forecasting, were calculated for secondary substations, considering the time of day, calendar data, and outdoor temperature. These profiles model hourly substation loads and their variance over the coming year.
- Load and grid data have been associated with each other, and a grid analysis has been done to evaluate the state of the grid in previous years of grid operation (e.g., 2022, 2023, 2024).

Based on the progress made so far and the remaining tasks, it is expected that by month 30, unit testing of critical grid algorithms including grid state forecast and CM planning are completed. The ideal case would be that the developed algorithms for edge are used as template for cloud subservices so that the development for cloud sub-services progress faster instead of starting from scratch. The integration and system testing of them will follow right after unit testing. The risk here is that the work on CM algorithm on the edge progress slower than expected and that affects the implementation of cloud systems. However, the mitigation method as mentioned above is to reuse the edge algorithms on the cloud with modifications as much as possible to reduce extra work.

4.4.5.1 Functionalities

Table 21 indicates the functionalities and their timeline to be completed for the predictive CM service.

TABLE 21 – PREDICTIVE CONGESTION MANAGEMENT FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Database	It stores static and dynamic data	0	M30
State forecast	It estimates the state variables of the grid	0	M30
CM planning	It makes CM decision depending on the grid state and aims to remove congestion	0	M30
Config publisher	It publishes parameters necessary for all components on the edge.	0	M36
Cloud-edge adaptor for using EDC	It sends and receives data to/from cloud	0	M36
Data fetcher	It received data from a third-party service provider (e.g., weather forecast)	0	M36
Market-based flexibility utilization	It participates in the LFM to procure flexibility in the day ahead to prevent congestion on real-time.	0	M36

4.4.5.2 Integration and dependencies

As shown in Figure 32, the predictive CM that is located on the cloud, at the enterprise level, must connect to the edge server where real-time CM is located through EDC. Therefore, the edge and cloud servers are dependent on each other, although the dependency has to be mitigated as much as possible in case of unavailability of EDC (N-1 principle), so those two servers can ideally work independently if situations demand that. In addition, the predictive CM service is dependent on third-party service providers for weather forecast and local flexibility market (LFM). The following section describes how the predictive CM service integrates with other services in the Finnish pilot:

1. [Service 3.5] Real time CM on the edge

- a. Sends ambient temperature data.
- b. Sends electricity market price data

2. EDC Connector

- a. Mainly publishes and possibly receives data to/from EDC

3. Third-party service provider

- a. Fetches the weather forecast data
- b. Interacts with LFM concerning DERs' flexibility

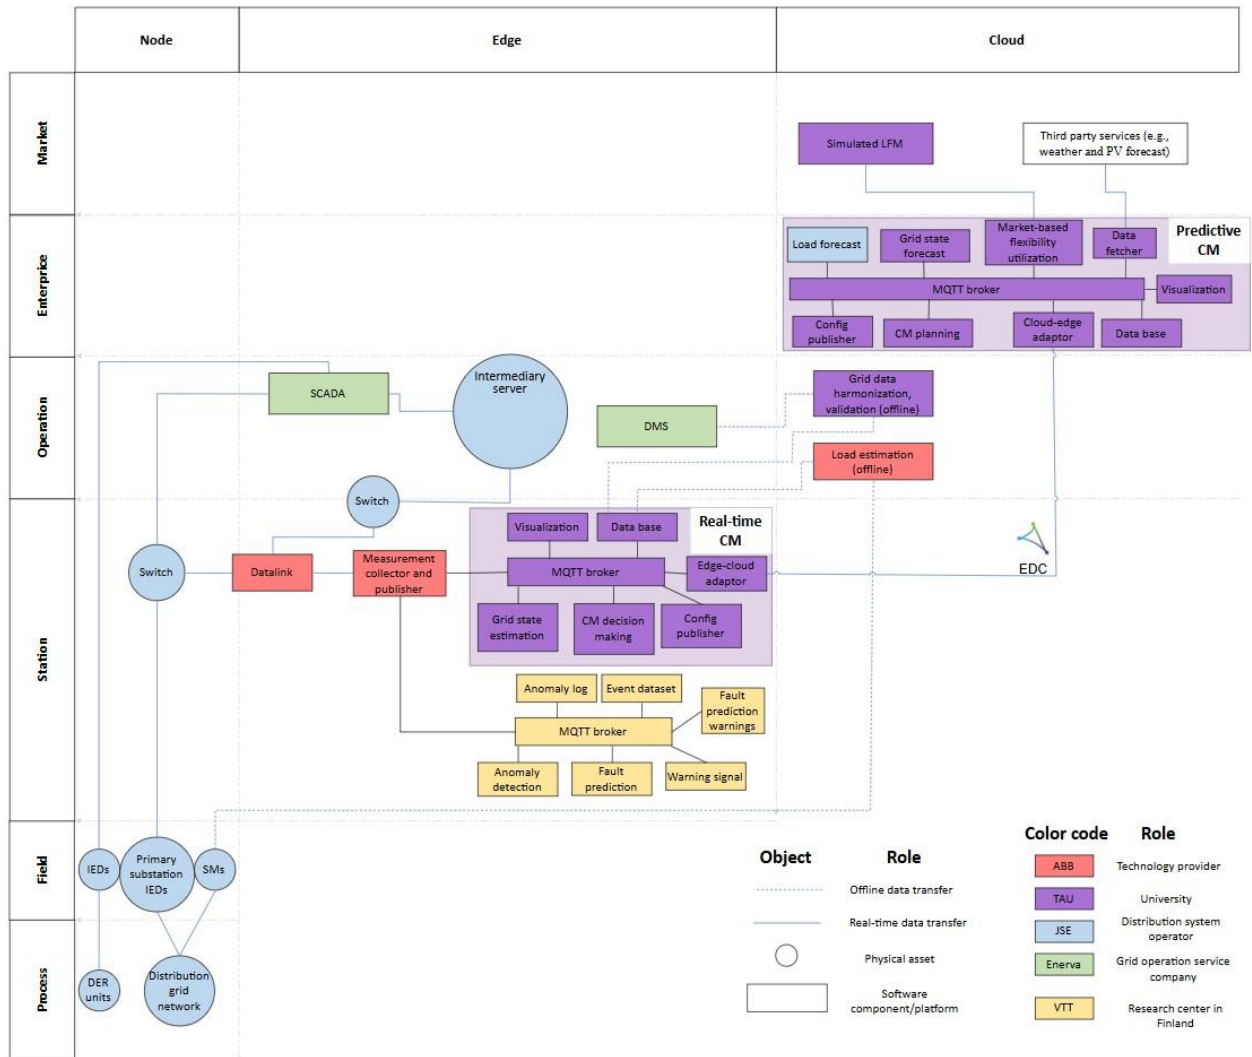


FIGURE 32 – PREDICTIVE CONGESTION MANAGEMENT PILOT SGAM DIAGRAM – TAU

PROVISION OF FLEXIBILITY FOR FREQUENCY RESTORATION RESERVE

As mentioned in the previous deliverable (D3.3), 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](#).

COMFORT OF RESIDENTIAL USERS

The service will be able to receive comfort restrictions from the end users, related to each asset, and integrate those restrictions into the optimization algorithms. The algorithm will take them into account, together with the needs of the TSO to form a bid that will not be harmful to the user's comfort but will still allow them to participate effortlessly in the tertiary reserve market.

4.5.3 Service data

This service receives most of its data via RESTful API and stores relevant in local databases. Stored data is stored in an encrypted manner and is only shared with outside partners via RESTful API with rule-based access.

The data sets stored locally on this service are shown in Figure 34 and described below:

- **Users and Assets:** data regarding users, their assets and pre-qualification status.
- **Network access tariffs:** the price to access the network via liberalized market (changes every 6 months or so).
- **Smart meter energy consumption measurements:** the energy consumption of the EC users.
- **Flexibility needs:** the needs provided by the TSO for each market time unit (MTU). Can be updated every 15 minutes.
- **Flexibility bids:** the EC aggregated bids provided to the market.
- **Activation signals:** the activation signals for the EC, after-market clearing.
- **Settlement prices:** the settlement values for the EC, aggregated and disaggregated.

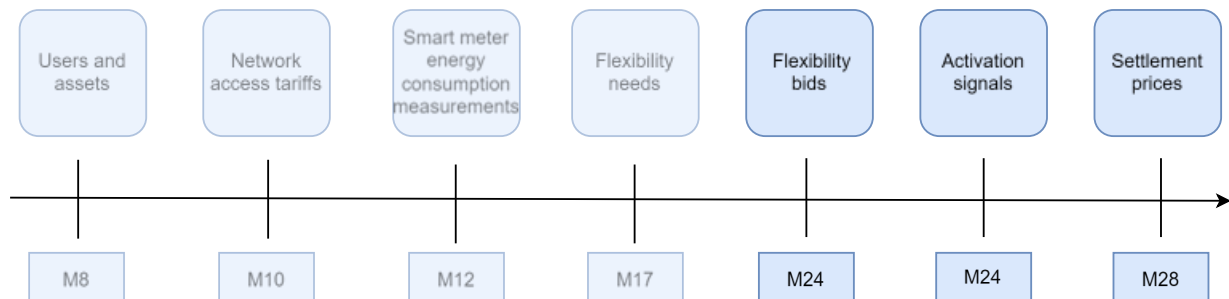


FIGURE 34 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE
 TIMELINE OF DATASETS – INESC

4.5.4 Integration with HEDGE-IoT interoperability framework

Currently, nothing has changed regarding the plan, detailed on D3.3, for the integration of this service with the interoperability framework of the project. The service will integrate with the EdgeConnect middleware (described on service 4.1), which provides an interface that will connect with HEDGE-IoT's data space, via the EDC connector. Furthermore, it will integrate with the App Store of the project, by being catalogued and packaged into it.

The integration with the EdgeConnect middleware has already started and is progressing according to the expected timeline.

4.5.5 Implementation details

At this moment, the platform is established at TRL 6, since it is part of a pilot in another European project (BeFlexible¹⁷). However, the service dedicated to the frequency restoration reserve is at TRL 4, since we have now started testing it in a controlled environment, using real IoT data from clients and example flexibility needs provided by PT partners responsible for the market simulator service (R&D NESTER).

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 (service 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.

4.5.5.1 Functionalities

Table 22 shows the list of functionalities comprehended by this service.

TABLE 22 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Integration of users and assets with pre-qualification	Creates and associates' users to an EC. Associates assets to the users and provides tests to pre-qualify the assets	100%	M13
Energy Management module	Optimizes the energy management internal to the community. Disaggregates and sends the activation signals to the assets.	80%	M30

¹⁷ <https://beflexible.eu/>

Transactions module	Identifies, calculates and routes transactions between EC users, based on energy allocation	100%	M13
Flexibility provisioning module	Performs the optimization of flexibility based on TSO needs and user comfort options. Aggregates a bid and sends it to market.	70%	M26
Settlement module	Calculates and disaggregates the settlement according to the data received from the market	50%	M30
Data storage	Stores the data from the EC's processes in a modular, secure and decoupled manner	100%	M8

4.5.5.2 Integration and dependencies

This section describes how the Energy Community Management service integrates with other services in the Portuguese Pilot.

1. [Service 4.1] EdgeConnect:

- a. Middleware to communicate with the market platform and the TSO
- b. Receives flexibility needs, settlement data and activation signals, previously relayed by the TSO (Transmission System Operator)
- c. Sends baselines, flexibility bids and metering data for the settlement process

2. CEVE Metering

- a. Receive aggregated energy consumption data from the smart meters of the DSO (CEVE)

3. [Service 2.2] FL Residential Forecasting Service

- a. Receives forecasts for day ahead consumption, based on real-time data collected on each household

4. App Store Integration:

- a. Packaging and deployment to the app store

Figure 35 shows a SGAM diagram that maps the integrations and dependencies of the service with the remaining PT pilot services, as well as showcases the status of the implementation of each integration.

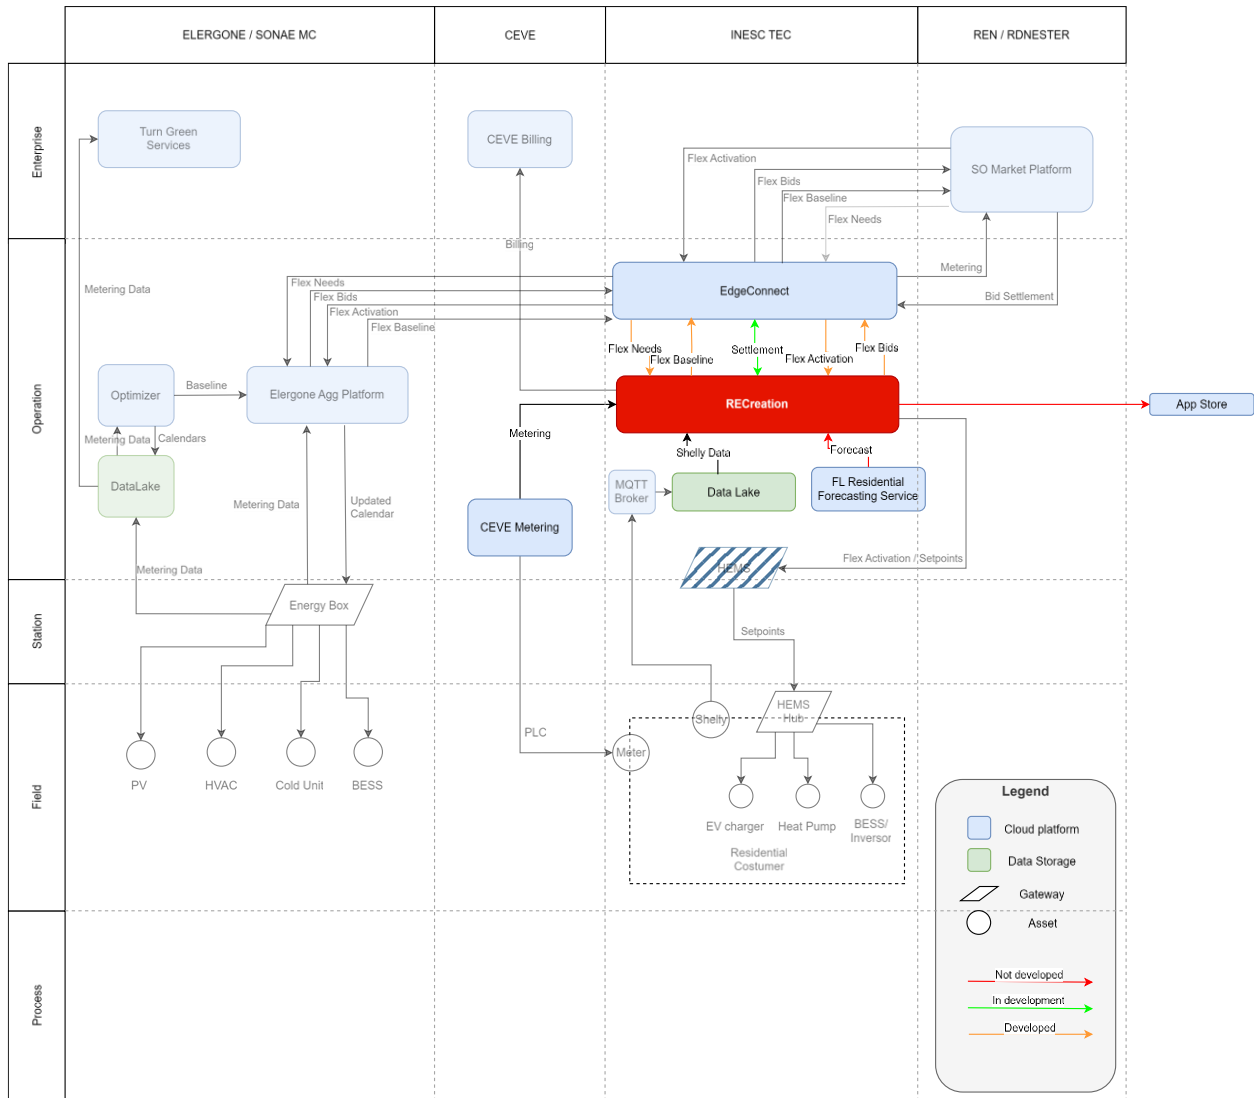


FIGURE 35 – ENERGY COMMUNITY MANAGEMENT SERVICE FOR FREQUENCY RESTORATION RESERVE PILOT SGAM DIAGRAM – INESC

4.6 LOCAL FLEXIBILITY MARKET PLATFORM

4.6.1 Description of technology enabler

The Local Flexibility Market (LFM) Platform (Figure 37) is a new technological enabler introduced in this deliverable. Its primary purpose is to enable and manage flexibility trading between Flexibility Service Providers (FSPs), Market Operators (MOs), and System Operators (SOs).

The platform is designed to support grid stability and enhance market efficiency by facilitating a structured flexibility market framework. The core processes supported by the LFM Platform include:

- The submission of flexibility requests and offers.
- Feasibility assessment of potential trades.
- Market clearing procedures.
- Settlement processes.

These functionalities are supported by robust information exchange and coordination mechanisms between all involved actors, ensuring transparency and responsiveness in market operations.

The current implementation of the LFM platform focuses on delivering the following core functionalities:

- **Role-based registration and pre-qualification:** A secure process for different market actors to register and be validated.
- **Asset Registration and prequalification:** A dedicated process for FSPs to register their flexibility assets for market participation.
- **Portfolio Management:** Allowing FSPs to create flexibility portfolios and assign their registered assets to them.
- **Market Time Unit (MTU) Management:** Daily generation and configuration of MTUs and their corresponding trading windows.
- **Trading Gates Operation:** Accurate and automated operation of trading gates according to the defined market schedule.
- **Market Clearing:** Execution of the flexibility trading resolution algorithm to clear the market, with validated outcomes that reflect market conditions.

Regarding the technical implementation, the LFM Platform's architecture (Figure 36) is structured as follows:

- **Deployment Environment:** The entire application is hosted on **Amazon Web Services (AWS)**, running on a dedicated **EC2 instance**.
- **Architecture Model:** It follows a modular, **containerized approach using Docker**, which enhances scalability and maintainability. The system is comprised of four distinct service containers:
 - **Frontend:** The container responsible for the user interface (UI), through which all market actors interact with the platform.

- **Backend:** This container encapsulates the core business logic, including user management, market processes, and the execution of the market resolution algorithms.
- **Database:** A **PostgreSQL** container that provides persistent data storage for all platform data, such as user information, assets, bids, and market results.
- **Message Broker:** A **RabbitMQ** container that facilitates asynchronous communication between services, primarily intended for handling real-time user notifications.
- **Communication Flow:** The Frontend communicates with the Backend via synchronous **REST API calls**. The Backend, in turn, orchestrates all operations and interacts directly with the PostgreSQL database.

