HEDGE-IoT

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

D2.4
HEDGE-IoT Reference Architecture (Final Release)

Co-funded by the European Union

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

DOCUMENT CONTROL SHEET

PROJECT INFORMATION

Project Number	101136216		
Project Acronym	HEDGE-IoT		
Project Full title	Holistic Approach towards Empowerment of the Digitalisation of the Energy Ecosystem through adoption of IoT solutions		
Project Start Date	01 January 2024		
Project Duration	42 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.	D2.4							
Deliverable Title	HEDGE-IoT Reference Architecture (Final Release)							
Work-Package No.	WP2							
Work-Package Title	Stakeholders' Requirements and System Specifications							
Lead Beneficiary	ED							
Main Authors	ED, TRIALOG							
Other Authors	RWTH							
Due date	30/04/2026							
Deliverable Type	X	Document, Report (R)		Data management plan (DMP)		Websites, press & media action (DEC)		Other
Dissemination Level	X	Public (PU)		Sensitive (SEN)		Classified		
	PU: Public, fully open SEN: Sensitive, limited under the conditions of the Grant Agreement Classified R-UE/EU-R – EU RESTRICTED under the Commission Decision No2015/444 Classified C-UE/EU-C – EU CONFIDENTIAL under the Commission Decision No2015/444 Classified S-UE/EU-S – EU SECRET under the Commission Decision No2015/444							

Co-funded by the European Union

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

DOCUMENT REVISION HISTORY

Version	Date	Description of change	List of contributors(s)
0,1	05/01/2026	Table of Contents of Document	ED, TRIALOG, RWTH
0.3	31/03/2026	First Draft ready	ED, TRIALOG, RWTH
0.4	05/05/2026	Final Version Ready for Review	ED, TRIALOG, RWTH
1.0	13/05/2026	Final reviewed Version Ready for Submission	ED

Co-funded by the European Union

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

PARTNERS

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

Co-funded by the European Union

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

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

Co-funded by the European Union

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

42	KONCAR – DIGITAL DOO ZA DIGITALNE USLUGE	KONC	HR
43	DS TECH SRL	DST	IT
44	CYBERSOCIAL LAB S.R.L.	CSL	IT
45	ACEA ENERGY MANAGEMENT	AEMA	IT

Co-funded by the European Union

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

DISCLAIMER

Funded by the European Union. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union. The European Commission is not liable for any use that may be made of the information contained herein.

COPYRIGHT NOTICE

Co-funded by the European Union

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

EXECUTIVE SUMMARY

HEDGE-IoT is a Horizon Europe Innovation Action that supports the digital transformation of the European energy system through an interoperable IoT-enabled edge-to-cloud framework. The project addresses energy-system digitalisation across multiple levels, from behind-the-meter assets to distribution and transmission system operation, with the aim of improving flexibility, resilience, observability and intelligence. Deliverable D2.4, "HEDGE-IoT Reference Architecture — Final Release", consolidates the final WP2 architectural baseline. It is an architectural consolidation deliverable, not the project's final implementation or validation report.

D2.4 builds on the requirements, business use cases, system use cases, transversal use cases and functional specifications developed in WP2. It consolidates 14 Business Use Cases, 39 System Use Cases and three Transversal Use Cases addressing data interoperability through the HEDGE-IoT Dataspace, computational interoperability through cloud-edge orchestration, and functional interoperability through the HEDGE-IoT App Store. These elements provide the functional basis for the final architecture and reflect common needs across the six pilot environments.

The final HEDGE-IoT Reference Architecture is organised into four principal layers: the Physical Layer, the Dataspace Layer, the Semantic Interoperability Layer and the Application Layer. The Physical Layer represents pilot infrastructures, IoT devices, local platforms, digital twins, data sources and edge nodes. The Dataspace Layer acts as the operational core of the architecture, enabling governed, sovereign and federated data and service exchange through connectors, catalogues, identity management, policy enforcement and contract negotiation. It also includes the App Store for service publication and reuse, and the Orchestrator for coordinating workloads across edge, fog and cloud resources.

The Semantic Interoperability Layer ensures that exchanged data, metadata and service descriptions can be interpreted consistently across heterogeneous systems and pilots. It includes semantic enablers, ontologies, vocabularies, validation mechanisms, model-management functions and knowledge-engine components. The Application Layer exposes HEDGE-IoT capabilities to users and stakeholders through dashboards, analytics, forecasting, optimisation, federated learning, fog/cloud services and interfaces to pilot-specific tools and operational platforms.

The architecture is aligned with recognised European and international frameworks, including SGAM, BRIDGE DERA, IDS-RAM, IDSA, Gaia-X, SAREF, IEC CIM and FIWARE-compatible approaches. This alignment supports interoperability, replicability and consistency with the broader development of European energy data spaces. D2.4 also defines interoperability profiles for key exchange points, including the dataspace connector interface, computational orchestrator interface, App Store interface and flexibility-market interfaces.

Security, privacy and AI trustworthiness are treated as cross-cutting concerns. The architecture incorporates identity management, role-based access, secure communication, policy enforcement, contract-based exchange and data-sovereignty principles. These

Co-funded by the European Union

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

mechanisms are complemented by the project's Cross-Cutting Characteristics Plan activities, which address cybersecurity, data privacy, AI trustworthiness and pilot-level maturity improvement.

Finally, the component-to-requirements traceability perspective links the architecture back to the WP2 functional requirements baseline. This traceability supports implementation, integration and validation by showing how architectural packages and service components respond to the capabilities required by the project's use cases.

The deliverable concludes by identifying the main architecture decisions, known limitations and future roadmap. Key open issues include variable implementation maturity across pilots, the continued need for local adapters and semantic refinement, heterogeneous flexibility-market interfaces, further validation of computational orchestration, and continued follow-up on cybersecurity, privacy and AI trustworthiness. Overall, D2.4 provides a stable, standards-aware and implementation-informed reference architecture for the remaining HEDGE-IoT integration, validation, exploitation and replication activities.

Co-funded by the European Union

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

TABLE OF CONTENTS

Co-funded by the European Union

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

Co-funded by the European Union

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

Co-funded by the European Union

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

LIST OF TABLES

Co-funded by the European Union

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

LIST OF FIGURES

Co-funded by the European Union

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

ABBREVIATIONS

AAS	Asset Administration Shell
ABAC	Attribute-Based Access Control
ADF	Architecture Description Framework
ADL	Architecture Description Language
aFRR	Automatic Frequency Restoration Reserve
AI	Artificial Intelligence
AIOTI	Alliance for Internet of Things Innovation
ALTAI	Assessment List for Trustworthy Artificial Intelligence
AMQP	Advanced Message Queuing Protocol
API	Application Programming Interface
BEMS	Building Energy Management System
BESS	Battery Energy Storage System
BDVA	Big Data Value Association
BMS	Building Management System
BSP	Balancing Service Provider
BTM	Behind-the-Meter
BUC	Business Use Case
CA	Consortium Agreement
CEM	Customer Energy Manager
CGMES	Common Grid Model Exchange Standard
CI/CD	Continuous Integration / Continuous Deployment
CIM	Common Information Model
CM	Congestion Management
CNN	Convolutional Neural Network
CRA	Cyber Resilience Act
CRUD	Create, Read, Update, Delete
CSA	Coordination and Support Action
CSV	Comma-Separated Values
DAPS	Dynamic Attribute Provisioning Service
DCAT	Data Catalog Vocabulary
DEMS	Distributed Energy Management System
DEP	Data Exchange Platform
DER	Distributed Energy Resource
DERA	Data Exchange Reference Architecture
DESAP	Digitalising the Energy System Action Plan
DID	Decentralised Identifier
DLR	Dynamic Line Rating
DM	Data Management

Co-funded by the European Union

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

DMP	Data Management Plan
DMS	Distribution Management System
DoA	Description of the Action
DPP	Digital Platform Provider
DRL	Deep Reinforcement Learning
DSC	Dataspace Connector
DSFM	Demand-Side Flexibility Management
DSO	Distribution System Operator
DSP	Dataspace Protocol
DSSC	Data Spaces Support Centre
DTR	Dynamic Thermal Rating
EaaS	Energy-as-a-Service
EC	European Commission
EDC	Eclipse Dataspace Components
eIDAS	Electronic Identification, Authentication and Trust Services
EMS	Energy Management System
EnC	Energy Community
ESCO	Energy Service Company
ESPR	Ecodesign for Sustainable Products Regulation
ETSI	European Telecommunications Standards Institute
EU	European Union
EV	Electric Vehicle
FIWARE	Future Internet Ware
FM	Flexibility Management
FR	Functional Requirement
FSP	Flexibility Service Provider
GA	Grant Agreement
GDPR	General Data Protection Regulation
GIS	Geographic Information System
GMC	Grid Monitoring and Control
GUI	Graphical User Interface
HEDGE-IoT	Holistic Approach towards Empowerment of the DiGitalisation of the Energy Ecosystem through adoption of IoT solutions
HEMRM	Harmonised Electricity Market Role Model
HEMS	Home Energy Management System
HLA	High-Level Architecture
HLEG	High-Level Expert Group
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
HVAC	Heating, Ventilation and Air Conditioning
IdP	Identity Provider

Co-funded by the European Union

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

IdPs	Identity Providers
IDS	International Data Spaces
IDSA	International Data Spaces Association
IDS-RAM	International Data Spaces Reference Architecture Model
IEC	International Electrotechnical Commission
IED	Intelligent Electronic Device
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IOP	Interoperability
IoT	Internet of Things
IoT-EPI	IoT European Platforms Initiative
ISO	International Organization for Standardization
IT	Information Technology
CoC ESA	Code of Conduct for Energy Smart Appliances
JSON	JavaScript Object Notation
JSON-LD	JavaScript Object Notation for Linked Data
LCIM	Levels of Conceptual Interoperability Model
LFM	Local Flexibility Market
LPWAN	Low-Power Wide-Area Network
LV	Low Voltage
M2M	Machine-to-Machine
MED	Main External Data
mFRR	Manual Frequency Restoration Reserve
ML	Machine Learning
MO	Market Operator
MQTT	Message Queuing Telemetry Transport
MQTTS	Message Queuing Telemetry Transport Secure
MV	Medium Voltage
NEMO	Nominated Electricity Market Operator
NGSI	Next Generation Service Interface
NGSI-LD	Next Generation Service Interface – Linked Data
NIS2	Network and Information Security Directive 2
NIST	National Institute of Standards and Technology
O&M	Operations and Maintenance
OAuth 2.0	Open Authorisation 2.0
OBJ	Objective
ODC	Open Data Connector
ODRL	Open Digital Rights Language
OCF	Open Connectivity Foundation
OF	Optimisation and Forecasting
OPC-UA	Open Platform Communications Unified Architecture

Co-funded by the European Union

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

OpenADR	Open Automated Demand Response
OWL	Web Ontology Language
PEP	Policy Enforcement Point
PF	Process Flow
PGUI	Power Grid User Interface
PMH	Project Management Handbook
PV	Photovoltaic
QoS	Quality of Service
RA	Reference Architecture
RAM	Reference Architecture Model
RBAC	Role-Based Access Control
RDF	Resource Description Framework
RES	Renewable Energy Sources
REST	Representational State Transfer
RR	Reserve Resource
RTU	Remote Terminal Unit
SaaS	Software-as-a-Service
SAREF	Smart Applications REFerence Ontology
SAREF4ENER	SAREF extension for energy
SAREF4GRID	SAREF extension for smart grids
SCADA	Supervisory Control and Data Acquisition
SD	Self-Description
SERV	Service Identifier
SG	Smart Grid
SGAM	Smart Grid Architecture Model
SIF	Semantic Interoperability Framework
SME	Small and Medium-sized Enterprise
SO	System Operator
SRI	Smart Readiness Indicator
SRV	Services Management
SUC	System Use Case
T&D	Transmission and Distribution
TLS	Transport Layer Security
TSO	Transmission System Operator
TUC	Transversal Use Case
UC	Use Case
UI	User Interface
UML	Unified Modelling Language
V2G	Vehicle-to-Grid
WG	Working Group
WP	Work Package

Co-funded by the European Union

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

X-CCP	Cross-Cutting Characteristics Plan
XML	Extensible Markup Language
YAML	YAML Ain't Markup Language

Co-funded by the European Union

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

1 INTRODUCTION

Deliverable D2.4, 'HEDGE-IoT Reference Architecture (Final Release)', consolidates the final WP2 architectural baseline for the operational framework and interoperable EDS-aligned Reference Architecture and is explicitly linked to Tasks 2.6 and 2.7. D2.4 is therefore an architectural consolidation deliverable, not the project's final implementation or validation report.

1.1. HEDGE-IOT PROJECT INTRODUCTION AND SUMMARY

HEDGE-IoT (Holistic Approach towards Empowerment of the Digitalisation of the Energy Ecosystem through adoption of IoT solutions) is a Horizon Europe Innovation Action (Grant Agreement No. 101136216) aimed at advancing the digital transformation of the European energy system through an interoperable IoT enabled edge-to-cloud framework. The project addresses the deployment of IoT assets across different levels of the energy value chain, from behind-the-meter environments to distribution and transmission system operation, with the objective of enhancing intelligence, flexibility, resilience, and observability.

At the core of the project is the development of a federated edge-to-cloud architecture capable of supporting real-time data processing, advanced analytics, AI/ML services, federated learning, and computational orchestration. In this way, HEDGE-IoT establishes a scalable and interoperable environment in which data, services, and computation can be shared securely across heterogeneous technical platforms and stakeholder domains.

The project is structured around four complementary pillars: Technology Facilitator, which strengthens intelligence at the edge through computational sharing; Interoperability, which addresses semantic and technical compatibility through Dataspace compliant approaches; Standardisation, which promotes the adoption and reuse of harmonised frameworks and ontologies such as SAREF; and Ecosystem Enabler, which supports trust, stakeholder engagement, and ethical digitalisation. HEDGE-IoT will be validated through six national pilots and extended through Open Calls targeting SMEs and innovators, thereby supporting wider replication and uptake in the European energy ecosystem.

1.2. D2.4 - SCOPE AND OBJECTIVES

D2.4 provides the final WP2 release of the HEDGE-IoT Reference Architecture (RA). Thus, it provides the updated and final version of the operational framework and of the interoperable EDS-aligned Reference Architecture, as a blueprint for software design, based on Tasks 2.6 and 2.7. Accordingly, D2.4 does not serve as the final implementation report of the project, nor as the final demonstration report. Instead, it consolidates the architectural baseline.

Within that scope, D2.4 has four main objectives. First, it provides its architectural principal layers, core components, interfaces, and viewpoints. Second, it documents how the requirements and functional specifications developed in WP2 are translated into an operational reference architecture that supports interoperability, secure data sharing, semantic alignment, and coordinated cloud-to-edge execution. Third, it incorporates the latest project evidence available at month 28, notably from the first and intermediate WP3 and WP4 releases and from the demo preparation and demo-phase reports available by that

Co-funded by the European Union

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

date, without overstating the maturity of deliverables scheduled beyond month 28. Fourth, it establishes a stable architectural reference that supports the remaining implementation, integration, validation, exploitation, and replication activities of the project.

1.3. WP2 – STAKEHOLDERS REQUIREMENTS AND SYSTEM SPECIFICATIONS

WP2, "Stakeholders Requirements and System Specifications", runs from month 1 to month 28 of the project and delivers the requirement, functional-specification, and reference-architecture baseline of HEDGE-IoT through D2.1, D2.2, D2.3, and D2.4. Within WP2, Task 2.6 is responsible for defining the functional specifications for interoperability and standardised data integration based on the BUCs and SUCs, while Task 2.7 is responsible for designing the operational framework and the interoperable EDS-aligned Reference Architecture. D2.4 is therefore the final architectural output of the WP2 workstream rather than a standalone document detached from prior WP2 results.

D2.4 consolidates the results of the WP2 inputs already available, namely D2.1 "Requirements on an IoT Cloud/Edge System for the Energy Ecosystem", D2.2 "Functional Specifications of the HEDGE-IoT system", and D2.3 "HEDGE-IoT Reference Architecture (First Release)". It also considers the project outputs that are already available by that date and are architecturally relevant: D3.1, D3.2, D3.3, D3.4, D4.1, D4.2, D5.1, D5.2, and D7.4. These documents form the validated evidence base that D2.4 draws upon, for its architectural consolidation at the conclusion of WP2.

The relationship with WP3 is therefore implementation-informed but also time-bounded. D3.1 maps the proprietary digital interfaces, platforms, and tools engaged in the pilots, while D3.3 and D3.4 provide the first and intermediate releases of the technological enablers. These outputs support the architectural refinement of D2.4, but the WP3 final release, D3.5, remains a downstream milestone due after D2.4.

The relationship with WP4 is similarly important and should be described with the same informative manner. D4.1 and D4.2 provide the first and intermediate releases of the interoperability framework and integrated solution, including the Open Services Catalogue, Middleware, Data Connector, IoT Edge/Cloud Integration, and Cybersecurity and AI Safety. These outputs are part of the evidence base for D2.4. However, D4.3 is due at a later stage of the project and must therefore be treated in D2.4 as a downstream implementation milestone rather than as an input already incorporated.

The relationship with WP5 is to provide the validation context available. D5.1 establishes the baseline analysis, evaluation criteria, and harmonisation guidelines for the demos; D5.2 reports on the pre-demo phase. These outputs provide the latest demonstrator evidence that can be reflected in D2.4 at the end of WP2. Later pilot reports (D5.3) and the final cross-demo synthesis remain outside the evidence window of this deliverable.

Accordingly, D2.4 should be read as the architectural synthesis of the requirements and specifications established in WP2 and of the implementation and demonstration evidence available up to now.

Co-funded by the European Union

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

1.4. REFERENCE AND APPLICABLE DOCUMENTS

The general guidelines for the project implementation are defined in the HEDGE-IoT Grant Agreement (GA), the Consortium Agreement (CA), the Project Management Handbook (D1.1 – PMH), and the Data Management Plan (D1.4 – DMP).

This deliverable does not replace or supersede any of these governing documents. In case of inconsistency, the following order of precedence applies:

- Grant Agreement
- Consortium Agreement
- D1.1 "Project Management Handbook" [44]
- D1.4 "Data Management Plan" [45]

To avoid ambiguity, this section distinguishes between:

i. project documents used as inputs or reference material in the present deliverable; and

ii. downstream deliverables that are mentioned only as future references and are not used as evidence for completed results in D2.4.

PROJECT DOCUMENTS USED AS INPUTS OR REFERENCE MATERIAL IN D2.4

The following project deliverables are available within the D2.4 evidence window and are used, where relevant, as inputs or reference material for this deliverable:

- D2.1 "Requirements on an IoT Cloud/Edge System for the Energy Ecosystem" [46]
- D2.2 "Functional Specifications of the HEDGE-IoT system" [47]
- D2.3 "HEDGE-IoT Reference Architecture (First Release)" [60]
- D3.1 "HEDGE-IoT Interfaces and Tools for Interoperability" [48]
- D3.2 "HEDGE-IoT Interfaces and Tools for Interoperability 2" [116]
- D3.3 "HEDGE-IoT Technological Enablers (First Release)" [49]
- D3.4 "HEDGE-IoT Technological Enablers (Intermediate Release)" [61]
- D4.1 "HEDGE-IoT Interoperability Framework and Integrated Solution (First release)" [62]
- D4.2 "HEDGE-IoT Interoperability Framework and Integrated Solution (Intermediate release)" [63]
- D5.1 "Guidelines for Demo Preparation" [64]
- D5.2 "Pre-Demo Phase Report" [65]

Co-funded by the European Union

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

- D7.4 "Dissemination, Exploitation and Market Exploration, Standardisation, and Community Building (Intermediate Release)" [66]

DOWNSTREAM DELIVERABLES REFERENCED (FUTURE ALIGNMENT)

The following deliverables are downstream or parallel outputs and may be mentioned in the text only to indicate planned follow-up work, later validation, final implementation evidence, or future dissemination and exploitation activities:

- D3.5 "HEDGE-IoT Technological Enablers (Final Release)" [72]
- D4.3 "HEDGE-IoT Interoperability Framework and Integrated Solution (Final Release)" [73]
- D5.3 "Full Demo Phase Report" [74]
- D7.5 "Dissemination, Exploitation and Market Exploration, Standardisation, and Community Building (Final Release)" [75]

1.5. STRUCTURE OF THE DOCUMENT

After this introductory chapter, the document is structured as follows:

Chapter 2 presents the methodological basis and consolidation logic of the reference architecture, including the main design principles and the rationale used to move from requirements and prior architectural work toward the final WP2 architecture baseline.

Chapter 3 summarises the final functional specification delta relevant to D2.4, including the updated view of BUCs, SUCs, and TUCs and the interoperability profile summary.

Chapter 4 contains the final HEDGE-IoT Reference Architecture itself, including its principal concepts, layers, mappings, and architectural viewpoints, and addresses security architecture concerns.

Chapter 5 provides the components-to-requirements traceability perspective that links architectural building blocks back to the requirements baseline.

Finally, **Chapter 6** concludes the deliverable by summarising the key architecture decisions, identifying known limitations and open issues, and outlining the roadmap toward the remaining implementation, integration, validation, exploitation, and replication phases of the project.

Co-funded by the European Union

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

2 REFERENCE ARCHITECTURE METHODOLOGY

This chapter presents the methodological framework used to evolve and consolidate the HEDGE-IoT RA from its initial version in D2.3 to the final WP2 release in D2.4. The approach integrates architectural description principles, stakeholder-driven design, use-case analysis, alignment with the technological landscape, and implementation feedback from related technical work packages. Rather than reiterating the detailed landscape analysis provided in D2.3, this chapter explains how those findings—together with the requirements and functional specifications developed within WP2—have informed and shaped the final architecture. It therefore focuses on the guiding design principles, architectural models, interoperability considerations, and convergence logic underpinning the resulting reference architecture.

2.1. REFERENCE ARCHITECTURE DESIGN APPROACH AND PRINCIPLES

The HEDGE-IoT RA has been developed to support a functional, interoperable, and scalable cloud-to-edge ecosystem for the energy sector. Its design is grounded in the need to combine secure data sharing, decentralised service execution, semantic interoperability, and cross-platform integration within a coherent operational framework. In practical terms, this means that the architecture must not only accommodate heterogeneous IoT devices, local platforms, and cloud services, but also enable these elements to interact under explicit governance, identity, and policy rules.

The architectural approach is based on a convergence between functional specifications and software architecture. The final RA does not emerge in isolation, but from the progressive integration of stakeholder requirements, business and system use cases, technical specifications, pilot commonalities, and implementation constraints. This convergence is essential for ensuring that architectural decisions remain anchored in real operational needs and that the resulting framework can support deployment across diverse pilot environments.

The design also follows a set of stable principles that remain valid from D2.3 to D2.4. These include:

- modularity, so that components can evolve without disrupting the whole system;
- interoperability, so that heterogeneous assets and services can interact meaningfully;
- data sovereignty, so that data providers retain control over access and usage;
- decentralisation, so that intelligence and computation can be distributed across edge and cloud resources; and
- regulatory awareness, so that the architecture remains compatible with European policy and data governance directions.

A further design principle is progressive refinement. The architecture has been elaborated iteratively, moving from a conceptual and high-level model towards a more implementation-oriented and pilot-grounded structure. This has allowed D2.4 to retain the core architectural

Co-funded by the European Union

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

rationale of D2.3 while simplifying its presentation and strengthening its relevance to the final integrated solution.

2.2. ISO/IEC/IEEE 42010:2022 MODEL

ISO 42010:2022 [1] has been used in HEDGE-IoT as a methodological framework for structuring the architectural description and for linking architectural decisions to stakeholder concerns. Its value in the context of HEDGE-IoT lies in the fact that the project addresses a complex socio-technical system that spans multiple layers of the energy system, multiple types of actors, and multiple interacting technologies, including IoT devices, local platforms, cloud services, AI/ML tools, and Dataspace components.

The framework has supported the finalisation of the RA in three main ways. First, it reinforced the distinction between the architecture itself and the architecture description, which is particularly important in a project where the same architectural logic must be communicated to researchers, software developers, system integrators, pilot operators, and external stakeholders. Second, it encouraged the explicit treatment of stakeholder concerns such as interoperability, scalability, trust, security, governance, and deployment feasibility. Third, it provided a disciplined way of organising viewpoints and architectural reasoning across the evolution from the first release to the final release.

In D2.4, the use of ISO 42010:2022 remains methodological rather than diagrammatic. As already clarified in D2.3, the standard was not used to dictate the visual form of the architecture, but to support the architecture design process through a sequence of structured activities. In the final release, this is reflected in the way the architecture description integrates requirements, viewpoints, and implementation evidence into a coherent final narrative rather than into a single formal meta-model representation.

More concretely, the finalisation process benefited from ISO 42010 through the continued use of correspondences between architectural elements, stakeholder concerns, and project artefacts. This allowed the final architecture to be refined without losing traceability to the concerns that originally motivated it, including semantic interoperability, secure data sharing, cross-layer orchestration, and the portability of services across pilots.

2.3. 4+1 ARCHITECTURAL VIEW MODEL

The 4+1 Architectural View Model [2] has been used in HEDGE-IoT as a practical structuring aid to capture the different dimensions of the system without forcing the final architecture into a rigid single-view representation. Its relevance lies in the fact that the HEDGE-IoT ecosystem must simultaneously address logical structuring, runtime processes, software decomposition, physical deployment, and operational scenarios.

In the HEDGE-IoT context, the logical view has helped frame the main architectural components and their relationships, especially across edge assets, dataspace services, semantic enablers, and application-layer services. The process view has informed the treatment of data flows, orchestration logic, negotiation processes, and service execution across the cloud-edge continuum. The development view has supported the structuring of software building blocks, middleware capabilities, interfaces, and interoperable modules. The

Co-funded by the European Union

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

physical view has been reflected in the representation of pilot-specific IoT-edge nodes, local platforms, and distributed deployment settings. Finally, the Scenario view has remained closely linked to the business and system use cases, as well as to the transversal use cases that connect the pilots through common interoperability mechanisms.

In line with the approach already adopted in D2.3, HEDGE-IoT did not produce five fully separate architectural descriptions. Instead, it borrowed the most relevant properties from the different views and consolidated them into a single RA that remains readable, implementation-aware, and grounded in operational scenarios. In D2.4, this application becomes more explicit because the final architecture clearly reflects the integration of logical layering, physical pilot representation, runtime orchestration, and use-case-driven validation in one consolidated model.

2.4. LANDSCAPE ANALYSIS: STANDARDS, INITIATIVES AND RELEVANT PROJECTS

The detailed landscape analysis was already carried out in D2.3 and is therefore not repeated here in full. Nevertheless, the final Reference Architecture remains directly informed by the principal initiatives, standards, and related projects reviewed in that chapter, since these continue to provide the main external architectural anchors for HEDGE-IoT.

The final architecture remains aligned with the most relevant European and international initiatives reviewed in D2.3, namely BRIDGE and the DERA framework [3], SGAM [4], the European AI Alliance and its trustworthy AI orientation [5][6][33], FIWARE [34][35], AIOTI [8][9][37], and IoT-EPI [10]. At the Dataspace and governance level, the most relevant reference points remain IDSA [11] [12], BDVA [15], and Gaia-X [16][17].

In terms of related projects, the final architecture continues to benefit from the architectural lessons identified in ATTEST [19][20], BRIGHT [54][55], Enershare [40], I-NERGY [21], OneNet [22], PLATONE [41], Resonance [42], and Synergy [43]. These projects were particularly useful in clarifying reusable patterns around data exchange, service portability, edge-cloud integration, flexibility management, marketplace functionalities, and interoperable digital energy services.

For D2.4, the role of the landscape analysis is therefore not to reopen the comparative review, but to confirm that the final HEDGE-IoT Reference Architecture remains standards-aware, policy-aligned, and consistent with the broader European evolution toward Dataspaces, semantic interoperability, open service ecosystems, and trustworthy AI. The reader should refer to D2.3 for the detailed rationale and full analytical treatment of each initiative and project.

2.5. IOT INTEROPERABILITY - CHALLENGES AND SOLUTIONS

IoT interoperability remains one of the central design drivers of the HEDGE-IoT architecture. As already discussed in D2.3, the main challenges arise from the lack of common standards across hardware components, communication protocols, and data formats; the persistence of proprietary technologies and closed ecosystems; increased security exposure due to interconnection complexity; the resource limitations of edge devices; and the governance

Co-funded by the European Union

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

difficulties associated with large-scale data management and ownership in distributed environments.

These challenges are particularly relevant in the energy sector, where critical digital infrastructures combine legacy operational technologies, new IoT devices, distributed energy resources, market-facing platforms, and cross-organisational data sharing. In such settings, interoperability cannot be reduced to simple connectivity. It must extend across protocol compatibility, semantic consistency, policy enforcement, trust, and operational integration. These challenges are revisited in Section 3.4.2, where they are translated into concrete interoperability challenges and architectural responses.

The solutions retained in the HEDGE-IoT final architecture follow the same direction as in D2.3 but are now reflected more concretely in the architecture itself. These solutions include the use of standardised protocols and frameworks, openness and collaboration across ecosystem participants, structured testing and validation approaches, interoperability platforms and gateways, and the distributed processing capabilities enabled by edge computing. Standardisation efforts by organisations such as IETF [50] and IEEE [51], communication patterns such as MQTT [52], and validation practices aligned with NIST-style approaches [53] remain highly relevant to the architectural choices adopted in HEDGE-IoT.

Accordingly, the final HEDGE-IoT RA responds to the interoperability challenge by combining dataspace connectors, semantic interoperability enablers, open catalogues, contract-based exchange, policy-governed access and cloud-edge orchestration within one coherent stack. This allows interoperability to be treated as a system property rather than as an isolated technical interface issue. The specific interoperability challenges addressed by the architecture, together with the corresponding architectural responses, are further detailed in Section 3.4.2.

2.6. KEY LESSONS LEARNED AND ARCHITECTURAL INFLUENCE

The detailed review of nine related European projects in D2.3 provided architectural insights relevant to the design of the HEDGE-IoT Reference Architecture. This section consolidates the main lessons extracted from ATTEST, BRIGHT, Enershare, I-NERGY, MATRYCS, OneNet, PLATONE, Resonance, and SYNERGY, and identifies the areas where they align with HEDGE-IoT's architectural choices.

A first cross-cutting observation concerns the role of dataspace governance in enabling multi-stakeholder data exchange. Enershare's reference architecture builds on the IDSA RAM and specialises it for energy and non-energy services, incorporating Dataspace Connectors that integrate local systems with the horizontal dataspace domain. Based on the Enershare experience, the analysis in D2.3 recommended that the adopted connector should be compliant with the Dataspace Protocol to ensure post-project sustainability of the developed services. SYNERGY complements this perspective through its Core Cloud Platform, which combines a Security and Authorisation Engine, an API Gateway, and a Data Handling Manager, ensuring that only authorised entities can access data. These approaches are consistent with the HEDGE-IoT Dataspace Layer, which adopts Eclipse Dataspace Component (EDC) Connectors, a federated data catalogue for metadata publication and discoverability, and an Identity Provider for secure and policy-compliant access.

Co-funded by the European Union

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

A second lesson relates to the decoupling of services from the underlying data infrastructure. I-NERGY organised its architecture into a Data Service Layer, an AI Trained Models Layer, and an Application Layer, supported by a Marketplace for developing and deploying AI-based energy services. PLATONE followed a different approach, structuring its Open Framework around a two-layer blockchain architecture (Access Layer and Service Layer) linked by a Shared Customer Database that stores and makes flexibility-related data available to authorised platforms and stakeholders. Both approaches illustrate the value of separating service logic from data exchange mechanisms, a principle reflected in the HEDGE-IoT Dataspace Layer, where the open service catalogue and the App Store coexist as distinct functional groupings alongside the dataspace core and the Orchestrator, each serving a different aspect of service publication, discovery, and reuse.

Third, semantic interoperability emerged as a necessary condition for cross-platform integration. OneNet adopted the FIWARE Context Broker as the unifying middleware component of its decentralised architecture, enabling different platforms from the Network of Platforms to communicate. Resonance contributed a complementary perspective through its DSFM framework, which provides catalogues of standardised software services for resource managers, customer energy managers, and aggregation platforms. These findings are aligned with the decision to elevate semantic interoperability into a dedicated architectural layer in the final HEDGE-IoT RA. The Semantic Interoperability Layer includes the SIF Knowledge Engine, ontologies and vocabularies, verification and validation tooling, and model management functions, making cross-pilot portability of data and services a direct architectural concern.

Fourth, several projects demonstrated the value of combining cloud-based analytics with distributed processing closer to the data source. ATTEST developed a joint ICT platform that integrates planning, operation, and asset management modules for TSO-DSO coordination, supported by a Data Access Layer, a Converter Layer, and an Orchestration Layer. BRIGHT deployed a set of technological enablers across different levels: B-EMHC for electricity metering and smart home automation, B-DT for digital twins at consumer and community level, and B-FLEX for AI-driven flexibility management at end-user, community, and system level. MATRYCS structured its reference architecture around three interconnected layers (Governance, Processing, and Analytics) addressing the buildings domain. These patterns are relevant to the final HEDGE-IoT architecture, where fog/cloud services and federated learning are placed in the Application Layer, while the Orchestrator in the Dataspace Layer coordinates workload distribution across the cloud-edge continuum.

Finally, scalability and replicability were recurrent themes across the project landscape. ATTEST committed to a modular, open-source toolbox using open data formats (MATPOWER, JSON, CSV, Excel). OneNet organised its demonstrators in four clusters covering countries across all European regions, validating its solutions in multiple regulatory environments. BRIGHT engaged approximately 1,000 predominantly residential consumers across four demonstration sites in diverse community configurations including Local Energy Communities, Citizen Energy Communities, and Virtual Energy Communities. These experiences support the HEDGE-IoT approach of designing a Reference Architecture that can accommodate six demonstration sites across different countries, grid topologies, and stakeholder configurations.

Co-funded by the European Union

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

In summary, the related projects collectively highlighted five areas of architectural relevance: (i) full-stack dataspace governance with IDSA-aligned principles, (ii) decoupled and marketplace-oriented service design, (iii) semantic interoperability as a dedicated architectural layer, (iv) distributed computation across fog/cloud and edge, and (v) architectural openness for cross-site scalability.

2.7. LANDSCAPE ALIGNMENT SUMMARY

This section maps the final HEDGE-IoT RA against major European and international initiatives, standards and frameworks. The alignment demonstrates that the RA was not developed in isolation, but rather as a natural evolution of established architectural frameworks, interoperability principles, and standardisation trends that are shaping the European digital data ecosystem. The landscape alignment serves three critical purposes: (i) it validates that HEDGE-IoT architectural decisions remain consistent with broader European policy directions such as the Green Deal, the Energy Transition, and the Digitalisation of Energy System Action Plan (DESAP); (ii) it ensures technical coherence with established reference architectures and interoperability frameworks such as SGAM and BRIDGE DERA; and (iii) it confirms that the project's semantic, governance, and technical choices are based on recognized standards and best practices.

2.7.1 Key Alignment Dimensions

- **Policy and Regulatory Alignment:** Clean Energy Package (EU 2019/944), GDPR, NIS2[1] Directive, Data Act[2], AI Act[3], Cyber Resilience Act[4] (CRA), JRC Code of Conduct for Energy Smart Appliances (CoC ESA), Smart Readiness Indicator (SRI), Ecodesign for Sustainable Products Regulation (ESPR) and DESAP
- **Reference Architecture Alignment:** SGAM (Smart Grid Architecture Model), BRIDGE DERA (Data Exchange Reference Architecture), IDS-RAM (International Data Spaces Reference Architecture Model)
- **Semantic and Standardisation Alignment:** SAREF (Smart Applications Reference Ontology), IEC CIM (Common Information Model), IEC 61850, FIWARE
- **Governance and Interoperability Alignment:** IDSA (International Dataspaces Association), BDVA (Big Data Value Association), Gaia-X
- **Initiative and Project Alignment:** AIOTI (Alliance for Internet of Things Innovation), IoT-EPI, related projects (ATTEST, BRIGHT, Enershare, I-NERGY, OneNet, PLATONE, Resonance, SYNERGY)

[1] NIS2 entered into force in January 2023; Member States had until 17 October 2024 to transpose it, and NIS1 was repealed from 18 October 2024. https://digital-strategy.ec.europa.eu/en/policies/nis2-directive
[2] The Data Act is applicable from 12 September 2025.
[3] The AI Act entered into force on 1 August 2024 and applies in phases. https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en
[4] The Cyber Resilience Act entered into force on 10 December 2024, with reporting obligations applying from 11 September 2026 and main provisions from 11 December 2027. https://digital-strategy.ec.europa.eu/en/policies/data-act

Co-funded by the European Union

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

TABLE 1 – DIMENSION MAPPING & ALIGNMENT LEVEL

Framework/ Standard	Primary Purpose	HEDGE-IoT Layers	Components/Elements	Alignment Level
SGAM	Energy System Interoperability	Physical, Dataspace, Application	Five interoperability layers + five zones	**Strong**
BRIDGE DERA	Energy Data Exchange Architecture	Physical, Dataspace, Semantic	14 services + semantic enablers	**Strong**
IDSA	Data Sovereignty & Governance	Dataspace Layer	EDC, Dataspace Protocol, Policy Enforcement	**Strong**
SAREF	Semantic Interoperability	Semantic Layer	Ontologies, Knowledge Engine	**Core**
IEC Standards	Power System Data Models	Physical, Semantic	CIM, 61850, 62559-2	**Essential**
FIWARE	Smart City/Energy Platforms	Application Layer	NGSI-LD, Smart Data Models	**Supported**
Gaia-X	Federated Cloud Infrastructure	Dataspace, Application	Sovereignty, Interoperability, Trust	**Aligned**
ISO 42010	Architecture Methodology	All Layers	Viewpoints, Stakeholder Concerns	**Methodological**
4+1 View Model	Multi-perspective Architecture	All Layers	Logical, Process, Development, Physical, Scenarios	**Structuring**

2.7.2 Mapping Explanations

SGAM (SMART GRID ARCHITECTURE MODEL)

SGAM is a five-layer interoperability model (Business, Functional, Information, Communication, Component) spanning five operational zones (Process, Field/Station, Operation, Enterprise, Market). HEDGE-IoT maps as follows:

- **Business Layer:** HEDGE-IoT's business use cases, governance logic, and market interactions (flexibility services, local flexibility markets, reserve participation)
- **Functional Layer:** Edge intelligence, cloud-edge orchestration, Dataspace services, AI/ML models, and policy enforcement
- **Information Layer:** SAREF, CIM, PowerCIM, semantic interoperability layer, and knowledge engine

Co-funded by the European Union

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

- **Communication Layer:** REST API, MQTT, HTTP/HTTPS, EDC connectors, and Dataspace Protocol
- **Component Layer:** IoT nodes, edge gateways, sensors, App Store, metadata broker, EDC connectors

BRIDGE DERA (DATA EXCHANGE REFERENCE ARCHITECTURE)

BRIDGE DERA provides a reference architecture model for energy data exchange across five layers: Business Actors & Ecosystems, Innovative Data Analytics Services, Data Interoperability (Information), Data Interoperability (Communication), and Data Sources & Components. HEDGE-IoT alignment:

- **Business Layer:** Regulatory frameworks (Clean Energy Package, GDPR, NIS2, AI Act), stakeholder roles (TSOs, DSOs, DER operators, aggregators, consumers, market actors)
- **Function Layer:** 14 energy services (forecasting, anomaly detection, congestion management, flexibility orchestration, AI/ML services, orchestration framework)
- **Information Layer:** IDS-RAM, SAREF/SAREF4ENER, CIM/CGMES, DCAT, RDF/OWL, semantic enablers (Semantic Treehouse, PowerCIM, Knowledge Engine)
- **Communication Layer:** JSON, RDF, Parquet, XML, MQTT, REST API, TLS, OAuth 2.0, HTTP/HTTPS, Kafka
- **Component Layer:** EDC, App Store, Metadata Broker, Identity Hub, SCADA, EMS, BEMS, meters, IoT devices, PV+BESS systems, smart appliances

IDSA (INTERNATIONAL DATA SPACES ASSOCIATION)

IDSA provides governance principles and technical specifications for sovereign data exchange. HEDGE-IoT incorporates:

- **IDS-RAM (4 layers):** Business, Function, Information, Communication, System
- **Dataspace Protocol (DSP: 2025-1):** Standardized negotiation, contract definition, and data exchange mechanisms
- **Eclipse Dataspace Connector:** Concrete implementation of sovereign data exchange
- **Data Sovereignty Principles:** Usage control, policy enforcement, data provider control

SAREF (SMART APPLICATIONS REFERENCE ONTOLOGY)

SAREF is the primary semantic interoperability enabler in HEDGE-IoT. The project adopts:

- **Core SAREF:** Device, service, property, and action definitions
- **SAREF4ENER:** Energy-specific extensions for electric appliances and energy management
- **SAREF4GRID:** Grid-specific concepts for distribution and transmission networks

Co-funded by the European Union

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

- **SAREF-CIM Bridge:** Mapping between SAREF and IEC CIM for harmonized semantic representation

IEC STANDARDS (CIM, 61850, 62559)

HEDGE-IoT incorporates key IEC standardisation:

- **IEC CIM:** Common Information Model for power system data exchange
- **IEC 61850:** Protocol and data model for substation automation
- **IEC 62559-2:** Use case methodology applied in D2.1, D2.2, and D2.4
- **IEC 62746 / OpenADR:** Demand-side flexibility and automated demand response

FIWARE AND NGSI-LD

HEDGE-IoT supports FIWARE-compatible approaches:

- **NGSI-LD:** Linked data API standard compatible with HEDGE-IoT data exchange
- **Smart Data Models:** Domain-specific data models aligned with FIWARE ecosystem
- **Context Broker patterns:** Semantic data aggregation and distribution mechanisms

GAIA-X FRAMEWORK

Gaia-X principles for federated sovereign cloud infrastructure:

- **Data Sovereignty:** Users retain control over their data and services
- **Interoperability:** Portability across multiple cloud and edge providers
- **Trustworthiness:** Transparency, security, and compliance certification

2.7.3 Key Alignment Findings

- **Standards Coherence:** HEDGE-IoT's layered architecture is closely aligned with both SGAM and BRIDGE DERA, demonstrating that the project's reference architecture is not a deviation from established models but a refined instantiation of them
- **Semantic Consensus:** The elevation of semantic interoperability to a dedicated architectural layer, using SAREF as the primary vocabulary and Knowledge Engine for federation, showcases a consensus in the energy transition community that semantic interoperability is non-negotiable for cross-platform integration
- **Governance Alignment:** HEDGE-IoT's adoption of IDSA principles, EDC technology, and Dataspace Protocol ensures post-project sustainability and alignment with European dataspace initiatives such as the Common European Energy Data Space
- **Policy Compliance:** The architecture's support for GDPR, NIS2, AI Act, and the Clean Energy Package demonstrates explicit alignment with European regulatory direction

Co-funded by the European Union

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

- **Project Learning:** Lessons from ATTEST, BRIGHT, Enershare, I-NERGY, OneNet, PLATONE, Resonance, and SYNERGY have directly influenced architectural decisions around dataspace governance, service decoupling, semantic layers, and computational orchestration

2.8. HEDGE-IOT REFERENCE ARCHITECTURE PATTERN

2.8.1 Architecture Development Path and Evolution

This section describes how the HEDGE-IoT Reference Architecture was developed, refined, and evolved from initial conceptual models through to its final release. The architecture did not emerge fully formed, but rather through a structured progression of design and review activities, stakeholder engagement, pilot feedback, and integration with emerging technological capabilities. Understanding this development path is essential for interpreting the final architecture, validating its soundness, and providing confidence in its replicability.

FIGURE 1 – HEDGE-IOT REFERENCE ARCHITECTURE PATTERN

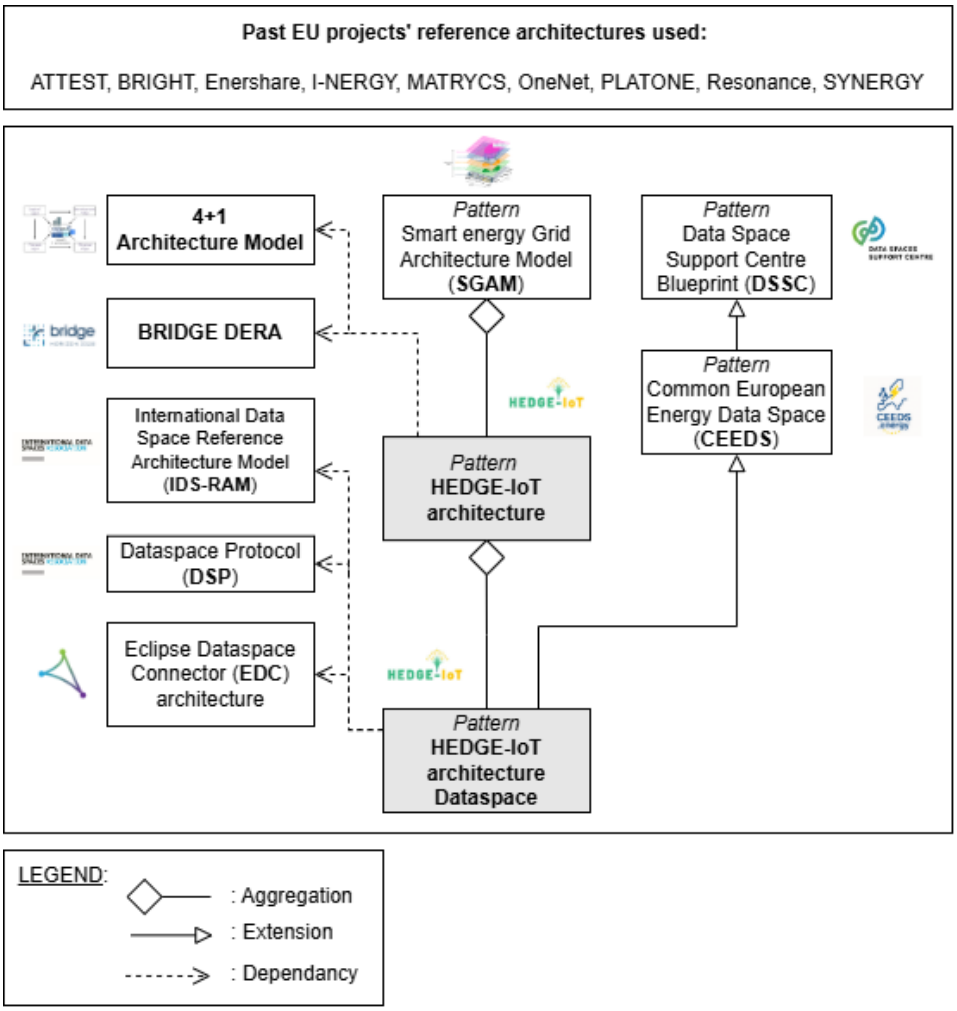

Co-funded by the European Union
This project has received funding from the European Union's Horizon Europe research and innovation programme under the Grant Agreement number 101136216. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European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

2.8.2 The HEDGE-IoT Architecture Development Paradigm

The HEDGE-IoT architecture follows a Convergence-Driven Design approach, integrating four streams of input as presented in Table 2:

TABLE 2 – ARCHITECTURE DEVELOPMENT PARADIGM

Input Stream	Key Artifacts	Architectural Influence	Outcome
Stakeholder Requirements	D2.1 (14 BUCs), D2.2 (39 SUCs), 3 TUCs	Defines functional scope, use case patterns, cross-cutting needs	Physical, Dataspace, Application layers
Landscape Analysis	Standards review, related projects, regulatory analysis	Ensures external alignment, informs component choices, validates feasibility	Dataspace core, semantic layer, governance mechanisms
Technical Implementation	D3.3, D3.4 (enablers), D4.1, D4.2 (interoperability)	Validates architectural assumptions, informs component architecture, enables refinement	App Store, Orchestrator, semantic tools
Pilot Validation	D5.1 demo prep, pilot site configurations	Validates architecture in operational reality identifies practical constraints	Physical layer specifics, deployment patterns

STAGE 1: LANDSCAPE ANALYSIS AND METHODOLOGICAL GROUNDING (D2.3)

Outputs:

- Comprehensive landscape analysis of 30+ standards, initiatives, and related projects
- Methodological framework selection (ISO 42010:2022, 4+1 View Model)
- First architectural model: Five-layer conceptual stack
- SGAM and BRIDGE DERA mapping to validate alignment

Design Principles Established:

- Modularity: Components can evolve independently
- Interoperability: Heterogeneous systems interact meaningfully
- Data Sovereignty: Data providers control access and usage
- Decentralisation: Intelligence and computation distributed across edge-cloud
- Regulatory Awareness: Compatible with European policy and data governance

Co-funded by the European Union

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

STAGE 2: REQUIREMENTS CONSOLIDATION (D2.1 & D2.2)

From Requirements to Architecture:

- **D2.1 (Requirements):** Established 14 Business Use Cases (BUCs) across six pilots
- **D2.2 (Specifications):** Defined 39 System Use Cases (SUCs) with functional requirements
- **Transversal Use Cases (TUCs):** Three cross-cutting TUCs (Data interoperability, Computational orchestration, App Store)
- **Functional Requirements:** 9 categories covering data management, interoperability, services, UI, optimisation, flexibility, grid monitoring, AI, and external data

Architectural Implications:

The BUCs across diverse pilot environments (Finnish, Greek, Italian, Dutch, Portuguese, Slovenian) revealed common architectural needs:

- Need for federated data exchange across organisational boundaries
- Requirement for semantic consistency across heterogeneous platforms
- Need for computational orchestration across edge-cloud continuum
- Requirement for service discovery and reuse across pilots
- Need for explicit governance, identity management, and policy enforcement

STAGE 3: FIRST ARCHITECTURAL CONSOLIDATION (D2.3)

Architecture Model Evolution:

D2.3 presented a five-layer reference architecture:

- **Layer 1 IoT-Edge Services:** Physical sensors, devices, local platforms, and edge processing
- **Layer 2 Dataspace Layer:** EDC connectors, Metadata broker, federated catalogue, contract negotiation
- **Layer 3 Semantic Interoperability:** SAREF, CIM, Knowledge Engine, ontology management
- **Layer 4 Application Services:** Forecasting, anomaly detection, optimisation, flexibility management
- **Layer 5 User Interfaces & Dashboards:** Role-based access, analytics, reporting

Key D2.3 Contributions:

- Established the Dataspace Layer as central to architecture

Co-funded by the European Union

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

- Introduced semantic interoperability as a dedicated concern
- Mapped to SGAM and BRIDGE DERA frameworks
- Identified interoperability profiles and points

STAGE 4: IMPLEMENTATION INTEGRATION AND REFINEMENT (WP3 & WP4 FEEDBACK)

Feedback Loops:

- **WP3 (Technological Enablers):** D3.3, D3.4 showed how AI/ML services, edge intelligence, and semantic tools could be implemented, informing Layer 4 refinements
- **WP4 (Interoperability Framework):** D4.1, D4.2 validated that EDC-based connectors, Open Service Catalogue, and App Store could operate as core architectural elements
- **WP5 (Pilot Validation):** D5.1 preparation guidelines informed on real deployment constraints and user needs

Architectural Refinements:

- App Store and Open Service Catalogue elevated to explicit architectural components
- Computational Orchestration (KubeEdge-based) defined as core Dataspace Layer function
- Federated Learning positioned in Application Layer with orchestration support
- Identity Provider and policy enforcement mechanisms made explicit

STAGE 5: FINAL ARCHITECTURE CONSOLIDATION (D2.4)

Simplification and Clarification:

D2.4 refined the five-layer D2.3 model into a cleaner four-layer structure:

- **Physical Layer:** Pilot infrastructure, IoT nodes, local platforms, digital twins
- **Dataspace Layer:** Dataspace core (EDC, catalogue, identity), App Store, Orchestrator
- **Semantic Interoperability Layer:** Knowledge Engine, ontologies, semantic enablers
- **Application Layer:** User interfaces, role-based access, federated learning, fog/cloud services

Key D2.4 Enhancements:

- Clearer layer separation: Physical, Dataspace, Semantic, Application
- Stronger visibility of dataspace as operational core

Co-funded by the European Union

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

- Elevated semantic interoperability to dedicated layer
- Explicit treatment of pilot-specific implementations
- Better representation of end-user interaction and application accessibility
- Introduction of Hourglass Model for stakeholders and capability visualisation

2.8.3 Key Design Pattern: The Convergence Principle

The HEDGE-IoT RA has been developed through explicit convergence of four streams: stakeholder requirements, landscape alignment, technical implementation, and pilot validation. This convergence principle distinguishes HEDGE-IoT from purely top-down or bottom-up architectural approaches. The architecture is neither dictated solely by external frameworks (top-down) nor emergent entirely from implementation experience (bottom-up), but rather the product of disciplined integration of multiple sources of truth. This approach ensures the architecture is both theoretically sound and practically grounded, both standards-aligned and pilot-validated, and both ambitious in scope yet achievable within project constraints.

2.8.4 Architecture Replicability and Future Evolution

The HEDGE-IoT RA pattern has been designed for replicability beyond the project context. The architecture's alignment with SGAM and BRIDGE DERA, its adoption of IDSA and Gaia-X principles, and its use of established semantic standards (SAREF, CIM) ensure that the framework can be extended and adapted in future energy digitalisation initiatives. The four-layer structure remains flexible enough to accommodate emerging technologies (new AI/ML models, evolving communication protocols, advanced edge hardware) while maintaining core architectural principles. Post-project, the architecture can be extended through community contributions, standardisation evolution, and new pilot implementations, while the landscape alignment chapter provides clear guidance on how extensions must maintain coherence with established frameworks.

Co-funded by the European Union

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

3 HEDGE-IOT FUNCTIONAL SPECIFICATIONS

This section presents the final version of the HEDGE-IoT functional specifications, which were defined in the form of System Use Cases (SUCs), Transversal (system) Use Cases (TUCs), functional requirements catalogue, and interoperability profiles. D2.2 and D2.3 serve as intermediate deliverables that provide the basis for this final version, and provide complementary information (e.g., methodologies, pilots' commonalities, and previous versions of this work). These results were used as input for the RA design, pilots' development, component development, and in general, as a basis for all project activities.

3.1. BUSINESS AND SYSTEM USE CASES

Business Use Cases (BUCs) and System Use Cases (SUCs) were defined respectively in deliverables D2.1 [46], and D2.2 [47]. They were defined using the IEC 62559 [25] template, and the methodology is described in these deliverables.

A use case (UC) describes the interactions of various actors in a system to achieve specific goals. A business use case (BUC) represents a business process, while a system use case (SUC) describes a function that supports one or more business processes. Both types of UCs are essential for the description of a system and can be used for the system architecture definition. In HEDGE-IoT, the BUCs were the starting point for the pilots to define SUCs and were also used for the definition of the RA. The project's functional requirements and non-functional requirements are specified in the form of SUCs.

Project pilots take the opportunity of this deliverable to update and refine their BUCs and SUCs when needed. It can be linked to changes in a pilot or the availability of new information.

To avoid duplicating extensive use-case documentation in the main body and the absence of major changes for the refined BUCs and SUCs of all pilots, these documents are compiled and uploaded to the project website. The compiled document will be publicly available, while the size of this deliverable is not compromised.

The following sections describe the main changes to highlight the refined BUCs and SUCs.

3.1.1 Business Use Cases

Fourteen (14) BUCs were identified and defined by the project pilots. Table 3 provides an overview of all the BUCs of the project. No changes to the BUCs have been introduced since D2.3, as the pilot-level business objectives remain unchanged.

TABLE 3 – BUCS OVERVIEW

Pilot	BUC ID	BUC name
1-Finnish	BUC-FI-01	Anomaly detection and fault forecasting to increase Medium Voltage (MV) distribution network resilience
	BUC-FI-02	Predictive and real-time congestion management (CM) to increase network hosting capacity

Co-funded by the European Union

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

2-Greek	BUC-GR-01	Flexibility management through active prosumers/consumers engagement
	BUC-GR-02	Leveraging data exchange and AI edge algorithms for energy forecasting and prevention of critical grid events
	BUC-GR-03	Flexibility trading platform for mitigating problems of the T&D networks
3-Italian	BUC-IT-01	Energy flow optimisation with dynamic grid limits
	BUC-IT-02	Flexibility provided by Energy Community to solve a local congestion
4-Dutch	BUC-NL-01	Energy Flexibility at business park
	BUC-NL-02	Enhance local grid resilience through detection & prevention
5-Portuguese	BUC-PT-01	GreenVale: Harnessing the potential of energy communities by leveraging Federated Learning strategies
	BUC-PT-02	Participation of industrial and residential energy communities in ancillary services market for the TSO
	BUC-PT-03	Flexibility aggregation at tertiary buildings
6-Slovenian	BUC-SI-01	Maximizing asset capacity for increased lifetime of DSO and TSO equipment
	BUC-SI-02	Enhanced Network Manageability and Observability

3.1.2 System Use Cases

This section presents the SUCs enabling the realisation of BUCs. Updates in this release are primarily slight or to ensure consistency with the pilot.

FINNISH PILOT SUCS

Table 4 maps the pilot BUCs to the corresponding SUC:

TABLE 4 - FINNISH PILOT SUCS

BUC ID & BUC name	SUC ID	SUC name
BUC-FI-01 Anomaly detection and fault forecasting to increase Medium Voltage (MV) distribution network resilience	SUC-FI-01.1	Data collection and anomaly detection
	SUC-FI-01.2	Fault forecasting
BUC-FI-02 Predictive and real-time congestion	SUC-FI-02.1	Congestion prediction in distribution grids

Co-funded by the European Union

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

management (CM) to increase network hosting capacity	SUC-FI-02.2	Congestion management planning in distribution grids
	SUC-FI-02.3	State monitoring of the distribution grid
	SUC-FI-02.4	Congestion management decision-making in real-time

The Finnish pilot's main SUCs refinements are:

1. Minor changes to reflect pilot real implementation.

GREEK PILOT SUCS

Table 5 links the BUCs and the SUCs of the pilot:

TABLE 5 - GREEK PILOT SUCS

BUC ID & BUC name	SUC ID	SUC name
BUC-GR-01 Flexibility management through active prosumers/consumers engagement	SUC-GR-01.01	Optimisation of Flexibility Distribution
	SUC-GR-01.02	Demand Forecasting
	SUC-GR-01.03	Production Forecasting
	SUC-GR-01.04	Edge Processing
	SUC-GR-01.05	User Interaction
BUC-GR-02 Leveraging data exchange and AI edge algorithms for energy forecasting and prevention of critical grid events	SUC-GR-02.01	Energy Grid Management using Forecasting Data
BUC-GR-03 Flexibility trading platform for mitigating problems of the T&D networks	SUC-GR-03.01	Registration & Prequalification on Local Flexibility Market
	SUC-GR-03.02	Flexibility Trading

The Greek pilot's main refinements are:

1. Minor textual refinements; no content changes.

ITALIAN PILOT SUCS

Table 6 links the BUCs and the SUCs of the pilot:

TABLE 6 - ITALIAN PILOT SUCS

Co-funded by the European Union

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

BUC ID & BUC name	SUC ID	SUC name
BUC-IT-01 Energy flow optimisation with dynamic grid limits	SUC-IT-01.1	Energy community power management
	SUC-IT-01.2	Energy community performance forecasting
BUC-IT-02 Flexibility provided by Energy Community to solve a local congestion	SUC-IT-02.1	Grid behaviour forecasting
	SUC-IT-02.2	Grid congestion computing
	SUC-IT-02.3	Localized weather forecast

No changes were introduced by the Italian pilots.

DUTCH PILOT SUCS

Table 7 links the BUCs and the SUCs of the pilot:

TABLE 7 - DUTCH PILOT SUCS

BUC ID & BUC name	SUC ID	SUC name
BUC-NL-01 Energy Flexibility at business park	SUC-NL-01.1	Monitor energy nodes and local grid & dashboard for data insights
	SUC-NL-01.2	Integrate energy nodes and EMS/BMS via semantics for control and explainability
	SUC-NL-01.3	Optimise energy production & consumption
	SUC-NL-01.4	Flexibility alignment
BUC-NL-02 Enhance local grid resilience through detection & prevention	SUC-NL-02.1	Anomaly and fault detection in the local grid
	SUC-NL-02.2	Predictive maintenance

The Dutch pilot's main changes are:

1. a stricter definition of the scope,
2. the removal of out-of-scope elements, and
3. an update of the UML diagrams to match the two previous elements.

PORTUGUESE PILOT SUCS

Table 8 links the BUCs and the SUCs of the pilot:

Co-funded by the European Union

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

TABLE 8 - PORTUGUESE PILOT SUCS

BUC ID & BUC name	SUC ID	SUC name
BUC-PT-01 GreenVale: Harnessing the potential of energy communities by leveraging Federated Learning strategies	SUC-PT-01.1	Connect flexibility providers across the DPP flexibility value chain
	SUC-PT-01.2	Enable Data Exchange via Dataspaces
	SUC-PT-01.3	Mobilizing Energy Flexibility
	SUC-PT-01.4	Activation of Energy Flexibility
BUC-PT-02 Participation of industrial and residential energy communities in ancillary services market for the TSO	SUC-PT-02.1	Bidding & Selection
	SUC-PT-02.2	aFRR/mFRR Activation
	SUC-PT-02.3	aFRR / mFRR Settlement
BUC-PT-03 Flexibility aggregation at tertiary buildings	SUC-PT-03.1	Integrate flexible assets from commercial buildings
	SUC-PT-03.2	Default valorisation scenario based on price hedging
	SUC-PT-03.3	TSO valorisation scenario

No changes were introduced by the Portuguese pilot.

SLOVENIAN PILOT SUCS

Table 9 links the BUCs and the SUCs of the pilot:

TABLE 9 - SLOVENIAN PILOT SUCS

BUC ID & BUC name	SUC ID	SUC name
BUC-SI-01 Maximizing asset capacity for increased lifetime of DSO and TSO equipment	SUC-SI-01.1	Dynamic Thermal Rating (DTR) edge calculation
	SUC-SI-01.2	Dynamic Line Rating (DLR) edge calculation
BUC-SI-02 Enhanced Network Manageability and Observability	SUC-SI-02.1	Semantic model of the substation
	SUC-SI-02.2	ML algorithm for enhanced network management and planning

No changes were introduced by the Slovenian pilot.

Co-funded by the European Union

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

3.2. TRANSVERSAL USE CASES

Following the definition of pilots' BUCs and SUCs, the project decided to define transversal system use cases to complete the project specifications with the aim to:

- address common needs across use cases,
- enable harmonisation and interoperability, and
- support transversal components specifications, implementation and integration (done by other tasks).

A transversal use case (TUC) is a SUC that refers to a cross-cutting use case that supports multiple BUCs and SUCs by addressing shared capabilities, services, or requirements that are common across different pilots. Unlike BUCs, which are driven by specific stakeholder goals, and pilots' SUCs, which describe vertical functional interactions, transversal use cases are designed to enable interoperability, reusability, and consistency across the entire ecosystem by specifying horizontal functional interactions.

The following sections provide an overview and a summary of the TUCs' content. Their complete documents can be found in Annex C of this deliverable.

Table 10 provides an overview of the selected HEDGE-IoT Transversal Use Cases (TUCs).

TABLE 10 - HEDGE-IOT TRANSVERSAL USE CASES

Transversal use case	Scenario	Description
[Dataspace] Data interoperability Data exchange through HEDGE-IoT Dataspace.	Sc1. Use of the dataspace by a data producer	It describes how a data producer makes its datasets or services available within the dataspace.
	Sc2. Use of the dataspace by a data customer	It describes how a data customer interacts with the dataspace to discover and access data shared by other parties.
	Sc3. Metadata Discovery and Planning	It describes how an organisation can discover which datasets or services are available, including their conditions of use, data formats, and applied semantic vocabularies.
	Sc4. Federated Service Chaining	It describes how a software component (e.g., an orchestrator or optimisation engine) uses the dataspace to access services or modules provided by other partners, in a dynamic and composable way.
[Computational Orchestration]	Sc1. Energy services orchestrations at edge geographic redundancy	It describes how the dynamic allocation of computational resources is done for energy services demand across edge, fog, and cloud

Co-funded by the European Union

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

Computational interoperability Automate coordination, management, and execution of HEDGE-IoT computing tasks across the distributed systems/services and cloud environment.		layers to maintain efficiency and responsiveness.
	Sc2. Federated AI services (hyperparameter tuning)	It describes how the minimisation of data transfer overhead in federated AI services is done by efficiently managing and optimizing the hyperparameters of the learning process.
	Sc3. Energy application rolling-up at Edge	It describes how the efficient update of energy services from cloud to edge avoiding execution disruption is ensured.
[App Store] Functional interoperability Use of the App Store as part of HEDGE-IoT.	Sc1. Publish a service/sub-service in the App Store	It describes how a service provider can publish a new service/sub-service in the HEDGE-IoT environment.
	Sc2. Reuse/Access a service/sub-service in the App Store	It describes how discovery mechanism for developers or users works to find available reusable services.
	Sc3. Interchangeable common services/sub-services (among pilots)	It describes how different providers publishing "functionally equivalent" services following the same data schemas can interchange services.

3.2.1 Transversal use cases summary and updates

TUC 1 - DATA INTEROPERABILITY DATA EXCHANGE THROUGH HEDGE-IOT DATASPACE [DATASPACE – DATA INTEROPERABILITY]

Short description: This use case describes how partners in the HEDGE-IoT project can exchange data across organisational boundaries using a shared dataspace infrastructure based on the Eclipse Dataspace Connector (EDC). A data provider exposes metadata and access policies for its resources, while a data consumer discovers and retrieves data through secure, policy-compliant mechanisms. The interaction ensures data sovereignty, access control, and interoperability. This pattern can be reused across different pilots and domains to enable secure and standardised data flows.

Changes: No changes implemented since D2.3.

TUC 2 – AUTOMATE COORDINATION, MANAGEMENT, AND EXECUTION OF HEDGE-IOT COMPUTING TASKS ACROSS THE DISTRIBUTED SYSTEMS/SERVICES AND CLOUD ENVIRONMENT [COMPUTATIONAL ORCHESTRATION - COMPUTATIONAL INTEROPERABILITY]

Short description: The use case focuses on enabling the automated coordination, management, and execution of computational tasks through a computational orchestrator. The orchestrator leverages swarm-based algorithms to optimise resource usage, ensuring both computational and communication efficiency. It integrates with non-AI Energy Services to manage deployment, coordination, and resource allocation, and with AI Federated Services

Co-funded by the European Union

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

to support hyperparameter tuning and training optimisation. Additionally, it enables services roll-up at the edge, allowing automated updates and deployment of new versions. Integration with the Eclipse Dataspace Connector ensures compliance with Dataspace standards and secure, interoperable data and service exchange.

Changes: Minor updates on the actor table and the KPIs.

TUC 3 – USE OF THE APP STORE AS PART OF HEDGE-IOT [APP STORE – FUNCTIONAL INTEROPERABILITY]

Short description: The HEDGE-IoT App Store is a repository for Software Applications that operates at least in one Dataspace configuration. Apps include/represent services/microservices enabling quick discovery, sharing, and reuse across different pilots and domains. It allows service owners to publish new service functionalities, while developers or other system components can request and acquire access. By ensuring semantic and technical interoperability, the App Store accelerates solution development and deployment, fosters collaboration, and ensures consistent service quality within the HEDGE-IoT framework. In line with the new Dataspace protocol (version >2), the App Store also promotes the possibility for dataspace compliant connectors to search for and adopt other versions of control and data planes available.

Changes: The TUC diagrams were enhanced to focus on clarity, aesthetics and alignment with the latest implementation activities. Also they are now more consistent with IDSA principles.

3.2.2 Pilot-Specific TUC Implementations

This section presents how the TUCs are covered by the different pilots of the project. A first version of the table was drafted in D2.3. Table 11 presents the latest available version aligned with pilot activities and implementations.

TABLE 11 – PILOTS' IMPLEMENTATIONS OF TRANSVERSAL USE CASES

Transversal use case	Scenario	FI	GR	IT	NL	PT	SI
[Dataspace] Data interoperability Data exchange through HEDGE-IoT Dataspace.	Sc1. Use of the dataspace by a data producer	✓	✓	✓	✓	✓	✓
	Sc2. Use of the dataspace by a data customer	✓	✓	✓	✓	✓	✓
	Sc3. Metadata Discovery and Planning		✓	✓	✓	✓	✓
	Sc4: Federated Service Chaining		✓	✓	✓		
[Computational Orchestration]	Sc1. Energy services orchestrations at edge geographic redundancy		✓ (Data preprocessing techniques, model inference on edge nodes)	✓			

Co-funded by the European Union

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

Computational interoperability Automate coordination, management, and execution of HEDGE-IoT computing tasks across the distributed systems/services and cloud environment.	Sc2. Federated AI services (hyperparameter tuning)		✓ (Model training, inference and weight updates between edge and cloud)			✓	
	Sc3. Energy application rolling-up at Edge	✓					
[App Store] Functional interoperability Use of the App Store as part of HEDGE-IoT.	Sc1. Publish a service/sub-service in the App Store	✓	✓ (mobile application, selection of services)	✓	✓	✓	✓
	Sc2. Reuse/Access a service/sub-service in the App Store			✓	✓	✓	✓
	Sc3. Interchangeable common services/sub-services (among pilots)		✓ (demand forecasting)		✓		

3.3. FUNCTIONAL REQUIREMENTS CATALOGUE

This section presents the final general HEDGE-IoT functional requirements, which were extracted from BUCs and SUCs. More information could be found in each BUCs, SUCs, and TUCs which present the use case specific functional and non-functional requirements.

A first version of the functional requirements was present in D2.2 and D2.3, defining the technical capabilities required for HEDGE-IoT. This first version was extracted from the SUCs provided by the pilots using the IEC 62559-2 template. This section aims to update and refine this list based on the progress of the project and pilots.

The HEDGE-IoT functional capabilities are split into 9 categories:

- Data management,
- Interoperability and data exchanges,
- Services management,
- User interfaces,
- Optimisation and forecasting,
- Flexibility management,
- Grid monitoring and control,
- Artificial intelligence, and
- Main external data.

Co-funded by the European Union

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

3.3.1 Data Management (DM)

- **FR-DM-01 | IoT data collection:** Collect and monitor real-time data from IoT devices, IEDs, and energy nodes.
- **FR-DM-02 | Data discoverability:** Data customers of the HEDGE-IoT Dataspace should be able to explore the Dataspace catalogue and the possible data from data producers.
- **FR-DM-03 | Computational orchestration:** Coordinating distributed computational tasks for energy services across edge-to-cloud systems, with goals of ensuring responsiveness and cost-effective data exchange, including minimised latency and bandwidth usage.
- **FR-DM-04 | Real-time data processing and aggregation:** Process and aggregate real-time data (e.g., DTR, DLR, DER). The data could be in the cloud or at the edge level, depending on each pilot.
- **FR-DM-05 | Data storage and access:** Store and access data at different levels of the project architecture (e.g., pilot, orchestrator, service catalogue, app store).
- **FR-DM-06 | Data validation and quality check:** Ensure data accuracy with quality checks.

3.3.2 Interoperability (IOP) and data exchanges

- **FR-IOP-01 | HEDGE-IoT middleware data exchange:** Ensure data exchange through HEDGE-IoT components using REST API libraries.
- **FR-IOP-02 | Interoperability among systems:** Ensure interoperability for communication among systems and grid components with a semantic interoperability layer.
- **FR-IOP-03 | Dataspace, dataspace connectors and dataspace protocol:** A shared dataspace infrastructure to facilitate secure, standardised, and scalable data exchange across the various components, stakeholders, and pilot sites involved in the HEDGE-IoT project. Specific connectors will be required for data producers and data customers.
- **FR-IOP-04 | Dataspace catalogue:** Provide a catalogue to make available data discoverable for data customers.
- **FR-IOP-05 | Pilot data exchanges:** A means of pilot data exchanges, such as a gateway and REST API.
- **FR-IOP-06 | Semantic enablers library:** Facilitating pilot semantic interoperability through enablers to cope with: integration of heterogeneous data sources, system and protocol fragmentation, data meaning and context, semantic ambiguity, and cross-domain interoperability.

3.3.3 Services management (SRV)

Co-funded by the European Union

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

- **FR-SRV-01 | App store for energy services:** A centralised repository for services/microservices enabling quick discovery, sharing, and reuse across different pilots and domains. It allows service owners to publish new functionalities, while developers or other system components can access. By ensuring semantic and technical interoperability, the App Store accelerates solution development and deployment, fosters collaboration, and ensures consistent service quality within the HEDGE-IoT framework.

3.3.4 User Interfaces (UI)

- **FR-UI-01 | User interfaces:** Provide user interfaces for DSOs, producers, consumers, flexibility operators, energy community users and operators.
- **FR-UI-02 | User interface configuration:** Configuration of user interfaces, such as preferences and parameter settings.
- **FR-UI-03 | Alerting, reporting & visualisation:** Generate security alerts, notifications, reports, and visualise relevant information.
- **FR-UI-04 | Dynamic tariffs interface:** Display dynamic tariffs for customer engagement and flexibility pricing.
- **FR-UI-05 | User registration and contract information:** Implement a user registration system that securely captures and stores user details and contract information, allowing for account creation, verification, and management.