The next and final iteration of this deliverable (D3.5), the development plan is focused on finalizing the platform's core trading lifecycle and preparing it for production use. Key objectives include:

- The full implementation of the **notification system** to inform users about bid status, trading results, and gate timings.
- The deployment and validation of the **settlement and clearing mechanisms** to finalize the trading lifecycle.
- Comprehensive **security hardening** of all exposed API endpoints, including robust authentication and permission layers.
- The establishment of a **user feedback loop** to collect insights and perform functional refinements.
- **Performance optimization** and scalability testing to ensure system responsiveness under realistic load scenarios.

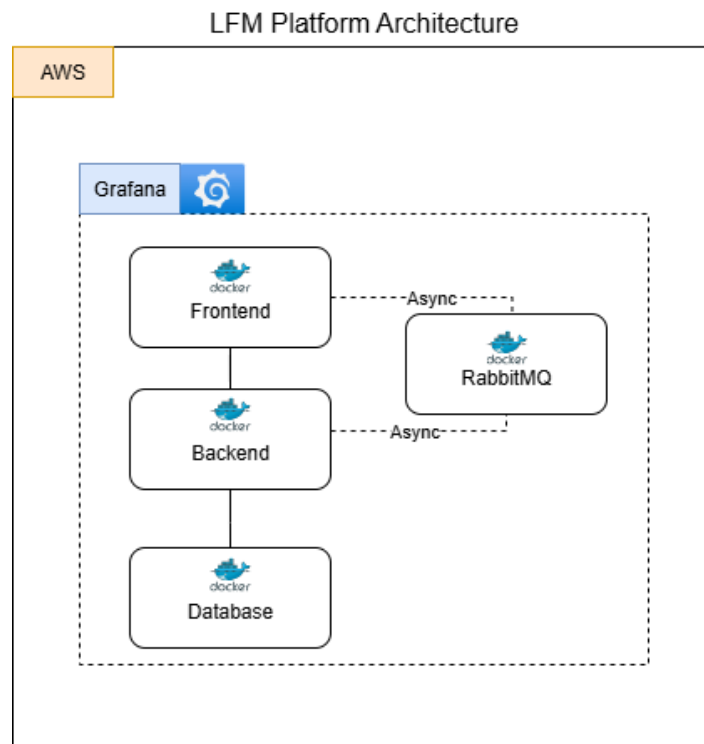


FIGURE 36 – LOCAL FLEXIBILITY MARKET PLATFORM ARCHITECTURE DIAGRAM – HENEX

4.6.2 Innovative aspects

This service pioneers the application of marginal pricing auction mechanisms within distribution-level electricity networks, addressing the emerging challenges of decentralized energy resources. The benefit of this approach is the establishment of **transparent and economically efficient price signals for flexibility**. This creates a level playing field for all market participants, from large aggregators to individual prosumers, and helps minimize the overall procurement cost for the System Operator.

The technological innovation lies in how the platform addresses the performance, scalability, and resilience challenges of a high-volume, distribution-level market. Instead of a traditional monolithic design, our approach leverages modern cloud-native principles, such as *containerization* and *microservice*, to ensure efficiency.

The core optimization algorithm is built as a self-contained, lightweight microservice. It is packaged as a Docker container and managed by an orchestration platform like Kubernetes.

This architecture introduces innovation by tightly coupling modern cloud-native infrastructure with the specific computational needs of the LFM platform's core algorithms—such as matching, market clearing, and settlement.

- **On-demand Scalability:** The platform is designed using container orchestration (e.g., Kubernetes), allowing the computational layer executing market-clearing and settlement algorithms to scale dynamically. During market peaks (e.g., bid submission deadlines or settlement windows), the system can spin up parallel instances of the algorithm engine, reducing latency and ensuring timely responses. For example, the settlement process benefits from this by distributing heavy computations (e.g., reward allocation, constraint resolution) across nodes, minimizing delay and supporting near-real-time execution.
- **Dependability Mechanisms:** The platform leverages cloud-native fault-tolerance features—such as automatic failover, health checks, and service mesh-level observability—to ensure continuity and correctness. If a compute node fails mid-process, orchestration tools reassign the task without human intervention, preventing market disruption or inconsistent results. This resilience enhances the reliability of time-critical processes like gate closure or price publication.
- **Maintainability and Continuous Delivery:** The modular design allows individual services (e.g., the pricing engine or settlement module) to be updated independently. This supports rapid iterations and bug fixes, reducing downtime and aligning with DevOps KPIs such as *Mean Time to Recovery (MTTR)* and *Deployment Frequency*.

These innovations directly support key performance indicators (KPIs) defined for the LFM service, such as:

- **System Availability ≥ 99.9%**
- **Maximum Market Algorithm Solving Latency ≤ 10 seconds**
- **Successful Auction Execution = 100%**
- **Scalability to ≥ 10,000 active assets with no performance degradation**

Overall, this architecture does not merely support the service—it actively enhances the performance, resilience, and reliability of critical LFM operations, for example system availability is critical when solving market auctions.

4.6.3 Service data

The LFM Platform processes several key datasets to facilitate the operation of the flexibility market. The data is stored locally within the platform's database and is managed according to the defined roles and permissions of the market actors. The three primary data groups are:

- **Company and User Information:** This dataset includes registration details for all participating entities (FSPs, SOs, MOs) and their respective users, including authentication credentials and roles.
- **Asset and Portfolio Data:** Contains detailed technical specifications of the flexibility assets registered by FSPs, along with information on how these assets are grouped into portfolios for market participation.
- **Trades and Results Data:** This includes all data related to market operations, such as the bids and offers submitted by participants, the market clearing results for each MTU, and final settlement information.

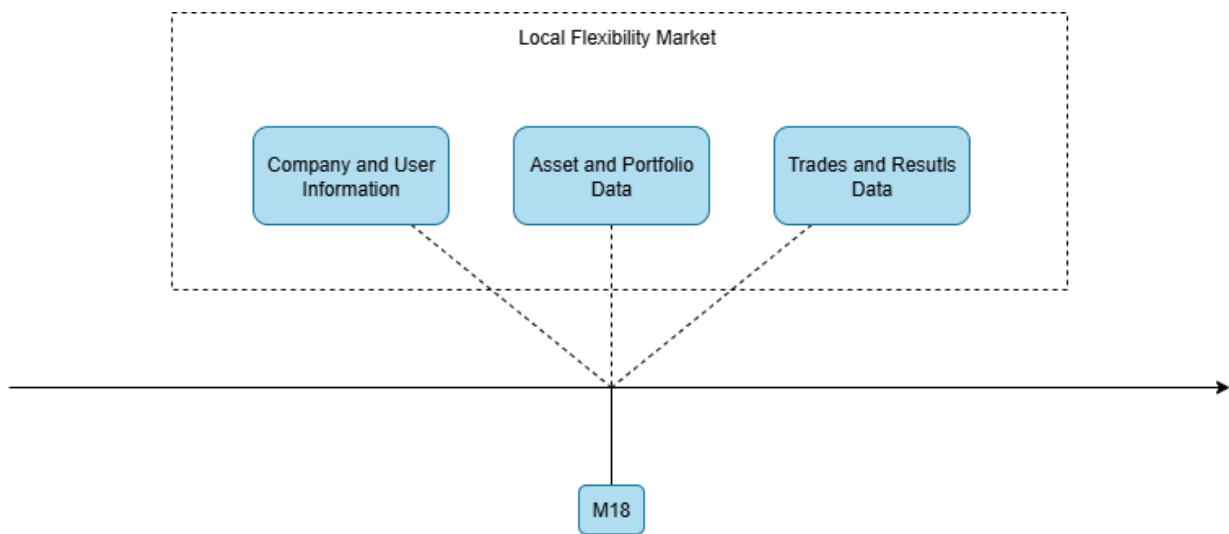


FIGURE 37 – LOCAL FLEXIBILITY MARKET PLATFORM TIMELINE OF DATASETS – HENEX

4.6.4 Integration with HEDGE-IoT interoperability framework

The LFM (Local Flexibility Market) Platform, in its current design and implementation phase, operates as a standalone service. However, the LFM platform is designed with interoperability in mind, in alignment with the broader objectives of the HEDGE-IoT framework. Specifically, the platform architecture anticipates future data exchange and integration scenarios and thus considers compatibility with relevant smart grid and IoT standards. These include:

- **IEC 61850:** While not directly implemented at this stage, the platform acknowledges IEC 61850 as a key standard for communication networks and systems in substations and distributed energy resources, particularly relevant for potential integration with field devices or DSO infrastructure.
- **IEC 61970 / 61968 (CIM - Common Information Model):** These standards are relevant for information exchange with grid management systems and may guide future extensions of the platform's data models.
- **Web-based standards (REST, OAuth 2.0, JSON-LD, etc.):** The LFM platform adopts modern web standards that support API-based integration and secure service interoperability, facilitating future interconnection with external services in a standardized manner.

As a result, while the SGAM diagram at this phase would depict the LFM platform in isolation, its architecture is aligned with interoperability best practices to enable smooth integration with other HEDGE-IoT components or third-party systems in subsequent phases of development.

4.6.5 Implementation details

The initial set of core functionalities has been implemented and is now in a continuous validation and debugging phase. These implemented functionalities include:

- Role-based registration and prequalification.
- Asset Registration and prequalification for FSPs.
- Portfolio Management.
- Market Time Unit (MTU) Management and trading window configuration.
- Automated operation of trading gates.
- Execution of the market clearing algorithm.

The expectation for M30 (for the D3.5 deliverable) is to have the service fully implemented and stable, ready for large-scale pilot activities. This status will be achieved through the completion of the remaining key features, which are:

- The full implementation of the **notification system**.
- The deployment and validation of the **settlement and clearing mechanisms**.
- Comprehensive **security hardening** of all exposed API endpoints.
- The establishment of a **user feedback loop** for functional refinements.
- **Performance optimization** and scalability testing.

4.6.5.1 Functionalities

Table 23 shows the list of functionalities and tasks comprehended by this service.

TABLE 23 – LOCAL FLEXIBILITY MARKET PLATFORM FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Role-based registration and prequalification	A secure process for different market actors to register and be validated.	90%	M20
Asset Registration and prequalification	A dedicated process for FSPs to register their flexibility assets for market participation.	90%	M20
Portfolio Management	Allowing FSPs to create flexibility portfolios and assign their registered assets to them.	90%	M20
Market Time Unit (MTU) Management	Daily generation and configuration of MTUs and their corresponding trading windows.	100%	M18
Automated operation of trading gates	Accurate and automated operation of trading gates according to the defined market schedule.	100%	M18
Execution of the market clearing algorithm	Execution of the flexibility trading resolution algorithm to clear the market.	100%	M18
Implementation of the notification system	To inform users about bid status, trading results, and gate timings.	10%	M32
Deployment and validation of the settlement and clearing mechanisms	Finalizing the trading lifecycle and enabling financial reconciliation.	10%	M32

4.6.5.2 Integration and dependencies

The HenEx Local Flexibility Market (LFM) platform plays a pivotal role in enabling intra-day flexibility trading within the Greek Pilot. The Flexibility Optimization Service on the aggregator side is tightly integrated with several services and endpoints on the LFM platform, forming a closed-loop system that supports bid submission, trading and clearing.

LFM Platform Integration

Through this interface:

- The DSO submits flexibility requests, quantifying the needed flexibility in terms of both upward and downward capacity.

- The aggregator submit flexibility bids as the price-quantity pairs, that can be dispatched based on current and forecasted grid conditions.

These submissions are triggered by the Flexibility Dispatch and Bid Formulation Service, which relies on inputs from demand forecasting, production forecasting, and NILM disaggregation.

Prequalification and Compliance Mechanisms

To participate in the LFM, each aggregator must undergo a prequalification process facilitated by HenEx. This includes validating the technical readiness of assets, verifying historical data, and ensuring that baseline and bid submissions meet regulatory requirements. The aggregator’s backend system integrates with LFM compliance checks via structured APIs, ensuring data integrity and eligibility at every bidding cycle.

Trading & Settlement Integration

The DSO flexibility requests, and aggregator flexibility bids are collected by the Trading and Settlement service, which then executes a market auction which calculates the market clearing prices and accepted bid quantities.

Coordination with DSO/TSO Services

The LFM platform also facilitates bid validation and system coordination via integration with the DSO & TSO Grid Management Service (Figure 38). This integration enables:

- The submission of baseline consumption/generation profile for each aggregator portfolio.
- Coordination with national grid needs, particularly during stress or high-demand scenarios.

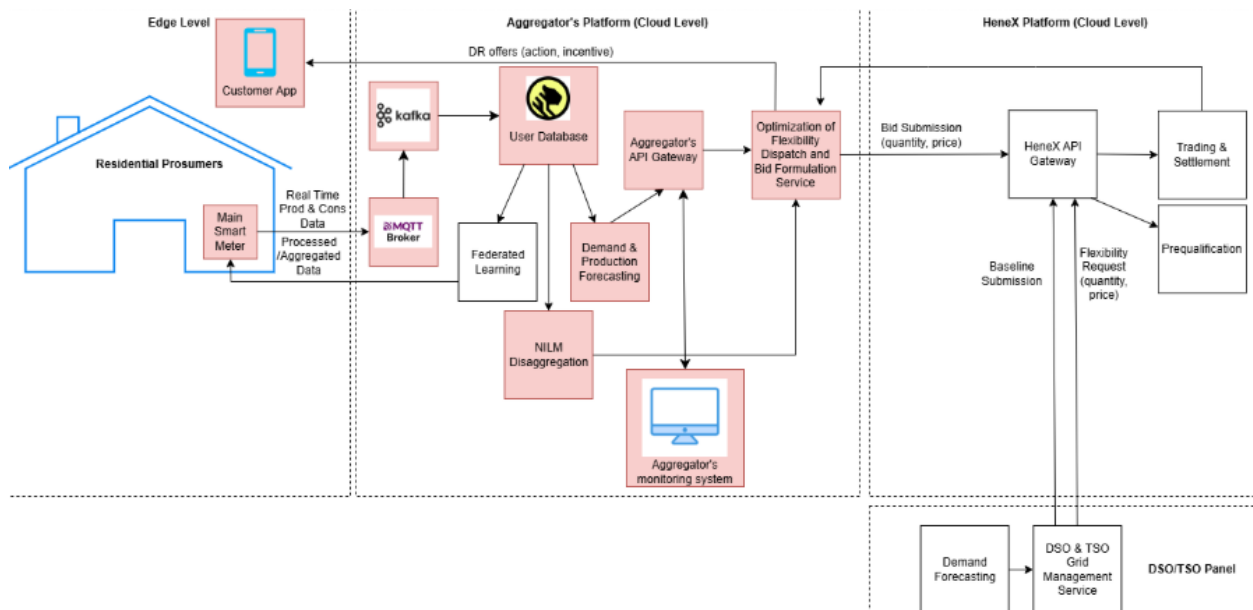


FIGURE 38 – LOCAL FLEXIBILITY MARKET PLATFORM PILOT SGAM DIAGRAM – HENEX

4.7 ENERGY COMMUNITY PLATFORM

4.7.1 Description of technology enabler

Apio Balance is a fully featured Distributed Energy Resources (DERs) management platform (Figure 39) with first class support for Renewable Energy Communities (RECs). It supports the integration of heterogeneous DERs, from EVs to BESS, allowing aggregators to manage them, configure them as Virtual Power Plants (VPPs) and exploit their flexibility maximizing both energy and economic performances.

D3.3 described the Apio IoT Platform and outlined its integration with the PGUI (Power Grid User Interface) for obtaining data from the field and ML methods to forecast load and production data.

The same IoT Platform is used within Balance (on a separate instance, because the users are not the same) to power its integration with field EMS, Edge Devices and IoT Devices.

The Same ML forecasting methods are used within Balance to forecast production and load baselines to assist the user in formulating the best possible flexibility offer.

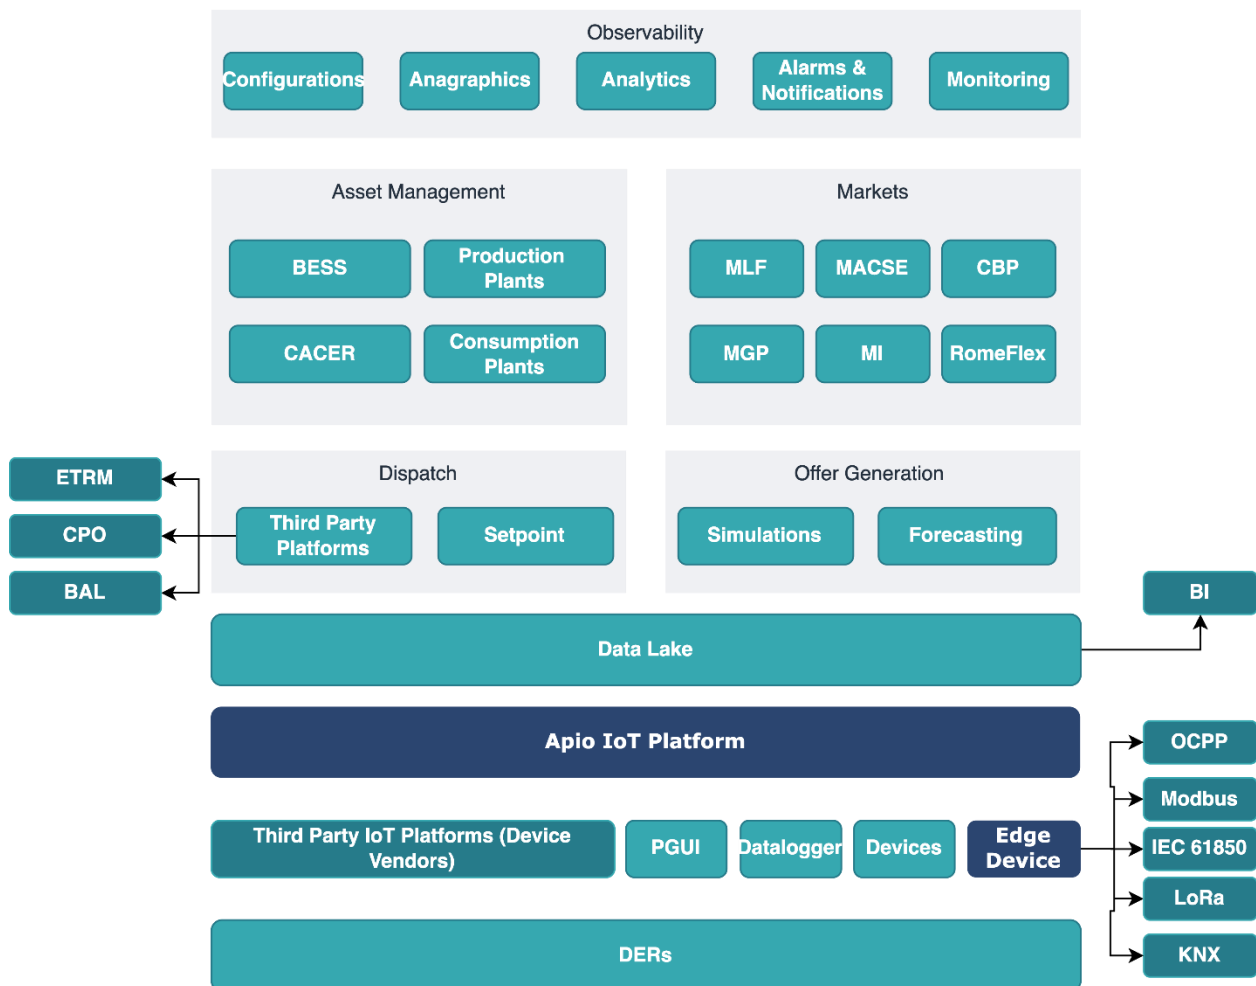


FIGURE 39 – ENERGY COMMUNITY PLATFORM ARCHITECTURE DIAGRAM - APIO

4.7.2 Innovative aspects

REC INTEGRATION

The system is designed to handle complex scenarios in which DERs participate in multiple market structures simultaneously – such as the energy market combined with REC (Renewable Energy Community) incentives. In such cases, a cost-benefit analysis is required to compare standard market offers with local incentive schemes. To address this, the platform employs a **multi-stage approach**: it first establishes a baseline scenario incorporating REC incentives and subsequently evaluates additional flexibility services. This allows for the optimization of REC’s energy costs while maximizing revenue from flexibility provision.

DYNAMIC AGGREGATION AND OPTIMIZATION

The platform offers advanced **dynamic DER aggregation strategies**, enabling manual and real-time optimization across heterogeneous pools of DERs. By functioning as a VPP platform, it overcomes the limitations of traditional, invitation-based VPP models, where DERs are statically assigned. This dynamic approach provides significant advantages also to REC administrators, allowing them to plan and optimize the performance of distributed self-consumption configurations.

HOLISTIC, MULTI-SYSTEM APPROACH

The platform adopts a holistic and multi-system perspective, enabling the **integration of heterogeneous DERs within the same asset portfolio**. From photovoltaic plants to battery energy storage systems, from electric vehicles to smart charging stations – any resource that can be profiled and dispatched for flexibility provision is seamlessly aggregated. This comprehensive approach allows for the coordinated management of diverse technologies, each characterized by distinct technical constraints and operational profiles (e.g., ramping limits, activation times, availability windows).

4.7.3 Service Data

The platform uses information about DERs and measurements, to compute baselines, forecasts, and optimal flexibility offers, which timeline is described on Figure 40.

- REC Members information profiles: technical information such as maximum absorbed/injected active power, maximum increase/decrease in absorption/injection, service timings such as ramp timings, minimum/maximum activation time.
- REC Members historical consumption/injection: time series of energy (and power) consumption and injection.
- REC Members consumption/injection forecasts: time series of forecasted energy (and power) consumption and injection
- REC Members measured consumption/injection and other available data, depending on the field equipment (TBD) provided as time series of energy (and power) consumption and injection

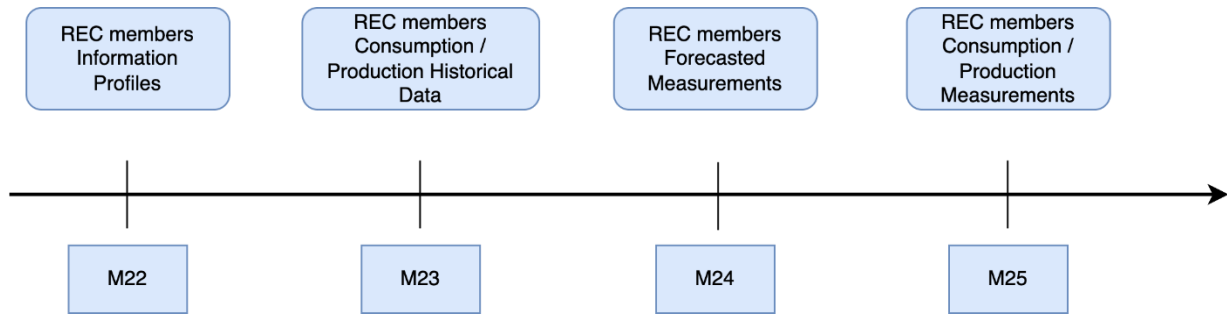


FIGURE 40 – ENERGY COMMUNITY PLATFORM TIMELINE OF DATASETS – APIO

4.7.4 Integration with HEDGE-IoT interoperability framework

In the previous iteration of this deliverable (D3.3), it was described that this service would integrate with the HEDGE-IoT interoperability framework by describing the PGUI data flows through a SAREF based ontology.

An extended adoption of SAREF is being considered for some of the data flows of the platform, with the aim of introducing interoperability in the interface with both markets and the field (as illustrated in the SGAM architecture, including flows to the Market Interface Platform and to Edge Devices).

4.7.5 Implementation details

In the previous iteration, only the forecasting capability of the platform was described as the main area of application of ML methods.

The implementation of the basic workflow of the Apio Balance platform was completed by M18, including integration with the Market, with the Flexibility Register, the ability to compute the best offer considering both market and REC incentives, DERs management, basic analytics functionalities, forecasting and the capability to receive activation commands from the DSO to dispatch grid limits. (BUC 2).

As for the writing of this deliverable, all features planned for M18 are already completed and in testing. Since RECs are not yet enrolled, all the data flows originate from simulated data.

At the beginning of the full demo operation phase (M21), REC members are planned to be available and connected to the local flexibility market, this will allow to test all the features of the SUCs.

At M33, with the large-scale demo operation phase, it is planned to be able to test interoperability features and have integrations with other HEDGE-IoT Pilots components.

4.7.5.1 Functionalities

TABLE 24 shows the list of functionalities comprehended by this service.

TABLE 24 – ENERGY COMMUNITY PLATFORM FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
---------------	-------------	-------------	----------------------------------

Baseline Computation	P2P procedure where data is shared and partially encrypted by all peers.	100%	M18
Flexibility Offer Optimization	P2P procedure where each agent trains its model locally and then shares encrypted coefficients with the remaining peers. Additional optimization algorithms may be introduced later.	100%	M18
Energy community power management optimization	Service registered and packaged in the HEDGE-IoT App Store	100%	M18
Integration with the Local Flexibility Market	Integration with the Market Interface Platform and the Flexibility Register in order to implement all the required data flows (requests, offers, measurements, baselines)	100%	M18
Prices Fetching	Fetching real time energy prices	0%	M21
Load Forecasting	ML Based and ARIMA based methods for Load forecasting	100%	M18
Production Forecasting	ML Based and ARIMA based methods for Production forecasting	100%	M18

4.7.5.2 Integration and dependencies

This section describes how the ECP platform integrates with other services in the Italian Pilot (Figure 41).

1. Market Interface Platform

- a. Receives Flexibility Offers
- b. Sends Flexibility Requests

2. Flexibility Register

- a. Sends Baselines
- b. Sends Measurements
- c. Sends resource registrations

3. Apio IoT Platform

- a. Provides behind the meter data, or more in general, field data.
- b. Provides the activation order generated by the DSO

4. Blockchain Access Layer

- a. Provides the Activation Order from the DSO

5. Field Devices and Edge Devices

- a. Receives measurements relative to behind-the-meter devices
- b. Sends activations to EMS

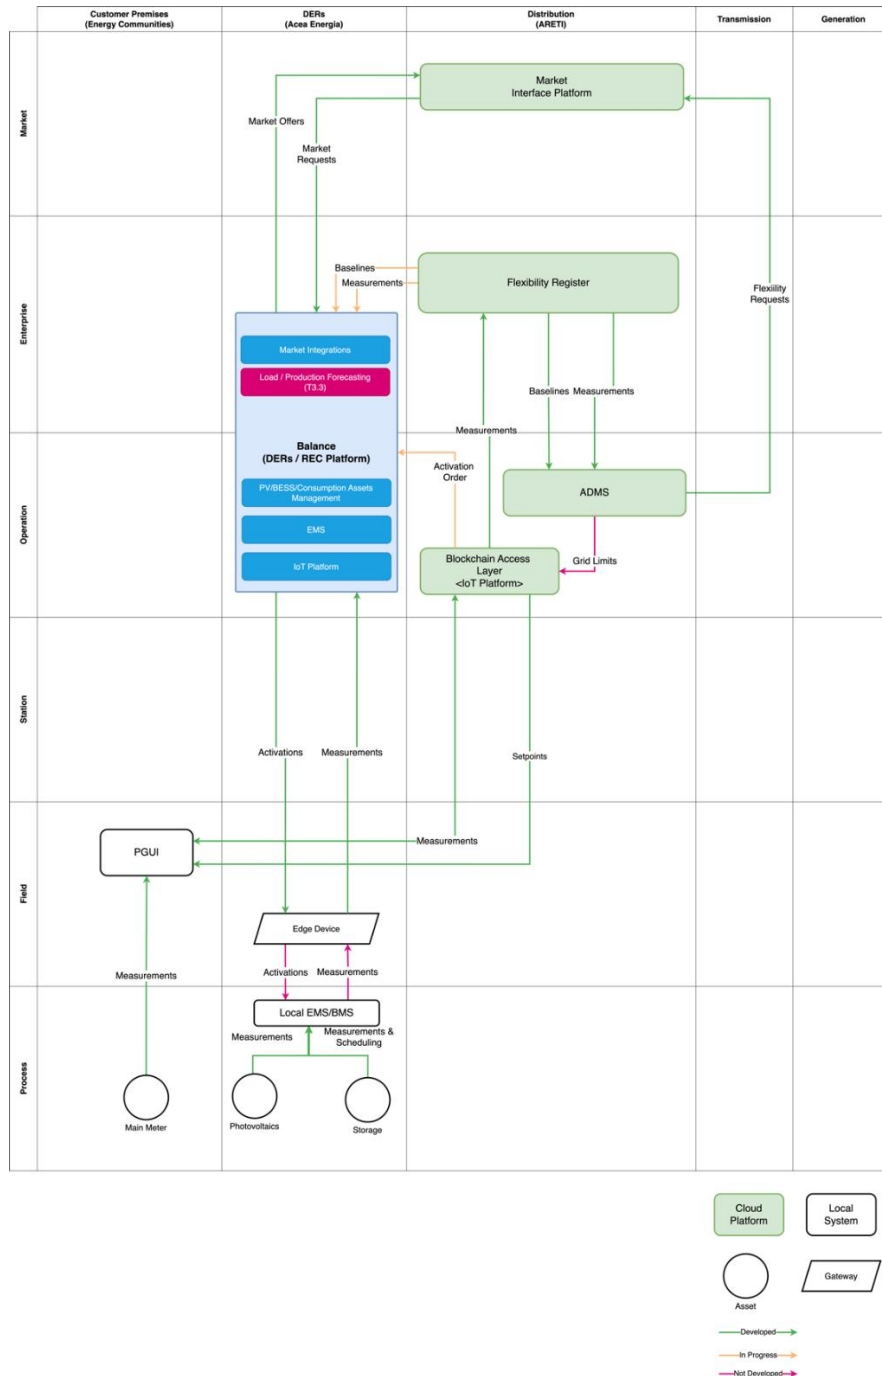


FIGURE 41 – ENERGY COMMUNITY PLATFORM PILOT SGAM DIAGRAM – APIO

4.8 TURNGREEN – OPTIFLEX

4.8.1 Description of technology enabler

OptiFlex is a Flexibility Optimization Service (Figure 42) designed and operated by Elergone/Sonae, focused on the aggregation, monitoring, and control of flexible assets in tertiary buildings within the Portuguese HEDGE-IoT pilot. The service is fully integrated with the Sonae MC infrastructure through the EnergyBox gateway, which enables secure connection of multiple assets (such as HVAC, industrial cooling, PV, and storage) via a dedicated VLAN.

OptiFlex combines rule-based automation and AI-powered forecasting algorithms to optimize the operation of these assets, aiming to maximize economic value while improving overall energy efficiency. The AI algorithms are trained in historical, weather, and market data, supporting predictive scheduling, baseline setting, and flexibility bids. The service operates continuously to adjust asset operation in response to price signals, flexibility requests, and grid conditions. Also, it will consider all user inputs and system constraints, such as comfort requirements and operation limits.

Additionally, OptiFlex enables cross-sector flexibility by orchestrating different types of loads and storage assets within the building environment, thereby unlocking synergies between heating/cooling, renewable generation, and storage. The platform is designed to reduce energy consumption during peak periods and to support the integration of renewable energy sources, contributing to both operational efficiency and sustainability goals.

OptiFlex allows automated and real-time management of these assets, maximizing economic valorization through both default price hedging scenarios and participation in external TSO requests (via EdgeConnect). The platform operates under strict data privacy and GDPR compliance, as required by EU guidelines and the HEDGE-IoT Grant Agreement.

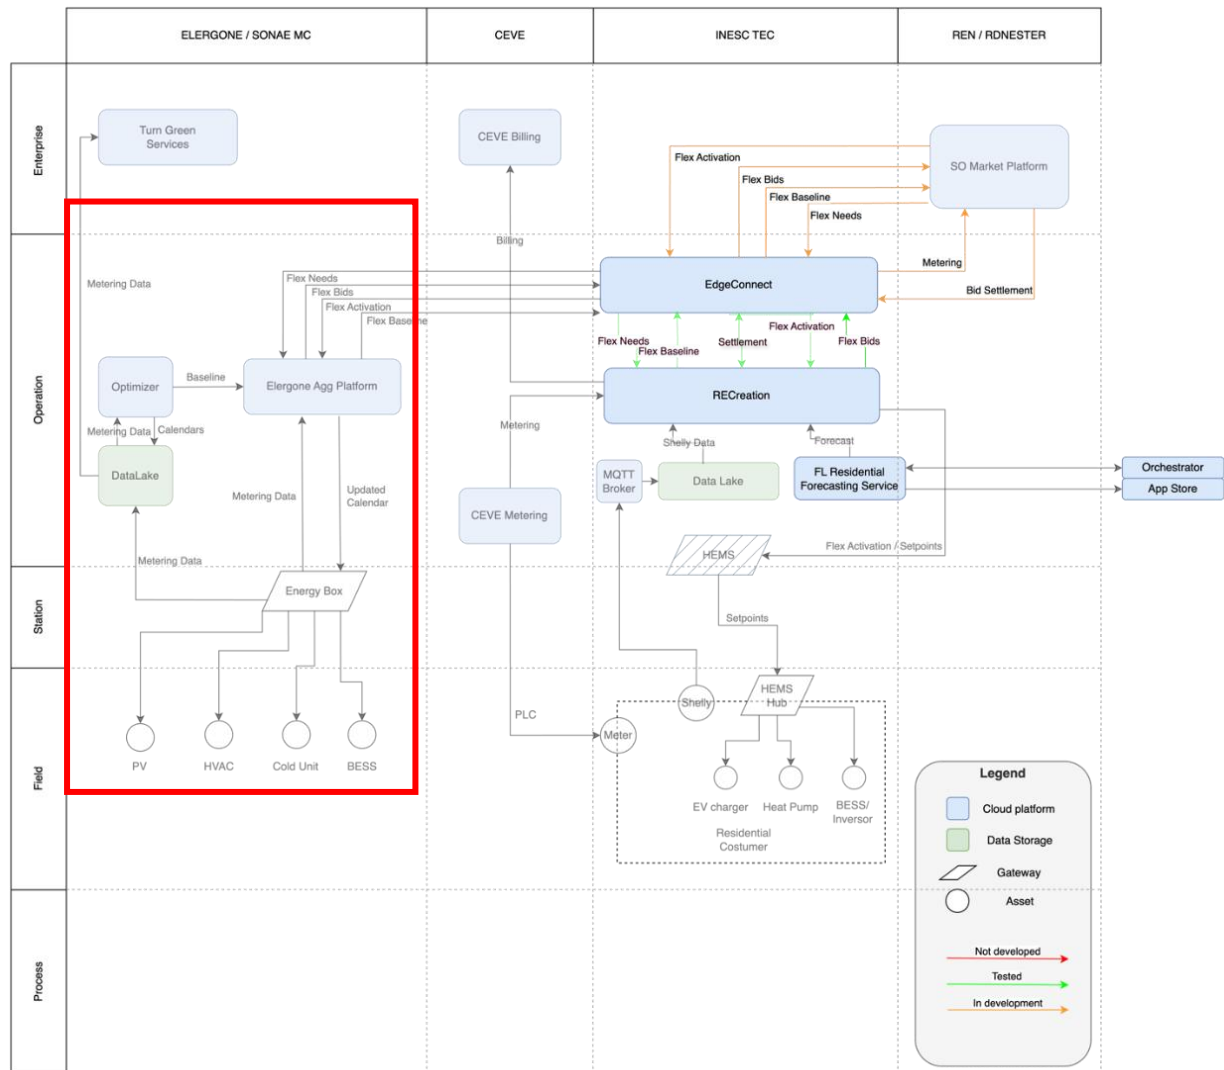


FIGURE 42 – TURNGREEN – OPTIFLEX ARCHITECTURE DIAGRAM – SONAE/ELERG

4.8.2 Innovative aspects

DYNAMIC ASSET OPTIMIZATION

OptiFlex enables the real-time and scheduled optimization of diverse energy assets, adapting asset operation (charging, discharging, modes such as ECO/BOOST) to external signals, such as price curves or grid flexibility requests.

ADVANCED BUSINESS LOGIC & AGGREGATION

The platform incorporates advanced business rules, including price and surplus arbitrage, portfolio-wide aggregation and disaggregation, and default bidding rules for participation in TSO markets (e.g., automated bid calculation as a function of forecasted savings).

MULTI-LAYER INTEGRATION

Supports native integration (API/MQTT), local industrial protocols (BACnet, Modbus), and seamless communication with external platforms and aggregators (EdgeConnect, the TSO interface).

USER-CENTRIC PLATFORM

OptiFlex provides a flexible user management framework, enabling customized access and dashboards for different stakeholder profiles (aggregators, portfolio managers, building operators, O&M providers).

AI-BASED FORECASTS AND DECISION SUPPORT

Internally developed algorithms leverage available datasets (consumption, price, meteo, PV forecasts) to generate optimized schedules and economic proposals, improving the competitiveness and automation level of flexibility bids.

4.8.3 Service Data

OptiFlex collects, processes, and manages multiple data types, including:

- **Asset operational data:** Real-time and scheduled measurements (e.g., consumption, state, capacity).
- **External signals:** Market prices, TSO requests, weather and PV forecasts.
- **Network and infrastructure data:** Asset registration, VLAN connectivity, and EnergyBox status.
- **KPIs and compliance data:** Asset availability, activation rates, GDPR compliance logs.

4.8.4 Integration with HEDGE-IoT interoperability framework

OptiFlex is designed for seamless interoperability with the HEDGE-IoT ecosystem, primarily through the EnergyBox and open API endpoints. Integration efforts focus on:

- Secure data exchange with other project components (e.g., EdgeConnect for market/TSO interaction).
- Compliance with data space and interoperability standards (such as Eclipse Dataspace Connector), as required in the project.
- Ongoing validation of GDPR and data protection practices in collaboration with Sonae MC teams.

The focus remains strictly on commercial buildings and their assets under Sonae/Elergone management.