3.3.5 Optimisation and Forecasting (OF)

- **FR-OF-01 | Optimisation:** Manage and optimise consumption, flexibility, congestion and energy management system (EMS) with real-time adjustments.
- **FR-OF-02 | Forecasting:** Predict production, consumption, grid limits, demand, weather, PV production, etc.
- **FR-OF-03 | Anomaly/fault detection and prediction:** Identify and forecast faults or anomalies in the grid.
- **FR-OF-04 | Anomaly/fault assessment and resolution:** Analyse and make decisions to solve or minimise anomaly/fault quickly to maintain operational continuity.
- **FR-OF-05 | Congestion prediction and management planning:** Anticipate and prevent grid congestion.
- **FR-OF-06 | Performance analysis:** Analyse system performance (e.g., photovoltaic (PV) production).
- **FR-OF-07 | Maintenance prediction:** Predict and manage maintenance to ensure grid reliability.

3.3.6 Flexibility Management (FM)

Co-funded by the European Union

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

- FR-FM-01 | **Registration, prequalification and resources enrolment:** Register and prequalify organisation and flexibility resources for integration.
- FR-FM-02 | **Flexibility offer handling:** Send, receive, accept, or reject flexibility offers; check feasibility and propose incentives.
- FR-FM-03 | **Market price tracking and forecasting:** Track and predict energy and flexibility market prices.
- FR-FM-04 | **Activation & planning:** Execute and plan flexibility actions at both grid and market levels.
- FR-FM-05 | **Flexibility estimation:** Estimate and calculate the required flexibility
- FR-FM-06 | **Real flexibility provided calculation:** Calculation of required/agreed flexibility and provided flexibility.
- FR-FM-07 | **Settlement & payments:** Manage flexibility settlements, payments, and penalties.
- FR-FM-08 | **Vulnerable user identification:** Identify users who have major constraints related to their infrastructures' power supply.

3.3.7 Grid Monitoring and Control (GMC)

- FR-GMC-01 | **Grid state monitoring:** Monitor and estimate grid status in real time.
- FR-GMC-02 | **Grid and energy node configuration:** Configure energy nodes and integrate new ones.

3.3.8 Artificial Intelligence (AI)

- FR-AI-01 | **AI trustworthiness:** Provide the assurance that AI systems used are trustable, including aspects like security, reliability, safety, ethics, integrity and accuracy.
- FR-AI-02 | **AI explainability:** Provide transparent AI-driven decisions for trust.
- FR-AI-03 | **AI maintenance**: Provide an AI maintenance system that continuously monitors model **performance** and facilitates updates and further developments.

3.3.9 Main External Data (MED)

- FR-MED-01 | **Grid historical data:** Access grid historical data as one major input to the system.
- FR-MED-02 | **Weather data:** Incorporate weather forecasts to enhance grid operation decisions.
- FR-MED-03 | **Geographic information system (GIS) model:** Access geographic information systems for environmental data.

Co-funded by the European Union

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

3.4. INTEROPERABILITY PROFILES

This section complements the TUCs and SUCs, specifications and architecture of the project by listing the main IOP challenges, IOP points, and related IOP profiles encountered by HEDGE-IoT and its pilots. In addition, it supports the European Commission's understanding of the interoperability issues encountered in energy innovation projects. These elements will contribute as well to outcomes on policy, roadmap, and standardisation. In the long term, the contributions of several European projects to this action, started in the OPEN DEI [84] and Int:net [85] European Coordination and Support Actions (CSA), will give insights into the definition of a smart grid interoperability profile.

3.4.1 Interoperability (IOP) profile principles

An interoperability profile gathers information and guidance that can be used to create interoperable systems.

Following ISO/IEC 19941:2017 (Information technology - Cloud computing - Interoperability and portability) [67] and ISO/IEC 21823-1:2019 (Internet of things (IoT) - Interoperability for internet of things systems - Part 1: Framework) [68], the interoperability between two (or more) interacting systems can be described by the 5-facets (or layers) model illustrated on (Figure 2).

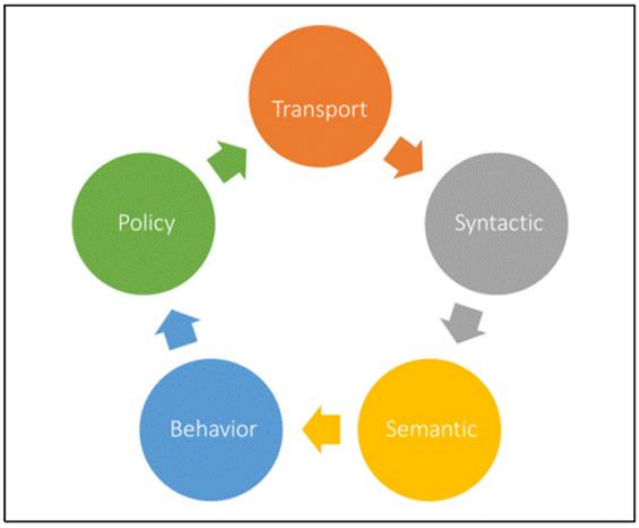

FIGURE 2 – REPRESENTATION OF THE 5 FACETS INTEROPERABILITY MODEL FROM ISO/IEC 19941:2017 [67] AND ISO/IEC 21823-1:2019 [68]

Table 12 describes these 5 interoperability facets and presents a correspondence with the European Interoperability Framework.

TABLE 12 - DESCRIPTION OF THE ISO/IEC 21823-1:2019 INTEROPERABILITY MODEL, AND COMPARISON WITH EIF (EUROPEAN INTEROPERABILITY FRAMEWORK), TABLE ADAPTED FROM ISO/IEC 21823-1:2019

Co-funded by the European Union

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

Facets from ISO/IEC 21823-1:2019	Description	Correspondence with EIF
Transport interoperability	Deals with data delivery	Technical interoperability
Syntactic interoperability	Allows reading the data in a known format and grammar	
Semantic interoperability	Responsible for the meaning, enabling the unambiguous interpretation and understanding of data	Semantic interoperability
Behavioural interoperability	Refers to the way in which business processes, responsibilities and expectations are aligned to achieve commonly agreed and mutually beneficial goals	Organisational interoperability
Policy interoperability	Ensures that organisations operating under different legal frameworks, policies and strategies can work together	Legal interoperability

TABLE 13 – INTEROPERABILITY PROFILE EXAMPLE

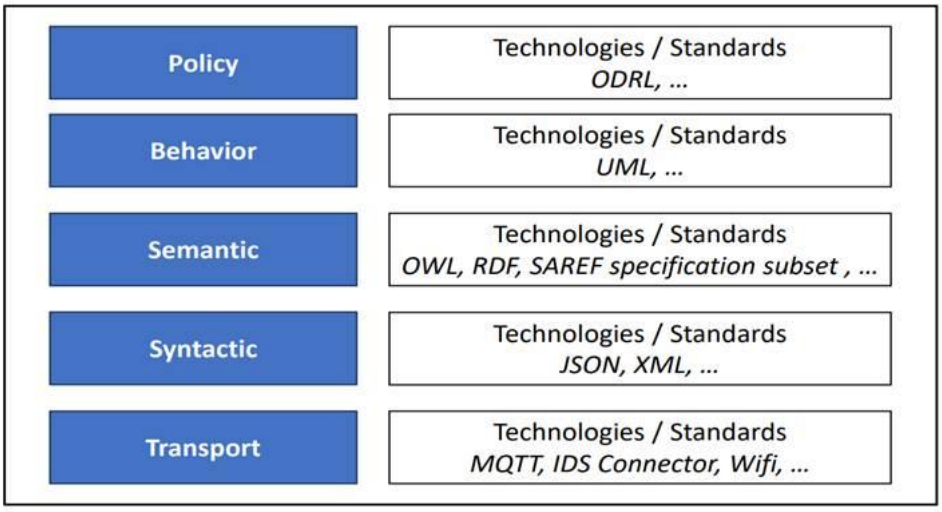

The European Coordination and Support Action OPEN DEI has defined the terms Interoperability point and interoperability case as follows:

- **An interoperability point,** i.e., a location in the overall system where data is exchanged according to an agreed interoperability specification (e.g., where interoperability takes place in a specific context).
- **An interoperability case,** i.e., a documented justification and agreement on an interoperability point (i.e., why interoperability is needed).

Co-funded by the European Union

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

The combination and availability of these two elements lead to the publication of an interoperability profile. As already mentioned, an interoperability profile gathers the information and guidance that can be used to create interoperable systems. It can be considered as the specification for the implementation of the interaction model taking place at an interoperability point.

Additional key notions:

- **An interoperability scenario**, i.e., a description of how interoperability is achieved at a specific interoperability point
- **An interoperability profile**, i.e., a documented set of technical specifications, standards, protocols, and constraints that define the requirements and capabilities needed to ensure interoperability for a defined interoperability scenario.

3.4.2 Interoperability Challenges

The energy ecosystem and its stakeholders continue to face interoperability challenges, and this is also the case for HEDGE-IoT. Some of them are listed below as generic IOP challenges:

- **Integration of heterogeneous data sources:** The diversity of data sources introduces significant variability in data formats, quality, and granularity, creating challenges for consistent interpretation and integration.
- **System and protocol fragmentation:** The use of legacy systems, diverse protocols, proprietary platforms, and non-standardised APIs create critical barriers to seamless data exchange and real-time coordination across the energy ecosystem.
- **Data meaning and context, semantic ambiguity:** Usage of the same terms but with different definitions or context
- **Governance and policy differences:** The variation in data sharing policies, licensing, and legal frameworks create barriers to data sharing, stakeholders' collaboration, and scalable interoperability.
- **Scalability and performance bottlenecks:** large-scale and real-time energy services demand high-performance infrastructure (e.g., AI-driven infrastructure) and an optimised interoperability to avoid latency and data overload that can reduce performance, limit scalability, or crash a system. This is particularly relevant for AI-driven applications.
- **Cross-domain interoperability:** Enabling seamless interaction and data exchange across different application domains or traditionally siloed energy domains (e.g., grids, markets, prosumers, mobility) remains a challenge.
- **Interoperability challenges linked to AI-based systems | AI-based continuous evolution and learning issue:** Issues with the continuous evolution and learning of AI-based systems that impact their behavior, and stability between interoperable systems.

The RA of HEDGE-IoT was designed to eliminate or reduce these IOP challenges. To this end, the choices made for the reference architecture to cope with these challenges are:

Co-funded by the European Union

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

- **Adhere to the Dataspace paradigm**, especially with the selection and the implementation of the Eclipse Dataspace Component (EDC) connector to enable secure, sovereign, and interoperable data sharing across heterogeneous stakeholders, while ensuring compliance with FAIR principles and EU data governance frameworks (e.g., GAIA-X, GDPR). The dataspace allows data interoperability.
- **To establish a specific reference architecture layer for semantic interoperability, including a semantic enablers library**, and using mainly SAREF ontology, CIM, ensuring consistent data interpretation across systems and stakeholders.
- **To implement a swarm-based computation orchestration framework** for the coordination, management, and execution of energy services across the computational continuum. Built on KubeEdge, the framework extends Kubernetes capabilities to edge environments, incorporating swarm-based heuristics to optimise resource allocation. It ensures a streamline, homogenous and efficient cloud-to-edge computational effort (e.g., energy services, AI training).

3.4.3 Selected Interoperability Points

The interoperability profiles for the HEDGE-IoT architecture were selected based on the TUCs, the main HEDGE-IoT API. Additionally, interoperability challenges were collected from project participants and especially pilots, to validate this selection. A fourth interoperability point was selected based on the concerns of the partners related to the data exchange with the flexibility market.

TABLE 14 - MAIN INTEROPERABILITY POINTS IDENTIFIED BASED ON IOP CHALLENGES

Targeted interoperability points	Related Transversal UCs	Description
Dataspace connector interface [Data interoperability]	TUC 1 – Data exchange through HEDGE-IoT Dataspace	Interface between the pilots and the Dataspace Connector through its API.
Computational orchestrator interface [Computational interoperability]	TUC 2 - Automate coordination, management, and execution of HEDGE-IoT computing tasks across the distributed systems/services and cloud environment	Interface between the pilots and distributed resources with the Computational Orchestrator through its API.
App Store interface [Functional interoperability]	TUC 3 - Use of the App Store as part of HEDGE-IoT	Interface between the pilots and the App Store through its API.

Co-funded by the European Union

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

DSO-aggregator Flexibility market interface	/	Interface between the DSO/pilot and the flexibility market.

Co-funded by the European Union

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

3.4.4 Interoperability profiles of the selected IOP points

This section details for each selected interoperability point the interoperability scenario, case, and profile. For the "DSO-aggregator Flexibility market interface" interoperability point, it also specifies the interoperability profile of each pilot.

DATASPACE CONNECTOR INTERFACE

TABLE 15 - DATASPACE CONNECTOR INTERFACE

Interoperability point	Dataspace connector interface
Interoperability case	This IOP point enables data exchange throughout the HEDGE-IoT architecture across organisational boundaries (e.g., data producer, data customer). Directly link to TUC 1 - Data interoperability Data exchange through HEDGE-IoT Dataspace [Dataspace – Data interoperability].
Interoperability scenario	The dataspace connector supports data exchange across organisational boundaries using a shared dataspace infrastructure based on the Eclipse Dataspace Connector (EDC) and the Dataspace Protocol (DSP). A data provider exposes metadata and access policies for its resources, while a data consumer discovers and retrieves data through secure, policy-compliant mechanisms. The interaction ensures data sovereignty, access control, and interoperability. This pattern can be reused across different pilots and domains to enable secure and standardised data flows. Key interoperability challenges: • Governance and policy differences • Data meaning and context, semantic ambiguity
HEDGE-IoT Interoperability profile	**Transport facet:** • HTTP/HTTPS (REST-based communication) is the primary protocol used by EDC • Additionally, the Dataspace Protocol (DSP) operates over HTTP for data exchange and negotiation processes. **Syntactic facet:** • JSON is the main format used in EDC APIs

Co-funded by the European Union

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

	• In some cases, JSON-LD may also be used, especially when linked data representations are required. API structures are typically defined via OpenAPI/Swagger specifications. **Semantic facet:** The semantic layer is not enforced by EDC itself but depends on the data models adopted by the participants. HEDGE-IoT especially uses SAREF when possible, for the semantic facet. Relevant standards that can be used include: • RDF / OWL • NGSI-LD • DCAT • Asset Administration Shell (AAS) • Domain-specific models such as SAREF, IEC 61850, OpenADR **Behavior facet:** The behavioral aspect is mainly defined by the interaction patterns and workflows implemented by EDC: • Compliance with IDSA principles and IDS-RAM • Alignment with the Dataspace Protocol (DSP), which defines negotiation, contract agreement, and data exchange flows • Control Plane / Data Plane architecture of EDC • Optional modelling support via BPMN/UML where workflows are documented • PWI JTC1-SC41-8 – Behavioural and policy interoperability • ISO/IEC PWI 25850 - Information technology — Cloud computing — Use cases for dataspace **Policy facet:** Policy interoperability is ensured through: • IDSA Rulebook (data sovereignty and usage control principles) • Gaia-X compliance framework • EDC policy framework (contract definitions, access policies, and usage control mechanisms) Please note that some elements (e.g., domain-specific semantic models) may vary depending on the specific use cases and pilots involved in the project.

Co-funded by the European Union

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

COMPUTATIONAL ORCHESTRATOR INTERFACE

TABLE 16 - COMPUTATIONAL ORCHESTRATOR INTERFACE

Interoperability point	Computational orchestrator interface
Interoperability case	This IOP point enables data exchange between pilots and distributed computational resources with the HEDGE-IoT Orchestrator. Directly link to TUC 2 - Automate coordination, management, and execution of HEDGE-IoT computing tasks across the distributed systems/services and cloud environment [Computational Orchestration - Computational interoperability].
Interoperability scenario	The orchestrator leverages swarm-based algorithms to optimise resource usage, ensuring both computational and communication efficiency. It integrates with non-AI Energy Services to manage deployment, coordination, and resource allocation, and with AI Federated Services to support hyperparameter tuning and training optimisation. Additionally, it enables services roll-up at the edge, allowing automated updates and deployment of new versions. Integration with the Eclipse Dataspace Connector ensures compliance with Dataspace standards and secure, interoperable data and service exchange.
Interoperability profile	**Transport facet:** • HTTP REST – communication with energy services • AMQP – grid events **Syntactic facet:** • JSON – data, information and events • JSON-LD – Dataspace artifacts • YAML – deployment configuration files for orchestrator components and energy services **Semantic facet:** • SAREF4MIND - Extension of SAREF4SYST for modelling distributed AI processes • DCAT – Dataspace catalogue vocabulary

Co-funded by the European Union

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

	Behavior facet: • compliance with the principles and standards established by the International Dataspaces Association (IDSA) • aligns with the Dataspace Protocol (DSP), a specification that defines the handling of transactions in data ecosystems to foster interoperability and trust among stakeholders. • Swagger API endpoints documentation **Policy facet:** • Dataspace policy and contracts definition using ODRL

APP STORE API INTERFACE

TABLE 17 – APP STORE API INTERFACE

Interoperability point	App Store API interface
Interoperability case	This IOP point enables data exchange between the App Store component and App Store users (i.e., pilot).
Interoperability scenario	The App Store is a repository for Software Applications that operate at least in one Dataspace configuration. Apps include/represent services/microservices enabling quick discovery, sharing, and reuse across different pilots and domains. It allows service owners to publish new service functionalities, while developers or other system components can request and acquire access. The App Store integrates into the IDS ecosystem as one of the building blocks. It interfaces with the IDS Connector and enables Data Apps to be distributed within the Dataspace.
Interoperability profile	**Transport facet:** The App serves as the link to distribute and deploy services that operate side-by-side with a connector. Thus, it supports the link of an IDS connector through its HTTPS interface. Moreover, the App Store also exposes its own API over HTTPS transport. **Syntactic facet:** The App contains service representations as JSON-LD asset representations in the data space. Albeit not implemented, as no use case for the use of the App Store

Co-funded by the European Union

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

	requires it, the JSON-LD schemas could be exported as RDF representations by including an internal model converter. **Semantic facet:** The metadata supporting the service catalogue extracts from the JSON-LD schema representations the corresponding ontological annotations. Some represent SAREF concepts. **Behavior facet:** In line with the IDS RAM, the App store exposes its RESTfull API via a OpenAPI/Swagger documentation system for its control API. The App store opted for a direct integration with one IDS connector, instead of directly adopting the DSP. Nonetheless, the semantic representations of assets are kept. **Policy facet:** The App Store registers new services as dataspace assets. Thus, when adding a new asset, the corresponding data policies are also mapped in ODRL. Beyond the IDS RAM, this is kept in line with the EU's Data Act, where the data owner is the role accountable for setting the sharing permissions of its data. GDPR is applied to user accounts only. Thus, exposure to GDPR related risks is limited.

FLEXIBILITY MARKET INTERFACE

TABLE 18 – FLEXIBILITY MARKET INTERFACE

Interoperability point		**Flexibility market interface**
Interoperability case		This IOP point enables data exchange between DSO/pilot and their related flexibility market for interactions like flexibility request, bids, and activations.
Interoperability scenario		Because this interface is not standardised across pilots, the IOP profiles below are described separately by pilot and illustrate different implementation patterns.
Interoperability profile	Finnish pilot	**Transport facet:** Over public internet using HTTP **Syntactic facet:** JSON **Semantic facet:** Nodes specific (custom semantic model used by Nodes) **Behavior facet:** Nodes specific (each local flexibility market (LFM), has its own way of interacting with flex buyers and sellers)

Co-funded by the European Union

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

		Policy facet: Clean energy package (because the Nodes allows DSOs to access flex from aggregators something that is encouraged by clean energy package)
	Greek pilot	**Transport facet:** REST API (HTTP(S)) **Syntactic facet:** JSON **Semantic facet:** SAREF **Behavior facet:** Hybrid model combining request-based data submission with event/time-driven processing, supporting the flexibility market lifecycle (bids, auctions, results) **Policy facet:** EU/National regulation (e.g., Clean Energy Package, Network Codes, GDPR)
	Italian pilot	**Transport facet:** • Kafka (any version) • HTTPS **Syntactic facet:** • JSON • XML **Semantic facet:** The semantic layer is defined by the market specifications and data models adopted by GME. As the Italian flexibility market relies on a single market operator, no cross-platform semantic interoperability based on generic standards (e.g. IEC CIM, ETSI SAREF, OpenADR) is currently required for the pilot **Behavior facet:** DSO–market interaction workflows and mechanisms are defined in dedicated project documentation, describing roles, processes and information exchanges across the flexibility market lifecycle. **Policy facet:** The Italian pilot is compliant with Directive (EU) 2019/944 (implemented by D.lgs 210/2021) and with the New Electricity Market Design Regulation (EU) 2024/1711, implemented in Italy by D.lgs No. 3 of 24/01/2026. RomeFlex is also aligned with the current draft status of the Network Code on Demand Connection and Congestion Regulation (NCDCR), with which the project is currently compliant.

Co-funded by the European Union

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

	Dutch pilot	**Transport facet:** • Ethernet • REST API (HTTP(S)) **Syntactic facet:** • JSON • XML **Semantic facet:** • QUDT ontology • SAREF4ENER ontology • SAREF ontology • Time ontology • NOTE: we do not interact with the flexibility market, we only get day-ahead prices via the ENTSO-E transparency platform, so we do not use any specific standard for energy markets (SAREF/4ENER concepts are sufficient for us) **Behavior facet:** None (we use UML sequence diagrams to guide the behavior, as we specified in our SUC-NL-01.04 Flexibility Alignment) **Policy facet:** None (as mentioned above, for the Dutch pilot we only need day ahead prices that we get via the ENTSO-E transparency platform)
	Portuguese pilot	**Transport facet:** SOAP over HTTP **Syntactic facet:** XML **Semantic facet:** • SAREF • IEC CIM 62325-2 **Behavior facet:** The most important documents are the bid structuring documents, that cover the possible data fields within the payloads. Namely, the bid structure, the bidding periods/intervals, and the units of measure therein. **Policy facet:** ENTSO-E and TSO/DSO guidelines
	Slovenian pilot	Not in scope

Co-funded by the European Union

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

4 HEDGE-IOT REFERENCE ARCHITECTURE — FINAL VERSION

This chapter presents the final HEDGE-IoT Reference Architecture. It translates the methodological principles, functional specifications, transversal use cases, interoperability profiles and implementation feedback described in the previous chapters into a coherent architecture model. The chapter first defines the core vocabulary and IoT-edge node concept, then introduces the dataspace framework, maps the architecture to SGAM and BRIDGE DERA, and finally presents the final four-layer architecture and its 4+1 architectural interpretation. The objective is to provide a stable, standards-aware and implementation-grounded reference model that can guide pilot integration, technical validation and future replication.

4.1. ARCHITECTURE VOCABULARY AND CONCEPTS

This section consolidates the key vocabulary required to interpret the final Reference Architecture. It focuses on terms that are used directly in the final architecture model, the SGAM and BRIDGE DERA mappings, and the component-to-requirements traceability matrix.

TABLE 19 - HEDGE-IOT REFERENCE ARCHITECTURE VOCABULARY

Term	Definition
Reference Architecture	A reusable architectural template that organises the main layers, components, roles and interactions of the HEDGE-IoT ecosystem.
Architecture Description	The documented representation of the architecture, including its views, concepts, mappings, rationale and traceability links.
Actor	A human, organisation, system or service that interacts with the HEDGE-IoT architecture by providing data, consuming data, offering services, using services or performing operational control.
Component	A modular architectural building block that provides a defined function and exposes interfaces to other parts of the system.
Platform	An integrated technical environment that hosts, manages or exposes data, services, applications or operational capabilities.
Physical Layer	The architecture layer represents pilot infrastructures, field assets, IoT devices, edge nodes, local platforms, data sources and digital twins.
IoT-Edge Node	A local physical and computational entry point where data is collected, processed or exposed to the wider HEDGE-IoT ecosystem.
IoT-Edge Service	A service executed close to physical assets or data sources to support local processing, low-latency analytics, monitoring or control-related functions.
Local Platform	A pilot-specific platform, such as SCADA, EMS, BEMS or another operational system, that hosts or exposes local energy data and control capabilities.
Digital Twin	A digital representation of a physical asset, system or operational context used for monitoring, forecasting, optimisation or decision support.

Co-funded by the European Union

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

Dataspace Layer	The architecture layer that enables governed, sovereign and federated data and service exchange across organisational boundaries.
Dataspace	A decentralised data-sharing environment in which participants exchange data and services under agreed governance, identity, policy and contract rules.
Eclipse Dataspace Components (EDC)	The open-source framework used in HEDGE-IoT to support dataspace-based, sovereign and policy-driven data exchange.
EDC Connector	The EDC component that enables a participant to publish, discover, negotiate and exchange data or services within the dataspace.
Dataspace Protocol	The protocol used to support catalogue access, contract negotiation and controlled data exchange between dataspace participants.
Data Provider	An actor or system that makes data assets available through the dataspace under defined access and usage conditions.
Data Consumer	An actor or system that discovers, negotiates access to and uses data assets made available through the dataspace.
Federated Data Catalogue	A catalogue that exposes metadata about distributed data assets without centralising the underlying data.
Open Service Catalogue	A catalogue that maintains metadata about reusable services and supports their discovery across the HEDGE-IoT ecosystem.
App Store	The architectural component that supports the publication, discovery, access and reuse of applications, services and sub-services.
Orchestrator	The component responsible for coordinating workloads, services and computational resources across edge, fog and cloud environments.
Computational Orchestration	The coordination of distributed computation, service deployment and workload placement across the cloud-edge continuum.
Semantic Interoperability Layer	The architecture layer that ensures data, metadata and service descriptions preserve consistent meaning across heterogeneous systems and pilots.
Semantic Interoperability	The capability to exchange and interpret information consistently through shared ontologies, vocabularies, models, mappings and validation mechanisms.
Semantic Enabler	A tool, service or model that supports semantic alignment, validation, transformation or interpretation of data and metadata.
Knowledge Engine	A semantic component that supports knowledge representation, semantic validation, model alignment and interoperable data interpretation.
Ontology	A formal representation of concepts, relationships and meanings within a domain, used to support semantic interoperability.
Application Layer	The architecture layer where user-facing services, dashboards, analytics, federated learning, optimisation and operational applications are provided.
Fog/Cloud Service	A service executed above the local edge level to support analytics, orchestration, data processing or cross-site digital service delivery.
Federated Learning	A machine-learning approach in which models are trained across distributed nodes without requiring centralisation of raw data.

Co-funded by the European Union

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

Interoperability Profile	A structured specification describing how interoperability is achieved at a selected interoperability point, using transport, syntactic, semantic, behavioural and policy facets.
Interoperability Point	A point in the system where data, services or control information are exchanged between systems, components or organisations.
Policy Enforcement	The application of access, usage, security or governance rules during data and service exchange.
Contract Negotiation	The process through which a data provider and data consumer agree on access and usage conditions before data exchange.
Identity Provider	A component that supports authentication, authorisation and identity management for users, services, systems or organisations.
Functional Requirement	A required system capability derived from BUCs, SUCs or TUCs and used to guide architecture, implementation and validation.
Component-to-Requirements Traceability Matrix	A matrix linking functional requirements to architecture packages or components to show how the requirements baseline is supported.
Energy Stakeholder	An actor in the energy ecosystem, including DSOs, TSOs, aggregators, flexibility service providers, market actors, energy communities, data users and pilot operators.

4.2. IOT-EDGE NODE ARCHITECTURE

4.2.1 IoT-Edge Node Reference Model

IoT-edge nodes constitute the distributed, field-level computational layer of the HEDGE-IoT ecosystem, acting as the interface between physical energy assets and the federated edge–cloud environment. They are responsible for acquiring, processing, and exposing data and services close to the source, enabling low-latency operations and context-aware intelligence. Architecturally, IoT-edge nodes bridge the physical and digital domains by connecting pilot-specific infrastructures—such as sensors, metering systems, and local control platforms—with higher-layer services, including interoperability frameworks, dataspace mechanisms, semantic services, and orchestration components.

The concept is intentionally technology-agnostic and deployment-flexible, allowing each pilot to implement edge nodes according to its operational constraints while adhering to a common architectural pattern. This abstraction ensures interoperability and scalability across heterogeneous environments. An IoT-edge node may encompass a combination of hardware and software elements, including sensing and actuation devices, embedded controllers, edge gateways, local compute resources, digital twins, and edge-native services. Regardless of implementation specifics, all nodes share a core responsibility: to generate or ingest operational data and expose standardised interfaces to the HEDGE-IoT ecosystem.

Functionally, IoT-edge nodes support:

- Real-time observability of field assets
- Local data processing and event-driven analytics
- Secure and governed data/service exposure

Co-funded by the European Union

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

- Bidirectional interaction with orchestration and higher-level intelligence

4.2.2 IoT-Edge Node baseline (mandatory requirements)

The following functional requirements constitute the operational baseline of a HEDGE-IoT architecture IoT-Edge node:

1. **Data Acquisition Capability**
   - Ability to collect or ingest data from physical assets or local systems
   - Support for relevant industrial/IoT protocols (e.g. MQTT, OPC-UA, Modbus where applicable)
2. **Standardised Data Exposure**
   - Provision of data through interoperable interfaces aligned with the HEDGE-IoT dataspace
   - Use of agreed data models and semantic annotations (e.g. aligned with project ontologies)
3. **Secure Communication**
   - Encrypted communication channels (e.g. TLS)
   - Authentication and authorization mechanisms for all external interactions
4. **Interoperability Integration**
   - Compatibility with HEDGE-IoT interoperability components (e.g. connectors, middleware)
   - Ability to register and be discoverable within the ecosystem
5. **Basic Local Processing**
   - Capability to perform minimal data filtering, aggregation, or transformation
   - Support for event-driven data handling
6. **Remote Management Support**
   - Ability to receive configuration updates or software deployments
   - Monitoring of node health and status

4.2.3 IoT-Edge Node optional requirements

These requirements enhance functionality but greatly depend on pilot needs and infrastructure maturity:

1. **Edge Analytics & AI/ML Execution**
   - Deployment of machine learning models or advanced analytics at the edge
   - Support for real-time inference or anomaly detection
2. **Federated Learning Participation**
   - Capability to participate in distributed training workflows
   - Local model training with privacy-preserving data handling
3. **Digital Twin Integration**
   - Hosting or interfacing with digital representations of physical assets
   - Synchronisation between physical and virtual states
4. **Autonomous Decision-Making**
   - Local control logic for automated responses (e.g., demand response actions)
   - Operation under intermittent connectivity

Co-funded by the European Union

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

5. **Data Caching and Buffering**
    - Temporary storage to handle connectivity loss or bandwidth constraints
    - Store-and-forward mechanisms
6. **Service Hosting Environment**
    - Containerised or virtualised runtime (e.g. Docker, Kubernetes at the edge)
    - Support for deploying reusable edge services
7. **Advanced Security Features**
    - Hardware-based security (e.g. TPM, secure enclaves)
    - Intrusion detection or anomaly monitoring
8. **Context Awareness**
    - Integration of environmental or operational context (e.g., weather, grid state)
    - Adaptive behavior based on local conditions

4.2.4 Pilot-Specific IoT-Edge Configurations

Table 20 summarises the main pilot-specific IoT-edge configurations retained in the final architecture narrative.

TABLE 20 – MAIN PILOT-SPECIFIC IOT-EDGE CONFIGURATIONS RETAINED IN THE FINAL ARCHITECTURE NARRATIVE.

PILOT	MAIN EDGE ASSETS	MAIN EDGE/LOCAL FUNCTIONS	ARCHITECTURAL OBJECTIVE LINK
Finnish pilot	IEDs, RTUs, smart meters, edge server, data storage	Local processing of substation data; anomaly detection; fault forecasting	Increase network resilience through local monitoring and intelligence
Greek pilot	Submetering IoT devices, smart meters, LV nodes, PVs, batteries, EVs and chargers	Real-time data gathering; smart building and flexibility modelling at edge and cloud; demand and production forecasting	Support grid management and grid-flexibility calculation
Italian pilot	Grid sensors, IEDs, behind-the-meter assets, weather stations, pilot IoT platform	Collection of metering, sensor, DER and weather data; local exposure of data to upper layers; forecasting of load and production	Support energy community management and congestion-related calculations
Dutch pilot	Metering devices, submetering IoT devices, PVs, batteries, heat pumps, EVs and chargers	Semantic interoperability moved closer to the edge; AI-based building and flexibility modelling; integration with EMS/BMS	Optimise energy use and support flexibility alignment in the business park
Portuguese pilot	Smart meters, commercial-building assets, heat pumps,	Integration of building and community asset data through	Exploit flexibility from buildings and energy communities

Co-funded by the European Union

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

	HVAC, batteries, controllers	controllers; forecasting; optimal dispatch calculations	
Slovenian pilot	DTR and DLR edge devices, power quality meters, weather stations, temperature sensors	Edge-side calculation of DTR and DLR; transfer of outputs to cloud and semantic substation model; use in ML solutions	Maximise asset capacity and improve planning and operation of distribution networks

4.3. DATASPACE FRAMEWORK - ECLIPSE EDC

The Eclipse Dataspace Components (EDC) is an open-source framework (available under Apache 2.0 license) that provides a foundational set of features, both functional and non-functional, for building and customizing dataspace solutions. Managed by the Eclipse Foundation, EDC offers developers a structured architecture, concept, and codebase that they can reuse and extend to ensure seamless interoperability through well-defined APIs. It is built on the principles of the Gaia-X AISBL Trust Framework and the IDSA Dataspace protocol. EDC is geared toward developers who need a standards-based foundation for creating dataspace implementations. They can use EDC to build and customize data-sharing services tailored to their customers' needs. The framework includes several core components essential for dataspace construction: the Connector, Federated Catalog, Identity Hub, Registration Service, and Data Dashboard. Its communication protocol is the Dataspace protocol 2024-01 [117] [121]. EDC is designed to enable decentralised data spaces, adhering to IDSA standards for identity management. It integrates tools such as the Metadata Broker, Dynamic Attribute Provisioning Service (DAPS), and supports identity management using Decentralised Identifiers and Federated Catalogs [118] [119] [120].

4.4. ARCHITECTURE ALIGNMENT WITH KEY FRAMEWORKS

4.4.1 SGAM Mapping of the HEDGE-IoT Reference Architecture

Figure 3 presents the mapping of the HEDGE-IoT RA against the Smart Grid Architecture Model (SGAM). The mapping confirms that the final HEDGE-IoT architecture can be consistently interpreted through the five SGAM interoperability layers, namely Business, Functional, Information, Communication, and Component, while spanning the relevant SGAM zones from Process and Station/Field up to Operation, Enterprise, and Market. This is consistent with the HEDGE-IoT project vision, which explicitly aims to deploy IoT assets across different levels of the energy system, from behind-the-meter up to the TSO level, add intelligence at the edge and cloud layers, and bridge the cloud-edge continuum through interoperable and federated applications.

Co-funded by the European Union

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

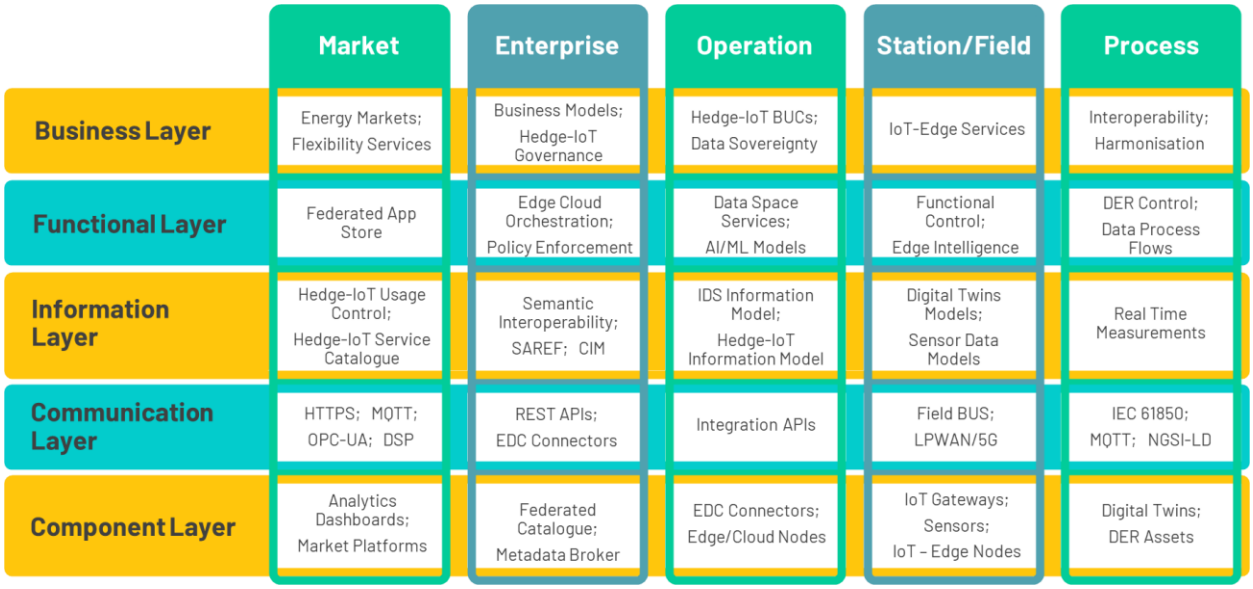

FIGURE 3 - SGAM MAPPING OF THE HEDGE-IOT REFERENCE ARCHITECTURE – FINAL VERSION

In the proposed HEDGE-IoT SGAM view, the horizontal layers describe the different interoperability abstractions of the architecture, while the vertical columns show how these abstractions are realised across the operational zones of the electricity system. In this way, the SGAM mapping demonstrates that HEDGE-IoT is not limited to one system level or one organisational layer but instead provides a structured digital framework connecting process-level assets, field and station intelligence, operational services, enterprise-level data and orchestration, and market-facing flexibility and service interactions.

BUSINESS LAYER

The Business Layer describes the business objectives, governance logic, market interactions, and operational roles supported by HEDGE-IoT across the different SGAM zones. In the Market zone, HEDGE-IoT addresses energy markets and flexibility services, reflecting the project's explicit focus on local flexibility markets, reserve market participation, and value creation through data-driven energy services. In the Enterprise zone, this layer is represented by business models and HEDGE-IoT governance, highlighting the organisational structures required for service operation, platform ownership, and coordinated exploitation of interoperable energy services. In the Operation zone, the business layer is expressed through HEDGE-IoT business use cases and data sovereignty considerations, which connect operational responsibilities with trusted data sharing and policy-aware cooperation. At the Station/Field level, the business layer is reflected through IoT-edge services, while at the Process level it is linked to interoperability and harmonisation objectives.

This mapping reflects the fact that HEDGE-IoT is not merely a technical platform, but a digital energy ecosystem enabler. The project description explicitly frames HEDGE-IoT as a multi-dimensional framework that supports resilience, interoperability, Dataspaces, and stakeholder engagement, while also contributing to market opportunities, business uptake, and the wider Digitalisation of Energy Action Plan. The Business Layer therefore captures the

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

institutional and value-chain logic within which the project's services are created, governed, and adopted

FUNCTIONAL LAYER

The Functional Layer captures the main digital and operational functions performed by HEDGE-IoT across the system. In the Market zone, this layer is represented by the federated service marketplace, which corresponds to the discovery, exposure, and use of interoperable services across the ecosystem. In the Enterprise zone, it includes edge-cloud orchestration and policy enforcement, reflecting the project's strong emphasis on computational orchestration, workload offloading, and controlled data and service usage. In the Operation zone, the functional layer includes Dataspace services and AI/ML models, which provide the operational intelligence required for forecasting, condition assessment, congestion management, and flexibility activation. At the Station/Field level, the layer is realised through functional control and edge intelligence, while at the Process level it is associated with DER control and data process flows.

This is directly aligned with the HEDGE-IoT work plan. WP3 defines the AI/ML tools and services for edge and cloud levels, including federated learning and computational orchestration, while WP4 establishes the Open Services Catalogue, App Store, interoperability middleware, and system integration framework. As a result, the Functional Layer in the SGAM mapping acts as the bridge between business needs and deployable technical mechanisms, showing how HEDGE-IoT moves from stakeholder requirements to executable services and coordinated operational logic.

INFORMATION LAYER

The Information Layer describes the information models, semantic structures, and knowledge representations that allow data to be interpreted consistently across systems, actors, and services. In the Market zone, this layer is represented by HEDGE-IoT usage control and the service catalogue, which provide structured metadata and service descriptions for market-facing interactions. In the Enterprise zone, the Information Layer includes semantic interoperability and the core semantic artefacts used by HEDGE-IoT, in particular SAREF and CIM. In the Operation zone, it includes IDS-based information models and the HEDGE-IoT information model. At the Station/Field level, the layer is expressed through digital twin models and sensor data models, while at the Process level it maps to real-time measurements.

This is one of the most important layers in the HEDGE-IoT architecture. The project explicitly builds on semantic interoperability through SAREF, IEC CIM, PowerCIM, and the Knowledge Engine, while WP4.3 defines semantic interoperability as a dedicated activity for ontology-based knowledge federation and decentralised exchange. The Information Layer therefore demonstrates how HEDGE-IoT addresses one of the most persistent barriers in digital energy systems, namely the lack of common semantics across heterogeneous devices, platforms, and market actors. Through this layer, raw data is transformed into interoperable information assets that can be reused consistently across pilots and services.

Co-funded by the European Union

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

COMMUNICATION LAYER

The Communication Layer captures the protocols, APIs, connector interfaces, and communication channels that enable data and service exchange between the different HEDGE-IoT components. In the Market zone, the layer includes protocols such as HTTPS, MQTT, OPC-UA, and the Dataspace protocol. In the Enterprise zone, it is represented by REST APIs and EDC connectors. In the Operation zone, the layer includes integration APIs, while at the Station/Field level it incorporates field bus communication and LPWAN/5G technologies. At the Process level, the mapping refers to protocols such as IEC 61850, MQTT, and NGSI-LD.

This communication view is coherent with the HEDGE-IoT interoperability framework. The project specifies that the Open Data Connector and interoperability middleware must provide secure, seamless, decentralised exchange based on standard interfaces and standardised data models. It also highlights the use of MQTT, AMQP, NGSI-LD, IEC 61850, and related messaging and API technologies as part of the project's interoperability and standardisation pillar. Accordingly, the Communication Layer in the SGAM mapping shows how HEDGE-IoT implements protocol-level interoperability between field infrastructure, operational platforms, enterprise services, and market-level applications.

COMPONENT LAYER

The Component Layer represents the actual physical and software assets that instantiate the HEDGE-IoT architecture. In the Market zone, this includes analytics dashboards and market platforms. In the Enterprise zone, it includes the federated catalogue and metadata broker. In the Operation zone, it includes EDC connectors and edge/cloud nodes. At the Station/Field level, the component layer is realised through IoT gateways, sensors, and IoT-edge nodes. At the Process level, it comprises digital twins and DER assets.

This interpretation is aligned with the HEDGE-IoT deployment logic. The project aims to integrate IoT devices, gateways, local platforms, digital twins, edge services, and cloud services into one coherent framework. It also includes concrete components such as the App Store, Open Service Catalogue, interoperability middleware, IDS connectors, Identity Hub, Semantic Treehouse, PowerCIM, and the Knowledge Engine. The Component Layer therefore provides the physical and platform grounding of the SGAM mapping, showing how the architectural abstractions above are realised in deployable artefacts across real pilot environments.

VERTICAL INTERPRETATION ACROSS SGAM ZONES

Beyond the horizontal layer interpretation, the HEDGE-IoT SGAM mapping also demonstrates clear vertical alignment across zones. In the Process and Station/Field zones, the architecture captures IoT-enabled sensing, device control, real-time measurements, digital twins, and local intelligence. In the Operation zone, it supports data space services, AI/ML models, grid operation functions, and operational interoperability. In the Enterprise zone, it provides semantic interoperability, orchestration, policy enforcement, metadata handling, and governance-supporting platform functions. In the Market zone, it connects these lower-level capabilities to flexibility services, marketplaces, and federated service publication and consumption. This vertical continuity is essential because it shows that HEDGE-IoT can link

Co-funded by the European Union

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

field-level observability and device intelligence with enterprise coordination and market-facing service activation within one coherent digital framework.

4.4.2 BRIDGE DERA Mapping of the HEDGE-IoT Reference Architecture

Figure 4 presents the mapping of the HEDGE-IoT Reference Architecture against the BRIDGE Data Exchange Reference Architecture (DERA). BRIDGE DERA provides a reference model for interoperable energy data exchange and has evolved from earlier BRIDGE Data Management Working Group outputs towards DERA 3.0 and 3.1, with a strong emphasis on data governance, interoperability, semantic alignment, security, and the emergence of European energy data spaces [3] [23] [31] [56]. The HEDGE-IoT mapping shows that the final project architecture is strongly aligned with the BRIDGE layered view, while specialising it for IoT-enabled energy systems, semantic interoperability, trusted data space exchange, and cloud-edge computational orchestration.

This alignment is consistent with the HEDGE-IoT objective of deploying IoT assets across the energy system, from behind-the-meter environments up to DSO and TSO levels, adding intelligence at the edge and cloud layers, and bridging the cloud-edge continuum through federated applications and orchestration mechanisms. The mapping is structured across the five BRIDGE DERA horizontal layers: Business, Function, Information, Communication and Component [23]. In the HEDGE-IoT specialisation, these layers are populated by project-specific building blocks, service groupings, standards, protocols, semantic enablers and field components derived from the final HEDGE-IoT Reference Architecture and the latest available project consolidation.

BUSINESS ACTORS AND ECOSYSTEMS LAYER

The upper horizontal layer corresponds to the Business layer of BRIDGE DERA and captures the regulatory, organisational and stakeholder context within which the HEDGE-IoT framework operates. This layer includes European policy and regulatory frameworks relevant to data exchange, energy digitalisation, cybersecurity, data governance and AI-enabled energy services. These include the Clean Energy Package and Directive (EU) 2019/944 [24], the European Green Deal [76], the European Strategy for Data [77], the Digitalising the Energy System Action Plan [7], eIDAS [78], the AI Act [79], the Data Act [80], GDPR [81], and NIS2[82].

The same layer also includes the principal associations, standardisation bodies and European initiatives shaping the interoperability and data space landscape. These include ENTSO-E and the Harmonised Electricity Market Role Model [26], E.DSO [83], IDSA and IDS-RAM [11] [13], BRIDGE and the BRIDGE Data Management Working Group [31], AIOTI [8], CEN-CENELEC-ETSI and the Smart Grid Architecture Model [4], [32], the Common European Energy Data Space and related blueprint work [59] [84], and Gaia-X [16].

Co-funded by the European Union

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

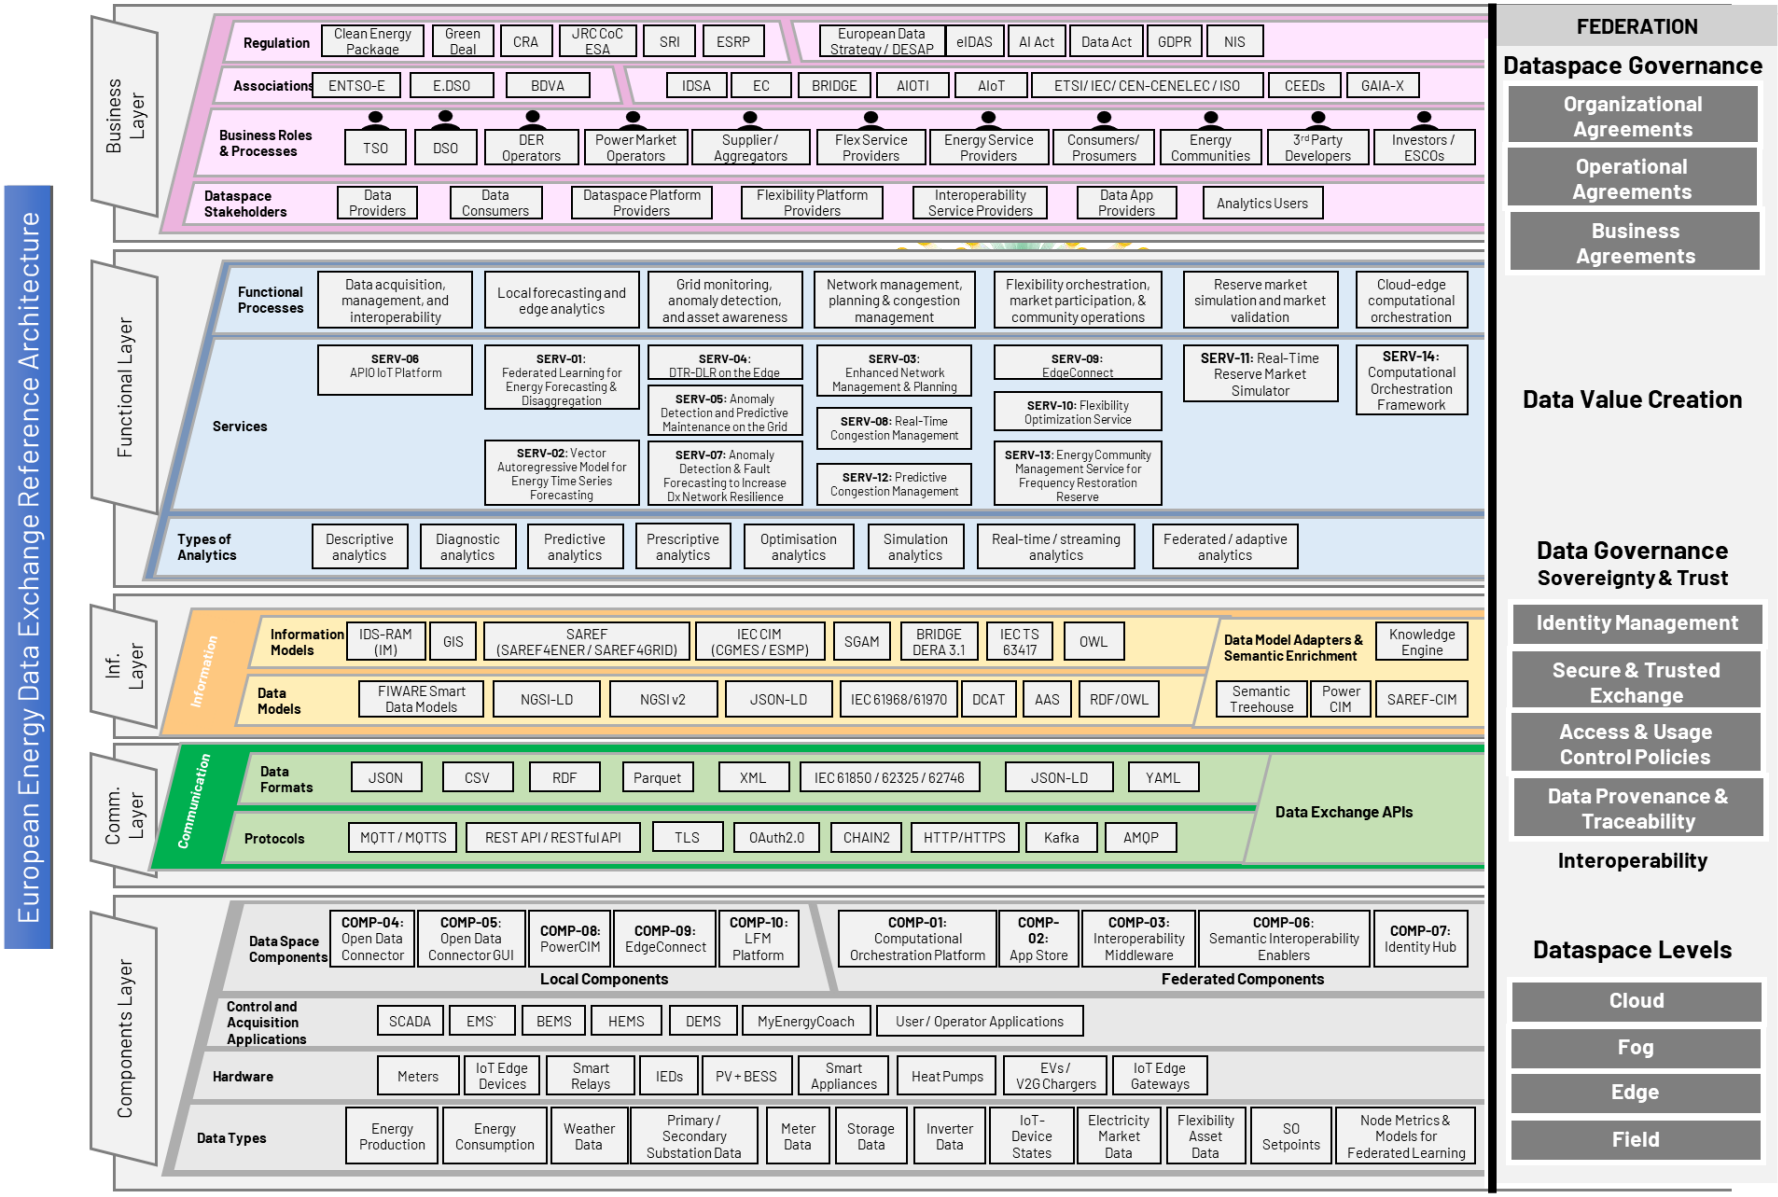

FIGURE 4 - BRIDGE DERA MAPPING OF THE HEDGE-IOT REFERENCE ARCHITECTURE – FINAL VERSION

Co-funded by the European Union

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

At stakeholder level, this layer represents the actors that participate in the HEDGE-IoT ecosystem both as energy-system roles and as data space participants. These include TSOs, DSOs, DER operators, suppliers, aggregators, flexibility service providers, power market operators, energy service providers, consumers, prosumers, third-party developers, investors and ESCOs. In parallel, the mapping identifies data-space-specific roles such as data providers, data consumers, data space platform providers, flexibility platform providers, interoperability service providers, data app providers and analytics users. This dual representation is important because HEDGE-IoT does not treat interoperability as a purely technical matter. Instead, it places data exchange, service publication and service consumption within a governed business ecosystem involving operational, market, technical and regulatory actors.

INNOVATIVE DATA ANALYTICS SERVICES LAYER

The second horizontal layer corresponds to the Function layer of BRIDGE DERA. In the HEDGE-IoT mapping, this layer is represented by the "Innovative Data Analytics Services" block and organises the system according to four internal dimensions: functional processes, service families, analytics types and processing capabilities. It therefore captures both the business-facing operational functions and the digital services that realise them.

At the top of this layer, the functional processes express the main operational domains addressed by HEDGE-IoT. These are: data acquisition, management and interoperability; local forecasting and edge analytics; grid monitoring, anomaly detection and asset awareness; network management, planning and congestion management; flexibility orchestration, market participation and community operations; reserve market simulation and market validation; and cloud-edge computational orchestration. Together, these process areas reflect the intended scope of HEDGE-IoT services across real-time monitoring, forecasting, system operation, flexibility activation, market participation and distributed execution across the cloud-edge continuum [61] [63] [65].

Within these functional processes, the HEDGE-IoT services are grouped into seven service families [61]:

- data acquisition, management and interoperability, represented by SERV-06 APIO IoT Platform, whose core role is to ingest, manage, secure and expose energy data across edge and cloud environments;
- forecasting and disaggregation analytics, including SERV-01 Federated Learning for Energy Forecasting and Disaggregation and SERV-02 Vector Autoregressive Model for Energy Time Series Forecasting;
- grid monitoring, anomaly detection and asset awareness, including SERV-04 DTR-DLR on the Edge, SERV-05 Anomaly Detection and Predictive Maintenance on the Grid, and SERV-07 Anomaly Detection and Fault Forecasting to Increase Distribution Network Resilience;
- network management, planning and congestion management, comprising SERV-03 Enhanced Network Management and Planning, SERV-08 Real-Time Congestion Management, and SERV-12 Predictive Congestion Management;

Co-funded by the European Union

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

- flexibility orchestration, market participation and community operations, including SERV-09 EdgeConnect, SERV-10 Flexibility Optimisation Service, and SERV-13 Energy Community Management Service for Frequency Restoration Reserve;
- reserve market simulation and market validation, represented by SERV-11 Real-Time Reserve Market Simulator;
- cloud-edge computational orchestration, represented by SERV-14 Computational Orchestration Framework.

This layer also identifies the analytics types supported by the HEDGE-IoT services. These include descriptive, diagnostic, predictive, prescriptive and optimisation analytics. In the broader service interpretation, the architecture also accommodates simulation analytics, real-time and streaming analytics, and federated or adaptive analytics. This confirms that the HEDGE-IoT Function layer is not restricted to monitoring but supports a wider range of data-driven reasoning and operational decision support, from condition assessment and forecasting to optimisation, market simulation and federated execution across the cloud-edge continuum.

INFORMATION LAYER

The third horizontal layer maps to the Information layer of BRIDGE DERA. Its purpose is to provide shared meaning, semantic consistency and model alignment across heterogeneous platforms, devices and services. In HEDGE-IoT, this layer is expressed through information models, data models, data-model adaptation and semantic enrichment functions.

The information models included in the mapping comprise IDS-RAM [13], SAREF and its energy-related extensions, including SAREF4ENER and SAREF4GRID [85] [86] [87], IEC CIM, including IEC 61970, IEC 61968, CGMES and ESMP-related profiles [88] [89] [90] [91], SGAM [4], and IEC SRD 63417 on smart energy ontologies [92]. These represent the main conceptual and semantic foundations for cross-platform and cross-stakeholder interoperability.

The data models and semantic technologies include FIWARE Smart Data Models [34], [36], [93], NGSI-LD [94], NGSI v2 [95], JSON-LD [96], IEC 61968/61970 [88] [89], DCAT [97], AAS [98], RDF [99], and OWL [100]. Through these artefacts, the Information layer supports both domain-specific and platform-oriented modelling needs, enabling the transformation of raw device, platform and market data into exchangeable, semantically interpretable information assets.

A central element of this layer is the Data Model Adapters and Semantic Enrichment block, where HEDGE-IoT-specific semantic enablers are positioned. These include the Semantic Treehouse, PowerCIM, the Knowledge Engine and the SAREF-CIM bridge. This reflects the project's explicit semantic-interoperability work, which leverages SAREF-based ontologies, IEC CIM-related models and knowledge-federation mechanisms to support decentralised semantic exchange among AI-IoT edge and cloud nodes [61], [63]. Accordingly, the Information layer is one of the strongest differentiators of HEDGE-IoT compared with a generic data space architecture: it ensures that data is not only exchanged, but also understood consistently across pilots, stakeholders and service contexts.

Co-funded by the European Union

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

COMMUNICATION LAYER

The fourth horizontal layer corresponds to the Communication layer of BRIDGE DERA and captures how data and services are exchanged between HEDGE-IoT components. In the mapping, this layer contains the data formats, protocols and data exchange APIs used to interconnect field devices, platforms, middleware components and higher-level services.

The data formats shown in the mapping include JSON [101], CSV [102], RDF [99], Parquet [103], XML [104], JSON-LD [96], YAML [105], and standards-related payload structures associated with IEC 61850, [106], IEC 62325 [90], [91], and IEC 62746 / OpenADR-related demand response interfaces [107][108]. These formats support the exchange of operational, semantic, market and flexibility-related information across heterogeneous energy platforms.

The protocols and communication mechanisms include MQTT and MQTTS [52], [109], REST APIs and HTTP/HTTPS [110], TLS [109], OAuth 2.0 [111], CHAIN2 [112], Kafka [113], AMQP [114], and data space protocol mechanisms [27][28]. These technologies collectively represent the communication mechanisms through which HEDGE-IoT implements data and service exchange between distributed actors, field systems, middleware components and data space participants.

This Communication layer is closely aligned with the HEDGE-IoT Open Data Connector and interoperability middleware defined in WP4. In particular, the project specifies that the connector must support secure, decentralised end-to-end data exchange, user-friendly interaction through APIs and graphical interfaces, and standardised data models such as SAREF and NGSI-LD [63], [85], [94]. Therefore, the Communication layer in the DERA mapping is not an isolated networking block; it is the operational bridge that connects component-level assets with semantically enriched, policy-aware and service-enabled data space interactions.

DATA SOURCES AND COMPONENTS LAYER

The bottom horizontal layer corresponds to the Component layer of BRIDGE DERA and represents the actual field, platform and software infrastructure from which HEDGE-IoT derives data and on which it deploys services. In the current mapping, this layer is structured into data space components, control and acquisition applications, hardware and data types. It therefore provides the physical and software deployment basis for the upper interoperability and service layers.

The data space components include the IDS Connector, Eclipse Dataspace Connector, App Store, Interoperability Middleware, Open Services Catalogue, Identity Hub, Eclipse GUI, Semantic Treehouse, PowerCIM and the Knowledge Engine [12] [27] [29] [30] [61] [63] [115]. Even though some of these components also contribute to upper-layer functionality, their inclusion in the Component layer reflects the fact that they are instantiated as deployable artefacts within the HEDGE-IoT infrastructure stack.

The layer also includes control and acquisition applications such as SCADA, EMS, BEMS, HEMS, DEMS, GIS and MyEnergyCoach, which act as operational platforms through which data is acquired, processed and made available to services. At hardware level, the mapping includes meters, IoT edge devices, smart relays, IEDs, PV and BESS systems, smart appliances,

Co-funded by the European Union

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

heat pumps, EVs, V2G chargers and IoT edge gateways. This is aligned with the HEDGE-IoT project scope, which aims to deploy IoT assets at multiple levels of the energy system and validate the framework through heterogeneous demonstrators operating across buildings, prosumer assets, substations, distribution grids and transmission-related contexts.

Finally, the data-types row identifies the categories of operational data handled by the architecture, including energy production, energy consumption, weather data, primary and secondary substation data, meter data, storage data, inverter data, IoT-device states, electricity market data, flexibility asset data, system-operator setpoints, and node metrics and models for federated learning. This confirms that the HEDGE-IoT Component layer is not limited to raw device connectivity but is structured around the information assets needed to support grid observability, forecasting, flexibility activation, market participation and cloud-edge orchestration.

Taken together, the five horizontal layers show how the HEDGE-IoT RA specialises BRIDGE DERA into a concrete European energy data space architecture. The resulting mapping demonstrates that HEDGE-IoT is conceptually aligned with BRIDGE DERA and provides an implementation-oriented interpretation tailored to IoT-enabled flexibility, semantic interoperability, sovereign data exchange and cloud-edge intelligence in the European energy system.

4.5. HEDGE-IOT REFERENCE ARCHITECTURE — FINAL VERSION

4.5.1 Architecture Overview and Design Rationale

Figure 5 presents the final HEDGE-IoT Reference Architecture adopted in D2.4. The final model consolidates the project into four major layers: (i) Physical Layer, (ii) Dataspace Layer, (iii) Semantic Interoperability Layer, and (iv) Application Layer. These layers are connected through dataspace connectors and supported by cross-layer governance, orchestration, and service reuse logic.

PHYSICAL LAYER

The Physical Layer represents the pilot infrastructures and the real-world environments in which HEDGE-IoT is deployed. It includes the local platforms, IoT physical services, digital twins, data sources, and edge services that exist within each pilot node. In practical terms, this is the layer where operational grid data, flexibility information, asset status, and environmental measurements originate.

The layer is also the point at which local control-relevant outputs can return to the field. In other words, the HEDGE-IoT architecture is not limited to upward data collection; it also supports the downward application of optimised set-points, orchestration decisions, and service outcomes where the pilot implementation permits such interactions.

In the final architecture, the Physical Layer is deliberately represented as a replicated pilot pattern. This means that the same architectural logic can be instantiated across different pilot sites even though the exact assets differ. Each pilot exposes its data and services to

Co-funded by the European Union

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

the broader HEDGE-IoT ecosystem through a dataspace connector, while computational resources can also be engaged to support distributed processing across sites.

DATASPACE LAYER

Above the Physical Layer sits the Dataspace Layer, which acts as the operational core for secure and sovereign data and computation sharing. This layer transforms multiple pilot environments into a federated Dataspace in which assets remain distributed but become discoverable, governable, and reusable through common mechanisms.

The Dataspace Layer can be read through three functional groupings. The first grouping is the Dataspace core, which includes the federated data catalogue, identity provider, Dataspace services, contract negotiation, connector configuration, policies, open service catalogue, and the Dataspace graphical user interface. Together, these components support participant onboarding, metadata publication, secure access control, contract-based exchange, and day-to-day interaction with the dataspace.

The second grouping is the App Store. In the final architecture, the App Store is not a peripheral add-on, but a central mechanism for the registration, publication, discovery, and reuse of applications and services. Through its GUI, registry functions, metadata handling, and connector-based integration, it enables modular digital capabilities to be exposed and reused across pilots and stakeholders.

The third grouping is the Orchestrator. This is the mechanism through which HEDGE-IoT addresses computation sharing across the cloud-edge continuum. The Orchestrator includes dedicated user interfaces, computational orchestration functions, infrastructure operation and planning capabilities, and application metadata handling. Its role is to plan, deploy, and coordinate workloads across distributed environments so that services can execute where they are most effective from operational, latency, and privacy perspectives.

In this final architecture, the Dataspace Layer therefore does more than exchange data. It provides the environment in which both data assets and executable digital capabilities can be governed, discovered, negotiated, orchestrated, and reused.

SEMANTIC INTEROPERABILITY LAYER

The Semantic Interoperability Layer provides the harmonisation logic required for data and services to remain portable across heterogeneous pilots and platforms. By elevating semantics into a dedicated architectural layer, the final RA makes explicit that interoperability in HEDGE-IoT is not only technical, but also conceptual and model-driven.

The layer includes the SIF Knowledge Engine, Semantic Treehouse, ODC Tester, PowerCIM, ontologies and vocabularies, and standards. It is further supported by ontology design and support, verification and validation, data sharing and orchestration, and model management and integration functions. Collectively, these elements provide the semantic tooling needed to define, align, validate, and operationalise common representations across the project.

A central role of this layer is to harmonise different models and vocabularies - including, for example, PowerCIM and SAREF-oriented perspectives - so that data originating from

Co-funded by the European Union

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

different pilots can be interpreted and reused consistently. It also supports the alignment of metadata, service descriptions, and model-management processes, thereby reducing the semantic fragmentation that would otherwise arise in a multi-pilot, multi-platform architecture.

In practical terms, this layer makes it possible for the same data-driven service or analytics logic to be transferred from one pilot context to another without losing meaning. It is therefore a foundational enabler for reuse, portability, and cross-pilot scaling.

APPLICATION LAYER

The Application Layer is the point at which the architectural capabilities of HEDGE-IoT become accessible to end users and operational stakeholders. It includes identity management, role-based access, visualisation and analytics, federated learning, data management, fog/cloud services, and interfaces to proprietary tools and operational platforms.

This layer is important because it expresses HEDGE-IoT as a usable digital environment rather than as a collection of isolated technical components. Operators, market actors, service providers, and third-party users experience the system through the services made available at this level. The presence of role-based access and identity management also ensures that this usability remains compatible with the governance and trust requirements established in the lower layers.

The inclusion of federated learning and fog/cloud services in the Application Layer highlights that advanced digital functions are expected to operate across distributed computational settings rather than only in central cloud platforms. In this way, the final architecture supports both user-facing digital services and advanced data-driven capabilities that depend on coordinated use of edge and cloud resources.

CROSS-LAYER INTERPRETATION

Read end-to-end, the final RA expresses a complete digital pathway from pilot infrastructure to user-facing services. Data originates in the Physical Layer, becomes discoverable and governable in the Dataspace Layer, is semantically aligned and validated in the Semantic Interoperability Layer, and is then consumed by applications, dashboards, digital tools, or AI-driven services in the Application Layer.

The same pathway also supports reverse operational influence. Services discovered or orchestrated in the upper layers can be deployed closer to the edge, and results produced by higher-level analytics or orchestration mechanisms can inform decisions and actions in the pilot environments. The architecture is therefore both data-driven and service-driven, and it supports a two-way relationship between operational infrastructures and digital intelligence.

Overall, the final HEDGE-IoT architecture can be summarised as a standards-aligned, semantics-driven, dataspace-centric edge-cloud stack that allows pilots to share data with sovereignty, orchestrate computation across distributed sites, and reuse digital services

Co-funded by the European Union

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

through an App Store and open service catalogue. This is one of the core architectural results of the project.

MAIN DIFFERENCES COMPARED WITH THE D2.3 REFERENCE ARCHITECTURE

Compared with the first release presented in D2.3, the D2.4 final architecture is more explicit, more layered, and more operationally structured. The first release evolved from a conceptual three-layer model to an intermediate implementation-oriented model and then to a first version. The D2.4 architecture consolidates this evolution into a clearer four-layer structure in which the Physical Layer, Dataspace Layer, Semantic Interoperability Layer, and Application Layer are more distinctly separated.

A first major difference is the stronger centrality of the dataspace. While D2.3 already introduced connector-based exchange and catalogue logic, the D2.4 version gives a more mature and visible role to the Dataspace Layer as the main federation environment for both data and services. The App Store and the Orchestrator are no longer supporting details but explicit architectural pillars.

A second major difference is the elevation of semantic interoperability into a dedicated layer. In D2.3, semantics were already important, but in D2.4 the semantic tooling is more clearly isolated and expanded through the explicit inclusion of knowledge-engine, ontology, validation, and model-management functions. This makes the portability of data and services across pilots a direct architectural concern.

A third difference is the clearer representation of the pilot side. D2.4 makes the Physical Layer and pilot-specific nodes more visible, emphasising that the architecture is grounded in real infrastructures and distributed resources. This is complemented by a more explicit treatment of computational resources and cloud-edge orchestration between sites.

Finally, the D2.4 version better reflects how end users and external stakeholders interact with the system. The Application Layer now gives a clearer place to identity-aware access, analytics, federated learning, fog/cloud services, and proprietary operational platforms, thereby presenting HEDGE-IoT as a usable digital ecosystem rather than only as an integration architecture.

The complete evolution of the Reference architecture across the lifetime of the HEDGE-IoT project is presented in Annex A.