4.8.5 Implementation details

OptiFlex is deployed in three Sonae MC pilot stores, with EnergyBox installed and VLAN connectivity established (VLAN95). Key milestones include:

- Asset onboarding and integration (HVAC, industrial cooling, PV, storage) - Ongoing, expected completion by M28.
- Validation and mapping of variables into the EnergyBox and platform data models - Ongoing, expected completion by M28.
- Implementation of business rules for price hedging and TSO valorization scenarios - Ongoing, expected completion by M30.
- Testing of privacy, security, and asset availability requirements, in line with project and EU standards - Ongoing, expected completion by M30.

Since the previous reporting period, OptiFlex has progressed with the installation and validation of the EnergyBox in three pilot stores, establishment of dedicated connectivity, integration of key assets, and the implementation of optimization algorithms and business rules adapted to Sonae MC operational needs.

4.8.5.1 Functionalities

Table 25 shows the list of functionalities comprehended by this service.

TABLE 25 – TURNGREEN – OPTIFLEX FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Asset integration (EnergyBox/VLAN)	Connection and mapping of all flexible assets through EnergyBox and VLAN infrastructure.	60%	M28
Automated schedule optimization	Generation and deployment of optimized operational schedules for each asset.	80%	M28
TSO/Market request activation	Automated response to TSO flexibility requests via EdgeConnect.	40%	M30
Data privacy and GDPR compliance	Implementation and audit of privacy/security protocols and documentation.	70%	M29
KPI monitoring and reporting	Continuous tracking of project-relevant KPIs and reporting tools.	50%	M29

4.8.5.2 Integration and dependencies

The successful operation of OptiFlex relies on integration with several key components and stakeholders within the pilot environment (Figure 43). The main dependencies and integration points are outlined below:

- **EnergyBox:** Central gateway for asset connectivity, data collection, and control execution.
- **VLAN infrastructure:** Secure and dedicated asset connectivity in pilot stores.
- **EdgeConnect:** Interface for market/TSO signals and participation.
- **Sonae/Elergone IT teams:** Ongoing support for privacy, security, and operational integration.
- **Third-party asset vendors:** Support for technical integration and mapping of all asset variables.

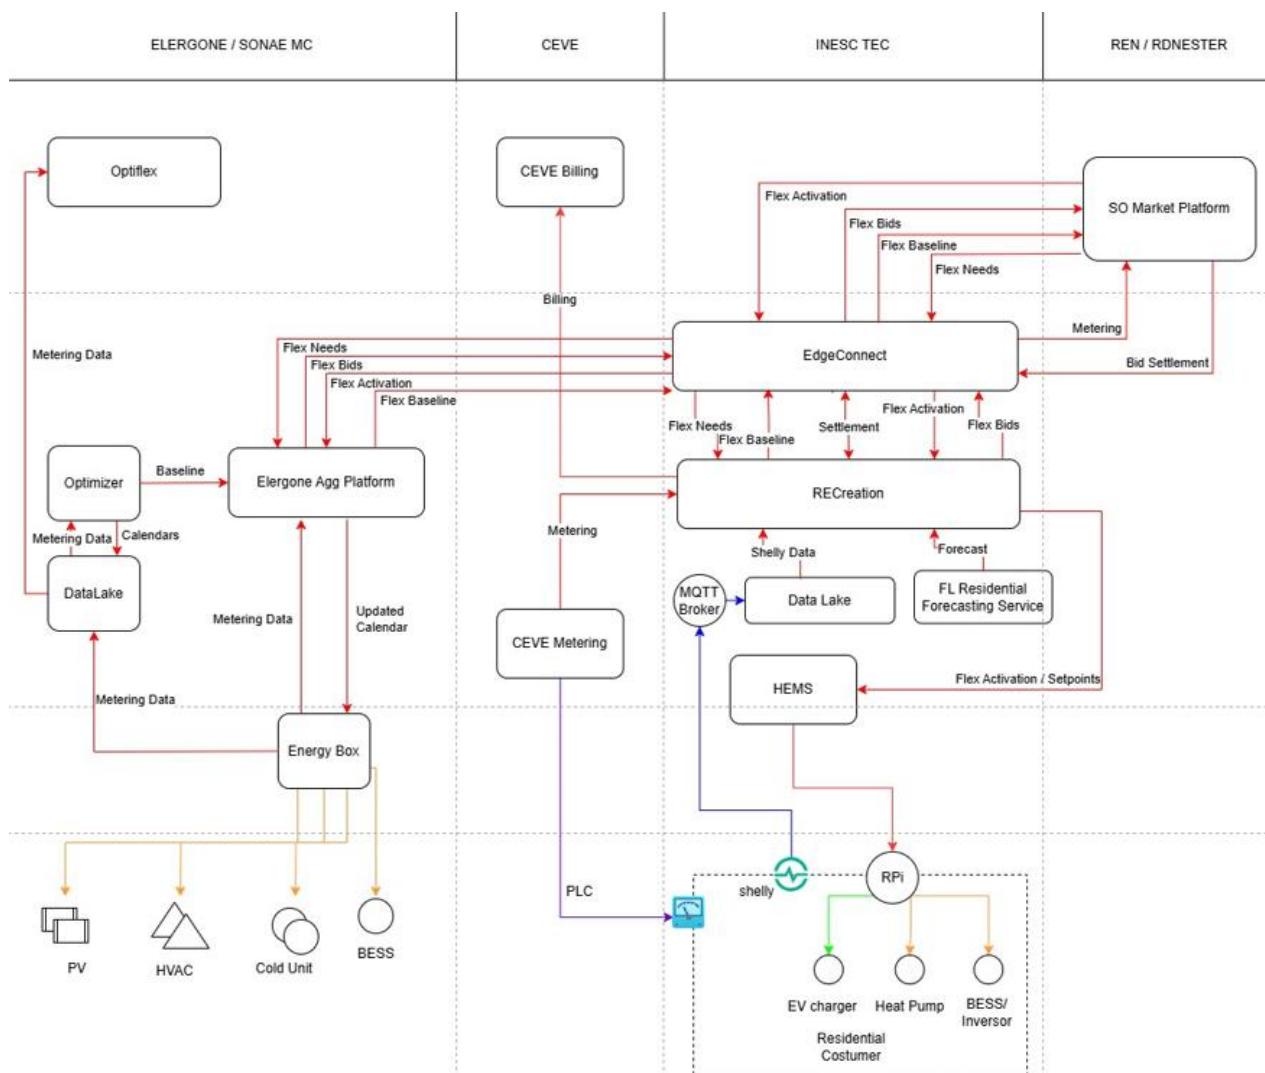


FIGURE 43 – TURNGREEN – OPTIFLEX PILOT SGAM ARCHITECTURE – SONAE/ELERG.

4.9 POWERCIM TOOL

4.9.1 Description of technology enabler

PowerCIM is a platform for storing and exchanging electrical grid models based on the IEC CIM (Common Information Model) standards. It enables the integration of data from various sources used by the distribution system operator, creating a unified, semantically aligned grid model. The platform supports versioned model management, semantic enrichment of data, and linking of telemetry time series. It consists of several modular components, including a core repository server, model viewer, standard CIM/XML adapters, and tools for metamodel and data flow management (Figure 44). Currently, it supports multiple CIM profiles such as EQ, SSH, SV, and GL.

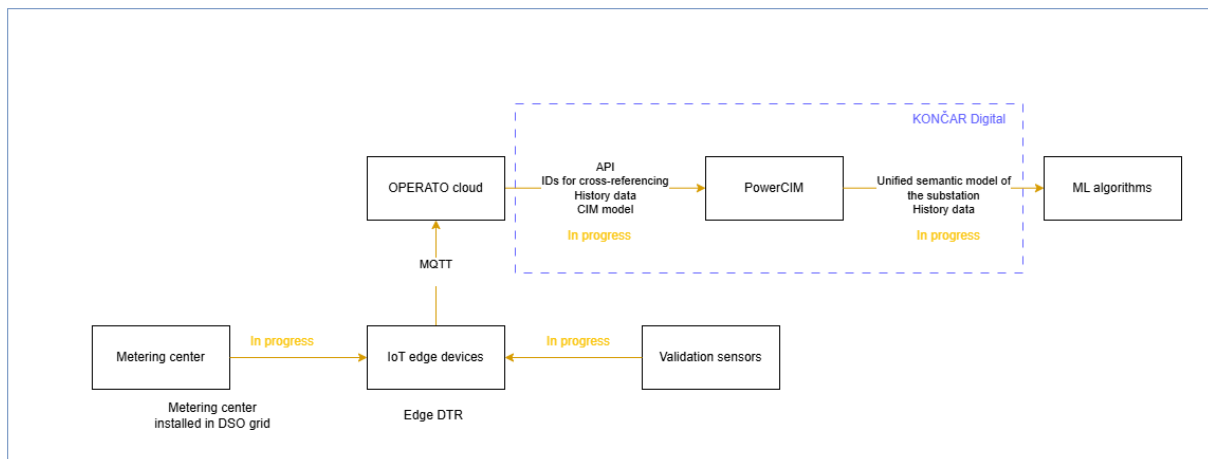


FIGURE 44 – POWERCIM TOOL PILOT SGAM DIAGRAM – KONCAR

4.9.2 Innovative aspects

The PowerCIM solution was not included in D3.3 but was described in D3.1 and D4.3. In this section, the innovative aspects of the service will be detailed.

EXTENSION OF THE SUBSTATION SEMANTIC MODEL WITH DTR DATA

The innovative aspect of this solution is the **extension of the substation’s semantic model with data from DTR (Dynamic Thermal Rating) calculations performed at the edge**. Following the IEC CIM standard, provides a standardized and integrated representation that unifies models from different systems. The model is further enhanced with weather data.

In this context, the use case successfully addresses the integration of telemetry data with CIM model, enabling their use in further analysis and advanced applications. The key contribution is **overcoming data silos and facilitating seamless data exchange within the distribution system operator’s domain**.

4.9.3 Service data

For that purpose, three datasets are used to create unified electrical grid models, as FIGURE 45 shows:

- **Preliminary data for one substation:** test data in CIM/XML and CSV formats for a single substation.
- **Deployment-phase data for one substation:** following the installation of edge devices performing DTR calculations, new data is utilized.
- **Deployment data for several substations:** the same process will be applied to multiple substations where DTR edge devices will be installed.

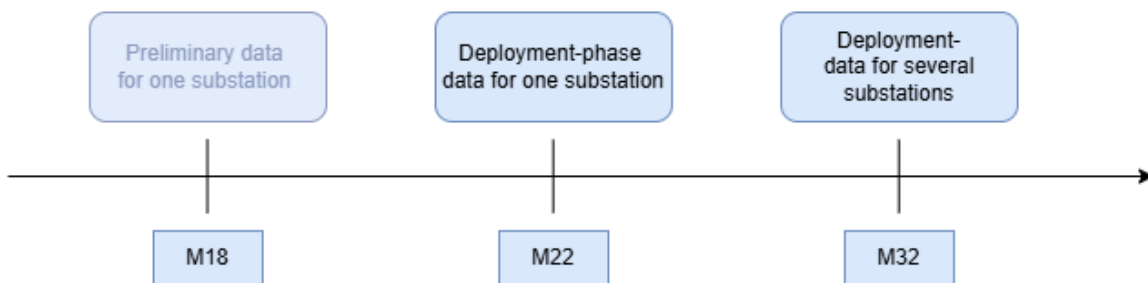


FIGURE 45 –POWERCIM TOOL TIMELINE OF DATASETS – KONC

4.9.4 Integration with HEDGE-IoT interoperability framework

PowerCIM within the HEDGE-IoT Interoperability Framework encompasses the distribution network, including transformer substations, and is based on the IEC CIM standard, which ensures semantic interoperability and data standardization. By integrating data from isolated systems that do not communicate with each other, PowerCIM enables a unified network model that serves as a foundation for advanced analytics, optimization, and application development. In this way, it addresses challenges related to network data management and supports system operators in operational control and planning. In addition to the Slovenian distribution system operator, the solution will also be used by the academic partner (UNIZG) for the development of advanced algorithms.

4.9.5 Implementation details

By Month 18 (M18), preliminary data from the GIS system was imported, followed by the creation of *AnalogValue* measurement point objects derived from time-series data. A new timeseries table was created and populated with sample data, and a cross-reference between measurement points and CIM Terminal elements was established manually. Substation measurement data and test datasets for DTR calculations were integrated. By Month 32 (M32), the integration of real operational data in PowerCIM is expected, following the installation of IoT devices in the substation, where DTR edge computations will be executed.

4.9.5.1 Functionalities

TABLE 26 shows the list of functionalities comprehended by this service.

TABLE 26 – POWERCIM TOOL FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Mapping measurement points to the substation CIM model	Each <i>AnalogValue</i> is linked to a specific terminal in the substation model through a cross-reference table (<i>AnalogValue</i> -> <i>Terminal</i>), enabling consistent mapping of measurements to the network topology.	30%	M32
CIM Extension for DTR Integration	Extended CIM transformer model with thermal parameters required for Dynamic Thermal Rating (DTR). Added support for dynamic temperature inputs and calculated thermal states, enabling integration of real-time DTR data into the substation model.	30%	M32
Functional-Asset Transformer Separation	Implemented separation between functional transformer entities (PowerTransformer) and physical assets (TransformerAsset) in the CIM model. Enables tracking of asset replacements without impacting the functional topology.	30%	M32
Frontend for fetching timeseries	Implemented a frontend functionality that fetches timeseries data for selected measurement points via API.	10%	M32

4.9.5.2 Integration and dependencies

1 [Service 3.1] Enhanced network management and planning

- a. This service uses output data in a standardized format, integrated with the CIM model, as the basis for performing calculations.

2 [Service 3.2] DTR-DLR on the edge:

- a. The results of the DTR calculation at the edge are used as input data for creating a semantic model of the substation.
- b. A database for telemetry data has been created in PowerCIM platform, which is linked to the CIM model

Figure 44 shows a diagram that maps the integrations and dependencies of the service with the remaining SL pilot services.

5 COMPUTATIONAL ORCHESTRATION TO ENSURE THE CLOUD-EDGE CONTINUUM

In the description of this section, we present the computational orchestration framework to ensure the cloud/edge continuum. In the previous iteration of this deliverable, a high-level design of the orchestrator was presented for two use-cases: edge offloading for low-latency data processing and federated learning orchestration.

In this version, we present the detailed design along with the current implementation.

Please note that this is not the final implementation, as further refinements may be made. Also, an additional use-case was identified and considered the application/federated learning models rolling out at edge ensuring seamless update of their versions at edge without service disruption.

Finally, we describe the orchestrator integration with the HEDGE-IoT framework focusing on available services, and data space connector.

5.1 DESCRIPTION OF THE ORCHESTRATION FRAMEWORK

The computational orchestration platform for smart grid services provides tailored functionalities across three different types of orchestration needs, identified in the pilots:

- **Energy services offloading at edge:** to improve data processing, reducing latency, responsiveness and geographic redundancy at the network's edge.
- **Federated AI-driven energy services:** orchestration of federated learning processes across the edge-fog-cloud continuum facilitating the efficient delivery of AI-driven applications and improvement of the training efficiency.
- **Energy services rolling out at edge:** enable automatic rolling out of new versions of the energy services through data space connector.

5.2 ENERGY SERVICES OFFLOADING AT EDGE

In the context of energy services offloading at the edge, the orchestration framework dynamically allocates computational resources across the edge, fog, and cloud layers based on real-time service demand. Its core functionality is to ensure that service requirements and service-level agreements are consistently met while optimizing resource utilization and maintaining responsiveness. By managing service placement, the framework enables efficient offloading of energy services to the edge, ensuring low latency and continuous availability.

5.2.1 Description

As described in the previous iteration (D3.3), for energy services offloading at edge, the orchestration platform leverages cluster technologies such as KubeEdge and Kube Prometheus to gather insights about the health and performance of the infrastructure components and services' state.

The updated version of the orchestration platform’s architecture is provided in Figure 46. The main contributions are regarding the scheduling and optimization components which are integrated into the platform together with a graph database used to capture the real-time state of the system.

The Eclipse Data Space Connector (EDC) has the role of enabling interoperable communication between distributed energy services. The connections between services must be preserved even after migration across different nodes. The orchestrator has its own EDC to communicate with the Open Service Catalogue and with the connectors of the energy services to access information about them and their requirements.

The blockchain network is used to ensure transparency and traceability during energy services placement and migration. When a service is placed and deployed on an edge node, an associated non-fungible token representing it will be transferred to the entity that represents the edge node on chain. Ownership of both services and edge nodes is tracked in smart contracts, and this enables also automatic payments when a resource is utilized.

Those components and their integration with the KubeEdge platform are presented in more detail in the following sections. Also, a more in-depth overview of blockchain technology and smart contracts will be presented.

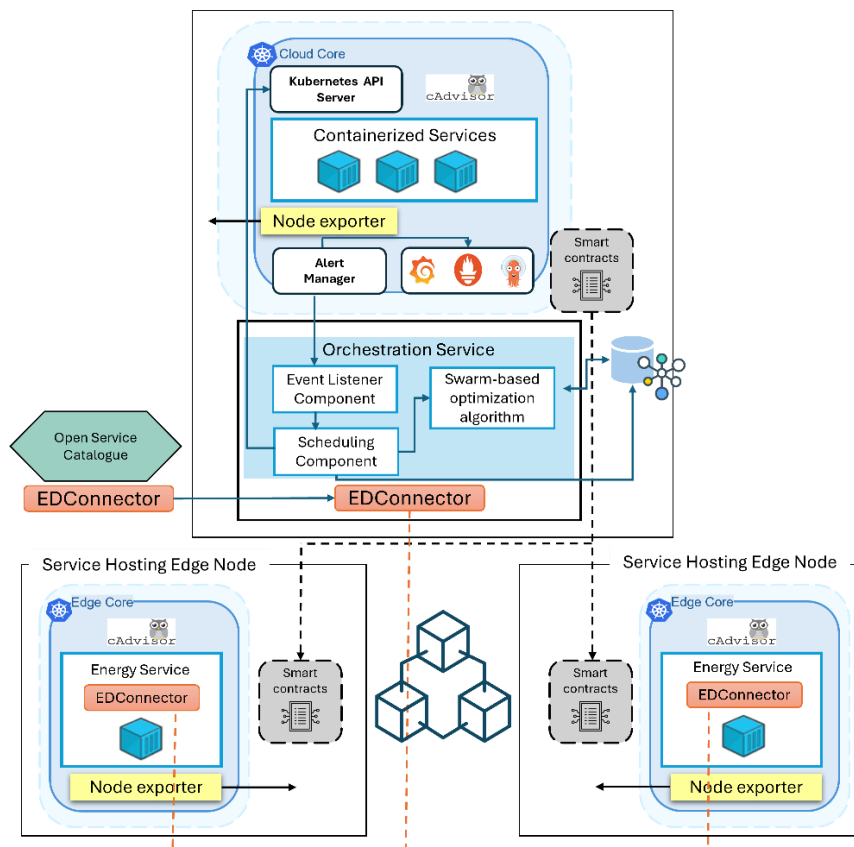


FIGURE 46 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR ENERGY SERVICES OFFLOADING AT EDGE – TUC

5.2.2 Implementation details

The orchestration service acts as a coordination layer, bringing together real-time monitoring, custom scheduling logic and SLA-driven rebalancing in a unified control loop. It consists of three components: event listener, swarm-based optimization algorithm and scheduling component.

Event listener

The event listener component is integrated with a monitoring stack and informs the scheduling component when certain events occur. This component triggers the optimization algorithm that generates a new planification of the resources and updates the state into the graph database. The service relocation is made through the Kubernetes API.

The **Monitoring Stack** used to gather all the insights needed by the computational orchestration service is presented in Figure 47.

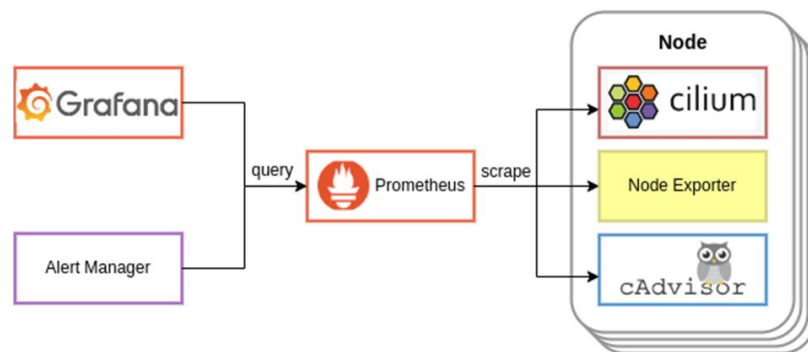


FIGURE 47 – MONITORING STACK ARCHITECTURE

Continuous monitoring of nodes and tasks within the infrastructure represent the foundation for effective workload placement and ongoing performance.

The **orchestrator service** is integrated with Kube-Prometheus stack to collect computational metrics at both the node and task levels. This monitoring layer is divided in two categories: Data Sources, deployed across all nodes in the cluster and Data Collectors, deployed on the cloud side that aggregate, store, query and act based on the collected metrics.

Although the Kube-Prometheus stack provides observability for node-level and task-level measurements it doesn't have directly integrated data sources to collect network related metrics. To address this problem, inter-node and pod-level network metrics were integrated through the cluster's CNI (container network interface). The CNI is responsible for providing network capabilities to pods and managing their communication inside and outside the infrastructure. It creates and assigns virtual network interfaces, manages unique IP allocation, handles routing and network policy enforcement. Cilium is used as the CNI, which, besides the standard functionalities, it also exposes network observability measurements between the nodes and their workloads.

The **data source components** are *cAdvisor*, *Node Exporter*, and *Cilium*. They are deployed across the cluster as DaemonSets, meaning they run on every node. These components periodically collect metrics and expose them via HTTP endpoints, which are then scraped by Prometheus. The *cAdvisor*, embedded within the kubelet on each node, gathers container-level metrics such as CPU usage, memory consumption, I/O statistics, and lifecycle events (e.g., container start and stop). Node

Exporter, also running on every node—whether in the cloud or at the edge—collects node-level metrics, including CPU and memory usage, disk utilization, filesystem statistics, load average, and hardware-related metrics like interrupts. Cilium, running as a Cilium Agent on each node, collects node-to-node network metrics, providing data on latency, drop rates, and packet loss between pairs of nodes. **Data Collection components** are *Prometheus, Grafana and Alert Manager*. They are deployed on the Cloud side to aggregate, analyse, and trigger actions based on infrastructural state.

Orchestration Service

The **Event Listener** component uses Kubernetes Client library and subscribes to Kubernetes API server events; thus, it can react fast to infrastructure and workload changes and trigger the scheduling component in case of pod lifecycle (*Pending* pods) or node status events (transitions from *Ready* to *NotReady*). Additionally, it exposes an API that is used by Alert Manager to trigger rebalancing when SLA rules are breached.

The **Scheduling Component** replaces the default Kubernetes scheduler, and it is designed to better support distributed environments and energy services requirements. It is triggered by the event listener component whenever pod lifecycle events or alerts are detected. The actual resource task placement logic has been delegated to a dedicated Swarm-based optimization algorithm, which makes decisions based on infrastructure metrics gathered from the graph database. It uses Prometheus Client to gather both computational and network metrics and after grouping them they are fed into the graph database.

Ant Colony Optimization Algorithm

The **swarm-based optimization algorithm** is responsible for assigning the available resources to running energy services (tasks). A custom Ant Colony Optimization (ACO) is developed for task offloading. The optimization objectives are to efficiently assign computational task (energy services) to edge whilst minimizes execution, communication, and energy costs. The algorithm retrieves the state of the system (edge nodes and running tasks) from the graph database. Each artificial ant builds a solution by probabilistically selecting execution nodes for tasks, guided by pheromone trails and a heuristic favoring high bandwidth and low latency paths. Only edge-task mappings that respect the resource requirements and service agreements are considered. Pheromone values are updated after each iteration to reinforce good solutions while enabling exploration. Over multiple iterations, the algorithm converges and provides a mapping between edge nodes and tasks.

Graph database: maintaining the state of edge nodes/services

A graph database is used to provide a unified representation of both the infrastructure and workload information (Neo4j). The graph database model can be seen in FIGURE 48, it captures computational resources information (both available and in use) for each node, along with the connectivity between nodes. Links include attributes such as latency and available bandwidth, enabling a comprehensive view of the system's state.

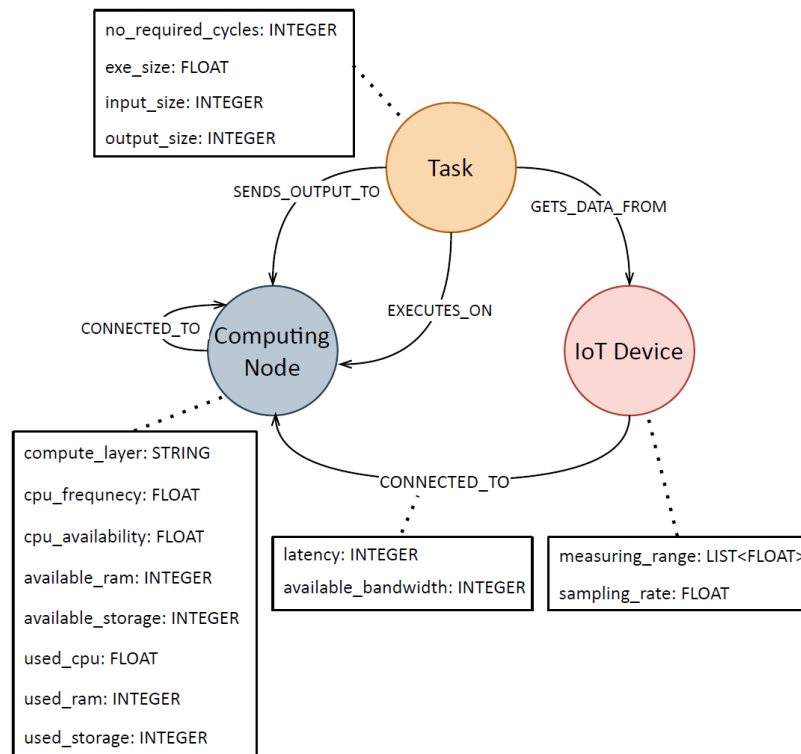


FIGURE 48 – GRAPH MODEL FOR UNIFIED TASK AND INFRASTRUCTURE REPRESENTATION

The orchestration workflow for energy services offloading at edge is presented in FIGURE 49. Whenever a new event or alert is identified by the event listener, it notifies the scheduler component. The scheduler queries the infrastructure information and sends it to the graph db. If the occurring events require a reallocation of resources, the swarm-based optimization algorithm is triggered that computes the placement plan and updates the graph database. The scheduler receives the new planification and interacts with the Kubernetes platform to move the services whilst maintaining their availability through scaling-up and down.

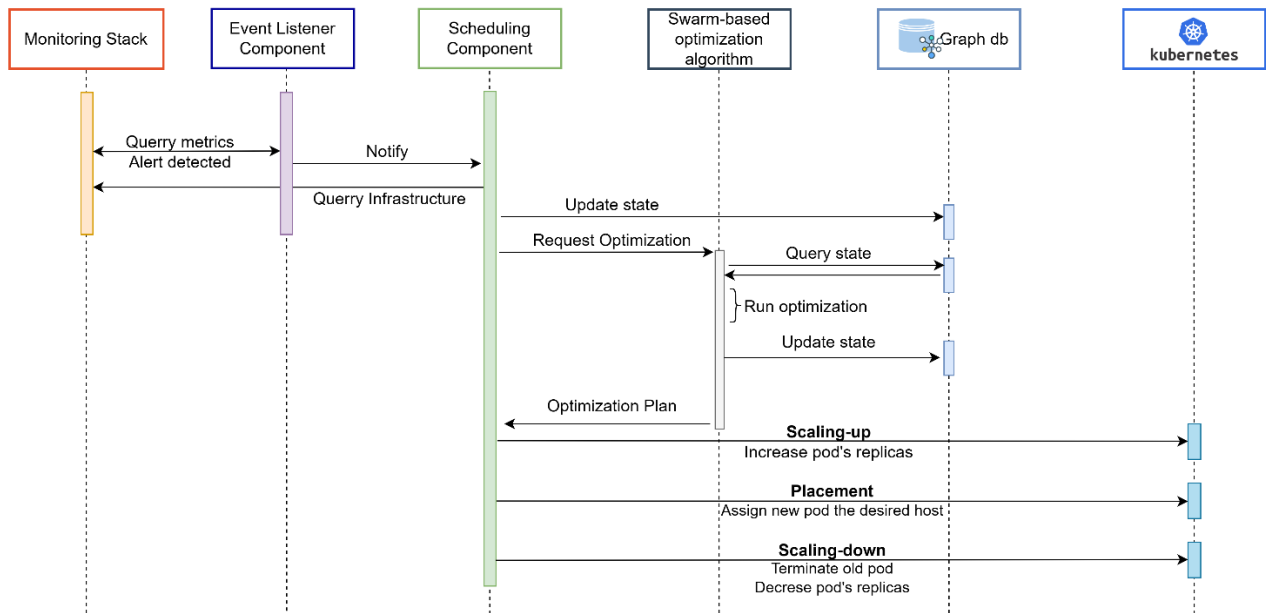


FIGURE 49 – ORCHESTRATION OF ENERGY SERVICES AT EDGE WORKFLOW

Blockchain and Smart Contract Integration

Blockchain technology is utilized to manage and track the location of tasks deployed across edge computing infrastructure. It provides transparency, traceability, and verifiability by recording all relevant interactions and enabling automated tracking and payments enforcement through smart contracts. Blockchain ensures that resource usage is notarized opening the space for further incentivization for edge resource utilization. To support this functionality, a multi-type token is defined by extending the ERC1155 token standard. Two types of tokens are considered: fungible (for automated payments) and non-fungible to represent the tasks (running energy services). The non-fungible tokens provide a globally unique identifier for the workload and serve as reference across task deployments, even when they are redeployed or migrated.

The blockchain integration with the orchestration framework is presented in FIGURE 50. When a task is deployed, a corresponding non-fungible token is minted and sent to the smart contract representing the edge node where the task is running. When a pod starts its execution the Event Listener Component detects this event and triggers a blockchain transaction, which using the unique token id, identifies the task (mints the token if necessary) and transfers the token to the node hosting it. This token transfer links the blockchain tokens to the real-world resources (CPU, RAM, HDD) consumed by this task.

When a pod finishes its execution (either because the task finished its operations or because it is migrated to a more suitable node) the Event Listener Component triggers the settlement process. It is performed by querying Prometheus to verify if the trade was respected by both sides, and payment is handled automatically. This ensures task placement traceability and enforces automated payments for resource consumption. If the task finished its execution, the smart contract transaction that burns the token representing the task is initiated, finalizing the transaction and ensuring that the allocated resources are properly settled.

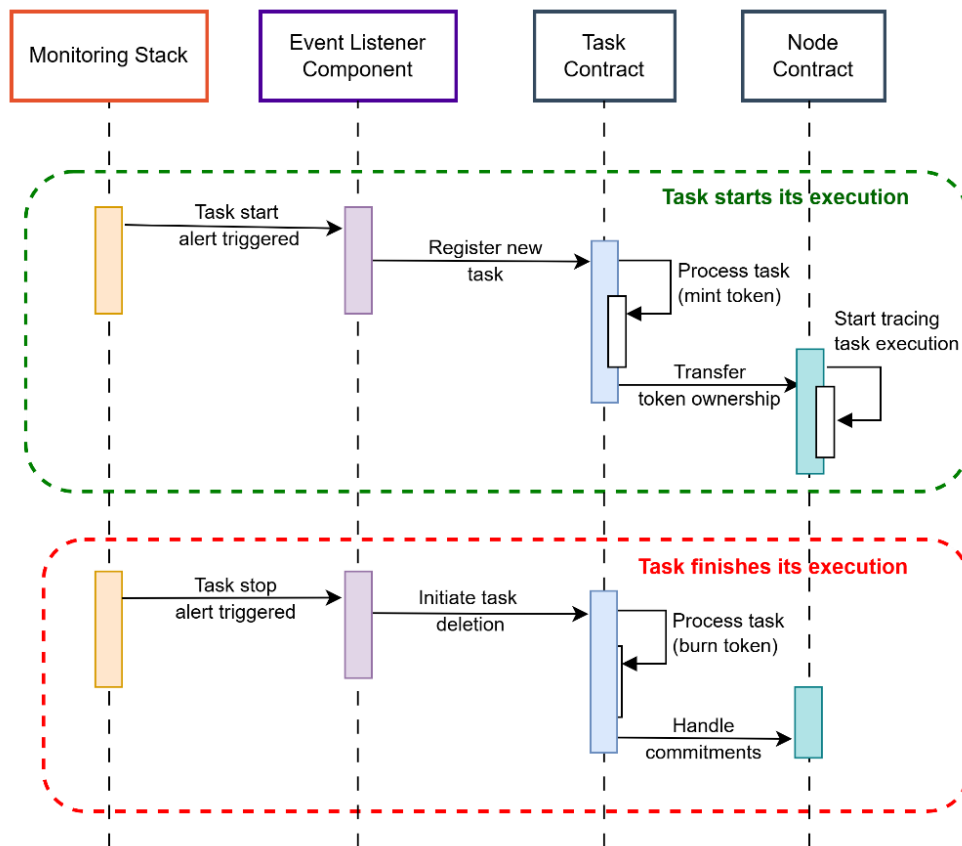


FIGURE 50 – BLOCKCHAIN INTEGRATION WORKFLOW

The fungible tokens are identified by the token id 0 whilst the non-fungible tokens are identified by structured IDs that encode metadata relevant to the task, making them identifiable and verifiable on-chain (FIGURE 51).

The non-fungible ERC1155 token ID is a 256-bit unsigned integer with a structured format:

- **Most significant bit:** represents the token type. Where value "0" is used for fungible tokens, further used for payments and "1" used for non-fungible tokens representing unique metadata for tasks. In this context the most significant bit always has value "1".
- **Next 160 bits:** contains the address of the account that initiated the token creation, the owner. Ensuring that tokens for a task can only be minted by the deploying entity.
- **Next 5 bits:** encodes resources requested by the task. Each bit represents a specific resource (CPU, RAM, HDD, bandwidth, latency), where "1" indicates the presence of the resource and "0" its absence.
- **Last 90 bits:** used to represent additional service metadata including resource requirements and a counter value that increments with each token or batch of token minted. This guarantees uniqueness for every token minted, even where the same address deploys multiple tasks with identical resource specifications.




FIGURE 51 – TOKEN ID STRUCTURE

5.3 ORCHESTRATION OF FEDERATED AI-DRIVEN ENERGY SERVICES

For Federated AI-driven services the orchestration framework manages the federated process by optimizing its hyperparameters to reduce communication overhead and increase training efficiency.

5.3.1 Description

In FIGURE 52 is represented the architecture of the orchestration framework for this scenario. Multiple edge nodes with private local data that train models locally are considered. On the cloud layer there is a synchronization process that coordinates the federated rounds. The weight aggregation of the global model can be either at cloud level (centralized/hierarchical) or at edge level by training peers' edge nodes (P2P). The orchestrator gathers information about available edge nodes and metadata about the federated process either by communicating with a central server or directly with the edge nodes. The clustering service can perform clustering of edge nodes based on available resources and provided metadata, or store clustering information by an external clustering mechanism. A scheduled synchronization thread is integrated into the orchestrator to assist in the synchronization mechanism of the federated process.

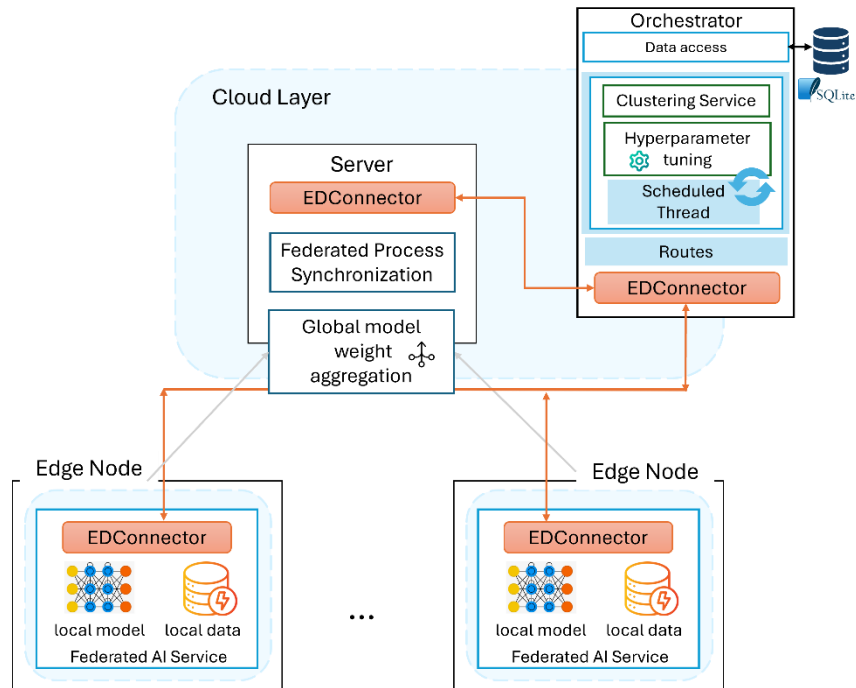


FIGURE 52 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR FEDERATED AI-DRIVEN SERVICES

5.3.2 Implementation details

Hierarchical Federated Learning Orchestration

In a hierarchical federated learning approach, training is organized in a multi-level structure. Our system considers four tiers: edge, fog and cloud nodes and the orchestrator and is presented in FIGURE 53.