Co-funded by the European Union

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

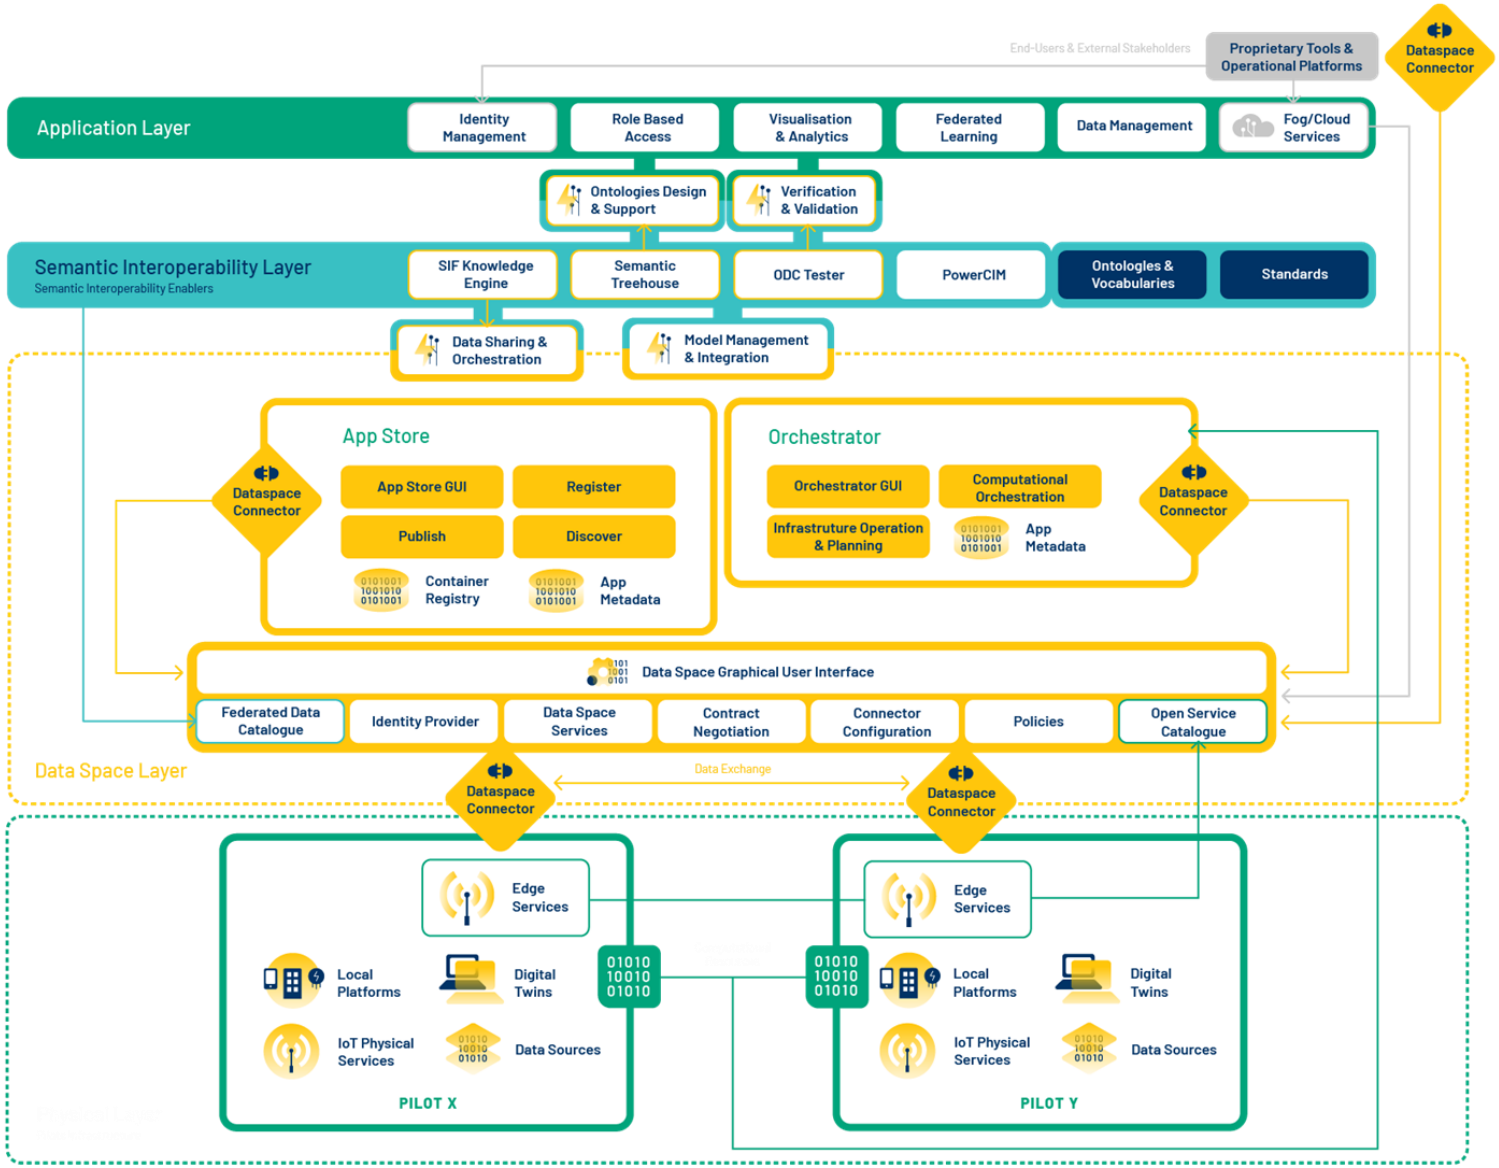

FIGURE 5 – HEDGE-IOT REFERENCE ARCHITECTURE – FINAL VERSION

Co-funded by the European Union

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

4.5.2 HEDGE-IoT 4+1 Architectural View Model

OVERVIEW AND METHODOLOGY

This Section describes the HEDGE-IoT RA using the 4+1 architectural view model, a well-established approach for capturing multiple perspectives of complex systems. This model organizes architectural descriptions into five distinct viewpoints: the Logical View, the Development View, the Process View, the Physical View, and the Scenarios View (the "+1"). Each view addresses specific actor's concerns and architectural dimensions, while together they provide a complete and comprehensive picture of the system. The 4+1 model is particularly valuable for HEDGE-IoT because it allows different stakeholders to understand the architecture from their respective viewpoints without requiring them to interpret a single monolithic model. This approach also ensures that no critical architectural concern is overlooked, as each view is explicitly designed to address specific quality attributes and stakeholder priorities.

Holistic System Understanding

The 4+1 views together provide a complete understanding of HEDGE-IoT:

- **Logical View** answers, "What are the main components and layers?": Four-layer architecture (Physical, Dataspace, Semantic, Application)
- **Development View** answers, "How is the software structured?": Components organized into packages; interfaces and dependencies defined
- **Process View** answers, "How do components interact at runtime?": Four key process flows; data flows through layers; concurrent operations supported
- **Physical View** answers, "How is the system deployed?": Distributed multi-site architecture; edge-fog-cloud tier model
- **Scenarios View** answers, "Can the architecture support required use cases?": TUCs and BUCs mapped to architectural layers; validation through concrete scenarios

TABLE 21 – VIEWPOINT SUMMARY

View	Primary Concern	Stakeholders	Key Questions	Main Artifacts
Logical	System structure & components	Architects, Senior Developers	What are main components? How do they relate?	Layer diagrams, Component relationships
Development	Software decomposition	Developers, Integrators	How is code organized? What are modules & interfaces?	Component hierarchy, Dependency graphs
Process	Runtime behavior & data flows	Performance Engineers	How do components interact? What are data flows?	Sequence diagrams, State machines

Co-funded by the European Union

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

Physical	Deployment & infrastructure	Deployment Engineers	Where does code run? How is it distributed?	Deployment diagrams, Network topology
Scenarios	Use case realisation	Business Analysts, All Stakeholders	Can architecture support required scenarios?	Scenario descriptions, Traceability matrix

LOGICAL VIEW

Purpose and Scope: The Logical View describes the main architectural elements of the HEDGE-IoT system and their relationships, independent of their physical deployment or runtime behavior. It presents the system as a collection of interrelated components and services organized into distinct layers. This view is essential for understanding the system's functional structure, how components interact, and how the system decomposition supports the stakeholder requirements and use cases defined in D2.1 and D2.2.

The HEDGE-IoT Logical View is organized into four distinct layers (also described in previous chapters), each with a specific role and set of responsibilities:

- **Layer 1 – Physical Layer:** Represents the physical infrastructure and data sources across pilot sites. Includes IoT devices, sensors, meters, edge computing nodes, local platforms (SCADA, EMS, BEMS), and data repositories. This layer is the source of all primary operational data (grid measurements, asset states, environmental conditions) and the point where control decisions are actualized in the field.
- **Layer 2 – Dataspace Layer:** Acts as the federated interoperable core enabling secure, policy-driven data and service exchange. Contains three main functional groupings: (a) Dataspace Core (EDC connectors, federated catalogue, identity provider (hub), contract negotiation, policies), (b) App Store (service registry, discovery, publication, and deployment), and (c) Orchestrator (computational resource allocation, workload scheduling, infrastructure management). This layer transforms isolated pilot data into discoverable, governable, and reusable assets.
- **Layer 3 – Semantic Interoperability Layer:** Provides the harmonisation mechanisms for consistent meaning across diverse data sources and services. Includes the SIF Knowledge Engine, Semantic Treehouse, semantic enablers (PowerCIM, ontology management), SAREF and CIM vocabularies, and verification/validation tooling. This layer ensures that data from different pilots, despite originating from heterogeneous systems, can be understood and reused consistently.
- **Layer 4 – Application Layer:** Delivers user-facing and analytics capabilities. Includes identity and access management, user interfaces and dashboards, federated learning services, fog/cloud analytics, visualisation tools, and adapters to operational platforms. This is where stakeholders (operators, analysts, decision-makers) engage with HEDGE-IoT services.

TABLE 22 – KEY LOGICAL COMPONENTS & THEIR RELATIONSHIPS

Layer	Key Components	Relationships	Responsibility
Physical	IoT devices, Edge nodes, Local platforms	Feed data to Dataspace Layer via	Operational data generation, local control execution

Co-funded by the European Union

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

	(SCADA/EMS/BEMS), Data repositories, Digital twins	connectors; receive control decisions	
Dataspace	EDC Connectors, Federated Catalogue, Identity Provider, App Store, Orchestrator, Policy Engine	Expose Physical data; coordinate with Semantic Layer; orchestrate Application execution	Secure data exchange, service discovery, computational orchestration
Semantic	Knowledge Engine, SAREF/CIM ontologies, Semantic Treehouse, PowerCIM, Validation tooling	Receive data from Dataspace; enrich and validate; serve semantically consistent data to Application	Semantic harmonisation, data validation, ontology management
Application	Analytics & Forecasting, Federated Learning, Optimisation services, User interfaces, External platform adapters	Consume semantically aligned data; invoke Orchestrator for deployment; provide user-facing services	End-user analytics, decision support, service delivery

Logical Data Flow Within Layers

Upward Flow (Data Acquisition to Services)

- Physical Layer: IoT devices and local platforms collect operational data
- Dataspace Layer: EDC connectors expose data through standardized APIs; federated catalogue makes data discoverable
- Semantic Layer: Knowledge Engine validates and enriches data semantics; mappings ensure consistent interpretation
- Application Layer: Services consume semantically aligned data to provide analytics, forecasting, optimisation, and control recommendations

Downward Flow (Services to Operations)

- Application Layer: Generates optimized setpoints, control decisions, and service recommendations
- Semantic Layer: Validates and contextualizes decisions before transmission
- Dataspace Layer: Orchestrator determines optimal deployment and coordinates execution
- Physical Layer: Control decisions are actualized in field assets through local platforms

DEVELOPMENT VIEW (COMPONENT ARCHITECTURE)

Purpose and Scope: The Development View describes the software decomposition of HEDGE-IoT into modules, components, and packages. It shows how the system's software structure supports modularity, reusability, and independent evolution. This view is critical for software developers, system integrators, and architects responsible for implementing the system, as it defines the actual building blocks, their interfaces, dependencies, and organisational structure.

HEDGE-IoT components are organized into logical packages aligned with the four-layer architecture as depicted in Table 23:

Co-funded by the European Union

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

TABLE 23 – HEDGE-IOT COMPONENTS

Package/Layer	Components	Type	Deployment
Physical Layer Package	IoT Gateway, Sensor Driver, Local Data Store, Controller Interface, Digital Twin	Hardware/Firmware	Edge nodes
Dataspace Core Package	EDC Connector, Federated Catalogue, Identity Hub, Contract Engine, Policy Engine	Microservice	Cloud/Fog
App Store Package	Service Registry, Discovery API, Deployment Manager, Metadata Handler	Microservice	Cloud
Orchestrator Package	Resource Allocator, Workload Scheduler, KubeEdge Adapter, Monitoring Agent	Microservice	Fog/Cloud
Semantic Layer Package	Semantic technological enabler APIs	Library/Microservice	Shared/Distributed
Application Layer Package	Forecasting Service, Optimisation Engine, Federated Learner, User Interfaces, Integration APIs	Microservice/SaaS	Cloud/Fog

Key Component Interfaces

Component interfaces define how software modules interact. Key interfaces include:

- **EDC Connector API:** REST-based interface for data providers and consumers to negotiate contracts and exchange data
- **Orchestrator API:** Interfaces for service deployment, resource allocation requests, and execution monitoring
- **App Store Registry:** API for publishing, discovering, and instantiating services
- **Knowledge Engine API:** Interface for semantic validation, ontology queries, and model transformation
- **Identity Provider Interface:** Decentralized identity management, leveraging verifiable credentials and contract-based communication for trust

Modularity and Dependency Management

Component architecture follows key modularity principles:

- **Layer Independence:** Each layer can be developed, tested, and deployed independently; loose coupling between layers through well-defined APIs
- **Pluggable Components:** Semantic enablers can be replaced or updated without affecting other layers
- **Service-Oriented Architecture:** Services in the Application Layer are designed as independent units discoverable and reusable via the App Store
- **Open Integration Points:** Standardized protocols allow integration with external systems

PROCESS VIEW (INTERACTION & DATA FLOWS)

Purpose and Scope: The Process View describes the dynamic behavior of HEDGE-IoT at runtime. It shows how components interact, how data flows through the system, and how services are orchestrated and executed. This view is essential for understanding the temporal aspects of the system, communication patterns, and the operational

Co-funded by the European Union

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

sequences that realise the use cases. It addresses concerns such as performance, responsiveness, scalability, and concurrency.

Four key process flows characterise HEDGE-IoT runtime behavior:

PF1: Data Providing and Registration Flow
- Data sources → Edge aggregation → EDC connector → Federated catalogue registration
- Timing: Continuous or event-triggered, depending on data source

PF2: Data Discovery and Consumption Flow
- Application/Service queries catalogue → Discovers available datasets and access policies à Initiates contract negotiation via EDC à Consumes data
- Timing: On demand, triggered by analytics or forecasting needs

PF3: Semantic Alignment and Validation Flow
- Incoming data → Knowledge Engine validation → Ontology mapping (SAREF ←→ CIM) → Semantic enrichment → Standardized output
- Timing: Asynchronous, processed before application consumption

PF4: Service Deployment and Orchestration Flow
- Service request (from user or automated trigger) → App Store lookup → Orchestrator receives deployment request → Resource availability analysis → Optimal placement decision (edge/fog/cloud) → Service instantiation à Monitoring and scaling
- Timing: Depending on placement optimisation

TABLE 24 – PROCESS FLOW CHARACTERISTICS

Flow Type	Source	Destination	Protocol	Frequency
Operational Data	IoT devices, smart meters	Edge aggregators, EDC	MQTT, Modbus, IEC 61850	Continuous/Real-time
Metadata Queries	Applications, Services	Federated Catalogue	REST API	On demand
Control Signals	Orchestrator, Applications	Local controllers	HTTP/HTTPS, MQTT	Event-triggered, Real-time
Semantic Mappings	Data in native format	Knowledge Engine	JSON, RDF	Asynchronous
Service Results	Analytics services	Application Layer, UI	REST API, WebSocket	Stream/Batch

Concurrency and Synchronisation

HEDGE-IoT supports concurrent operations:

- Multiple data providers simultaneously exposing datasets without mutual interference
- Parallel service execution across edge and cloud nodes
- Federated learning operations executing simultaneously across pilots
- Real-time data streams and batch processing coexisting without deadlock

Co-funded by the European Union

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

PHYSICAL/DEPLOYMENT VIEW

Purpose and Scope: The Physical View describes how the logical and software components are mapped onto actual hardware, network infrastructure, and distributed environments. It addresses deployment concerns including physical node distribution, network connectivity, cloud/edge resource allocation, failover mechanisms, and operational infrastructure. This view is essential for deployment engineers, system administrators, and infrastructure planners.

HEDGE-IoT follows a distributed, multi-site deployment model with six pilot environments as shown in Table 25:

TABLE 25 – DEPLOYMENT ARCHITECTURE

Pilot	Country	Primary Domain	Edge Deployment	Cloud Provider	Key Infrastructure
Finnish	Finland	Distribution grid resilience	Local substations (2-3 nodes)	Regional cloud	MQTT broker, IED gateways
Greek	Greece	Flexibility markets	Prosumer sites (10+ nodes)	National cloud	Smart meters, EMS
Italian	Italy	Energy communities	Community nodes (5-7 nodes)	Regional cloud	PV + battery systems
Dutch	Netherlands	Business Park optimisation	Park-local controller	Regional cloud	EMS, BMS integration
Portuguese	Portugal	Federated learning	Building controllers (8-10 nodes)	National cloud	HVAC, battery systems
Slovenian	Slovenia	Asset optimisation (DTR/DLR)	Substation edge server	Regional cloud	Temperature/ power quality sensors

Distributed Processing Architecture

HEDGE-IoT employs a three-tier computational model:

- **Edge Tier:** Local IoT edge nodes and gateways running lightweight services (data preprocessing, local anomaly detection, time-sensitive control). Low latency limited computational resources.
- **Fog Tier:** Pilot-local cloud infrastructure (regional data centers) running medium-complexity services (forecasting, optimisation). KubeEdge managed; supports 100ms-level latency.
- **Cloud Tier:** Centralized cloud services (federated learning, complex analytics, dataspace core). Supports complex computations; acceptable latency for non-real-time services.

Network Connectivity

Key connectivity patterns:

Co-funded by the European Union

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

- **Intra-Pilot (Local):** Direct connections via LAN/WAN; protocols: MQTT, HTTP/HTTPS, IEC 61850
- **Inter-Pilot (Federated):** Through EDC connectors over public internet; protocols: HTTPS, Dataspace Protocol
- **Cross-Layer (Orchestration):** Kubernetes API for edge-cloud coordination; REST API for service communication

Redundancy and Resilience

HEDGE-IoT deployment supports i) Geographic redundancy since data can be replicated across pilot sites, ii) Service failover and iii) Data sovereignty through EDC implementation (each pilot retains its own data locally).

SCENARIOS VIEW – TRACEABILITY TO TUCS

Purpose and Scope: The Scenarios View (the '+1' in the 4+1 model) describes concrete usage scenarios and use cases that validate the architecture. It demonstrates how the logical, development, process, and physical views come together to realise actual business and technical requirements. Scenarios are essential for validating the architecture that can support the operational needs defined in D2.1 and D2.2, and that it can implement the Transversal Use Cases (TUCs) that enable cross-pilot interoperability.

Scenarios in HEDGE-IoT are organized around the three Transversal Use Cases (TUCs) which represent cross-cutting capabilities applicable across all pilots:

TUC1: Data Interoperability — Data Exchange through HEDGE-IoT Dataspace
- **Scenario 1a (Data Producer Registration):** A pilot DSO registers its grid state data with the dataspace, specifying access policies and semantic models
- **Scenario 1b (Data Consumption):** A forecasting service discovers available flexibility asset data, negotiates access terms, and retrieves real-time consumption patterns
- **Scenario 1c (Metadata Discovery):** A developer queries the catalogue to find all datasets related to PV production, including their data formats and semantic vocabularies
- **Scenario 1d (Federated Service Chaining):** A service optimisation engine: retrieves load forecasts à applies grid constraints à invokes flexibility coordinator à sends control decisions

TUC2: Computational Orchestration — Automated Coordination of Computing Tasks
- **Scenario 2a (Edge Deployment):** An anomaly detection service is deployed on edge nodes near grid sensors for low-latency detection
- **Scenario 2b (Federated Learning):** A demand forecasting model is trained cooperatively across three pilot sites, with model updates synchronized without centralizing raw data
- **Scenario 2c (Service Rolling Update):** A new version of a forecasting service is deployed to edge nodes; old version gradually replaced; no service interruption

TUC3: App Store — Service Publication, Discovery, and Reuse
- **Scenario 3a (Service Publication):** A developer publishes a congestion prediction service with API specification, input/output formats, and deployment requirements
- **Scenario 3b (Service Discovery):** A grid operator searches the App Store for "flexibility estimation" services; finds 4 implementations; compares quality metrics and cost
- **Scenario 3c (Interchangeable Services):** Multiple pilots adopt a common "demand forecasting" service; all use the same SAREF data schema; service operates identically across sites

Each scenario maps to specific architectural layers and components:

Co-funded by the European Union

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

TABLE 26 – SCENARIO-TO-ARCHITECTURE TRACEABILITY

Scenario	Logical Layer	Dev Component	Process Flow	Physical Nodes
S1a: Data Producer Registration	Dataspace, Physical	EDC Connector, Federated Catalogue	PF1: Data Ingestion	Edge nodes + Cloud
S1b: Data Consumption	Dataspace, Application	EDC, Catalogue, Service	PF2: Discovery & Consumption	Cloud + Edge
S1d: Service Chaining	All four layers	Orchestrator, Services	PF4	All nodes
S2a: Edge Deployment	Dataspace, Physical	Orchestrator, KubeEdge	PF4: Orchestration	Edge nodes
S2b: Federated Learning	Application, Dataspace	Federated Learner, EDC	PF4 + Custom	All fog/cloud nodes
S3a: Service Publication	App Store, Dataspace	App Store Service Registry	Discovery API call	Cloud

Use Case Validation Through Scenarios

The Scenarios View demonstrates that HEDGE-IoT's four-layer architecture can successfully implement all TUCs. The vertical trace from Physical Layer (data collection) through Dataspace Layer (exposure and discovery), Semantic Layer (validation), and Application Layer (service execution) ensures that each scenario's objectives are architecturally supported. Furthermore, the distributed deployment across six pilots validates that the architecture scales horizontally without requiring modification of its core structure.

Scenarios Across Business Use Cases (BUCs)

While this detailed view focuses on the three TUCs (which represent horizontal cross-cutting capabilities), the architecture also supports the pilot-specific Business Use Cases. Each BUC leverages the TUCs infrastructure: for example, the Finnish pilot's "Anomaly Detection" BUC uses the Data Interoperability TUC to expose grid measurements, the Computational Orchestration TUC to deploy detection algorithms, and the App Store TUC to make the service discoverable. This pattern repeats across all pilots and BUCs, demonstrating the architecture's ability to serve diverse operational needs through a unified framework.

4.6. HEDGE-IOT HOURGLASS MODEL

This section introduces the Hourglass model, its principles, and presents the HEDGE-IoT Hourglass model (version 1.0).

4.6.1 Hourglass model principles

The Hourglass model was designed in the CEI-Sphere project [69]. According to CEI-Sphere article on the Hourglass model here is the definition:

Co-funded by the European Union

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

"The Hourglass Model is a conceptual framework that illustrates how different layers of digital capabilities and corresponding stakeholder groups interact to create intelligent, interoperable, and scalable digital ecosystems. Named for its distinctive shape — wide at the top and bottom with a narrow central layer — this model is especially relevant for complex systems such as smart cities, mobility platforms, industrial automation, and digital twins.

It shows how data flows from the physical world (sensors and devices) through enabling platforms and infrastructure, and ultimately up to applications and user interfaces. Each layer is structured to highlight a specific type of stakeholder, representing those responsible for building, enabling, or governing that part of the digital system. Opposite them are the key capabilities, the technological functions that make the system work.

Together, these dual aspects of the model help visualise the alignment between people, processes, and platforms." [70] [71]

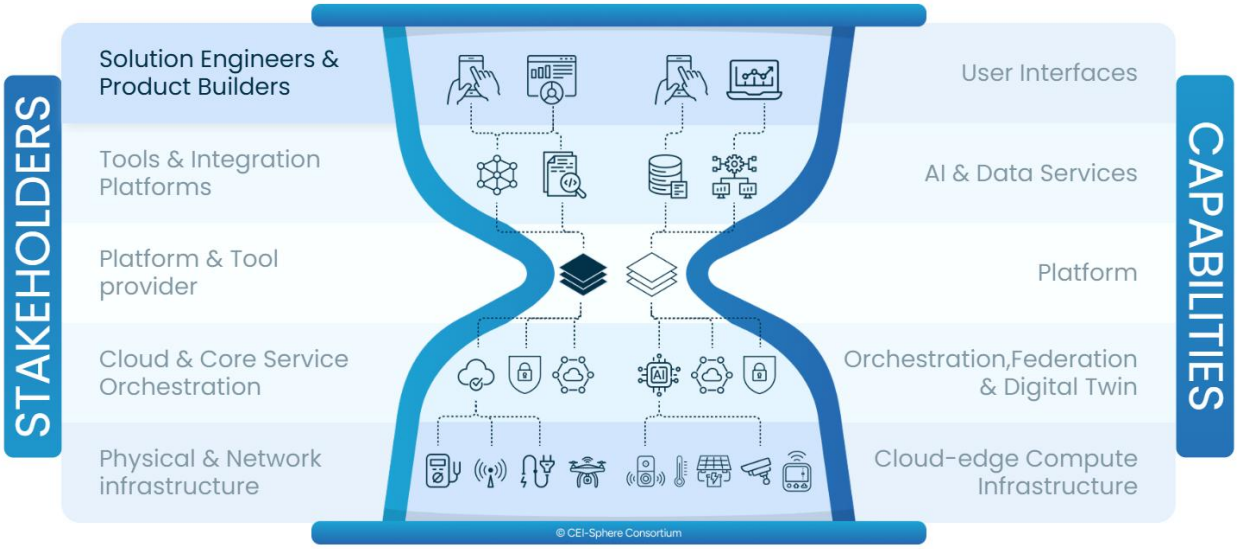

FIGURE 6 - HOURGLASS MODEL CONCEPTUAL FRAMEWORK [70]

This conceptual framework could be especially used to provide an overall picture to understand the concept of a project, stakeholders, capabilities, layers, and added values,

4.6.2 HEDGE-IoT Hourglass model

HEDGE-IoT Hourglass model (version 1.0) presented in Figure 7, was designed based on project content like use cases, specifications, main components, main assets, and stakeholders. Not all capabilities and stakeholders are shown on it to ensure the figure remains clear. This HEDGE-IoT hourglass model does not replace the project specifications or architecture. It is a tool for presenting the various aspects, issues, developments and challenges of the project.

It was decided to highlight as the key project result an AI-based services library for energy dataspace (at the middle of the hourglass model).

The stakeholders highlighted on the HEDGE-IoT Hourglass model are:

Co-funded by the European Union

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

- Grid operator (TSO, DSO),
- Energy community,
- Local Flexibility Markets,
- Grid data user for grid operation,
- Energy community user for operation,
- Local Flexibility Markets,
- Grid/energy/flexibility application capabilities developer,
- Flex service platform provider,
- Dataspace connector developer,
- Interoperability service provider,
- IDSA,
- Eclipse Foundation,
- Computational orchestration service provider,
- Grid data and system provider, and
- Smart Meter Data Sharing provider.

The capabilities highlighted on the HEDGE-IoT Hourglass model are:

- User interfaces for:
    - grid/grid infrastructures management,
    - energy/flexibility,
    - energy community
- Market interactions for flexibility,
- App store for grid/flexibility management services,
- Grid infrastructure, energy/flexibility CLOUD AI-based services,
- Interoperability layer for data exchange and connectivity,
- Dataspace compliance with IDSA, IDS-RAM, and the Dataspace Protocol (DSP),
- Semantic layer for semantic interoperability,
- Data catalogue,
- Orchestration of AI-based and infrastructure capabilities (incl. edge, cloud),
- Minimise data transfer for the services,
- AI-based services update,
- Data collection for grid infrastructure and energy/flexibility services,
- EDGE AI-based services for grid infrastructures, and energy/flexibility systems, and
- Edge computing infrastructure.

Co-funded by the European Union

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

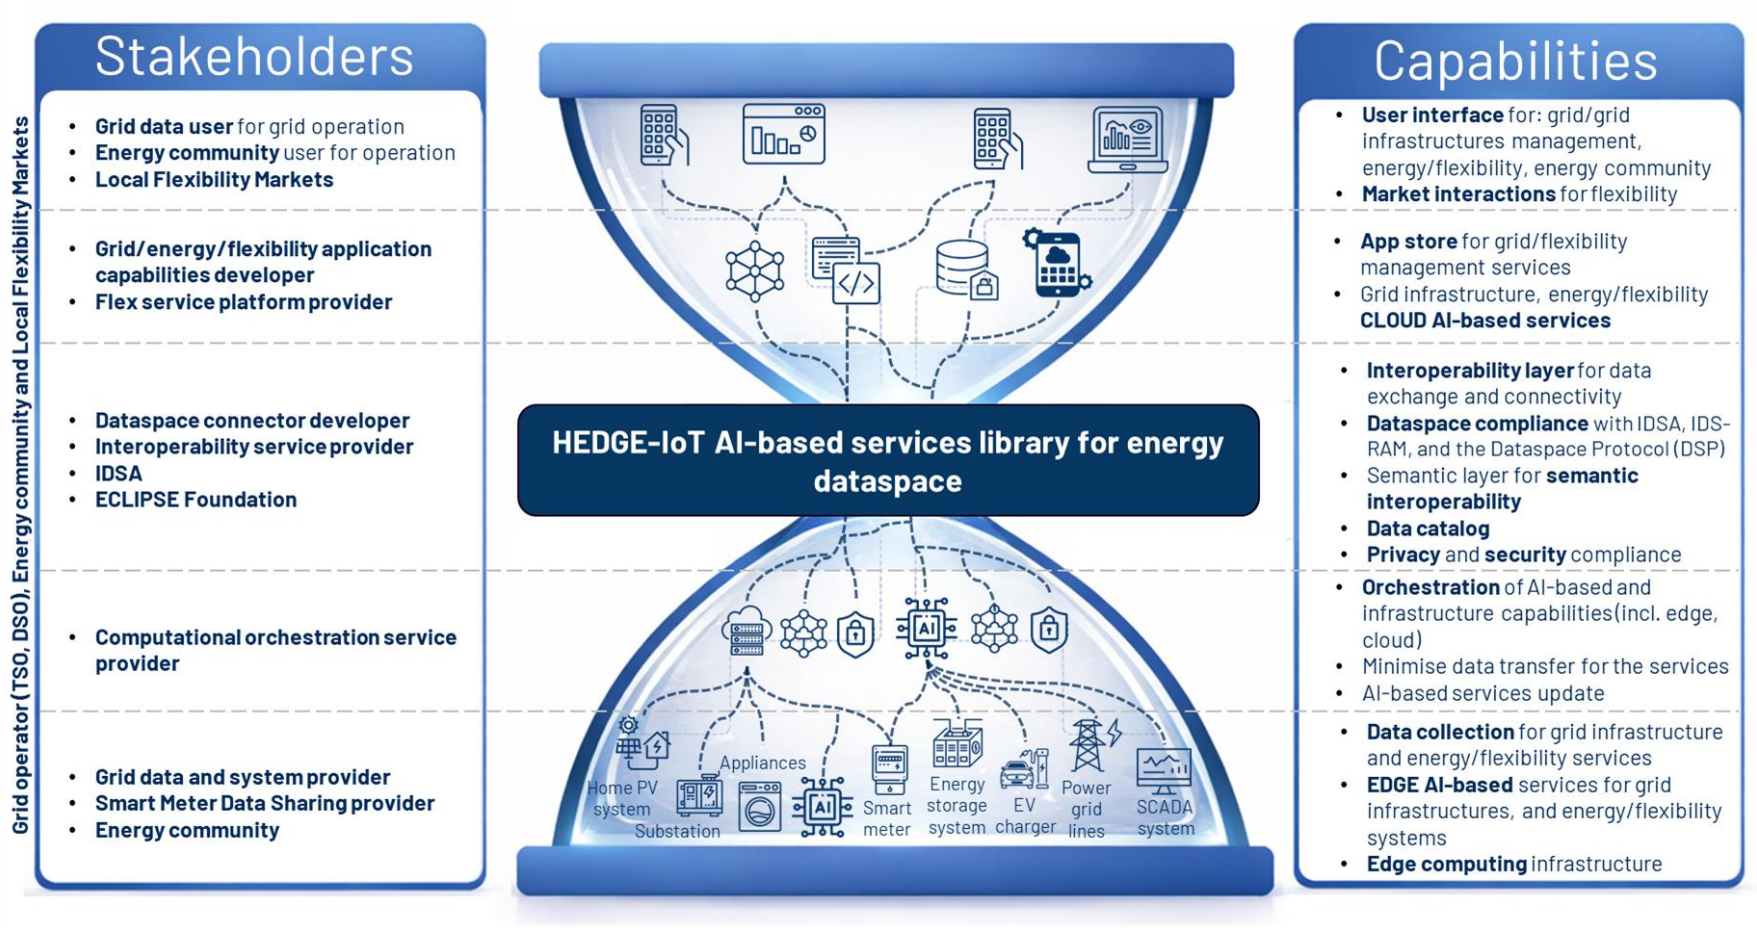

FIGURE 7 – HEDGE-IOT HOURGLASS MODEL (VERSION 1.0)

Co-funded by the European Union

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

4.7. SECURITY AND PRIVACY ARCHITECTURE

This chapter aims to provide an overview of project activities related to HEDGE-IoT trustworthiness, cybersecurity, and privacy considerations. It is based on the input coming from Task 4.5 – Cybersecurity considerations and AI safety.

The following trustworthiness aspects are covered by WP4 task 4.5:

- Cross-cutting characteristics plan (X-CCP) analysis throughout the pilots (main activity of T4.5),
- Compliance and regulatory considerations (e.g., GDPR, NIS2, CRA),
- HEDGE-IoT architecture security analysis and recommendations,
- Dataspaces cybersecurity and privacy, and
- AI explainability analysis.

CROSS-CUTTING CHARACTERISTICS PLAN (X-CCP) ANALYSIS THROUGHOUT THE PILOTS

- Objective:
  Ensure that cybersecurity, data privacy and AI trustworthiness is adequately managed in HEDGE-IoT demonstrators, system-of-interest and its associated ecosystem.
- Description:
  To achieve this objective, a Cross-Cutting Characteristics Plan (X-CCP) was adapted and developed according to the specific needs of HEDGE-IoT to address the identified targeted trustworthiness characteristics. This activity is the main one achieved within T4.5.

  The X-CCP covers the following characteristics:
  - AI trustworthiness,
  - data privacy,
  - cybersecurity, and
  - progress monitoring.

  This practice is based on an existing methodology developed by TRIALOG and successfully demonstrated in previous EU projects, e.g., InterConnect, PARMENIDES. This method was initially named Privacy and Security Plan (PSP) and changed over time to Cross-Cutting Characteristics Plan (X-CCP) to consider additional characteristics like AI trustworthiness in HEDGE-IoT. The practice is composed of 4 trainings, 4 workshops (including an analysis) for each pilot, and associated activities.

  The X-CCP practice expects to provide the following outputs to the project and pilots:
  - training sessions,
  - results analysis reports for each pilot on privacy, cybersecurity and AI trustworthiness,
  - pilots' maturity and security status and progress assessment,

Co-funded by the European Union

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

- action plan for improvement, and
- knowledge to tend to conformity based on the analysis reports.

Trustworthiness profiles will be explored.

- Content location:
  The method is described in D4.1, the first half of the results in D4.2, and the second half and trustworthiness profiles will be reported in D4.3.

COMPLIANCE AND REGULATORY CONSIDERATIONS (E.G., GDPR, NIS2, CRA)

- Objective: Raise awareness within the HEDGE-IoT consortium on compliance and regulation.
- Description: This activity is in the form of a state-of-the-art, awareness-raising in X-CCP sessions, and lists of regulation requirements and recommendations to support partners and their systems to move towards compliance.
- Content location:
  The state of the art on regulation and standardisation is available on D4.1.
  The lists of regulation requirements and recommendations are available in D4.1. and will be reported in D4.3.

HEDGE-IOT ARCHITECTURE SECURITY ANALYSIS AND RECOMMENDATIONS

- Objective:
  Make sure that cybersecurity is well managed throughout the HEDGE-IoT RA.
- Description:
  A cybersecurity analysis on the HEDGE-IoT architecture was performed, going through threats identification and recommendations issuance.
- Content location:
  The results will be reported in D4.3.

DATASPACES CYBERSECURITY AND PRIVACY CONSIDERATIONS

- Objective:
  Make sure that dataspace cybersecurity and privacy are well-known and managed within HEDGE-IoT, following standards and guidelines.
- Description: Summary of the different cybersecurity and privacy considerations (e.g., identity and trust management, data sovereignty and usage control).
- Content location:
  Results can be found in D2.3.

AI EXPLAINABILITY ANALYSIS

- Objective:
  Ensure that AI system developers and producers take AI explainability properly into account by raising their awareness of the concept.

Co-funded by the European Union

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

- Description:
  To go further on AI trustworthiness, an unplanned outlook was added to the task scope on AI explainability. This activity is composed of a training session, a study of some AI systems from the explainability point of view (supported by workshops).
- Content location:
  The method and the results will be reported in D4.3.

Co-funded by the European Union

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

5 COMPONENT-TO-REQUIREMENTS TRACEABILITY MATRIX

This chapter establishes traceability between the functional requirements defined in Section 3.3 and the components from the HEDGE-IoT RA described in Section 4.5.2. The objective is to demonstrate how the final HEDGE-IoT RA supports the WP2 requirements baseline and to provide a structured basis for implementation, integration and validation activities in subsequent work packages.

Table 27 maps each functional requirement (described in Section 3.3) to the HEDGE-IoT components grouped in packages/layers from Table 23 An "x" indicates that the requirement is supported by corresponding package / layer. The complete component-level traceability matrix, mapping each functional requirement to each individual component is presented in Annex BANNEX B — Detailed Component-to-Requirements Traceability Matrix.

Some functional requirements are not mapped directly to the common architectural packages in Table 27 because they are primarily realised through HEDGE-IoT services, pilot-specific applications, market-facing functions or service chains rather than through a reusable reference-architecture package alone. This is particularly the case for FR-UI-04 Dynamic tariffs interface and the Flexibility Management requirements FR-FM-01 to FR-FM-08. These requirements are not excluded from project traceability; instead, they are addressed at service level and through the relevant BUCs, SUCs and pilot implementations.

This distinction is consistent with the commonalities analysis presented in D2.3 (see D2.3 Table 11: COMMONALITIES - OBJECTIVES), where the normalised objectives breakdown by pilot identifies which functional areas are common across pilots and which ones are primarily implemented through pilot services. In particular, the D2.3 table shows that several objectives related to flexibility and user interaction are service-oriented rather than architecture-package-oriented. These objectives correspond directly to the requirements that remain blank in the architecture-package traceability matrix, especially FR-UI-04 and FR-FM-01 to FR-FM-08 and confirm that their primary implementation route is through HEDGE-IoT services, pilot workflows and market-facing applications rather than through the common reference-architecture packages alone.

Accordingly, Table 27 should be read as the common reference-architecture traceability view, while the complementary FR-to-Services mapping covers requirements whose fulfilment depends mainly on service-level capabilities. This includes dynamic-tariff interfaces, flexibility registration and prequalification, flexibility offer handling, market price tracking, activation and planning, flexibility estimation, settlement and payment functions, and vulnerable-user identification. Together, the architecture-package matrix and the FR-to-Services mapping preserve end-to-end traceability from functional requirements to architectural layers, technical components, HEDGE-IoT services, pilot-specific implementations and the normalised cross-pilot objectives identified in D2.3.

Co-funded by the European Union

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

TABLE 27 – COMPONENT-TO-REQUIREMENTS TRACEABILITY MATRIX

Functional Requirement ID	Physical Layer Package	Data Space Core Package	App Store Package	Orchestrator Package	Semantic Layer Package	Application Layer Package
FR-DM-01	x					
FR-DM-02		x				
FR-DM-03				x		
FR-DM-04	x					
FR-DM-05	x	x	x			x
FR-DM-06	x	x	x			x
FR-IOP-01		x			x	x
FR-IOP-02					x	x
FR-IOP-03		x				
FR-IOP-04		x				
FR-IOP-05	x	x			x	x
FR-IOP-06						
FR-SRV-01			x			
FR-UI-01	x	x	x	x	x	x
FR-UI-02	x	x				x
FR-UI-03	x	x	x	x		x
FR-UI-04						
FR-UI-05		x				x
FR-OF-01						x
FR-OF-02						x
FR-OF-03						x
FR-OF-04						x
FR-OF-05						x
FR-OF-06	x					x
FR-OF-07						x
FR-FM-01						
FR-FM-02						
FR-FM-03						
FR-FM-04						
FR-FM-05						
FR-FM-06						
FR-FM-07						
FR-FM-08						
FR-GMC-01	x					x
FR-GMC-02	x				x	
FR-AI-01						x
FR-AI-02						x
FR-AI-03						x

Co-funded by the European Union

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

FR-MED-01						x
FR-MED-02						x
FR-MED-03						x

Note on service-level traceability: *Requirements without direct mapping in Table 27 are not considered unsupported. They are handled through the FR-to-Services mapping, where service-oriented and pilot-specific requirements are linked to the relevant HEDGE-IoT services, BUCs, SUCs and implementation activities.*

Co-funded by the European Union

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

6 CONCLUSIONS

D2.4 consolidates the final WP2 RA for HEDGE-IoT. It brings together the methodological foundations, stakeholder requirements, business and system use cases, transversal use cases, functional requirements, interoperability profiles, landscape alignment and implementation feedback developed across WP2 and related technical work packages. The result is a coherent reference model for secure data sharing, semantic interoperability, cloud-edge orchestration and reusable energy services across heterogeneous pilot environments.

The final architecture should be understood as a reference model rather than as a single deployment blueprint. It defines the common structure, layers, component groupings and interoperability mechanisms that guide the project, while allowing pilot-specific deployment choices. This flexibility is necessary because the six HEDGE-IoT pilots operate in different national, technical and organisational contexts.

A central outcome of D2.4 is the stabilisation of the architecture into four main layers: the Physical Layer, the Dataspace Layer, the Semantic Interoperability Layer and the Application Layer. This structure clarifies the separation of responsibilities between field assets, governed data and service exchange, semantic harmonisation and user-facing digital capabilities.

The Dataspace Layer is the operational core of the architecture. It enables distributed participants to publish, discover, negotiate and exchange data and services under explicit governance and policy conditions. Through EDC-based principles, federated catalogues, identity management, contract negotiation and policy enforcement, the architecture supports data sovereignty while enabling cross-organisational interoperability.

Semantic interoperability is also established as a core architectural concern. By elevating semantics to a dedicated layer, D2.4 recognises that HEDGE-IoT must support not only data exchange, but also consistent interpretation of data, metadata and service descriptions across pilots, platforms and applications. Semantic enablers, ontologies, vocabularies, knowledge-engine functions and validation mechanisms provide the basis for cross-pilot reuse and portability.

The final architecture also strengthens service reuse and distributed computation. The App Store and Open Service Catalogue support publication, discovery and reuse of applications and services, while the Orchestrator supports workload deployment and execution across edge, fog and cloud resources. Together, these components allow HEDGE-IoT to coordinate data, services and computation according to operational needs, latency requirements, privacy constraints and infrastructure availability.

The SGAM and BRIDGE DERA mappings confirm that the HEDGE-IoT RA is aligned with recognised European energy-sector frameworks. This alignment improves the interpretability, replicability and long-term relevance of the architecture, and positions HEDGE-IoT within the broader evolution toward interoperable European energy data spaces.

Co-funded by the European Union

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

Security, privacy and AI trustworthiness remain cross-cutting concerns. The architecture incorporates identity, access control, policy enforcement, secure exchange, data sovereignty and governance principles as foundational mechanisms. In parallel, the X-CCP, cybersecurity, privacy and AI trustworthiness activities provide a basis for analysing risks and improving pilot maturity.

Finally, the component-to-requirements traceability matrix demonstrates how the WP2 requirements baseline is reflected in the final architecture. It provides a structured basis for implementation, integration and validation, and should remain a living artefact as pilot deployment and service validation continue.

6.1. SUMMARY OF ARCHITECTURE DECISIONS

The final HEDGE-IoT RA was stabilised as a four-layer model composed of the Physical Layer, Dataspace Layer, Semantic Interoperability Layer and Application Layer. This structure improves separation of concerns while preserving an end-to-end pathway from field assets and IoT-edge nodes to governed data exchange, semantic harmonisation and user-facing applications.

A dataspace-centric architecture was adopted as the main federation mechanism for secure and sovereign data sharing. The Dataspace Layer brings together connector-based exchange, federated catalogue functions, identity management, policy enforcement and contract negotiation, allowing data providers and data consumers to interact under explicit governance rules.

Semantic interoperability was elevated to a dedicated architectural layer. This decision reflects the project's need to preserve consistent meaning across heterogeneous pilots, data models, service descriptions and operational contexts. The Semantic Interoperability Layer supports harmonisation, validation, ontology management, model alignment and cross-pilot portability.

A generic IoT-edge node model was retained and complemented with pilot-specific configurations. This allows the architecture to represent the diversity of the Finnish, Greek, Italian, Dutch, Portuguese and Slovenian pilots while maintaining a common pattern for data acquisition, local processing, edge intelligence and integration with the wider HEDGE-IoT ecosystem.

The App Store and Open Service Catalogue were made explicit architectural building blocks to support service publication, discovery, reuse and future onboarding of third-party applications. This strengthens the modularity and exploitation potential of the HEDGE-IoT ecosystem.

Computational orchestration was positioned as a core capability of the Dataspace Layer. The Orchestrator supports workload placement, service deployment, federated-learning support, rolling updates and coordinated execution across edge, fog and cloud resources.

Identity management, trust, policy enforcement and contract-based access were treated as foundational cross-layer concerns. These mechanisms are essential because HEDGE-IoT operates across organisational boundaries and involves multiple actors, including system

Co-funded by the European Union

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

operators, energy communities, aggregators, service providers, data providers and data consumers.

The final architecture was aligned with SGAM, BRIDGE DERA, IDS-RAM, IDSA, Gaia-X, SAREF, IEC CIM and FIWARE-compatible approaches. This alignment improves consistency with European energy-system interoperability models and supports future replication.

The 4+1 architectural view model and ISO/IEC/IEEE 42010 were used as practical structuring aids. They supported the description of the architecture from logical, development, process, physical and scenario perspectives without imposing a rigid formalism.

Interoperability profiles were introduced to translate architectural principles into practical guidance for selected interoperability points. These profiles use transport, syntactic, semantic, behavioural and policy facets to bridge high-level architecture and implementation-level integration.

Overall, the final architecture reflects convergence between requirements, landscape analysis, implementation feedback and pilot validation. It is therefore neither purely top-down nor purely implementation-driven, but grounded in stakeholder needs, European frameworks, technical feasibility and pilot realities.

6.2. KNOWN LIMITATIONS AND OPEN ISSUES

Although the D2.4 RA provides a stable common model for HEDGE-IoT, implementation maturity still differs across pilots and components. Some elements are already supported by technical releases and pilot activities, while others remain under integration, validation or deployment. The architecture should therefore be read as the final WP2 reference baseline, not as evidence of uniform implementation maturity.

Full plug-and-play interoperability cannot yet be assumed across all environments. Several pilots still depend on local adapters, gateway functions, data transformations, legacy-system interfaces and pilot-specific integration choices. The reference architecture reduces fragmentation by defining common patterns, but practical integration effort remains necessary.

Semantic interoperability also remains an iterative activity. Shared ontologies, vocabularies, knowledge-engine functions and model-management tools provide a strong basis for harmonisation, but semantic alignment still depends on the quality, granularity and stability of pilot data models.

The dataspace approach depends on organisational and governance readiness as well as technical connector deployment. Contract negotiation, data-usage policies, identity management and participant onboarding require clear roles, responsibilities and operational procedures.

The App Store and Open Service Catalogue introduce important mechanisms for service reuse, but their long-term governance remains open. Decisions will be needed on ownership, maintenance, quality assurance, lifecycle management, licensing, publication rules and post-project operation.

Co-funded by the European Union

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

Computational orchestration requires further validation under realistic distributed conditions. Performance, latency, resilience, resource availability, failure handling and workload-placement decisions must continue to be tested through pilot integration and technical validation.

Flexibility-market interfaces remain heterogeneous across pilots. The architecture documents pilot-specific interoperability profiles, but it does not define a single harmonised flexibility-market interface for all contexts. Differences in national regulation, market design and platform maturity will continue to influence implementation choices.

Security, privacy and AI trustworthiness require continued follow-up. The architecture includes relevant mechanisms, but their effectiveness depends on pilot-level action plans, system hardening, risk analysis, AI explainability work and practical implementation of cybersecurity and privacy recommendations.

The component-to-requirements traceability matrix should remain a living artefact. Some requirements, especially those related to flexibility management, dynamic tariffs and market-facing services, are better addressed at service or pilot level than through common architectural packages alone.

Finally, scalability beyond the current pilot scope remains to be demonstrated. Larger ecosystems will place additional demands on metadata management, identity federation, service discovery, policy negotiation, semantic validation and orchestration. D2.4 provides the baseline for such scaling, while operational proof will depend on subsequent integration, validation, exploitation and replication activities.

Co-funded by the European Union

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

REFERENCES

[1] ISO, "ISO/IEC/IEEE 42010:2022 -Software, systems and enterprise — Architecture description," [Online]. Available: https://www.iso.org/standard/74393.html

[2] P. B. Kruchten, "The 4+1 View Model of architecture," in IEEE Software, vol. 12, no. 6, pp. 42-50, Nov. 1995, doi: 10.1109/52.469759.

[3] European Commission: Directorate-General for Energy, "European (energy) data exchange reference architecture 3.0," Publications Office of the European Union, 2023, https://op.europa.eu/en/publication-detail/-/publication/dc073847-4d35-11ee-9220-01aa75ed71a1/language-en.

[4] CEN – CENELEC – ETSI: Smart Grid Coordination Group, Smart Grid Reference Architecture Report V2.0, 2012, https://www.researchgate.net/publication/263264218_CEN_-CENELEC_-_ETSI_Smart_Grid_Coordination_Group_-_Smart_Grid_Reference_Architecture_Report_20

[5] European Commission, "The European AI Alliance," Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/policies/european-ai-alliance

[6] European Commission, "Ethics guidelines for trustworthy AI," Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai#:~:text=%2A%20Societal%20and%20environmental%20well,Moreover%2C%20adequate%20an%20accessible (accessed Jul. 01, 2025).

[7] European Commission, "Key actions for digitalising energy," Energy. https://energy.ec.europa.eu/topics/eus-energy-system/digitalisation-energy-system/key-actions-digitalising-energy_en#:~:text=,2024 (accessed Jul. 01, 2025).

[8] AIOTI – Alliance for AI, IoT and Edge Continuum Innovation. https://aioti.eu/

[9] Interconnect "D2.1 - Secure interoperable IoT smart home/building and smart energy system reference architecture," 2020. [Online] Available: https://interconnectproject.eu/wp-content/uploads/2022/03/D2.1-Secure-Interoperable-Smart-Home-Building-and-Smart-Energy-System-Reference-Architecture_FR_v2.pdf

[10] O. Vermesan, "Advancing IoT Platforms Interoperability," June 2018, [Online]. Available: doi: https://doi.org/10.13052/rp-9788770220057

[11] International Data Spaces Association, "Home - International Data Spaces," International Data Spaces, Jun. 21, 2025. https://internationaldataspaces.org/

[12] Giussani G., Steinbuss S., Data Connector Report, International Data Spaces Association, (6), 2024 https://doi.org/10.5281/zenodo.13838396

[13] B. Otto, S. Steinbuß, A. Teuscher, and S. Lohmann, "REFERENCE ARCHITECTURE MODEL: Version 3.0," International Data Spaces Association, Apr. 2019. [Online]. Available: https://internationaldataspaces.org/wp-content/uploads/IDS-Reference-Architecture-Model-3.0-2019.pdf

[14] Bader, S., et. al." The International Data Spaces Information Model – An Ontology for Sovereign Exchange of Digital content", 2020, https://doi.org/10.1007/978-3-030-62466-8_12

[15] "Task forces Archive - BDV Big Data Value Association," BDV Big Data Value Association. https://bdva.eu/task-forces/ "BDVA | Task Forces," [Online]. Available: https://bdva.eu/task-forces/.

Co-funded by the European Union

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

[16] Gaia-X European Association for Data and Cloud AISBL, "GAIA-X Framework - GAIA-X: a federated Secure data infrastructure," Gaia-X: A Federated Secure Data Infrastructure -, Mar. 24, 2023. GAIA-X Framework," [Online]. Available: https://gaia-x.eu/gaia-x-framework/.

[17] GAIA-X, "Policy Rules Document," Apr. 2022. [Online]. Available: https://gaia-x.eu/wp-content/uploads/2022/05/Gaia-X_Policy-Rules_Document_v22.04_Final.pdf

[18] GAIA-X, "Architecture Document," Apr. 2022. [Online]. Available: https://gaia-x.eu/wp-content/uploads/2022/06/Gaia-x-Architecture-Document-22.04-Release.pdf

[19] ATTEST, "ATTEST - home," Attest Project, Nov. 03, 2023. https://attest-project.eu/

[20] ATTEST "D2.2: Toolbox Specifications," 2020, [Online]. Available: https://attest-project.eu/wp-content/uploads/Attachment_0-4-1.pdf

[21] I-NERGY "D2.5: I-NERGY Architecture and I-NERGY–AI4EU Synergies, Section 2.2.4: Application Layer – Energy Analytics Applications" (p. 22)", 2022, [Onine] Availbale: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5f47f03f5&appId=PPGMS

[22]OneNET D5.2: "OneNet Reference Architecture" [Online] Available: https://onenet-project.eu/wp-content/uploads/2022/12/OneNet_D5.2_v1.0.pdf

[23]M. Couto, "BRIDGE - European (Energy) Data Exchange Reference Architecture 3.1," European Commission, Data Management Working Group, Oct. 2024. [Online]. Available: https://op.europa.eu/en/publication-detail/-/publication/6c3b1add-a0a7-11ef-85f0-01aa75ed71a1#

[24]E. Union, "DIRECTIVE (EU) 2019/944 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 5 June 2019 on common rules for the internal market for electricity and amending Directive 2012/27/EU," [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1593158348328&uri=CELEX:32019L0944.

[25]International Electrotechnical Commission, "IEC 62559-2: Use case methodology – Part 2: Definition of the templates for use cases, actor list and requirements list," IEC, Apr. 2015. [Online]. Available: https://cdn.standards.iteh.ai/samples/20300/c3c2f3905fd045cba1b81c615be6d3c3/IEC-62559-2-2015.pdf

[26] European Network of Transmission System Operators for Electricity, "The Harmonised Electricity Market Model," Nov. 2022. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/EDI/Library/HRM/Harmonised_Role_Model_2022-01.pdf .

[27]International Data Spaces Association, "Dataspace Protocol 2024-1," [Online]. Available: https://docs.internationaldataspaces.org/ids-knowledgebase/v/dataspace-protocol.

[28]International Data Spaces Association, "Advancing interoperability: the Dataspace Protocol," [Online]. Available: https://internationaldataspaces.org/offers/dataspace-protocol-overview/.

[29] International Data Spaces Association, "App Store and App Ecosystem," [Online]. Available: https://docs.internationaldataspaces.org/ids-knowledgebase/v/ids-ram-4/layers-of-the-reference-architecture-model/3-layers-of-the-reference-architecture-model/3_5_0_system_layer/3_5_3_app_store_and_data_apps#app-store-and-ids-apps.

[30] International Data Spaces Association, "Metadata Broker," [Online]. Available: https://docs.internationaldataspaces.org/ids-knowledgebase/v/ids-ram-4/layers-of-the-reference-architecture-model/3-layers-of-the-reference-architecture-

Co-funded by the European Union

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

model/3_5_0_system_layer/3_5_4_metadata_broker#metadata-broker.

[31] European Commission: Directorate-General for Energy, Data Management Working Group, "Directorate-General for Energy, Data Management Working Group," Publications Office of the European Union, 2023. https://bridge-smart-grid-storage-systems-digital-projects.ec.europa.eu/working-groups/data-management

[32] European Commission Directorate-General For Energy, "M/490 - Standardization mandate to European Standardisation Organisations (ESOS) to support European smart grid deployment.", 2021 [Online]. Available: https://energy.ec.europa.eu/publications/mandate-m490-smart-grids-march-2011_en

[33]European Commission, "Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment,"2020. [Online]. Available: https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment

[34] G. De Panfilis, "FIWARE - Open APIs for Open Minds," FIWARE,. https://www.fiware.org/

[35] A. Tejado, T. Sapia, and C. Pezuela, "FI-NEXT - D5.2: FIWARE Go-to-Market Y2," 2018. [Online]. Available: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5c20897fa&appId=PPGMS

[36] FIWARE, "FIWARE: THE OPEN SOURCE PLATFORM OF CHOICE FOR BUILDING SMART ENERGY SOLUTIONS." Accessed: Jul. 01, 2025. [Online]. Available: https://www.fiware.org/wp-content/directories/marketing-toolbox/material/FIWAREBrochure_SmartEnergy.pdf

[37]Alliance for Internet of Things Innovation "High Level Architecture (HLA), Release 5.0", December 2020, [Online] Available: https://aioti.eu/wp-content/uploads/2020/12/AIOTI_HLA_R5_201221_Published.pdf

[38] The GridWise Architecture Council, "GridWise Interoperability ContextSetting Framework," Mar. 2008. [Online]. Available: https://gridwiseac.org/pdfs/GridWise_Interoperability_Context_Setting_Framework.pdf

[39] Gaia -X, "GAIA-X: Technical Architecture Release," Federal Ministry for Economic Affairs and Energy (BMWi), Jun. 2020. [Online]. Available: https://www.bundeswirtschaftsministerium.de/Redaktion/EN/Publikationen/gaia-x-technical-architecture.pdf?__blob=publicationFile&v=7

[40] S. Jiménez, "Interoperability Framework in energy data spaces: Position Paper | Version 2.0," International Data Spaces Association, Mar. 2025. [Online]. Available: https://enershare.eu/wp-content/uploads/IDSA-Position-Paper-Interoperability-Framework-in-Energy-Data-Spaces-v2-2.pdf

[41] Platone "Platform for operation of distribution networks," Platone - Platform for Operation of Distribution Networks. https://www.platone-h2020.eu/

[42]"RESONANCE 'D2.1: Initial Requirements and Common System Architecture,'" 2023. [Online]. Available: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5fb270ea8&appId=PPGMS

[43] Lampathaki, F., Biliri, E., Tsitsanis, T., Tsatsakis, K., Miltiadou, D., Perakis, K. (2022). Toward an Energy Data Platform Design: Challenges and Perspectives from the SYNERGY Big Data Platform and AI Analytics Marketplace. In: Curry, E., Scerri, S., Tuikka, T. (eds) Data Spaces . Springer, Cham. https://doi.org/10.1007/978-3-030-98636-0_14

[44] HEDGE-IoT "D1.1: Project Management Handbook", 2024, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D1.1-Management-communication-and-quality-approaches-Data-management-and-IPR-protection-procedures.pdf

Co-funded by the European Union

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

[45] HEDGE-IoT "D1.4: Data Management Plan", 2024, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D1.4-Data-Management-Plan.pdf
[46] HEDGE-IoT "D2.1: Requirements on an IoT Cloud/Edge System for the Energy Ecosystem", 2024, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D2.1-Requirements-on-an-IoT-CloudEdge-System-for-the-Energy-Ecosystem.pdf
[47]HEDGE-IoT "D2.2: Functional Specifications of the HEDGE-IoT system", 2024, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D2.2-Functional-Specifications-of-the-HEDGE-IoT-system.pdf
[48] HEDGE-IoT "D3.1: HEDGE-IoT Interfaces and Tools for Interoperability", 2025, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D3.1-HEDGE-IoT-Interfaces-and-Tools-for-Interoperability.pdf
[49] HEDGE-IoT "D3.3: HEDGE-IoT Technological Enablers (First Release)", 2025, [Online] Available: https://hedgeiot.eu/wp-content/uploads/2025/03/D3.3-HEDGE-IoT-Technological-Enablers-First-Release.pdf
[50] IETF - https://www.ietf.org/
[51] IEEE - https://www.ieee.org/
[52] MQTT - https://mqtt.org/
[53] NIST - https://www.nist.gov/
[54] Cordis. Europa. "Boosting DR through increased communlty-level consumer engaGement by combining Data-driven and blockcHain technology Tools with social science approaches and multi-value service design," CORDIS | European Commission, Sep. 11, 2020. https://cordis.europa.eu/project/id/957816/reporting
[55] Bright "D2.5: Cross-Domain Data & Service Interoperability", 2022, [Online] Available: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5f511726d&appId=PPGMS
[56] European Commission, "European Energy Data Exchange Reference Architecture," BRIDGE - Data Management Working Group, 2020. [Online]. Available: https://energy.ec.europa.eu/system/files/2021-06/bridge_wg_data_management_eu_reference_architcture_report_2020-2021_0.pdf
[57]Cordis , "OPEN DEI: Aligning reference architectures, open platforms and large scale pilots in digitising European industry," CORDIS | European Commission, https://cordis.europa.eu/project/id/857065/reporting/es
[58] Cordis, "Int:NET: Interoperability Network for the energy Transition," CORDIS | European Commission, https://cordis.europa.eu/project/id/101070086
[59] Alberto Dognini, "Blueprint of the Common European Energy Data Space," Interoperability Network for the Energy Transition (int:net), Jul. 2024. doi: 10.5281/zenodo.12609569.
[60] HEDGE-IoT "D2.3: HEDGE-IoT Reference Architecture (First Release)", 2025
[61] HEDGE-IoT "D3.4: Technological Enablers (Intermediate Release)", 2025
[62] HEDGE-IoT D4.1 "HEDGE-IoT Interoperability Framework and Integrated Solution (First release)", 2025
[63] HEDGE-IoT D4.2 "HEDGE-IoT Interoperability Framework and Integrated Solution (Intermediate release)", 2025
[64] HEDGE-IoT D5.1 "Guidelines for Demo Preparation"
[65] HEDGE-IoT D5.2 "Pre-Demo Phase Report"
[66] HEDGE-IoT D7.4 "Dissemination, Exploitation and Market Exploration, Standardisation, and Community Building (Intermediate Release)"
[67]ISO/IEC 19941:2017, Available at: https://www.iso.org/standard/66639.html.
[68] International Electrotechnical Commission, "ISO/IEC 21823-1," IEC, Feb. 2019. [Online]. Available:

Co-funded by the European Union

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

https://cdn.standards.iteh.ai/samples/100715/7e7bb111e4a84829897e7d7b6137714b/ISO-IEC-21823-1-2019.pdf
[69] CEI-Sphere | Empowering Innovation in Cloud-Edge-IoT Available at:https://ceisphere.eu/.
[70] Hourglass Model 2025. https://ceisphere.eu/hourglass-model.
[71] Kung, A. et al. (2025) 'Hourglass Model 2025,' Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.16574644.
[72] HEDGE-IoT D3.5 "HEDGE-IoT Technological Enablers (Final Release)"
[73] HEDGE-IoT D4.3 "HEDGE-IoT Interoperability Framework and Integrated Solution (Final Release)"
[74] HEDGE-IoT D5.3 "Full Demo Phase Report"
[75] HEDGE-IoT D7.5 "Dissemination, Exploitation and Market Exploration, Standardisation, and Community Building (Final Release)"
[76] European Commission, "The European Green Deal," COM(2019) 640 final, 2019. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52019DC0640
[77] European Commission, "A European strategy for data," COM(2020) 66 final, 2020. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0066
[78] European Union, "Regulation (EU) No 910/2014 on electronic identification and trust services for electronic transactions in the internal market, as amended by Regulation (EU) 2024/1183," 2014/2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2014/910/oj
[79] European Union, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence," Official Journal of the European Union, 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj
[80] European Union, "Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data," Official Journal of the European Union, 2023. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2023/2854/oj
[81] European Union, "Regulation (EU) 2016/679 — General Data Protection Regulation," Official Journal of the European Union, 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj
[82] European Union, "Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union," Official Journal of the European Union, 2022. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2555/oj
[83] E.DSO, "European Distribution System Operators." [Online]. Available: https://www.edsoforsmartgrids.eu/
[84] European Commission, "European Commission advances development of Common European Energy Data Space (CEEDS)," 2024. [Online]. Available: https://digital-strategy.ec.europa.eu/en/library/european-commission-advances-development-common-european-energy-data-space-ceeds
[85] ETSI, "ETSI TS 103 264 V4.1.1 — SmartM2M; Smart Applications; Reference Ontology and oneM2M Mapping," 2025. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/103200_103299/103264/04.01.01_60/ts_103264v040101p.pdf
[86] ETSI, "SAREF4ENER: SAREF for Energy Flexibility." [Online]. Available: https://saref.etsi.org/saref4ener/
[87] ETSI, "SAREF4GRID: an extension of SAREF for the Smart Grid domain." [Online]. Available: https://saref.etsi.org/saref4grid/
[88] IEC, "IEC 61970-301:2020 — Energy management system application program interface (EMS-API) — Part 301: Common information model (CIM) base," 2020. [Online]. Available: https://webstore.iec.ch/en/publication/62698
[89] IEC, "IEC 61968-11:2013 — Application integration at electric utilities — System interfaces for distribution management — Part 11: Common information model (CIM) extensions for distribution," 2013. [Online]. Available: https://webstore.iec.ch/en/publication/6199

Co-funded by the European Union

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

[90] ENTSO-E, "Common Information Model (CIM) for Grid Models Exchange." [Online]. Available: https://www.entsoe.eu/digital/common-information-model/cim-for-grid-models-exchange/
[91] ENTSO-E, "Electronic Data Interchange (EDI) Library — IEC 62325 Framework for energy market communications and European Style Market Profile." [Online]. Available: https://www.entsoe.eu/publications/electronic-data-interchange-edi-library/
[92] IEC, "IEC SRD 63417:2025 — Guidance and plan to develop smart energy ontologies and other domain-based ontologies within smart energy," 2025. [Online]. Available: https://webstore.iec.ch/en/publication/69080
[93] FIWARE Foundation, "Smart Data Models." [Online]. Available: https://www.fiware.org/smart-data-models/
[94] ETSI, "ETSI GS CIM 009 V1.9.1 — Context Information Management (CIM); NGSI-LD API," 2025. [Online]. Available: https://www.etsi.org/deliver/etsi_gs/CIM/001_099/009/01.09.01_60/gs_CIM009v010901p.pdf
[95] FIWARE, "FIWARE-NGSI v2 Specification." [Online]. Available: https://fiware.github.io/specifications/ngsiv2/latest/
[96] W3C, "JSON-LD 1.1 — A JSON-based Serialization for Linked Data," W3C Recommendation, 2020. [Online]. Available: https://www.w3.org/TR/json-ld11/
[97]W3C, "Data Catalog Vocabulary (DCAT) — Version 3," W3C Recommendation, 2024. [Online]. Available: https://www.w3.org/TR/vocab-dcat-3/
[98] Industrial Digital Twin Association, "Specification of the Asset Administration Shell — Part 1: Metamodel," IDTA Number 01001. [Online]. Available: https://industrialdigitaltwin.org/en/content-hub/aasspecifications/specification-of-the-asset-administration-shell-part-1-metamodel-idta-number-01001
[99] W3C, "RDF 1.1 Concepts and Abstract Syntax," W3C Recommendation, 2014. [Online]. Available: https://www.w3.org/TR/rdf11-concepts/
[100] W3C, "OWL 2 Web Ontology Language Document Overview," W3C Recommendation, 2012. [Online]. Available: https://www.w3.org/TR/owl2-overview/
[101] IETF, "RFC 8259 — The JavaScript Object Notation (JSON) Data Interchange Format," 2017. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc8259
[102] IETF, "RFC 4180 — Common Format and MIME Type for Comma-Separated Values (CSV) Files," 2005. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc4180
[103] Apache Software Foundation, "Apache Parquet File Format." [Online]. Available: https://parquet.apache.org/docs/file-format/
[104] W3C, "Extensible Markup Language (XML) 1.0, Fifth Edition," W3C Recommendation, 2008. [Online]. Available: https://www.w3.org/TR/xml/
[105] YAML Language Development Team, "YAML Ain't Markup Language (YAML™) Version 1.2.2," 2021. [Online]. Available: https://yaml.org/spec/1.2.2/
[106] IEC, "IEC 61850 Series — Communication networks and systems for power utility automation." [Online]. Available: https://webstore.iec.ch/en/publication/6028
[107] IEC, "IEC 62746-10-1:2018 — Systems interface between customer energy management system and the power management system — Part 10-1," 2018. [Online]. Available: https://webstore.iec.ch/en/publication/26267
[108] OpenADR Alliance, "OpenADR 2.0 and 3 Specifications." [Online]. Available: https://www.openadr.org/specification
[109] IETF, "RFC 8446 — The Transport Layer Security (TLS) Protocol Version 1.3," 2018. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc8446
[110] IETF, "RFC 9110 — HTTP Semantics," 2022. [Online]. Available: https://www.rfc-editor.org/rfc/rfc9110.html
[111] IETF, "RFC 6749 — The OAuth 2.0 Authorization Framework," 2012. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc6749

Co-funded by the European Union

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

[112]CEI, "CEI TS 13-82 — Chain 2 user profile for second-generation smart meters," referenced in CEI 0-21 Annex X. [Online]. Available: https://static.ceinorme.it/strumenti-online/doc/EstrattoAllegatoX.pdf
[113] Apache Software Foundation, "Apache Kafka Documentation." [Online]. Available: https://kafka.apache.org/documentation/
[114] OASIS, "Advanced Message Queuing Protocol (AMQP) Version 1.0," OASIS Standard, 2012. [Online]. Available: https://www.oasis-open.org/standard/amqp/
[115] Eclipse Foundation, "Eclipse Dataspace Components (EDC)." [Online]. Available: https://projects.eclipse.org/projects/technology.edc
[116] HEDGE-IoT, "D3.2: HEDGE-IoT Interfaces and Tools for Interoperability 2", 2025
[117]"Dataspace Protocol 2024-1 | IDS Knowledge Base." Accessed: Oct. 24, 2024. [Online]. Available: https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol/
[118] T. Dam, L. D. Klausner, S. Neumaier, and T. Priebe, "A Survey of Dataspace Connector Implementations," Jan. 09, 2024, arXiv: arXiv:2309.11282. doi: 10.48550/arXiv.2309.11282.
[119] eclipse-edc/Connector. (Oct. 24, 2024). Java. Eclipse Dataspace Components. Accessed: Oct. 24, 2024. [Online]. Available: https://github.com/eclipse-edc/Connector
[120] M. T. Delgado, "Eclipse Dataspace Components | projects.eclipse.org." Accessed: Oct. 24, 2024. [Online]. Available: https://projects.eclipse.org/projects/technology.edc
[121] International Data Spaces Association. (2026, January 9). IDSA Data Space Connector Report - International Data Spaces. International Data Spaces. Available at: https://internationaldataspaces.org/idsa-data-space-connector-report/