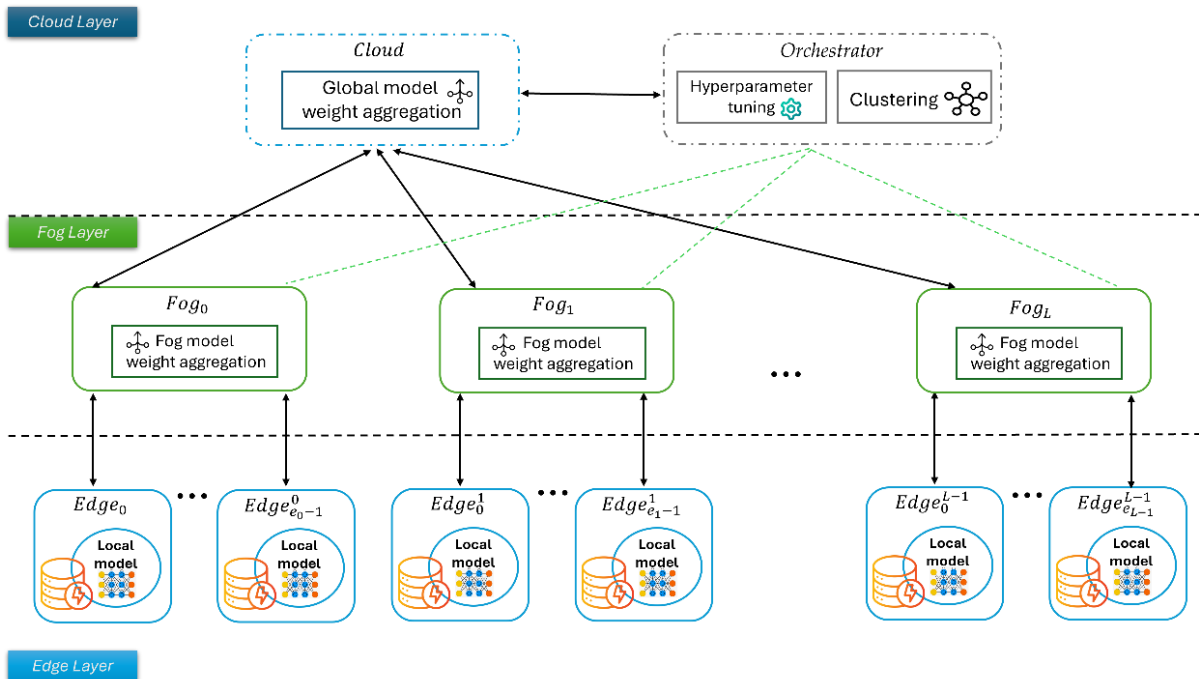


FIGURE 53 – ORCHESTRATION FRAMEWORK OVERVIEW FOR HIERARCHICAL FEDERATED LEARNING PROCESS

This structure reduces communication overhead and allows for tiered aggregations from the localized training rounds. In this setup, the Orchestrator coordinates the overall process and optimization. Cloud manages grid's commands and global aggregation. The rest of the nodes are organized in operational clusters. Fog Nodes are their cluster's management controller, handling communication between edge nodes and cloud. Edge nodes serve as workers within each cluster, performing private local training and evaluation tasks using sensitive energy consumption data. This design enables tiered model aggregation, local validation and centralized control over the hyperparameter optimization, while maintaining data privacy feature and reducing energy usage.

This design enables tiered model aggregation, local validation and centralized control over the hyperparameter optimization, while maintaining data privacy feature and reducing energy usage.

The following roles are identified:

- **Edge Nodes:** they are the private data owners that hold energy consumption metrics. Responsible for local training and evaluating models using their dataset, while keeping the data private. These nodes are split into two functional groups within each cluster: training nodes and validation node, both receive the current model weights and a hyperparameter configuration from their fog node. Training nodes update their local models, minimizing loss on their personal dataset. Validation node is selected by the cloud to evaluate the hyperparameter candidates during the genetic algorithm (GA) based tuning process. After local training or validation, each node sends updated weights and performance scores to the fog layer that aggregates and sends back to the cloud for further optimization.
- **Fog Nodes:** they facilitate synchronization between cloud layer and the clusters, acting as middleware controllers. They receive commands and configurations from the cloud. They forward these hyperparameter configurations and model weights to edge nodes, coordinate

local training and validation procedures and aggregate edge model updates before sending results to the cloud.

- **Cloud Node:** is the top-level aggregator for a grid. The cloud facilitates communication between Orchestrator and the clusters. It receives control commands and configurations from the orchestrator. It forwards hyperparameter configurations and model weights to the fog nodes, and aggregates fog model updates before sending results back to the Orchestrator.
- **Orchestrator:** organizes nodes into grid, grouping them based on clustering results, coordinates training and validation rounds, sends to the cloud the commands then forwards to each fog node the commands for their clusters. Also is orchestrating the hyperparameter tuning using a Genetic Algorithm (GA) for population evolution and Simulated Annealing (SA) for probabilistic decision-making.

P2P Federated Learning Orchestration

In a peer-to-peer federated learning approach, training is coordinated among a network of peer nodes. Edge nodes are both responsible to train on local available data, but also to communicate with other peers to obtain the global model.

In this scenario the complexity to manage the peer communication and computational overhead is higher. The orchestration service performs clustering in order to reduce communication between peers (each peer will send updates only to the peer nodes from the same cluster) and reduce the size of the model. A leader is selected by the orchestrator for direct communication between clusters. Additionally, the orchestrator facilitates synchronization to support reliable peer-to-peer federated learning. The orchestration framework overview for P2P Federated Learning services can be seen in FIGURE 54.

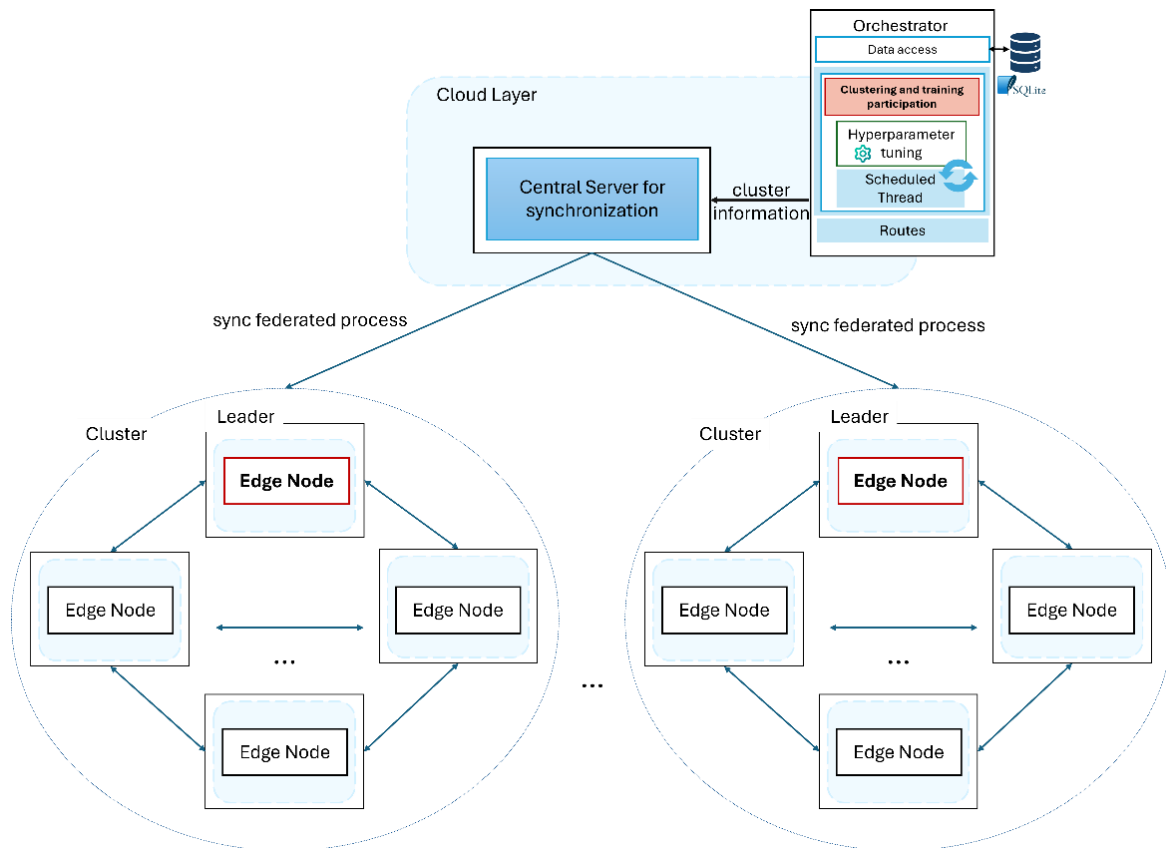


FIGURE 54 – ORCHESTRATION FRAMEWORK OVERVIEW FOR P2P FEDERATED LEARNING PROCESS

Three roles are identified:

Edge Node: performs local training on private data, ensuring that data remains on the device, without leaving the local environment. Upon joining the system, the node submits an initial set of performance metrics such as: CPU utilization, available RAM and time intervals associated with the locally stored data. These metrics are often updated and reported to better support clustering and orchestration processes. Edge Nodes may also act as leaders within their clusters. Leaders are the coordinating point of each cluster that handles the Central Server requests such as training or aggregation operations.

Central Server: facilitates synchronization between cluster leaders, it acts as a middleware between the Orchestrator and the Peer Nodes. It receives high-level commands from the Orchestrator such as initiating training rounds, performing aggregation or starting forecasting operations and coordinates the leaders to execute these operations. Besides leader communication, the Central Server communicates directly with all the nodes to inform members about their assigned cluster and their specific role.

Orchestrator: manages the entire training lifecycles across the system, this includes registering nodes at their request, collecting members metrics, grouping nodes into clusters using a spectral clustering approach, and assigning a leader to each cluster. It communicates with the Central Server to initiate training rounds or other operations. Additionally, it exposes a REST API that supports node registration, cluster related queries, training control commands and metric collection.

The clustering strategy aims to reduce energy consumption while maintaining models' performance and accuracy. The Orchestrator groups Edge Nodes into clusters based on their resource similarity and training compatibility. Ensuring nodes within a cluster share comparable computational capabilities in terms of CPU and memory availability minimizes the synchronization overhead and avoids performance bottlenecks caused by slower members. While, aligning nodes with similar data timeline intervals improves predictive performance. To form these clusters, the Orchestrator applies a spectral clustering approach. This method allows the system to detect grouping in a topology fashion, even when groupings are not linearly separable.

The entire clustering flow is depicted in FIGURE 55. This process requires a coordination among all participants to ensure the Peer Nodes are all available and ready to operate. Once clusters are formed, each node is notified about its membership and specific role before any training begins. In case of a node communication fail the system triggers a re-clustering process, temporarily excluding the unresponsive node.

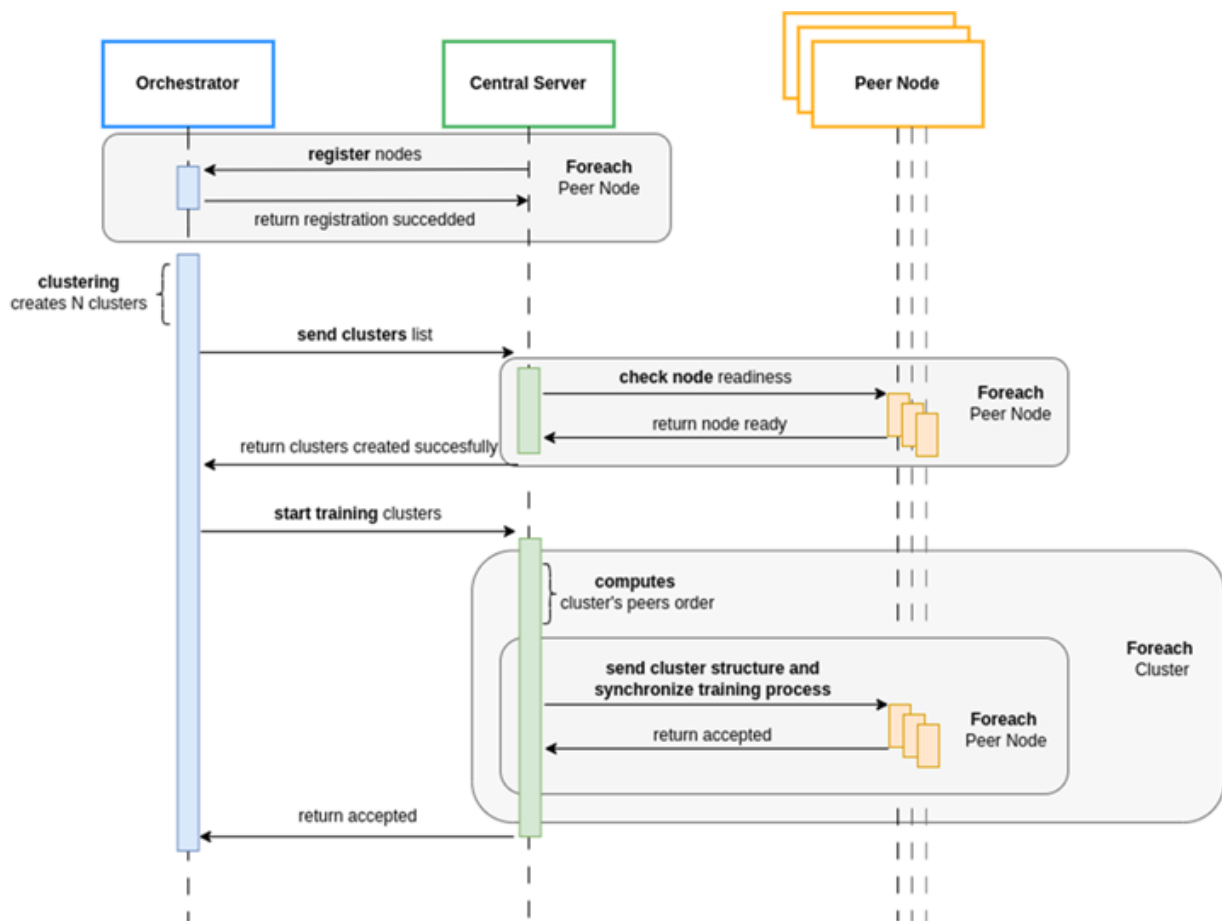


FIGURE 55 – ORCHESTRATION CLUSTERING FLOW

5.4 ENERGY SERVICES ROLLING OUT AT EDGE

In the context of energy services rolling out at edge, the orchestration framework handles automated updates of energy services by detecting updates for application components, manages versioning and handles service interruption during updates.

5.4.1 Description

FIGURE 56 represents an overview of the orchestration framework architecture for rolling out energy services. In this architecture, containerized energy applications are deployed and running across multiple edge nodes. The orchestrator requires a registration process to gather information about edge nodes and the applications currently running.

To ensure compatibility with edge environments, applications must either be designed as edge-native or adapted accordingly. The orchestrator is responsible for managing the lifecycle of the applications, including the automated rollout of updates when new service versions are available. To leverage the edge-cloud continuum, services should be designed or adapted to run distributed workloads (handling real-time, lightweight processing at the edge) and perform aggregation or coordination tasks in the cloud if necessary.

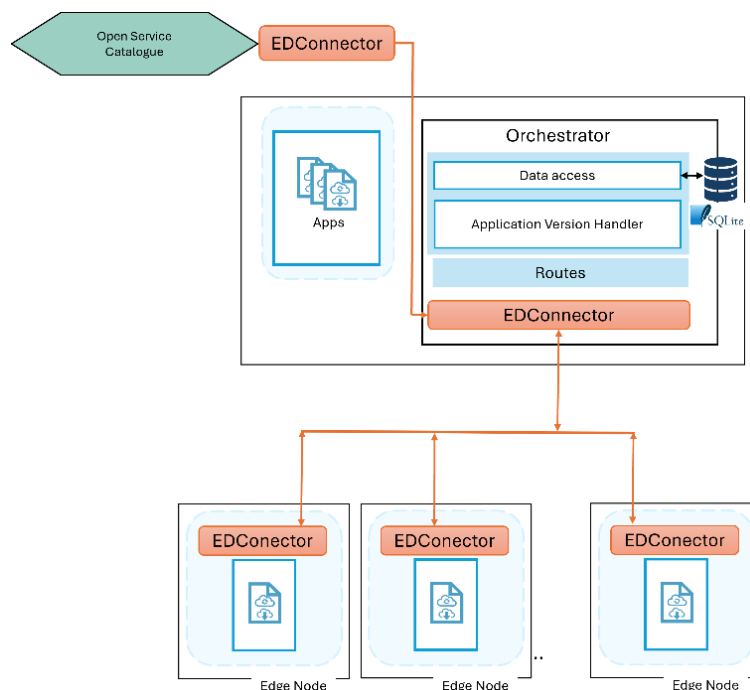


FIGURE 56 – ARCHITECTURE OF THE ORCHESTRATION FRAMEWORK FOR ENERGY SERVICES ROLLING OUT AT EDGE

5.4.2 Implementation details

After the registration phase, the orchestrator stores this information in a database (SQLite) and accesses it when needed. The Application Version Handler Component is connected to public registry containing application images and listens for updates. In this case, the application packages / images must be available in a docker image registry to which the orchestrator has access rights, or in the open services catalogue that provides update information (when the services are available and integrated with the service catalogue).

To ensure high availability and minimize service disruption during updates, the orchestrator supports parallel deployment of old and new service versions. When a new version of an application becomes available, it is deployed alongside the existing one on the same or different edge nodes.

Both versions can run concurrently for a defined transition period, allowing for updates without service interruption, or rollback in case of failures.

The architecture could also support communication through a Data Space Connector, enabling secure, standardized, and policy-governed exchange of application packages and updates across edge devices.

5.5 INTEGRATION WITH HEDGE-IOT FRAMEWORK

This section describes how the orchestrator will be integrated with the Hedge-IoT framework considering different types of services available in pilots and their computational orchestration requirements.

FIGURE 57 illustrates the general orchestration workflow required to integrate the orchestrator with available services and edge resources. The first step for integration is to gather information about the types of services, their computational requirements and the resources available. Thus, the orchestrator needs a discovery and registration phase for both services and edge nodes. For the services that will be available in the Open Service Catalogue, their information could be gathered from there by the orchestrator.

The second phase consists of executing the optimization algorithms and making the decisions or planification of computational resources. Finally, the last step involves sending the decisions to the part of the system that will ensure that the services' processes or the edge nodes' resources are executed and utilized accordingly. This aspect is handled by different components depending on the type of orchestration use case.