Co-funded by the European Union

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

ANNEX A — HEDGE-IOT REFERENCE ARCHITECTURE EVOLUTION

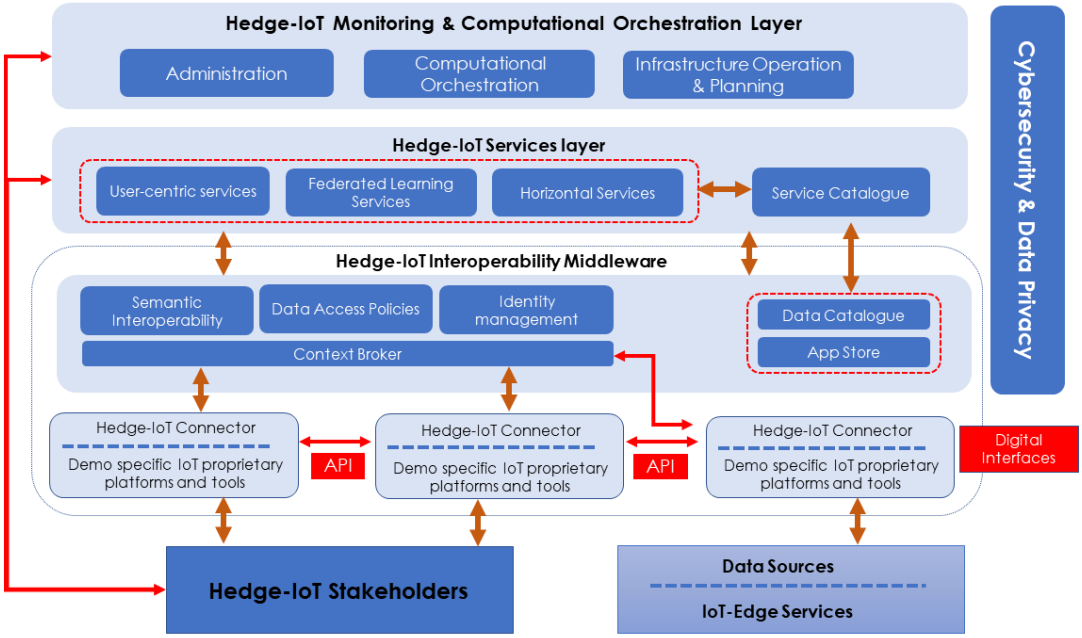

FIGURE 8 – HEDGE-IOT REFERENCE ARCHITECTURE INDICATIVE CONCEPT MODEL

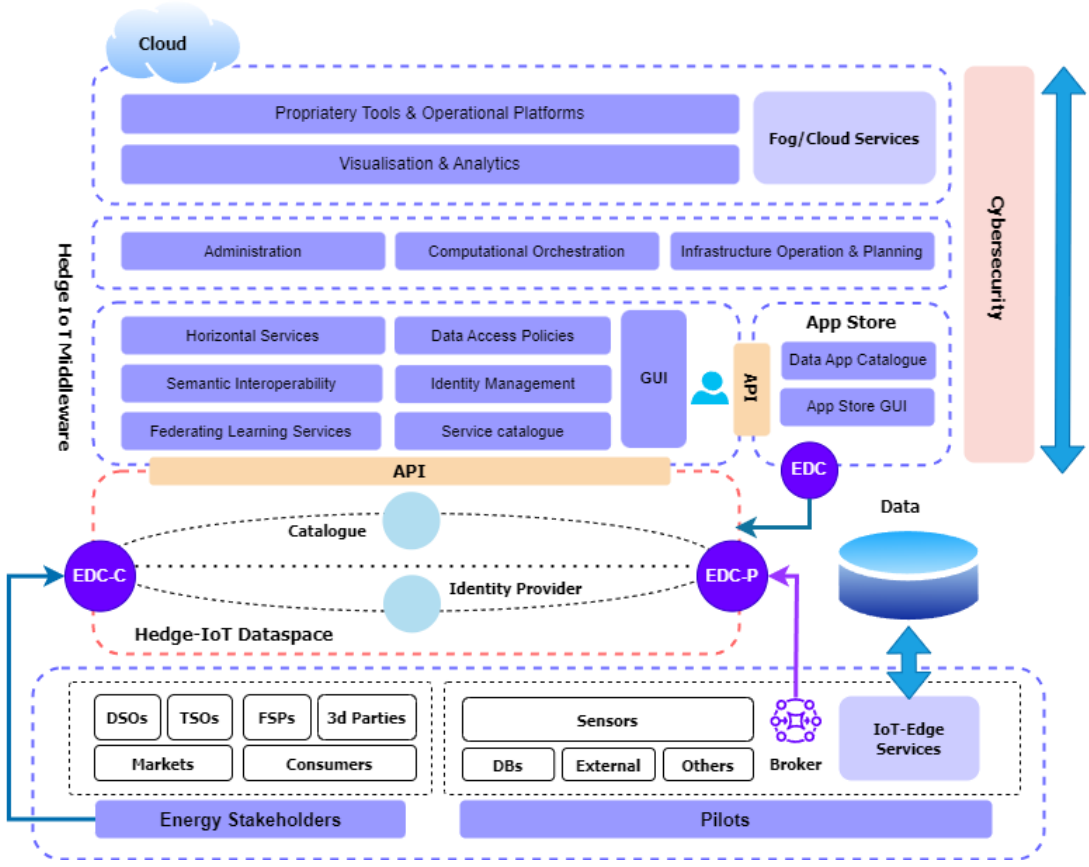

FIGURE 9 – HEDGE-IOT REFERENCE ARCHITECTURE (1ST RELEASE) - INTERMEDIATE VERSION

Co-funded by the European Union

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

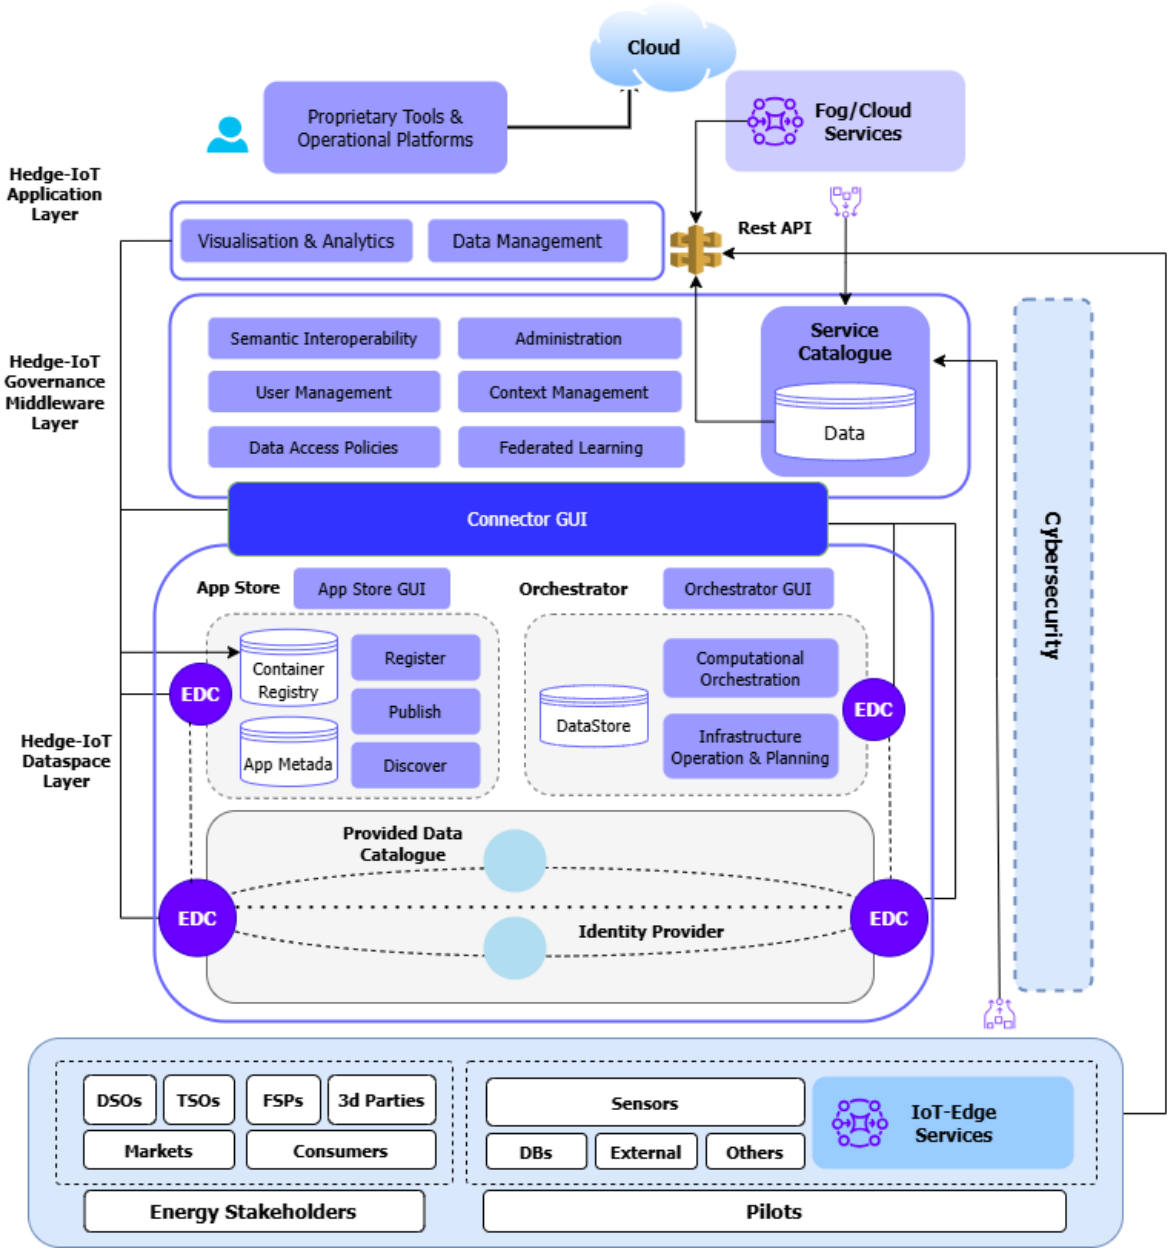

FIGURE 10 – HEDGE-IOT REFERENCE ARCHITECTURE (1ST RELEASE) – FINAL VERSION

Co-funded by the European Union

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

ANNEX B — DETAILED COMPONENT-TO-REQUIREMENTS TRACEABILITY MATRIX

The following section presents the complete components-to-requirements matrix which maps each functional requirement presented in Section 3.3 to each individual components from the architectural development view. Table 28 lists the HEDGE-IoT components and assign IDs for each of them, which are used in the traceability matrix in Table 29.

Table 28 – HEDGE-IoT Components

Package/Layer	Component ID	Component
Physical Layer Package	**C-01**	IoT Gateway
	C-02	Sensor Driver
	C-03	Local Data Store
	C-04	Controller Interface
	C-05	Digital Twin
Data Space Core Package	**C-06**	EDC Connector
	C-07	Federated Catalogue
	C-08	Identity Hub
	C-09	Contract Engine
	C-10	Policy Engine
App Store Package	**C-11**	Service Registry
	C-12	Discovery API
	C-13	Deployment Manager
	C-14	Metadata Handler
Orchestrator Package	**C-15**	Resource Allocator
	C-16	Workload Scheduler
	C-17	KubeEdge Adapter
	C-18	Monitoring Agent
Semantic Layer Package	**C-19**	Semantic technological enabler APIs
Application Layer Package	**C-20**	Forecasting Service
	C-21	Optimization Engine
	C-22	Federated Learner
	C-23	User Interfaces
	C-24	Integration APIs

Co-funded by the European Union

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

Table 29 – Detailed Components Traceability Matrix

Requirement ID	Physical Layer Package					Data Space Core Package					App Store Package				Orchestrator Package				Semantic Layer Package	Application Layer Package				
	C-01	C-02	C-03	C-04	C-05	C-06	C-07	C-08	C-09	C-10	C-11	C-12	C-13	C-14	C-15	C-16	C-17	C-18	C-19	C-20	C-21	C-22	C-23	C-24
FR-DM-01	x	x																						
FR-DM-02						x	x																	
FR-DM-03															x	x	x	x						
FR-DM-04	x		x		x																			
FR-DM-05			x				x	x			x			x								x		
FR-DM-06					x		x							x										x
FR-IOP-01						x													x					x
FR-IOP-02																			x					x
FR-IOP-03						x	x	x	x	x														
FR-IOP-04							x																	
FR-IOP-05	x					x													x					x
FR-IOP-06																								
FR-SRV-01											x	x	x	x										
FR-UI-01				x		x					x		x		x	x		x	x				x	
FR-UI-02				x		x																	x	
FR-UI-03				x		x							x					x		x	x	x	x	
FR-UI-05								x	x	x													x	
FR-OF-01																					x			
FR-OF-02																				x				
FR-OF-03																				x		x		
FR-OF-04																					x	x		
FR-OF-05																				x	x	x		
FR-OF-06					x															x	x	x		
FR-OF-07																				x		x		
FR-FM-01																								

Co-funded by the European Union

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

FR-FM-02																								
FR-FM-03																								
FR-FM-04																								
FR-FM-05																								
FR-FM-06																								
FR-FM-07																								
FR-FM-08																								
FR-GMC-01					X															X	X	X		
FR-GMC-02					X														X					
FR-AI-01																				X	X	X		
FR-AI-02																				X	X	X		
FR-AI-03																				X	X	X		
FR-MED-01																				X	X	X		
FR-MED-02																				X	X	X		
FR-MED-03																				X	X	X		

Co-funded by the European Union

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

ANNEX C – TRANSVERSAL USE-CASES DETAILED SPECIFICATIONS

Co-funded by the European Union

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

Transversal Use Case n°1:

Data exchange through HEDGE-IoT Dataspace

[Data interoperability]

1 Description of the use case

This use case describes how partners in the Hedge-IoT project can exchange data across organizational boundaries using a shared dataspace infrastructure based on the Eclipse Dataspace Connector (EDC). A data provider exposes metadata and access policies for its resources, while a data consumer discovers and retrieves data through secure, policy-compliant mechanisms. The interaction ensures data sovereignty, access control, and interoperability. This pattern can be reused across different pilots and domains to enable secure and standardized data flows.

1.1 Name of the use case

ID	*Area / Domain(s) / Zones(s)*	*Name of Use Case*
TUC-1	Data interoperability	Data exchange through HEDGE-IoT Dataspace.

1.2 Version management

Version Management			
Version No.	*Date*	*Name of Author(s)*	*Changes*
0.1	10/04/2025	Trialog	First structure based on project brainstorming.
0.2	07/05/2025	DST	1st Draft version
0.3	20/05/2025	Trialog	Feedback Provided by Trialog
0.4	23/05/2025	DST	Updates provided by DST
0.5	28/05/2025	Trialog	Feedback provided by Trialog
0.6	09/06/2025	DST	1st semi-final version
1.0	12/06/2025	Trialog	1st final version

1.3 Scope and objectives of use case

Scope and Objectives of Use Case	
Scope	The scope of this transversal use case is to leverage a shared dataspace infrastructure to facilitate secure, standardized, and scalable data exchange across the various components, stakeholders, and pilot sites involved in the Hedge-IoT project.
Objective(s)	The objectives that the use case is expected to achieve are to: • **Objective 1:** allow data exchange for an edge-cloud continuum • **Objective 2:** connect data provider and data customer
Related business case(s)	/

1.4 Narrative of use case

Narrative of Use Case
Short description
This use case describes how partners in the Hedge-IoT project can exchange data across organizational boundaries using a shared dataspace infrastructure based on the Eclipse Dataspace Connector (EDC). A data provider exposes metadata and access policies for its resources, while a data consumer discovers and retrieves data through secure, policy-compliant mechanisms. The interaction ensures data sovereignty, access control, and interoperability. This pattern can be reused across different pilots and domains to enable secure and standardized data flows.
Complete description

In the Hedge-IoT project, several partners collaborate to develop intelligent edge computing solutions for diverse sectors, including energy, mobility, and public services. These partners often need to exchange data across organizational and technical boundaries. However, sharing data between organizations raises concerns about security, control, and compliance with different regulations and usage agreements.

To address this, the project uses a shared **dataspace** infrastructure. This allows each organization to remain the owner of its data while making it available to others in a controlled and standardized way. The core of this infrastructure is the **Eclipse Dataspace Connector (EDC)**, a component that enables organizations to publish, find, and exchange data based on clearly defined policies.

Here's how it works in practice: a **data provider**, such as a company operating an edge service, wants to make a dataset available to other partners. The provider describes the dataset—what it is, how it can be used, and under which conditions—in a metadata format, and publishes it into a shared catalog managed by the dataspace. This information does not include the actual data, but tells potential users what is available and how they can request access.

A **data consumer**, for example a pilot partner developing a mobility application, browses the catalog and finds the dataset. If the dataset fits their needs, the consumer initiates a data access request. This triggers a **negotiation phase**, where the consumer's request is matched against the provider's policies (such as usage rights, contract terms, or allowed frequency). If both parties agree, the data is transferred securely using a trusted communication protocol.

The actual data never becomes publicly accessible—only those who are authorized through the dataspace infrastructure can retrieve it. Every interaction is logged and monitored to ensure compliance with the agreed rules.

This setup ensures **data sovereignty** (each partner controls how their data is used), **interoperability** (partners use common standards), and **security** (data is exchanged securely and only between trusted parties).

This kind of interaction is expected to be replicated in multiple use cases within the project—whether it's exchanging energy grid information, mobility patterns, or sensor data—and can also serve as a template for data exchange in future cross-domain projects.

This transversal use case could be split into different scenarios:

- Scenario 1: Use of the dataspace by a data producer
- Scenario 2: Use of the dataspace by a data customer
- Scenario 3: Metadata Discovery and Planning
- Scenario 4: Federated Service Chaining

Sc.1 Use of the dataspace by a data producer - Description:

- This scenario describes how a data provider makes its dataset or service available within the dataspace. The provider prepares the asset, defines the associated metadata and access policies, and publishes it through its local EDC connector. The asset becomes discoverable by other parties via the federated catalog, allowing compliant and secure access negotiations .Additionally, by exposing curated datasets through the dataspace, data producers contribute to cross-pilot AI training efforts, enabling other partners to discover and evaluate datasets suitable for model development.

Sc.2 Use of the dataspace by a data customer - Description:

- In this scenario, a data consumer interacts with the dataspace to discover and access data assets shared by other parties. The consumer queries the catalog, evaluates metadata and policy terms, and initiates a contract negotiation through its EDC connector. Upon agreement, the data is securely transferred according to the defined usage rules.

Sc.3 Metadata Discovery and Planning - Description:

- A partner uses the dataspace not to directly retrieve data, but to discover which datasets or services are available, including their conditions of use, data formats, and applied semantic vocabularies.

This scenario highlights the catalog and *resource discovery* capabilities of the dataspace, enabling informed planning, semantic mapping, and potential future agreements.
  - Particularly useful in the design or pre-integration phase.
  - Reduces effort in bilateral discussions, as the catalog serves as a shared point of reference.
  - Reinforces semantic interoperability goals (e.g., Task 4.3).

Sc.4 Federated Service Chaining - Description:

- A software component (e.g., an orchestrator or optimization engine) uses the dataspace to access services or modules provided by other partners, in a dynamic and composable way. For instance, an edge node may call a forecasting module hosted in the cloud by another partner, sending data via the dataspace and receiving a processed result in return.

1.5 Key performance indicators (KPI)

ID	*Name*	*Description*	*Reference to mentioned use case objectives*
KPI1	Adoption rate of the dataspace	Number of partners successfully integrated with the dataspace infrastructure	Supports the objective of creating a federated, reusable integration layer across pilots
KPI2	Data asset discoverability	Number of data assets made available and searchable via the federated catalog	Enables semantic discovery and contributes to interoperability across work packages
KPI3	Cross-WP or cross-pilot usage patterns	Number of use cases or pilots that rely on the dataspace to communicate and share data	Validates the transversal nature and reusability of the dataspace beyond vertical or isolated implementations

1.6 Use case conditions

Use case conditions
Assumptions
• All relevant partners have or will adopt a compatible version of the Eclipse Dataspace Connector (EDC), or are integrated via proxies. • All participating organizations agree to define and enforce data usage policies according to IDS principles. • Semantic vocabularies used across pilots (e.g., SAREF,IEC CIM) are sufficiently aligned to enable resource discoverability. • A common trust framework will be setup by DST for identity and access control is in place
Prerequisites
• EDC instances are deployed and reachable by partner systems. • Each partner has registered at least one data asset in its local catalog with appropriate metadata. • Policies for access control and data usage are defined and operational. • Communication between connectors is secured and authorized (e.g., via certificates or trusted endpoints). • At least two pilots or services require cross-organizational data exchange.

1.7 Further Information to the use case for classification / mapping

Classification Information
Relation to other use cases
Computational Orchestration
Level of depth
Prioritisation

Generic, regional or national relation
Italian Pilot, Greek Pilot, Dutch Pilot, Portuguese Pilot, Slovenian Pilot, Finnish Pilot
,Nature of the use case
System Use Case, Transversal Use Case
Further keywords for classification
Dataspace, Semantic Interoperability, Connector, Middleware

1.8 General Remarks

General Remarks

2 Diagrams of use case

Diagram(s) of use case

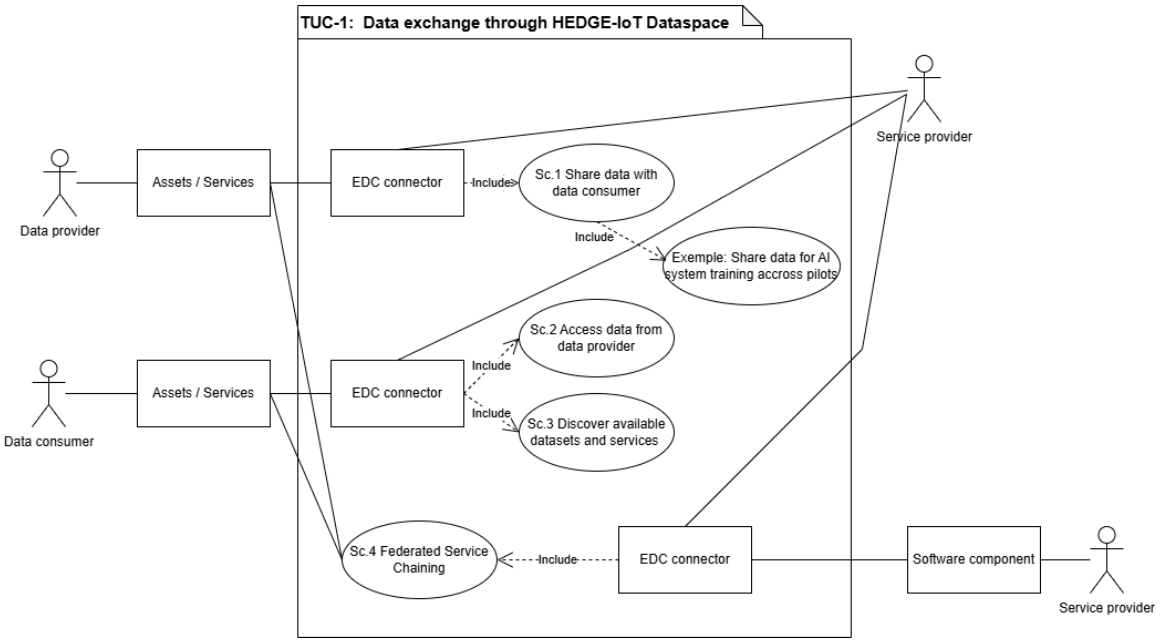

Figure 1: UML use case diagram

3 Technical details

3.1 Actors

Actors			
Actor Name	***Actor Type***	***Actor Description***	***Further information specific to this use case***
Data Provider	Business Actor	Any pilot or service that exposes data (e.g., sensor data, flexibility info, forecasts)	Publishes data assets and defines access policies. Initiates sharing via its EDC connector. {Pilot}
Data Consumer	Business Actor	Any pilot or service that requests data from other partners	Discovers assets, negotiates usage terms, and consumes data via its EDC connector. {Pilot,Service}

EDC Connector	Logical Actor	Eclipse Dataspace Connector instance deployed by each participant	Manages metadata, policies, negotiation, and transfer on behalf of the provider/consumer. {DST}
Dataset Catalog	Logical Actor	Central or distributed catalog for publishing and discovering data assets	Enables data discovery and lookup of published assets in the dataspace. {DST}

3.2 References

References						
No.	***Reference Type***	***Reference***	***Status***	***Impact on use case***	***Originator / organisation***	***Link***
1	GitHub Repository	Eclipse Dataspace Connector (EDC) Documentation	Online	Provides docs for EDC	Eclipse Foundation	https://github.com/eclipse-edc1

4 Step by step analysis of use case

4.1 Overview of scenarios

Scenario conditions						
No.	**Scenario name**	**Scenario description**	**Primary actor**	**Triggering event**	**Pre-condition**	**Post-condition**
1	Use of the dataspace by a Data Provider	A partner exposes a data asset in the dataspace with metadata and access policy	Data Provider	A dataset or service is ready to be shared	EDC connector is deployed, and metadata is defined	Data asset is published and available for discovery
2	Use of the dataspace by a Data Consumer	A partner discovers a data asset, negotiates usage, and retrieves it securely	Data Consumer	A need for external data arises	Catalog is populated and connectors are trusted	Data is transferred securely to the consumer
3	AI Training across Pilots	Data from one pilot is shared with another to support training of AI models	Data Consumer	A training pipeline needs external data	Relevant dataset is available, and policy permits reuse	AI component has access to shared training data
4	Federated Service Chaining	A service (e.g., orchestrator) triggers execution of remote services hosted by other partners via the dataspace	Orchestrator / Service	A workflow requires a remote component's output	Remote service is available and can be invoked via EDC	Result is returned to the orchestrating service
5	Metadata Discovery and Capability Advertising	A partner explores available assets to identify future collaboration or integration opportunities	Data Consumer	Integration planning or exploratory phase begins	Metadata has been published in the catalog	Consumer identifies useful assets or services

4.2 Steps – Scenarios

Scenario								
Scenario name:		Use of the Dataspace by a Data Provider						
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
1	New data available	Create data asset	The data provider creates or selects a dataset or service output intended for external sharing.	CREATE	Data Provider	Data Provider	Inf.01	DE.3
2	Asset identified	Define metadata	The provider defines metadata describing the asset (format, structure, purpose, etc.).	CREATE	Data Provider	Data Provider	Inf.02	MD.5
3	Access control needed	Configure access policy	The provider defines access policies such as allowed consumers, usage rights, and expiration terms.	CREATE	Data Provider	Data Provider	Inf.03	MD.1
4	Asset ready to publish	Register asset in connector	The data asset and metadata are registered in the local EDC connector catalog.	REPORT	Data Provider	EDC Connector	Inf.04	MD.6
5	Publication triggered	Publish asset to dataspace	The asset becomes discoverable in the federated dataspace via the connector's catalog endpoint.	REPORT	EDC Connector	Federation Catalog	Inf.05	MD.3
6	Idle	Await request from consumer	No further action until another actor (consumer) discovers and requests the asset.	TIMER	Data Provider	Data Provider	--	--

Scenario								
Scenario name:	Use of the Dataspace by a Data Consumer							
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
1	Need for external data arises	Discover data asset	The consumer searches the dataspace catalog for relevant data assets.	GET	Catalog	Data Consumer	Inf.06	MD.6
2	Matching asset found	Request asset access	The consumer selects a dataset and initiates a data usage request through its EDC connector.	EXECUTE	Data Consumer	Data Provider	Inf.07	DE.3
3	Negotiation starts	Negotiate contract	The EDC connectors negotiate a usage agreement (contract offer, response, confirmation).	EXECUTE	EDC Connector	EDC Connector	Inf.08	DE.1
4	Contract accepted	Authorize access	Access is granted according to policy and contract terms.	REPORT	EDC Connector	Data Consumer	Inf.09	CR.3
5	Transfer initialized	Retrieve data	The data is securely transferred from the provider to the consumer.	GET	Data Provider	Data Consumer	Inf.10	DE.3
6	Data received	Confirm transaction	The consumer confirms successful receipt and logs the transaction outcome.	REPORT	Data Consumer	EDC Connector	Inf.11	CR.2

Scenario								
Scenario name:	AI Training Across Pilots							
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
1	Training data needed	Discover training dataset	The consumer explores the dataspace catalog to find datasets suitable for model training.	GET	Catalog	Data Consumer	Inf.12	MD.6
2	Dataset selected	Request data access	The consumer requests access to the selected training dataset.	EXECUTE	Data Consumer	Data Provider	Inf.13	DE.2
3	Policy evaluation triggered	Negotiate usage agreement	The connectors negotiate the contract terms specific to AI training usage (e.g., retention, reuse).	EXECUTE	EDC Connector	EDC Connector	Inf.14	DE.1
4	Access granted	Transfer training data	Upon contract approval, the training data is transferred securely from provider to consumer.	GET	Data Provider	Data Consumer	Inf.15	DE.3
5	AI model development ongoing	Use data for model training	The consumer uses the dataset locally or in cloud environment to train its AI models.	EXECUTE	Data Consumer	Data Consumer	Inf.16	—
6	Optional feedback	Share model or metadata	Optionally, trained model metadata or feedback can be shared back with the provider or the platform.	REPORT	Data Consumer	Data Provider / Dataset Catalog	Inf.17	MD.4

Scenario								
Scenario name:		Federated Service Chaining						
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
1	Execution plan initialized	Discover available services	The orchestrator queries the dataspace to find remote services (e.g., optimizers, simulators).	GET	Dataset Catalog	Orchestrator	Inf.18	MD.6
2	Suitable service found	Request remote service execution	The orchestrator selects the target service and initiates execution via the dataspace.	EXECUTE	Orchestrator	Service Provider	Inf.19	—
3	Request received	Validate access and execute	The provider verifies the request and executes the requested computation or process.	EXECUTE	Service Provider	Service Provider	Inf.20	CR.1
4	Processing completed	Return execution result	The result of the remote computation is returned via the dataspace to the orchestrator.	REPORT	Service Provider	Orchestrator	Inf.21	—
5	Log transaction	Confirm transaction and record trace	Both parties log the operation for traceability and potential auditing.	REPORT	EDC Connectors	EDC Connectors	Inf.22	CR.2

Scenario								
Scenario name:		Metadata Discovery and Capability Advertising						
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
1	Integration or planning initiated	Explore available metadata	A partner accesses the dataspace catalog to explore available data assets and services.	GET	Dataset Catalog	Data Consumer / Planner	Inf.23	MD.6
2	Discovery of potential match	Analyze metadata	The consumer examines information such as data format, access policy, vocabularies used, etc.	GET	Dataset Catalog	Data Consumer	Inf.24	MD.5
3	Need for clarification	Request additional metadata	If necessary, the consumer contacts the provider (via connector) to clarify or retrieve extended info	EXECUTE	Data Consumer	Data Provider	Inf.25	—
4	Info provided	Share additional details	The provider responds with additional documentation or technical specifications	REPORT	Data Provider	Data Consumer	Inf.26	MD.4
5	Decision phase	Log discovery / plan integration	The consumer logs relevant metadata for planning future data usage or service integration	REPORT	Data Consumer	Data Consumer	Inf.27	—

5 Information exchanged

Information exchanged			
Information exchanged (ID)	***Name of information***	***Description of information exchanged***	***Requirement, R-IDs***
inf.01	Raw or processed dataset	Dataset created by a provider, may include time series, measurements, or computed results.	DE.3
inf.02	Asset metadata	Descriptive information about the data asset: type, format, update frequency, unit, etc.	MD.5
inf.03	Usage policy	Rules defined by the provider about who can access the data, under what conditions and purposes	MD.1
inf.04	Metadata + policy	The combination of asset metadata and usage policy used for registration in the local catalog.	MD.6
inf.05	Federated metadata	Metadata exposed to the federated dataspace catalog, discoverable by external consumers.	MD.1
inf.06	Asset list, metadata	List of available data assets retrieved by a consumer during discovery phase.	MD.6
inf.07	Access request, usage intent	A formal request to access a specific dataset, with declared purpose of use.	DE.3
inf.08	Contract offer, policy details	The contract terms proposed between provider and consumer, including allowed actions, pricing, duration, etc.	DE.1
inf.09	Contract confirmation, access token	Signed agreement and authentication material allowing access to the requested asset.	CR.3
inf.10	Dataset (payload)	The actual data transferred from provider to consumer, according to the contract.	DE.3
inf.11	Transfer status, acknowledgment	Confirmation of successful delivery and receipt of the dataset.	CR.2
inf.12	Metadata on training datasets	Information about datasets suitable for AI model training.	MD.6
inf.13	Access request, intended use	Consumer expresses interest in using dataset for AI training and provides justification.	DE.2
inf.14	Contract offer, policy details	Negotiated terms for use of dataset in training context, possibly stricter than general access.	DE.1
inf.15	Training dataset (input data)	Transferred data used for training AI models.	DE.3
inf.16	Internal model development	Not externally transferred, but part of consumer's local processing (e.g., neural network training).	—
inf.17	Trained model info, evaluation metrics	Feedback or metadata on trained models shared optionally with the original data provider.	MD.4
inf.18	Service metadata, capabilities	Catalogued information about available services and their invocation parameters.	MD.6

inf.19	Input parameters, invocation request	The parameters sent from an orchestrator to a remote service provider via the dataspace.	—
inf.20	Execution results	Output data generated by the remote service in response to the request.	CR.1
inf.21	Output data, execution status	Returned result of the invoked service, sent to the orchestrator.	—
inf.22	Transaction log, timestamps	Audit information logged by the connectors on both sides of the interaction.	CR.2
inf.23	Asset descriptions, metadata	High-level discovery information retrieved from the catalog during integration planning.	MD.6
inf.24	Semantic info, usage conditions	Includes ontologies, tags, units of measure, allowed use cases etc.	MD.5
inf.25	Clarification request, metadata query	A message from a consumer asking for more detail or technical info from the provider.	—
inf.26	Dataset schema, ontology references	Additional descriptive material provided by the data provider.	MD.4
inf.27	Internal planning data	Information recorded by the consumer for later analysis or integration planning.	—

6 Requirements

Data Exchange		
Categories ID	**Category name for Requirement**	**Category description**
DE	Data Exchange	Requirements for the exchange of data between connectors via the dataspace.
Requirement ID	**Requirement Name**	**Requirement description**
DE.1	Max negotiation attempt	Connectors must retry negotiation max 3 times before failing.
DE.2	Metadata must be discoverable	Metadata of published data must be indexed and retrievable via the Broker.
DE.3	Push and pull supported	Both push and pull mechanisms must be available for data transfers.

Metadata Requirements		
Categories ID	**Category name for requirements**	**Category description**
MD	Metadata Requirements	Requirements for metadata structure, discoverability, and lifecycle within the dataspace.
Requirement ID	**Requirement name**	**Requirement description**
MD.1	Metadata Publication Mandatory	Every dataset published by a provider must include a metadata description accessible via broker.
MD.4	Metadata Update Trigger	Metadata must be updated if the associated dataset is modified or deprecated.
MD.5	Metadata Minimum Attributes	Metadata must contain at least title, provider ID, data format, licensing, and update date.

MD.6	Metadata Retrieval Availability	Metadata must be queryable at all times through the dataspace discovery component.

Connector Requirements		
Categories ID	**Category name for requirements**	**Category description**
CR	Connector Requirements	Technical capabilities expected from each EDC connector instance.
Requirement ID	**Requirement name**	**Requirement description**
CR.1	DSP Compliance	The connector must implement the IDSA Data Space Protocol (DSP).
CR.2	Logging enabled	Every connector must log events and transactions for traceability.
CR.3	Identity-based auth	The connector must use the dataspace's Identity Provider for authorization.

7 Common Terms and Definitions

Common Terms and Definitions	
Term	**Definition**
EDC (Eclipse Dataspace Connector)	An open-source component that enables secure, policy-based data exchange across organizations in line with IDS and GAIA-X principles.
IDS (International Data Spaces)	A reference architecture model for trusted data exchange between entities, ensuring sovereignty and compliance.
Catalog	A list of available data assets or services, usually including metadata such as format, owner, and access policy.
Data Provider	An actor that owns a data asset and makes it available through the dataspace.
Data Consumer	An actor that requests and uses data assets shared by other parties in the dataspace.
Metadata	Descriptive information about a data asset, including format, purpose, owner, and usage conditions.
Usage Policy	A set of rules defined by the data provider to govern access and use of the shared asset.
Contract Negotiation	The automated process by which a consumer and provider agree on the terms for data access and usage.
Dataspace	A federated infrastructure for controlled data exchange, based on principles like sovereignty, traceability, and interoperability.
Orchestrator	A component that manages workflows or services by invoking remote or local functions based on data availability or triggers.

Transversal Use Case n°2:

Orchestrate the coordination, management, and execution of energy services across the computational continuum

[Computational interoperability]

1 Description of the use case

1.1 Name of the use case

ID	Area / Domain(s) / Zones(s)	Name of Use Case
TUC-2	Computational interoperability	Orchestrate the coordination, management, and execution of energy services across the computational continuum

1.2 Version management

Version Management			
Version No.	Date	Name of Author(s)	Changes
0.1	10/04/2025	Trialog	The first structure based on project brainstorming.
0.2	5/05/2025	TUC	Phase 1 Transversal Use Case Definition
0.3	26/05/2025	TUC	First complete version
0.4	6/06/2025	TUC	Updated version
1.0	13/06/2025	Trialog	1st final version
2.0	27/03/2026	TUC	2nd version Minor updates on the actor table and the KPIs.

1.3 Scope and objectives of use case

Scope and Objectives of Use Case	
Scope	Coordinating distributed computational tasks for energy services across edge-to-cloud systems, with goals of ensuring responsiveness and cost-effective data exchange, including minimized latency and bandwidth usage. We consider containerized energy services that are data space compliant through eclipse data connector
Objective(s)	The goals that the use case is expected to achieve are to: • **Objective 1:** Dynamically allocate computational resources based on energy services demand across edge, fog, and cloud layers to maintain efficiency and responsiveness. • **Objective 2:** Minimize data transfer overhead in federated AI services by efficiently managing and optimizing the hyperparameters of the learning process. • **Objective 3:** Ensure the efficient update of energy services from cloud to edge avoiding execution disruption
Related business case(s)	• Predictive and real-time congestion management (non-AI) • Forecast energy production and consumption for energy communities or residentials or commercial buildings (AI)

1.4 Narrative of use case

Narrative of Use Case
Short description
The use case focuses on enabling the automated coordination, management, and execution of computational tasks through a computational orchestrator. The orchestrator leverages swarm-based algorithms to optimize resource usage, ensuring both computational and communication efficiency. It integrates with non-AI Energy Services to manage deployment, coordination, and resource allocation, and with AI Federated Services to support hyperparameter tuning and training optimization. Additionally, it enables services rolling out at the edge, allowing automated updates and deployment of new versions. Integration with the Eclipse Data Space Connector ensures compliance with data space standards and secure, interoperable data and service exchange.
Complete description

This transversal use case could be split into three scenarios:

- **Scenario 1: Energy services orchestration at edge for responsiveness and geographic redundancy (Sc.1)**

Containerized energy services are deployed across edge-fog-cloud distributed infrastructure overlapping the smart grid. The services and the infrastructure available computing nodes are registered with the computational orchestrator. The orchestrator continuously monitors service locations, resource availability, and task assignments. An integrated Kubernetes component handles the initial task allocation based on current resource availability and predefined configurations. It also supports live monitoring of service status and system resources. When predefined events occur, such as violations of service level agreements or policy conditions, the orchestrator responds by executing a swarm-based optimization algorithm to determine which services should be migrated to other nodes. During service migration, the persistent state and data of each service must also be transferred, and their connectivity via the Eclipse Data Space Connector must be maintained to ensure secure and interoperable data exchange. Therefore, it ensures service responsiveness and geographic redundancy.

- **Scenario 2: Federated AI-driven energy services orchestration for cost-effective data exchanges (Sc.2)**

The orchestrator manages federated learning processes initiated by AI services deployed across edge, and fog/cloud nodes. The federated architecture may follow either a hierarchical or peer-to-peer model. All participating nodes are registered with the orchestrator, providing metadata on their computational capabilities and availability. Based on service-specific requirements, the orchestrator can cluster nodes to enhance training efficiency. It also performs hyperparameter tuning and training optimization using heuristic-based algorithms, ensuring efficient use of distributed resources at edge and minimizing data exchange overhead among edge and fog/cloud nodes.