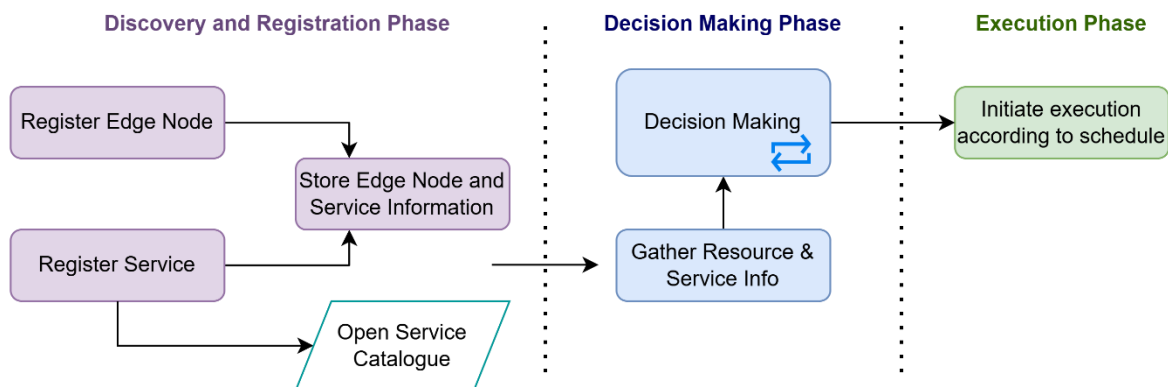


FIGURE 57 – SERVICE ORCHESTRATION WORKFLOW

5.5.1 Italian pilot

For the integration of the orchestration framework with the Italian Pilot, the congestion prediction and optimal power flow services are considered in the scenario of service offloading for geographical redundancy. The output of the congestion prediction is used by the optimal power flow service. Multiple servers located in different regions are considered for the validation of this scenario. The infrastructure is continuously monitored, and the services are replicated across

different locations. Income requests are redirected by a load balancer to the appropriate service instance to ensure high availability and minimize response time.

A KubeEdge platform is deployed to extend Kubernetes capabilities to edge locations. This setup allows for efficient service offloading and continuous monitoring of the physical resources and status of the services. Cilium is used for managing the networking layer, providing efficient, secure connectivity and it enables traffic control, service discovery, load balancing, and network monitoring.

5.5.2 Portuguese pilot

For the integration of the orchestration framework in the P2P federated learning scenario, the *Vector Autoregressive Model for Energy Time Series Forecasting* service from the Portuguese Pilot is considered. The federated learning process is decentralized: each peer communicates with other peers, while a central server manages synchronization of the federated process.

In this scenario, the computational orchestrator aims to reduce communication between peers and minimize the computational overhead caused by large model sizes, resulting in a more scalable federated solution. It receives analytics on available resources and statistical information about the time series data at each node and then performs a clustering algorithm to group edge nodes with similar characteristics. Based on the clustering results, the orchestrator generates a list of nodes for each cluster that includes identifying the peer nodes in each cluster, selecting a leader, and defining a queue order for the peer nodes. The peer nodes will only communicate with other peers from the same cluster.

As part of the initial integration of this service with the orchestrator, dedicated endpoints were introduced for edge, server, and the orchestrator APIs.

5.5.3 Finish pilot

In the Finnish Pilot the energy services rolling out at edge scenario is considered. The orchestrator handles version updates for the energy services that are deployed at edge with the objective of ensuring high availability and automated updates. The strategy considered is to run old and new versions in parallel until the new version is stable. The energy services are dockerized and their images will be pushed to a docker registry initially, then on the App Store. The orchestrator listens to updates on the registry and has information about where the services are deployed. It enables automated services updates whilst handling the running of previous version replicas of the service until the new version updates are completely rolled out at edge.

5.6 INTEGRATION WITH HEDGE-IOT INTEROPERABILITY FRAMEWORK

The orchestrator will incorporate open data connectors to enable secure and standardized communication with edge devices and energy services that will also be integrated with the connector.

Eclipse Dataspace Components (EDC) is an open-source framework that provides the necessary components for creating and operating data spaces in an extensible fashion to better suit different purposes. EDC architecture is divided into two roles: Data Provider that exposes data assets and Data Consumer that requests access to them. The same entity can act as both Data Provider and Consumer at the same time. Federated Catalogue enables data discovery by advertising the dataset's metadata. Identity Hub manages Decentralized Identifiers (DIDs) and the credentials. It ensures the validity of participants. EDC is designed for cloud-native deployment, and it supports containerization. Its microservice architecture permits each component (Control Plane, Data Plane) to be deployed and scaled independently. Persistent Volume Claims may be used to maintain contract states, DID documents and keys, while secrets may store credentials securely.

The EDC design is modular, extensible and has a decoupled architecture isolating failures and supporting retry operations. Also, it can be dockerized and deployed in different environments. In distributed infrastructures running tasks are often migrated to other nodes to better support the infrastructure and its workload. The connector can resume operations using the persistent storage layer. This ensures that active transfers are not disrupted, even when the Connector itself is moved to another host.

The edc-observability module is integrated to expose operational metrics and enable distributed tracing across transfer events. The observability setup can be configured to use Jaeger that provides real-time insights into communication, data availability and bottlenecks. This custom runtime is built using Gradle and packed as a Docker image. It can be deployed as a Kubernetes Pod or running on a simple Docker container.

In the Energy Services Orchestration context, data connectors can be used to ensure transfer continuity during service migrations, as the infrastructure nodes host services that continuously exchange operational data. This approach allows the orchestration framework to migrate services without disruption. If each of these nodes has a connector deployed, data access and exposure can be handled in a securely way ensuring that communication respects the contract agreements and policies defined.

In the Federated AI-driven services context, the orchestrator coordinates training cycles across distributed nodes. The connectors can act as Data Providers and Data Consumers at the same time on all levels: Edge, Fog and Cloud. The connectors can be used in exchanges such as: hyperparameters configurations distribution, model weights, updates, and performance metrics. This approach offers a secure, standardized and persistent communication layer decoupling the federated learning logic from security handling.

For the application lifecycle management, the EDC connector can be deployed on every participant node and acts as a Data Consumer. The nodes maintain a contract agreement for receiving application updates, handled as a non-finite, push-based transfer. Nodes receive updated

application images or configurations through the connector. This setup enables secure and reliable software distribution to ensure consistent updates for a dynamic distributed environment.

During the discovery and registration phase, the orchestrator service can rely on the open service catalogue, where energy services will be published along with relevant metadata. By querying this catalogue, the orchestrator will be able to automatically identify the services and coordinate deployments, federated processes or version rollouts.

5.7 Next steps

Action plan until the next iteration of this deliverable:

- (Energy Services Offloading at Edge) Integrate the meta-heuristic optimization algorithm with the orchestration framework and Kubernetes platform
- (P2P Federated Learning) Orchestrator API deployment and testing for P2P Federated Learning and clustering algorithm
- (Energy Applications Rolling Out at Edge) Identify the feasible services, establish the underlying infrastructure details (multiple edge nodes available) and develop additional deployment strategies and logic for pilot integration
- Complete the integration of blockchain platform with the orchestration framework and ensure that the smart contracts implementation aligns with the specification and requirements of the services
- Integration with Open Service Catalogue for gathering the specifications of the services to be offloaded
- Leverage the visualization tools to develop a dashboard for real-time metric monitoring and performance insights

6 CONCLUSION

This document delivers the intermediate Technology Enablers' (TEs) release that will be integrated into the HEDGE-IoT pilots and contribute to its multi-dimensional framework. A summary of results and status details of work package progress and objectives follows.

The first project objective is to **“add local/distributed intelligence leveraging IoT solutions at the edge/fog/cloud layers establishing the edge/cloud continuum through computational orchestration”**. The current technology enablers ecosystem addresses this objective in a straightforward way, by detailing federated learning approaches, edge-to-cloud services and a computational orchestration tool, which integrates with some of these technology enablers. As such, the following KERs can be observed and have been achieved or partially achieved:

- **KER1:** “A set of IoT solutions deployed across the energy system offering intelligence at the edge level” is considered to be **fully achieved**.
 - Designed a total of **10 TEs** (2 federated learning and 8 edge to cloud), which offer intelligent solutions for the energy system at the edge level, while exploiting **7 different types of IoT/edge devices**.
- **KER2:** “A service orchestrator to facilitate computational sharing between the cloud and edge levels” is considered to be **partially achieved**.
 - **Service orchestrator detailed in section 5** with implementation details for **3 different use cases** across pilots, which showcase how it facilitates the computational sharing between the cloud layer and edge nodes.

The second project objective is to **“design AI/ML tools for edge/fog/cloud services for increased flexibility, resilience and observability”**. The current technology enablers landscape addresses this objective by providing designs and specifications of data-driven services, which also leverage IoT/edge devices, to provide solutions for use cases such as congestion management, flexibility provision, demand response, manual frequency restoration reserve, predictive maintenance and residential user’s comfort. As such, the following KERs can be observed and validated:

- **KER3:** “A set of AI/ML tools for edge and cloud levels towards optimized planning, operation, resilience of interconnected assets” is considered to be **fully achieved**.
 - Designed **11 TEs** (6 edge-to-cloud and 5 cloud) that have the system operator as a stakeholder and work towards completing use cases related to optimized planning, operation and resilience of interconnected assets, such as congestion management, predictive maintenance and flexibility optimization.
- **KER4:** “A set of scalable data-driven energy and non-energy services for end-users (consumers, building occupants, system operators, etc.)” is considered to be **fully achieved**.
 - Designed **6 data-driven energy services for end-users**, addressing use cases like flexibility provision, demand response and home energy management.

- Designed **5 data-driven non-energy services for end-users**, addressing use cases like comfort of building occupants and advanced weather forecast.
- **KER5:** “An open app repository to populate edge/cloud and fog level AI/ML tools for energy stakeholders and SOs” is considered to be **partially achieved**.
 - Although a more in-depth specification of this repository will be done under T4.1, on work package 4, a repository has been created, and it will be publicly accessible by M30 (D3.5). It contains the specifications of the applications being developed under work package 3.

In conclusion, D3.4 showcases the intermediate release of designs and specifications for the technology enablers of the HEDGE-IoT project. Furthermore, it shows that the project and the work package are progressing well towards addressing its proposed challenges and objectives and that the technology enablers ecosystem is broad, increasing in maturity and ready to progress to the final stage of getting ready to integrate the pilots under the HEDGE-IoTs multi-dimensional framework.

REFERENCES

- [1] A. C. P. V. G. A. Daniel Rosendo, "Distributed intelligence on the Edge-to-Cloud Continuum: A systematic literature review," *Journal of Parallel and Distributed Computing*, pp. 71–94, 2022.
- [2] J. Z. Z. L. Y. e. a. Wen, "A survey on federated learning: challenges and applications," *International Journal of Machine Learning and Cybernetics*, vol. 14, no. 2, pp. 513–535, 2023.
- [3] W. Y. B. a. L. N. C. a. H. D. T. a. J. Y. a. L. Y.-C. a. Y. Q. a. N. D. a. M. C. Lim, "Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 2031–2063, 2020.
- [4] C. a. B. R. J. a. P. P. Gonçalves, "Privacy-Preserving Distributed Learning for Renewable Energy Forecasting," *IEEE Transactions on Sustainable Energy*, vol. 12, no. 3, pp. 1777–1787, 2021.
- [5] H. L. a. Z. W. a. T. H. a. M. A. Piette, "Energy flexibility of residential buildings: A systematic review of characterization and quantification methods and applications," *Advances in Applied Energy*, vol. 3, no. 1, 2021.
- [6] F. a. Y. Y. a. B. K. Ç. a. T. E. Yıldırım, "Comprehensive Review of Edge Computing for Power Systems: State of the Art, Architecture, and Applications," *Applied Sciences*, vol. 15, no. 8, 2025.
- [7] B. C. C. D. Gargya Gokhale, "Explainable Home Energy Management Systems based on Reinforcement Learning using Differentiable Decision Trees," in *e-Energy '24: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, Singapore, 2024.
- [8] L. a. M. J. a. G. K. a. S. R. a. V. J. Rodrigues, "Building Flexibility Bidding Curves for Energy Communities," in *2024 20th International Conference on the European Energy Market (EEM)*, Istanbul, 2024.

7 APPENDIX A

In the following sections, some tables can be found with some details relating to the services (x.y) described above, along the document.

(2.1) FEDERATED LEARNING SERVICE FOR ENERGY FORECASTING & DISAGGREGATION

TABLE 27 – INPUT AND OUTPUT DATA (SERVICE 2.1)

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	Weather data	Weather-related data, including temperature, humidity, etc.	°C, %, m/s	Structured, CSV	Numerical/Text
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

(2.2) VECTOR AUTOREGRESSIVE MODEL FOR ENERGY TIME SERIES FORECASTING

TABLE 28 – INPUT AND OUTPUT DATA (SERVICE 2.2)

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
		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

(3.1) ENHANCED NETWORK MANAGEMENT AND PLANNING

TABLE 29 - INPUT AND OUTPUT DATA (SERVICE 3.1)

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
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

(3.2) DTR-DLR ON THE EDGE

TABLE 30 - INPUT AND OUTPUT DATA (SERVICE 3.2)

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/m ²		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
	Wind direction	float	°		Intermediate output
	Ambient temperature	float	°C		Intermediate output
	Solar irradiance	float	W/m ²		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
Validation data	Time to overheat	float	s		Output
	Skin temperature	float	°C		Input
Validation data	Top oil temperature	float	°C		Input

(4.1) EDGECONNECT

TABLE 31 – INPUT AND OUTPUT DATA (SERVICE 4.1)

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

Bids	Divisible	Boolean	-	JSON	Input
	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

(4.2) FLEXIBILITY OPTIMIZATION SERVICE

TABLE 32 - INPUT AND OUTPUT DATA (SERVICE 4.2)

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	Numeric (€/kWh)	Real

		day-ahead and intraday market		
	Incentive_Limits	Maximum and minimum incentives allowed per customer	Numeric (€/kWh)	Real
	Historical_energy_consumption	Historical energy consumption data for the forecasting and NILM tasks	Numeric (kWh)	Real
	Historical_device_consumption	Historical energy consumption data from appliances for the NILM task	Numeric (kWh)	Real
	Historical_energy_production	Historical energy production data from the PVs for the forecasting task	Numeric (kWh)	Real
	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
	Disaggregation_output	Timeseries data per device, coming from NILM	Numeric (kWh)	Real
	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

(4.7) ENERGY COMMUNITY PLATFORM

TABLE 33 – INPUT VARIABLES (FROM D3.3) (SERVICE 4.7)

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
Sunmeter	SelfConsumptionEnergy	Wh	Integer
	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

TABLE 34 – OUTPUT VARIABLES (FROM D3.3) (SERVICE 4.7)

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

(4.7) PGUI, MARKET INTERFACE PLATFORM AND FLEXIBILITY REGISTER

TABLE 35 – CHARACTERIZATION OF VARIABLES (SERVICE 4.7)

Group	Variable Name	UoM	Data Type	Type
PGUI	Absorbed Active Energy	Wh	Integer	Output

Injected Active Energy	Wh	Integer	Output
Absorbed Reactive Energy	Varh	Integer	Output
Injected Reactive Energy	Varh	Integer	Output
Mean Absorbed Active Power	W	Integer	Output
Mean Injected Active Power	W	Integer	Output
Mean Absorbed Reactive Power	Var	Integer	Output
Mean Injected Reactive Power	Var	Integer	Output
P Up	W	Integer	Input
P Down	W	Integer	Input
Q Up	Var	Integer	Input
Q Down	Var	Integer	Input

Market Interface Platform

Flexibility Register

Measurements

Input

(4.9) POWERCIM PLATFORM

TABLE 36 – INPUT AND OUTPUT DATA (SERVICE 4.9)

Data Group	Variable	Variable description	Units	Format type	Type of data
Telemetry data	Wind speed		m/s	float	Input/Output
	Wind direction		°	float	Input/Output
	Ambient temperature		°C	float	Input/Output
	Solar irradiance		W/m ²	float	Input/Output
	Relative humidity		%	float	Input/Output
	Rain Intensity		mm/h	float	Input/Output
	Current		A	float	Input/Output
	Voltage		V	float	Input/Output
	Active power		W	float	Input/Output
	Reactive power		Var	float	Input/Output
	Apparent power		VA	float	Input/Output
	Power factor			float	Input/Output
	Ampacity		A	float	Input/Output
	Skin temperature		°C	float	Input/Output

Top oil temperature	°C	float	Input/Output
---------------------	----	-------	--------------

(APPENDIX B) USER INTERFACES

TABLE 37 – INPUT AND OUTPUT DATA (SERV. APPENDIX B)

Category of data	Variable name	Description	Data Format	Real or Synthetic
Input Data	Customer_id	The identifier of each customer	Numeric	Real
	House_id	The identifier of each house	Numeric	Real
	Aggregator_id	The identifier of each aggregator	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
	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
	Historical_energy_consumption	Historical energy consumption data for the forecasting and NILM tasks	Numeric (kWh)	Real
	Historical_device_consumption	Historical energy consumption data from appliances for the NILM task	Numeric (kWh)	Real
Historical_energy_production	Historical energy production data from the PVs for the forecasting task	Numeric (kWh)	Real	
Output data	Forecasted_energy_consumption	Timeseries of forecasted data for the main energy consumption and the energy	Numeric (kWh)	Synthetic

	consumption per device		
Forecasted_energy_production	Timeseries of forecasted data for the energy production and the energy consumption per device	Numeric (kWh)	Synthetic
Disaggregation_output	Timeseries data per device, coming from NILM	Numeric (kWh)	Real
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

8 APPENDIX B

As this service focuses on user interfaces and does not fit into any of the major categories of technology enablers of the project, it is presented here in the annexes.

8.1 USER INTERFACES (CLOUD SERVICES)

8.1.1 Description of Service

The User Interface layer of the system is composed of four distinct subservices, each designed to facilitate interaction between users, aggregators, and the underlying flexibility infrastructure. Each subservice addresses the needs of different stakeholders (residential users, aggregators, and system operators) by providing tailored tools for visualization, control, data interpretation, and secure information exchange.

8.1.1.1 Mobile & Web Application for Residential Users

To support active participation in demand-side flexibility, the system provides a mobile and web application tailored for residential users. These applications act as the main touchpoint for households, offering real-time visualizations of electricity consumption, solar production (if applicable), and flexibility opportunities. Users receive notifications about upcoming flexibility events and can view personalized incentives based on their past behavior and current energy usage. The app also helps users understand their consumption at a device level, promoting more energy-aware behavior. Its design prioritizes ease of use and accessibility, making complex energy interactions approachable and actionable for everyday users.

8.1.1.2 Web Application for Aggregator

For aggregators tasked with managing distributed energy resources, a web-based application serves as the operational backbone. It enables aggregators to monitor enrolled households, assess flexibility potential across their portfolio, and orchestrate participation in grid services such as demand response or local balancing. The interface includes forecasting tools, real-time dashboards, and bid preparation modules that integrate both historical and live data. This functionality helps aggregators optimize flexibility activation based on market signals, user behavior, and system constraints, ensuring both profitability and system reliability. It also supports compliance tracking and reporting for regulatory oversight.

8.1.1.3 Data Identification with Vision

A novel aspect of the system lies in its use of computer vision to enhance data identification and contextual tagging. By combining sensor data with visual or spatial information—such as floor plans, device images, or installation metadata—the system can more accurately associate measurements with specific appliances, rooms, or usage scenarios. This improves the granularity and accuracy of disaggregation, especially in settings with limited metadata or mixed loads. Vision-based tagging also enables faster system setup and more intuitive user interfaces, where users can confirm or adjust visual mappings of their devices, creating a richer foundation for personalization and optimization.

8.1.1.4 Data Sharing Platform

All components are connected through a secure, scalable data sharing platform that ensures interoperability, privacy, and real-time responsiveness. The platform supports standardized energy data models (e.g., OpenADR, SAREF), enabling seamless integration with external services and regulatory bodies. Fine-grained access control policies allow users and aggregators to define who can access what data, under what conditions, and for how long—supporting GDPR compliance and trust. The platform also includes synchronization mechanisms to align edge-collected data with cloud-based analytics and forecasts, enabling both fast local decisions and long-term learning. This infrastructure underpins the entire flexibility management ecosystem, ensuring that insights and actions flow smoothly across all layers.

8.1.2 Innovative Aspects

The system’s innovative aspects combine advanced AI and IoT integration, secure data handling, and user-focused design. These features work together to deliver accurate forecasting, personalized insights, and versatile access across devices, empowering users to effectively monitor and manage their household energy dynamics.

8.1.2.1 AI + IoT Integration

At the heart of the service lies the integration of IoT-enabled household monitoring with artificial intelligence. Real-time energy consumption/production data is collected from in-home smart meters and connected devices and fed into machine learning models designed for short and mid-term consumption/production forecasting. These models continuously learn from user behavior, time-of-day trends, and contextual data such as weather or occupancy patterns, enabling accurate predictions of household demand. This predictive intelligence powers an interactive user interface, where users can visualize upcoming consumption/production, receive personalized efficiency suggestions, and gain insights into how their daily habits affect energy use.

In parallel, a new energy disaggregation module is under active development, aiming to break down total consumption into individual appliance-level usage without requiring intrusive hardware. This future capability will further enhance transparency and control, allowing users to pinpoint inefficiencies and make more informed, device-specific decisions.

8.1.2.2 Integration Of Different Device Brands

The system is designed with interoperability in mind, supporting integration with a wide range of smart meters, IoT sensors, and appliances from different brands and manufacturers. This ensures that users can connect their existing hardware without being locked into a specific vendor ecosystem. Through standardized protocols and modular device drivers, the platform harmonizes data across heterogeneous sources, maintaining consistency in analysis and predictions. This brand-agnostic approach not only simplifies onboarding and scalability but also enhances user flexibility and future-proofs the system against evolving IoT landscapes.

8.1.2.3 Integration with Local Energy Providers Billing Systems

The platform integrates with local energy providers to access real-time and historical billing data, enabling accurate mapping between actual consumption and corresponding costs. This integration allows users to monitor their energy usage not only in kilowatt-hours but also in monetary terms,

based on the exact tariff structure applied to their household. By aligning metering data with provider billing systems, users can validate their utility bills, track spending patterns, and gain a transparent view of how their behavior translates into costs. This creates a foundation for both awareness and accountability, empowering users to make better-informed decisions—whether they're reviewing their monthly bill or adjusting daily habits.

8.1.2.4 User-Centric Personalization

Personalization is embedded throughout the user interface, making every interaction relevant and intuitive for each individual. Consumption forecasts and device-level insights generated by NILM are visually presented in a clear, accessible way, allowing users to easily understand their energy patterns. Incentive offers and notifications are dynamically tailored and delivered in real time via the UI, reflecting each user's unique habits, preferences, and engagement history. This personalized approach ensures that recommendations feel meaningful and actionable, enhancing user motivation and satisfaction while fostering more sustainable consumption behaviors.

8.1.2.5 Versatility Across Mobile and Desktop App

The service offers a versatile user experience, seamlessly accessible via both mobile and desktop applications. Whether users are on-the-go or at home, the UI adapts responsively to different screen sizes and interaction modes, ensuring intuitive navigation and consistent access to all features. This cross-platform design empowers users to monitor consumption, receive alerts, and manage settings anytime, anywhere —enhancing engagement and convenience across all devices.

8.1.2.6 Advanced Security & Privacy

Security and privacy are integral to the user experience, with the UI designed to transparently communicate data protection measures. Access to sensitive information such as personalized forecasts and appliance-level insights is strictly controlled through user authentication and role-based permissions, ensuring that only authorized users can view their data. All interactions with the system whether through mobile or desktop apps occur over encrypted connections to protect data in transit. Additionally, privacy-preserving techniques used in data processing, especially for NILM, are reflected in UI features that allow users to control data sharing preferences and review privacy settings easily, fostering trust and confidence in the platform.

8.1.3 Service Data

The User Interface aggregates and presents real-time and historical data from IoT devices, forecasting models, and user preferences to enable intuitive interaction. It supports personalized visualizations, notifications, and control features, facilitating user engagement and informed decision-making across multiple platforms.

- **Customer_id:** A unique identifier for each user, utilized by the UI to personalize data presentation, tailor notifications and recommendations, and associate NILM insights with the user profile, enabling a customized and targeted user experience.
- **House_id:** A unique identifier for each participating household, used by the UI to aggregate and present energy data, forecasts, and notifications specific to that household, enabling personalized and accurate monitoring of consumption and production.
- **Aggregator_id:** Unique identifier linking users to their energy aggregator. This enables the UI to tailor content, notifications, and control options based on aggregator-specific

programs or incentives, fostering personalized communication and streamlined interaction between end-users and service providers.

- **Main_energy_consumption:** Real-time total household energy demand (in kWh), directly ingested by the demand forecasting module for model inference, and also used by NILM to disaggregate appliance-level loads for finer control of flexibility assets.
- **Mtu (Market Time Unit):** A categorical timestamp (typically in 15-minute intervals) used to align real-time data across all subservices. It serves as the temporal backbone for forecasting horizons, bid construction, and incentive scheduling.
- **PV_Production:** Current solar generation data (in kWh), collected via IoT from PV inverters, and used by the production forecasting module to calibrate and validate short-term generation predictions for flexibility assessment.
- **Weather_data:** External meteorological parameters (e.g., temperature, irradiance, wind speed) retrieved from weather forecasting APIs¹⁸. This data is essential for production forecasting, as environmental conditions significantly influence load patterns and PV output.
- **Electricity_price:** Prices from the day-ahead and intraday wholesale markets (€/kWh), used by the bidding optimization engine to set economically efficient bid prices. These also serve as a benchmark to evaluate the cost-effectiveness of activating local flexibility.
- **Forecasted_energy_consumption:** Timeseries of predicted consumption values (in kWh), generated by the demand forecasting module and used by the optimization engine to identify periods of high load where flexibility activation would yield the most benefit.
- **Forecasted_energy_production:** Timeseries of predicted PV generation (in kWh), output by the production forecasting module. These forecasts are crucial inputs for planning dispatch and forming accurate, risk-aware bids in the local market.
- **Historical_energy_consumption:** Timeseries of historical household consumption values (in kWh), used by the demand forecasting module for model training.
- **Historical_energy_production:** Timeseries of historical PV generation (in kWh), used as input in the production forecasting module for model training.
- **Historical_device_consumption:** Timeseries of historical appliance consumption (in kWh), used as input in the NILM for model training.
- **Disaggregation_output:** Real-time timeseries of appliance consumption (in kWh), provided by NILM.
- **Flexibility_offers:** The flexibility offers to the customers (load shifting, incentive).
- **Bid_price:** The optimal bid price for the aggregator to submit to the LFM.
- **Bid_quantity:** The optimal bid quantity for the aggregator to submit to the LFM.

8.1.4 Integration with Hedge-IoT Interoperability Framework

Based on the ongoing discussions within WP4 and the initial set of requirements, this service is expected to be integrated with the Hedge-IoT Interoperability Framework. The details around the architecture and integration approach are still evolving. We anticipate having a clearer picture and more concrete information in the next deliverable.

¹⁸ <https://open-meteo.com/>

8.1.5 Implementation Details

8.1.5.1 Mobile

The mobile application is developed using Flutter¹⁹, enabling cross-platform deployment across Android and iOS with a unified codebase. It serves as the primary interface for residential users, allowing them to monitor their electricity consumption, PV production, and flexibility opportunities in real time. The UI is designed for ease of use, featuring clear visualizations and responsive interactions. Authentication is managed via Keycloak²⁰, ensuring secure and centralized identity management across all user-facing services. The app also integrates push notifications to inform users about flexibility events and incentive updates, enhancing engagement and responsiveness. Real-time data is visualized using native Flutter charting libraries, optimized for performance on low-power devices.

- **Framework:** Flutter
- **Authentication:** Keycloak
- **Data Visualization:** Native Flutter charting libraries
- **Platform Support:** Android & iOS
- **Communication:** Secure REST APIs

8.1.5.2 Residential Web App

The Residential Web App is implemented using React²¹, offering a responsive and lightweight interface tailored for desktop and tablet users. It provides similar functionalities as the mobile app, including real-time visualizations of energy usage and incentives, historical comparisons, and event notifications. For data visualization, the app uses Chart.js²², which allows for interactive and customizable charts. Authentication and session management are handled via Keycloak, ensuring consistent access control across platforms. The app communicates with backend services through secure APIs, enabling live updates and a personalized experience that aligns with the user's consumption profile and preferences.

- **Frontend Framework:** React
- **Data Visualization:** Chart.js
- **Authentication:** Keycloak
- **Communication:** RESTful APIs
- **Responsiveness:** Fully responsive design for tablet and desktop

8.1.5.3 Aggregator Web App

The Aggregator Web App is also built using React, with a focus on operational efficiency and oversight. It includes dashboards for monitoring household portfolios, bid preparation, and flexibility analytics. Chart.js is used to display forecasting outputs, device-level disaggregation, and bidding metrics in a visually interpretable format. Access is controlled via Keycloak, supporting role-based permissions for different aggregator staff roles. The app integrates closely with the optimization engine, NILM outputs, and forecasting modules, providing aggregators with the tools

¹⁹ <https://flutter.dev/>

²⁰ <https://www.keycloak.org/>

²¹ <https://react.dev/>

²² <https://www.chartjs.org/>

to assess flexibility in real time, prepare market bids, and manage user participation effectively. The interface is optimized for use on large displays and supports export of reports and visualizations for regulatory or business use.

- **Frontend Framework:** React
- **Data Visualization:** Chart.js
- **Authentication:** Keycloak (role-based access control)
- **Communication:** REST APIs to backend optimization and NILM services
- **UI Design:** Desktop-optimized layout with export and reporting support

8.1.5.4 Functionalities

TABLE 38 shows service functionalities and status of implementation.

TABLE 38 – USER INTERFACES FUNCTIONALITIES

Functionality	Description	Implemented	Expected Month of Implementation
Aggregated Energy Monitoring	Visualization and analytics of energy consumption/production across all monitored sites.	60%	M26
Individual Energy Monitoring	Per-customer energy usage/production tracking with dashboards and historical insights.	60%	M26
Aggregated Energy Forecasting	Short-term and 24-hour ahead forecasts of total energy consumption/production across prosumers.	50%	M26
Individual Energy Forecasting	Short-term and daily forecasts of each customer's energy consumption/production	50%	M26
NILM Disaggregation	Non-intrusive load monitoring to infer appliance-level usage from aggregate consumption.	25%	M32
Electricity Price Forecasts Integration	Integrate the electricity price of day-ahead market as a benchmark for the optimization algorithms	70%	M26
Bidding Optimization	Real-time continuous quantity & price bids for the aggregator to participate in the local flexibility market	50%	M26
IoT Data Ingestion	15-min load & PV readings streamed from prosumers' IoT meters for forecasting modules, 1-min load for NILM.	60%	M26
Integrate customer acceptance probabilities	Consider the probabilities of the customers to respond positively in flexibility offers	0%	M32

Functionality	Description	Implemented	Expected Month of Implementation
Integrate different levels of incentives	Consider different levels of incentives provided to the customers	0%	M32
Flexibility Offers	UI/API for defining and managing flexibility offers to customers, based on forecast and pricing.	10%	M32
Aggregated Health Monitoring	System-wide operational health checks, outage detection, and aggregated KPIs.	40%	M26
Individual Health Monitoring	Alerting and diagnostics for each prosumer and consumption anomalies.	30%	M26
Aggregated Energy Statistics	Aggregated KPIs and summaries across all customers for grid-level analysis.	75%	M26
Individual Energy Statistics	Detailed statistical summaries (avg, peak, load factor) for individual prosumers.	30%	M26
Authentication (Web)	Web-based user authentication and account management.	100%	-
Authentication (Mobile)	Mobile app login, session, and account management support.	100%	-

8.1.5.5 Integration and dependencies

The User Interfaces (UIs) in the Greek Pilot serve as the central point of interaction for operators, aggregators, and possibly prosumers. They integrate with multiple real-time backend services to support forecasting (demand & production), load disaggregation (NILM), optimization (incentives & bidding), flexibility offers, market trading (HenEx) and monitoring (health, energy statistics). Each backend component contributes data to the UI, ensuring that the user sees a synchronized, real-time view of the system state.

1. Demand Forecasting

The demand forecasting module is an essential backend service integrated with the Greek Pilot UI, designed to process both real-time and historical energy consumption data. These data streams, which originate from smart meters and are transmitted via MQTT and Kafka protocols, feed into machine learning models that produce short- and medium-term consumption forecasts. The results are presented within the UI for operator oversight and are crucial for downstream systems like the Flexibility Optimization Service. By accurately predicting demand, this module allows the platform to estimate the available flexibility of households and determine optimal timings for load shifting, making it a foundational component for both user awareness and grid-side decision-making.

2. Production Forecasting

Production forecasting plays a critical role in supporting prosumers with photovoltaic (PV) systems by predicting their solar generation. This service combines historical PV output data with external weather forecasts to estimate upcoming energy production. These forecasts are vital for determining surplus energy that can be dispatched or sold, and they directly influence the platform's bidding and dispatch logic. The UI provides visibility into individual

and aggregated production forecasts, allowing users and operators to assess energy self-sufficiency, manage storage, or contribute to flexibility markets. The close coupling of this forecast with weather services ensures responsive and accurate decision support.

3. NILM (Non-Intrusive Load Monitoring)

The NILM service is integrated within the aggregator's backend and interfaces with the UI to deliver appliance-level insights derived from total household energy usage. Utilizing advanced machine learning models, such as denoising autoencoders, this module disaggregates a home's energy signal to identify when and how much energy specific appliances are consuming. This granularity enhances the understanding of load profiles and helps classify loads as flexible or inflexible. In the UI, users or operators can view breakdowns of appliance usage, track behavioral patterns, and use these insights to craft more effective flexibility offers or design targeted incentives. The NILM service also feeds valuable features into forecasting and optimization models.

4. Edge Processing

Edge processing serves as the first stage in the data pipeline, enabling smart meters, PV inverters, and other IoT devices to stream real-time energy data into the system. This data includes both consumption and generation at high frequency (e.g., every 1 or 15 minutes). The edge-cloud architecture is designed to handle latency-sensitive computations close to the source, while the cloud aggregates and scales the information for broader analytics. In the UI, this edge data is reflected in near real-time dashboards, device health monitors, and anomaly alerts. This tight integration ensures that operators and prosumers receive up-to-date information for decision-making, while also enabling low-latency responses to grid events or flexibility activations.

5. Flexibility Optimization

Flexibility optimization is the core intelligence layer of the aggregator platform, responsible for determining the most efficient incentives to offer users in exchange for modifying their energy usage patterns. Leveraging inputs from forecasting, NILM, and user behaviour history, the optimization service computes the best mix of price, time, and customer targeting to achieve desired load shifting. The UI provides tools to simulate these outcomes, adjust assumptions, and review proposed flexibility offers. Additionally, the platform generates optimal bid curves for the flexibility market and allows operators to submit them directly. This end-to-end pipeline from data ingestion to bid deployment is orchestrated by the flexibility optimization engine and made fully accessible through the UI for oversight and control.

FIGURE 58 shows an architecture diagram that maps the integrations and dependencies of the UIs with the remaining Greek pilot services, as well as showcases the status of the implementation of each integration.

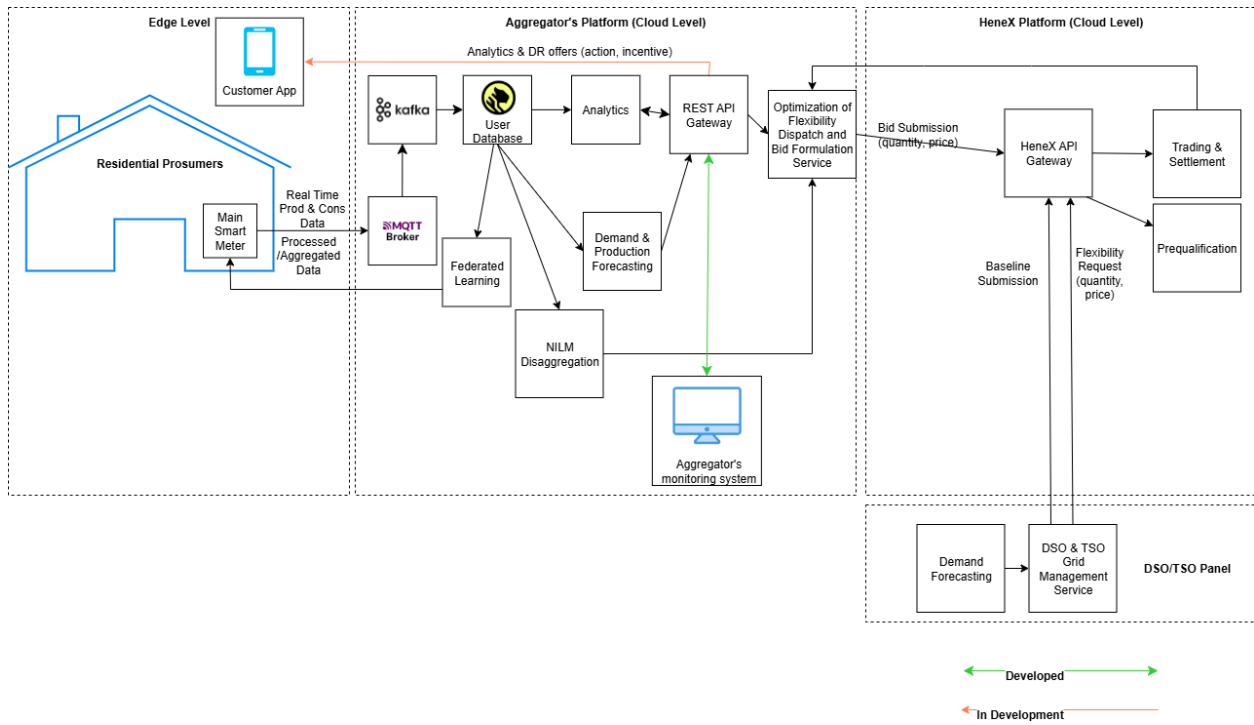


FIGURE 58 – USER INTERFACES PILOT SGAM DIAGRAM – ICCS



HEDGE-IoT