- **Scenario 3: Energy service rolling out at edge (Sc.3)**

The orchestrator can support service providers such as DSO to automatically roll out new versions for energy services for the consumers. The data models (packages or files) are sent through the Eclipse Data Space Connector. It detects available updates for application components, manages versioning and handles service interruption during updates.

1.5 Key performance indicators (KPI)

ID	*Name*	*Description*	*Reference to mentioned use case objectives*
KPI1	Decrease of data exchange between nodes	Measures the percentage reduction in communication data volume due to orchestrator optimization (%)	Objective 2
KPI2	Savings in network bandwidth and lower latency	Measures the percentage savings in network bandwidth usage and percentage reduction in service response latency through localized processing and optimized resource allocation (%)	Objective 1 and Objective 3

1.6 Use case conditions

Use case conditions
Assumptions
Sc.1 • The services are data space compliant • Edge nodes with sufficient computational resources are available at the deployment sites. **Sc.2** • Federated AI services for the energy sector are assumed to be in place, utilizing either hierarchical or peer-to-peer architectures. • Training nodes are assumed to have access to local datasets that are not shared with other nodes and must remain private due to data sovereignty or regulatory constraints.

• A communication infrastructure exists between nodes, using a data space connector or an equivalent secure communication protocol. • In scenarios involving hyperparameter tuning, each node can perform training using distinct hyperparameter configurations **Sc.3.** • All existing services are assumed to be registered in a federated catalog or app store to support service discovery and orchestration. • Federated models are assumed to be updatable through ongoing training or fine-tuning at participating nodes
Prerequisites

Sc. 1
• All services intended for deployment on edge nodes are assumed to be containerized (e.g., using Docker).
• An existing workload management platformed is assumed to be in place for managing edge nodes and monitoring deployed services (e.g. Kubernetes-based platform with KubeEdge).
Sc.2
• Each service exposes a communication interface and includes a data space connector, enabling the orchestrator to:
  ○ Retrieve metadata and status information about services and edge nodes
  ○ Configure optimization parameters such as hyperparameters, selected training nodes, and other task-specific settings.
Sc. 3
• Edge nodes that deploy applications from the catalog or use federated models are assumed to be registered with the orchestrator
• The orchestrator is assumed to have access to application images and federated models and is notified when updated versions become available.
• The orchestrator is assumed to have the capability and permission to manage updates to federate models.

1.7 Further Information to the use case for classification / mapping

Classification Information
Relation to other use cases
/
Level of depth
/
Prioritisation
/
Generic, regional or national relation
/
Nature of the use case
System Use Case, Transversal (system) Use Case
Further keywords for classification
/

1.8 General Remarks

General Remarks
/

2 Diagrams of use case

Diagram(s) of use case

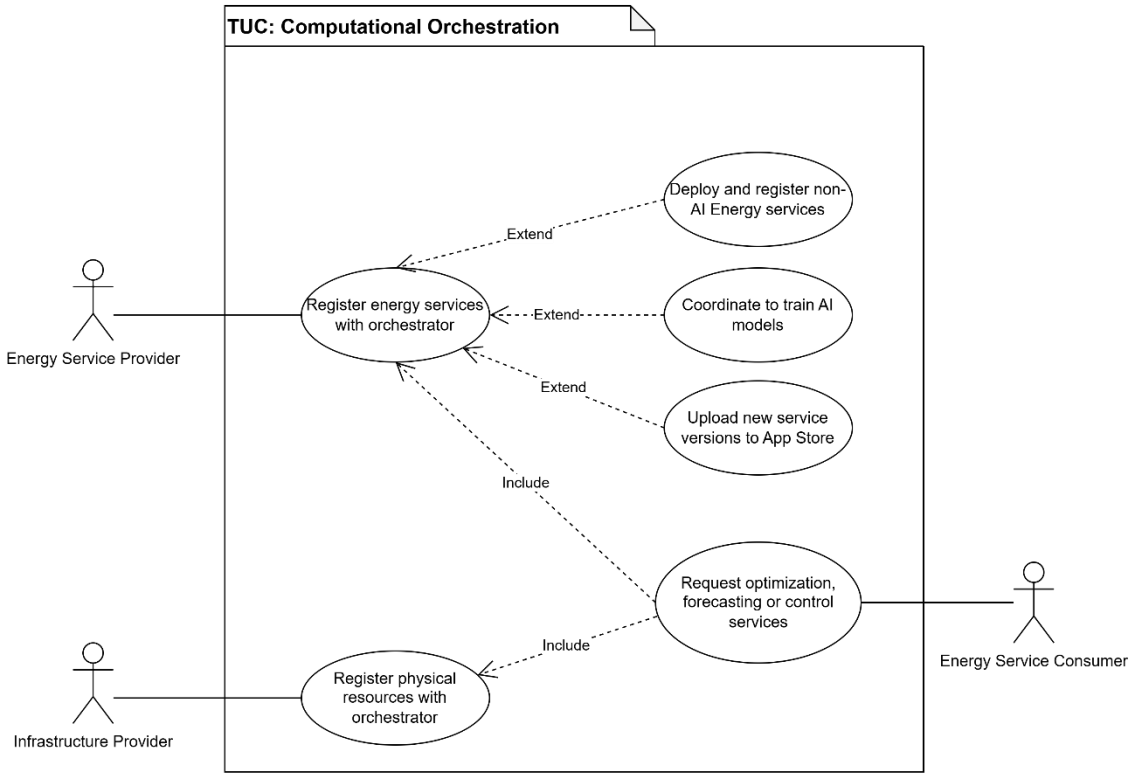

Figure 1: UML use case diagram

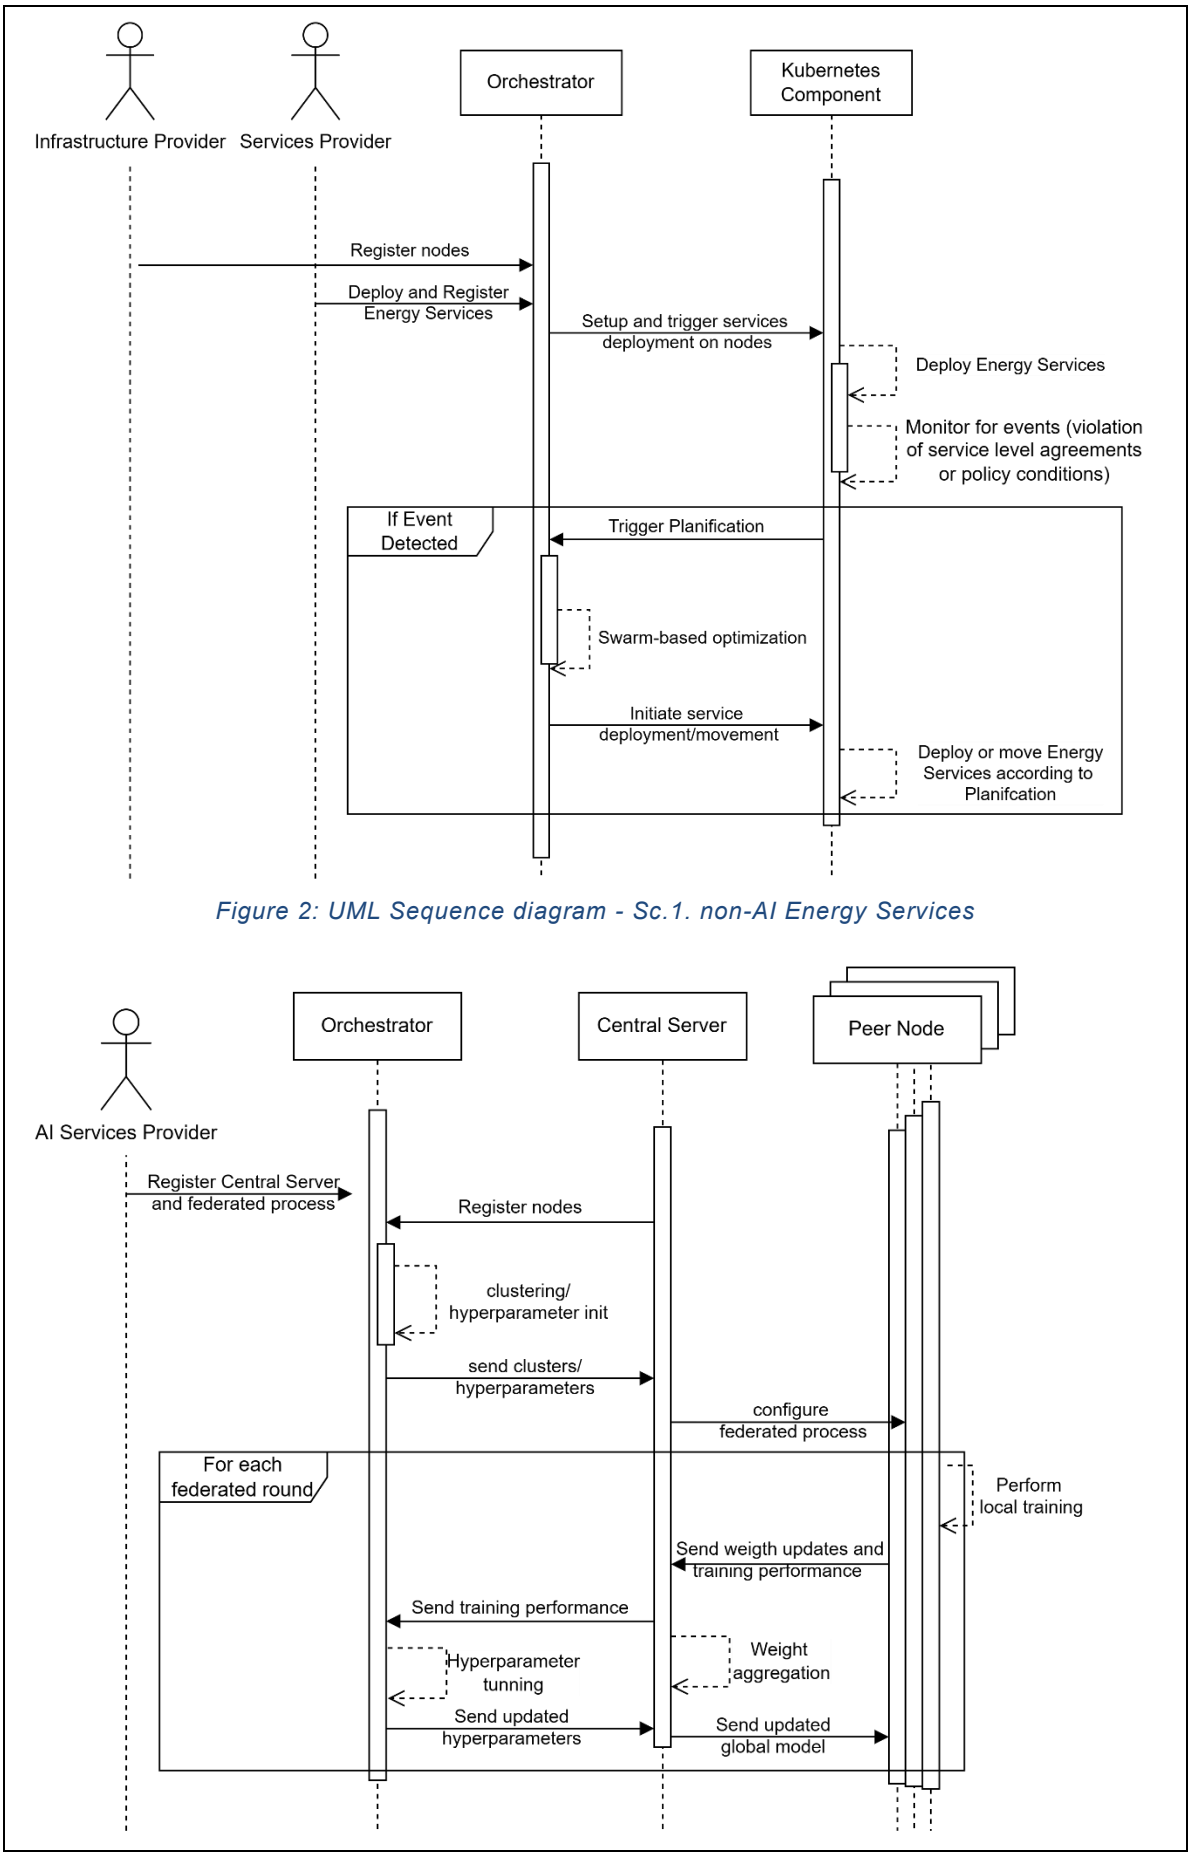

Figure 2: UML Sequence diagram - Sc.1. non-AI Energy Services

Figure 3: UML Sequence diagram - Sc.2. AI Energy Services

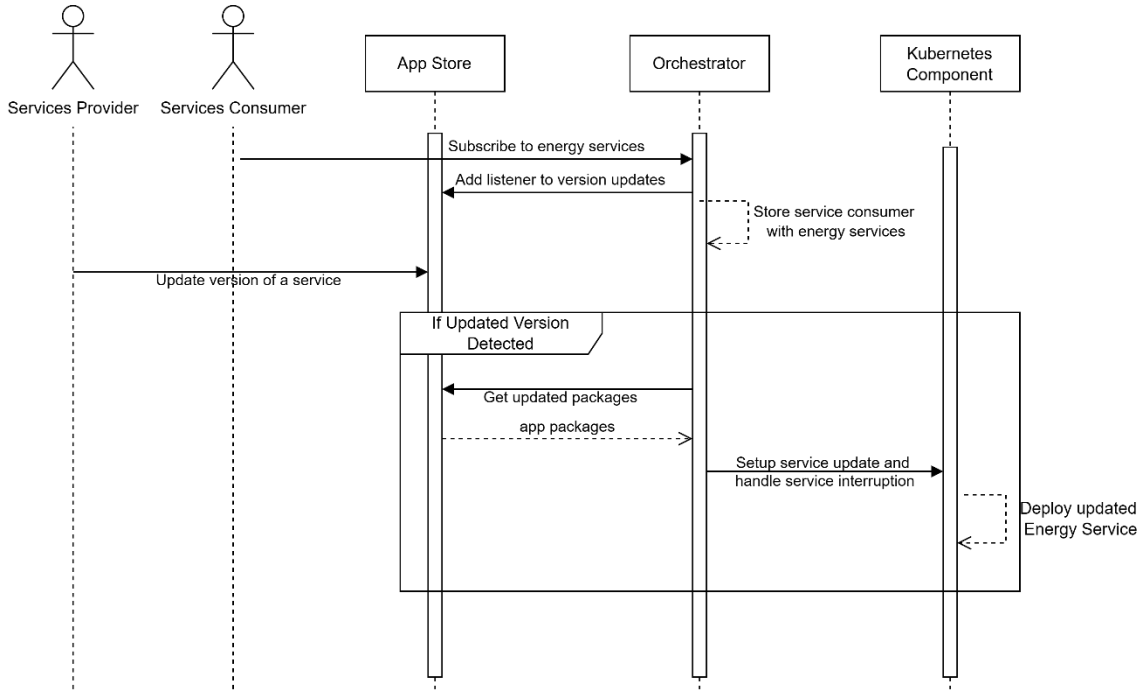

Figure 4: UML Sequence diagram - Sc. 3 Energy Services rolling out at edge

3 Technical details

3.1 Actors

Actors			
Actor Name	***Actor Type***	***Actor Description***	***Further information specific to this use case***
Energy Service Provider	Business actor	Provides energy forecasting, or congestion management services	- ARETI DSO (Italian Pilot – Sc.1) Congestion prediction and optimal power flow service - INESC Energy Community Service Provider (Portuguese Pilot – Sc.2) - TAU (Finland Pilot - Sc.3) State Estimation Service
Infrastructure Provider	Business actor	Owns or operates the computational resources and communication network infrastructure	- ARETI & DST (Italian Pilot Sc.1) - TUC & INESC (Portuguese Pilot - Sc. 2) - TUC & TAU (Finland Pilot – Sc. 3)
Energy Service Consumers	Operator	Needs the optimization, forecasting, or control services	Energy communities, prosumers, DSO, customers
Computational Orchestrator	Logical actor	Manages registration, scheduling, and coordination of services and nodes.	Orchestration logic
EDC Connector	Logical actor	Used for communication between services	Ensure connectivity and secure data exchange

3.2 References

References						
No.	*Reference Type*	*Reference*	*Status*	*Impact on use case*	*Originator / organisation*	*Link*

4 Step by step analysis of use case

4.1 Overview of scenarios

Scenario conditions						
No.	**Scenario name**	**Scenario description**	**Primary actor**	**Triggering event**	**Pre-condition**	**Post-condition**
Sc. 1	Energy services orchestration at edge for responsiveness and geographic redundancy	A distributed orchestration system leverages Kubernetes and swarm-based optimization to dynamically allocate and migrate containerized energy services across edge-fog-cloud infrastructure	Energy service provider	An energy services needs to be deployed/relocated due to violations of service level agreements or policy conditions	The service is containerized and data space compliant	SLA compliance, service is successfully relocated, communication links with other services are preserved
Sc. 2	Federated AI-driven energy services orchestration for cost-effective data exchanges	The orchestrator coordinates federated learning across edge, fog, and cloud nodes using hierarchical or peer-to-peer models, optimizing node clustering, hyperparameters, and training efficiency through heuristics while minimizing data exchange overhead.	Energy service provider	Initiate the model training for the federated AI service	The federated service exposes a communication interface, and the training process is configurable	Average data exchange is reduced compared with the baseline
Sc. 3	Energy service rolling out at edge	The orchestrator enables automated rollout of updated energy service versions for consumers on behalf of service providers	Energy service provider	A new version of an energy service is available	Service consumers are registered with the orchestrator; Orchestrator has access to updated versions through EDC	Updated service version is deployed

4.2 Steps – Scenarios

Scenario								
Scenario name:		**Sc. 1**						
Step No.	***Event***	***Name of process/ activity***	***Description of process/ activity***	***Service***	***Information producer (actor)***	***Information receiver (actor)***	***Information Exchanged (IDs)***	***Requirement, R-IDs***
St.1	New energy services available	Register new energy service	Service provider registers a service, including its policies and service level agreements	CREATE	Energy service provider	Orchestrator	Inf.01	D.1
St.2	New infrastructure resource available	Register new nodes	Infrastructure provider registers new available edge nodes	CREATE	Infrastructure provider	Orchestrator	Inf.02	D.2
St.3	New service or constraint violation	Trigger planification	Edge nodes and services are monitored; Orchestrator listens for new events.	REPORT	Kubernetes Component	Orchestrator	Inf.03	QoS.1 QoS.2
St.4	Planification triggered	Run swarm-based optimization	Orchestrator runs optimization algorithm based on monitoring and service information	EXECUTE	Orchestrator	-	-	QoS.1 QoS.2
St.5	Optimization completed	Initiate service deployment/movement	Orchestrator sends the optimization results to Kubernetes platform	CHANGE	Orchestrator	Kubernetes Component	Inf.04	Conf.2
St.6	Kubernetes Component receives new planification	Deploy/relocate services	Kubernetes component handles deployment of the services and preserves the links between services	EXECUTE	Kubernetes Component	Kubernetes Cluster	Inf.05	Conf.1 QoS.2

Scenario								
Scenario name:		**Sc. 2**						
Step No.	***Event***	***Name of process/ activity***	***Description of process/ activity***	***Service***	***Information producer (actor)***	***Information receiver (actor)***	***Information Exchanged (IDs)***	***Requirement, R-IDs***
St.7	New initiated federated process	Register new federated process	Service provider registers a federated process with its configuration and hyperparameters	CREATE	AI Energy service provider	Orchestrator	Inf.06	D.1 Conf.2
St.8	Federated process registered	Register edge nodes	Available training edge nodes are registered with orchestrator	CREATE	Central Server/Edge Node/ AI Service Provider	Orchestrator	Inf.07	D.2
St.9	Federated process triggered	Initialization of the federated process	Orchestrator performs clustering/initialization for optimization algorithm	EXECUTE	Orchestrator	-	-	Conf.2
St.10	Next federated round/ optimization needed	Orchestrator performs optimization	Orchestrator runs optimization algorithm for the configured hyperparameters	REPEAT (5, 6)	Orchestrator	-	-	Conf.2
St.11	Optimization completed	Sends optimization results	Orchestrator sends the optimization results to edge nodes	CHANGE	Orchestrator	Edge nodes	Inf.08	D.2 Conf.2
St.12	New hyperparameters/configuration available	Perform local training and send training performance	Edge nodes perform training based on the received configuration/hyperparameters and send performance results back	REPORT	Edge nodes	Orchestrator	Inf.09	D.2 Conf.2

Scenario name:		**Sc. 3**						
Step No.	***Event***	***Name of process/ activity***	***Description of process/ activity***	***Service***	***Information producer (actor)***	***Information receiver (actor)***	***Information Exchanged (IDs)***	***Requirement, R-IDs***
St.13	New subscriber to energy services	Register energy service consumer	Service consumer requests energy service	CREATE	Energy service consumer	AppStore / Orchestrator	Inf.02	D.1 D.2
St.14	New version of a service / federated model available	Detect new updated version	Orchestrator checks the consumers of the service and gets the updated package through data space connector	GET	App Store / Federated Catalogue	Orchestrator	Inf.10	D.1 D.3 Conf. 2
St.15	New app package / federated model received	Handle version update	Handle possible service interruption and send update command to the Kubernetes component	CHANGE	Orchestrator	Kubernetes Cluster / Edge Nodes	Inf.11	QoS.1 Conf.1 Conf.2
St.16	New update command (if the updates are made through the Kubernetes component)	Update energy services/federated model	Update the service version at edge nodes	CHANGE	Kubernetes Cluster	Edge Nodes	Inf.12	QoS.1 Conf.1

5 Information exchanged

Information exchanged			
Information exchanged (ID)	***Name of information***	***Description of information exchanged***	***Requirement, R-IDs***
Inf.1	Energy service metadata	Computational requirements, links with other services, policies and service level agreements Communication protocol: HTTP Format: Json	*D.1*
Inf.2	Edge node metadata	Computational resources, location Communication protocol: HTTP Format: Json	*D.2*
Inf.3	Service monitored data	Data monitored in real time by Kubernetes component and alerts for policy/agreement violations - Kubernetes API	*Conf. 1*
Inf.4	Service planification	The results of the swarm-based optimization run by orchestrator Communication protocol: HTTP Format: Json	*QoS.1* *QoS.2* *Conf.1* *Conf.2*
Inf.5	Deployment information	Container image, deployment description, configuration - JSON or YMAL format for specification, docker image	*Conf.1*
Inf.6	Federated process hyperparameters	Training nodes, number of rounds Communication protocol: HTTP Format: Json	*D.1*
Inf.7	Descriptive Data Profile	Statistics on energy data stored on the edge nodes Communication protocol: HTTP Format: Json	*D.1* *D.2*
Inf.8	Federated process configuration	Optimization/clustering result computed by the orchestrator Communication protocol: HTTP Format: Json	*D.2* *Conf.2*
Inf.9	Training performance data	Training performance of the edge nodes for the configuration given in the current round Communication protocol: HTTP Format: Json	*D.2*
Inf.10	New version notification	Produced when a new version is available for a service Communication protocol: HTTP or MQTT	*D.1*
Inf.11	App package / federated model parameters	Updated version of the app package or federated model Binary package or Docker image or JSON	*Conf.1*
Inf.12	Deployment update information	Updated image, configuration, rollout strategy - Kubernetes API in JSON or YMAL format	*QoS.1* *Conf.1*

6 Requirements

Quality of Service Requirements		
Categories ID	***Category name for requirements***	***Category description***
QoS	Quality of Service	Generic properties that service/SUC should provide – quality attributes.
Requirement ID	***Requirement name***	***Requirement description***
QoS.1	Service availability	The system must ensure availability of the services after relocation or updating
QoS.2	Services inter-connection	The connections between services are maintained after relocation

Security Requirements		
Categories ID	***Category name for requirements***	***Category description***
Sec	Security	Authentication of user, confidentiality, integrity, prevention of denial of service, non-repudiation or accountability, error management.
Requirement ID	***Requirement name***	***Requirement description***
Sec.1	Secure data exchange	Data exchange between services and orchestrator is secured (provided by Dataspace Connector)
Sec.2	Controlled access	Provided by the Dataspace Connector that ensures authorized access to data

Data Management Requirements		
Categories ID	***Category name for requirements***	***Category description***
D	Data Management	Type of source of data, correctness or validity of data, timeliness or time stamping of data, volume of data, synchronization, or consistency of data across systems, timely access to data, validation of data across organizational boundaries, transaction management, data naming, identification, formats across disparate systems, maintenance of data and databases.
Requirement ID	***Requirement name***	***Requirement description***
D.1	Management of service data exchange between provider and orchestrator	The orchestrator must access, store, and manage metadata and deployment information about energy services provided by the service provider
D.2	Management of data between orchestrator and edge nodes	The orchestrator must receive and store information about edge node computational resources, location and it sends configurations, updated app packages
D.3	Compliance with Eclipse Data Space Connector	Data exchanges between services should be made using Eclipse Data Space Connector

Discovery and Configuration Requirements		
Categories ID	***Category name for requirements***	***Category description***
Conf	Configuration	Locations, distances, communication layout, commonly used communication protocol media, network bandwidth, existing protocols, number of devices, systems, volume of data items, expected growth, etc.
Requirement ID	***Requirement name***	***Requirement description***
Conf.1	Deployment automation using Kubernetes	The services deployment should be automated using Kubernetes component
Conf.2	Service configurability	Services should support external configuration, and the orchestrator generates and manages these configuration

7 Common Terms and Definitions

Common Terms and Definitions	
Term	**Definition**
EDC	Eclipse Dataspace Component
API	Application programming interface
JSON	JavaScript Object Notation
YMAL	“YAML Ain’t Markup Language” or “Yet Another Markup Language”
MQTT	Message Queuing Telemetry Transport
HTTP	Hypertext Transfer Protocol

Transversal Use Case n°3:

Use of the App Store as part of HEDGE-IoT

[Functional interoperability]

1 Description of the use case

1.1 Name of the use case

ID	*Area / Domain(s) / Zones(s)*	*Name of Use Case*
TUC-3	Functional interoperability	Use of the App Store as part of HEDGE-IoT.

1.2 Version management

Version Management			
Version No.	*Date*	*Name of Author(s)*	*Changes*
0.1	10/04/2025	Trialog	First structure based on project brainstorming.
0.2	05/05/2025	ED	1st Draft: Sections 1-3
0.3	26/05/2025	ED	Sequence Diagrams and corrections on comments
0.4	30/05/2025	ED	Completed Scenario Tables
0.5	11/06/2025	ED	Updated Sequence Diagrams
0.6	18/06/2025	INESC	Updated Scenario Descriptions
0.7	19/06/2025	ED	Updates based on partner feedback
0.8	20/06/2025	ED	Finalised 1st Draft Version
1.0	27/06/2025	Trialog	1st final version
2.0	07/05/2026	ED	2nd final version (with refinements)

1.3 Scope and objectives of use case

Scope and Objectives of Use Case	
Scope	This use case describes the HEDGE-IoT App Store, a component of the HEDGE-IoT digital middleware, and how it enables publishing, discovering, using and creating new data apps (edge/cloud services) in a trusted, semantically interoperable energy IoT ecosystem. It covers the full lifecycle of data-driven services at the edge or cloud: from App Providers publishing services, to App Users discovering and deploying them, to composing new services from existing ones, and ensuring services are interchangeable in a data space, namely through to semantic standards.
Objective(s)	The main objective is to deploy an App Store where third-party developers can publish data apps, certified them, and where energy stakeholders can discover and deploy these apps on their HEDGE-IoT Connectors at the edge or in the cloud. The objectives that the use case is expected to achieve are to: • **Objective 1:** Facilitate publication of new services/sub-services in the HEDGE-IoT environment. • **Objective 2:** Provide a discovery mechanism so developers or users can find available reusable services. • **Objective 3:** Enable reusability and interoperability among different HEDGE-IoT components, pilot projects or datasets. • **Objective 4:** Promote semantic interoperability (services whose data assets can be used across different pilot contexts).
Related business case(s)	/

1.4 Narrative of use case

Narrative of Use Case
Short description
The HEDGE-IoT App Store is a repository for Software Applications that operate at least in one data space configuration. Apps include/represent services/microservices enabling quick discovery, sharing, and reuse across different pilots and domains. It allows service owners to publish new service functionalities, while developers or other system components can request and acquire access. By ensuring semantic and technical interoperability, the App Store accelerates solution development and deployment, fosters collaboration, and ensures consistent service quality within the HEDGE-IoT framework. In line with the new data space protocol (version >2), the App Store also promotes the possibility for dataspace compliant connectors to search adopt other versions of control and data planes available.
Complete description
This transversal use case is therefore split into four scenarios: • **Scenario 1:** Publish a service/sub-service in the App Store A service provider develops or containerizes a new IoT/energy service and registers it within the HEDGE-IoT App Store, supplying descriptive metadata (inputs, outputs, resource requirements, license) along with the container image. An automated certification process checks the connector's compatibility, security, and semantic conformance. Once approved, the service becomes discoverable in the App Store catalogue, ready for other stakeholders to reuse. • **Scenario 2:** Find/Retrieve/Reuse/Access a service/sub-service in the App Store An application developer or system component browses or queries the App Store to locate a suitable service based on functionality or usage terms. After accepting any licensing agreements or data usage policies, the user's edge/cloud environment retrieves the service container from the App Store's registry. The service is then deployed at the selected node(s), where it can securely process data under the HEDGE-IoT dataspace policies. • **Scenario 3:** Interchangeable common services/sub-services *(semantic interoperability across data consumers)* Different providers publish "functionally equivalent" services following the same data schemas. Thanks to uniform semantic definitions, a user can seamlessly swap one service with another without having to modify underlying workflows or data pipelines. This ensures plug-and-play upgrades, vendor-neutral deployments, and consistent service behaviour across diverse HEDGE-IoT environments.

1.5 Key performance indicators (KPI)

ID	Name	Description	Reference to mentioned use case objectives
KPI1	Open source released developments related to data connector implementation	Target: > 2 Y1, > 3 Y2, > 5 Y3 This KPI provides a counter of the number of Apps/services available in the App Store. It expects more than 2 services in year 1 of the project, more than 3 in year 2 and more than 5 in year 3. The counter is cumulative with past years. (*KPI5*).	Objective 1

1.6 Use case conditions

Use case conditions
Assumptions
1. The App Store must be part of the HEDGE-IoT Dataspace and operating via a functioning, compatible Connector. 2. There is a functioning HEDGE-IoT Connector environment for both App Providers and App Consumers (including identity management etc.). 3. Each participants Connector must be configured with valid identities and certificates, and in some cases equipped with an App Execution Environment to run the downloaded apps. 4. There is an established semantic model, ontology or set of ontologies so that published services can be semantically described. 5. The certification process is established for connector's that require it. 6. The HEDGE-IoT dataspace is operational 7. Apps are packaged as container images and include metadata according to the IDS Information Model and project-specific schemas. 8. The App Store has the necessary container registry and metadata database in place.
Prerequisites
• App Provider has corrected access rights and credentials to access the dataspace and to publish container images. • App Consumer is onboarded to the HEDGE-IoT environment, with a certified connector that can pull and deploy container images.

1.7 Further Information to the use case for classification / mapping

Classification Information
Relation to other use cases
Ties in with all pilot-specific SUCs/BUCs that need data processing or specialized analytics.
Level of depth
Prioritisation
High priority (since the App Store is a fundamental piece for the entire HEDGE-IoT solution).
Generic, regional or national relation

Generic
Nature of the use case
System use case, Transversal use case Technical/system *use case* that supports business processes across different pilot use cases.
Further keywords for classification
App Store, Services, Interoperability, Data App, Connector

1.8 General Remarks

General Remarks
• HEDGE-IoT App Store design is based on the IDSA guidelines (RAM v4.0) • HEDGE-IoT Connector design is based on EDC Framework

2 Diagrams of use case

Diagram(s) of use case

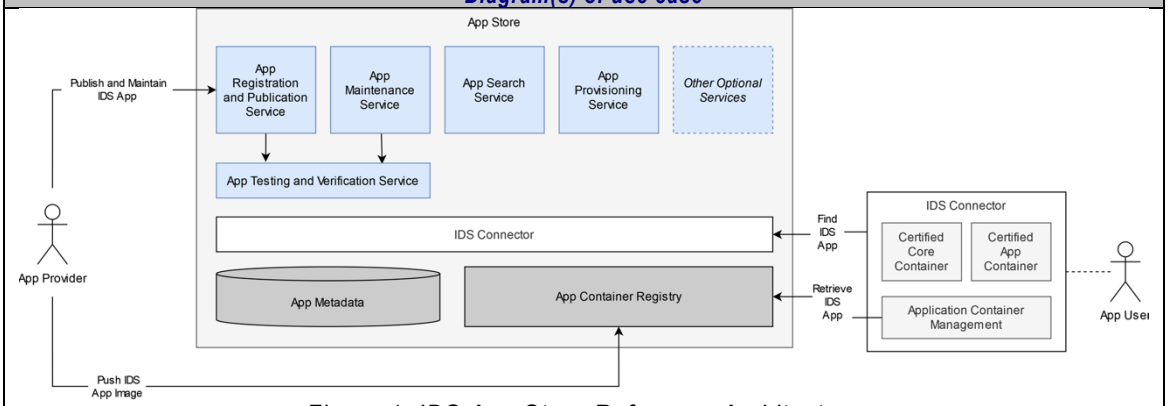

Figure 1: IDS App Store Reference Architecture

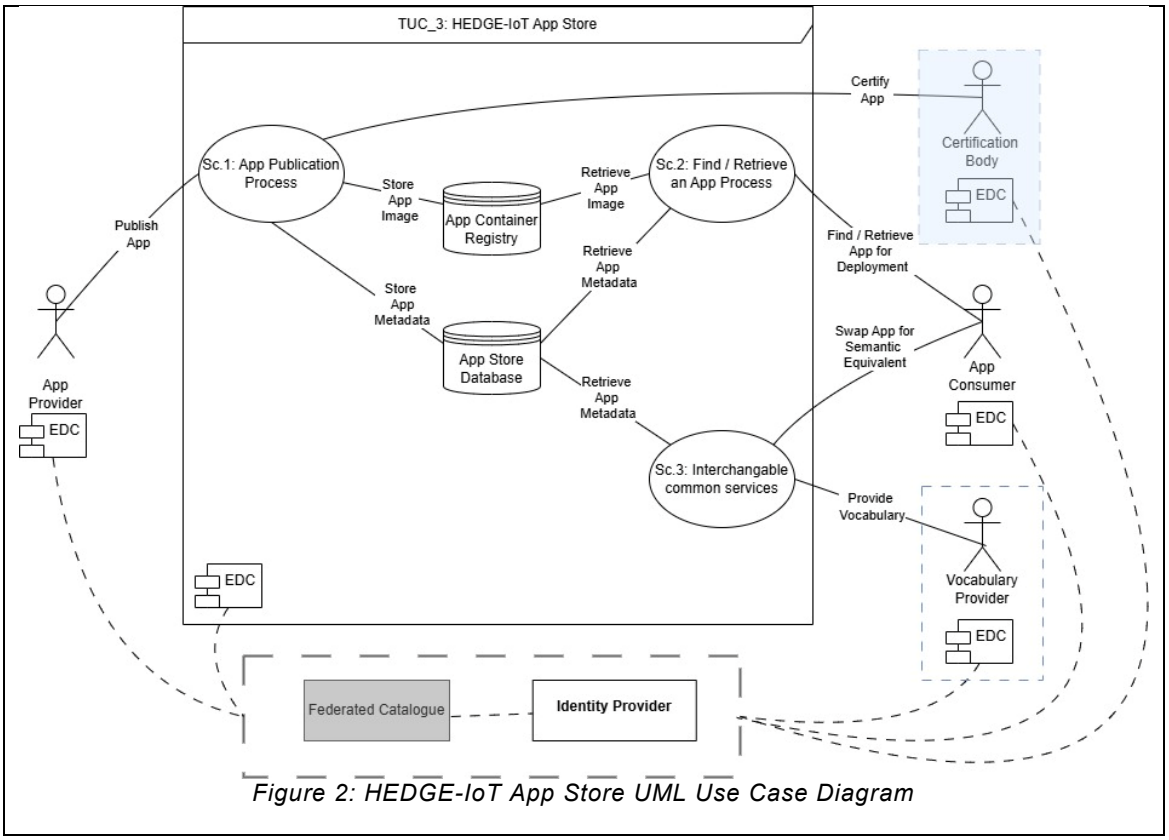

Figure 2: HEDGE-IoT App Store UML Use Case Diagram

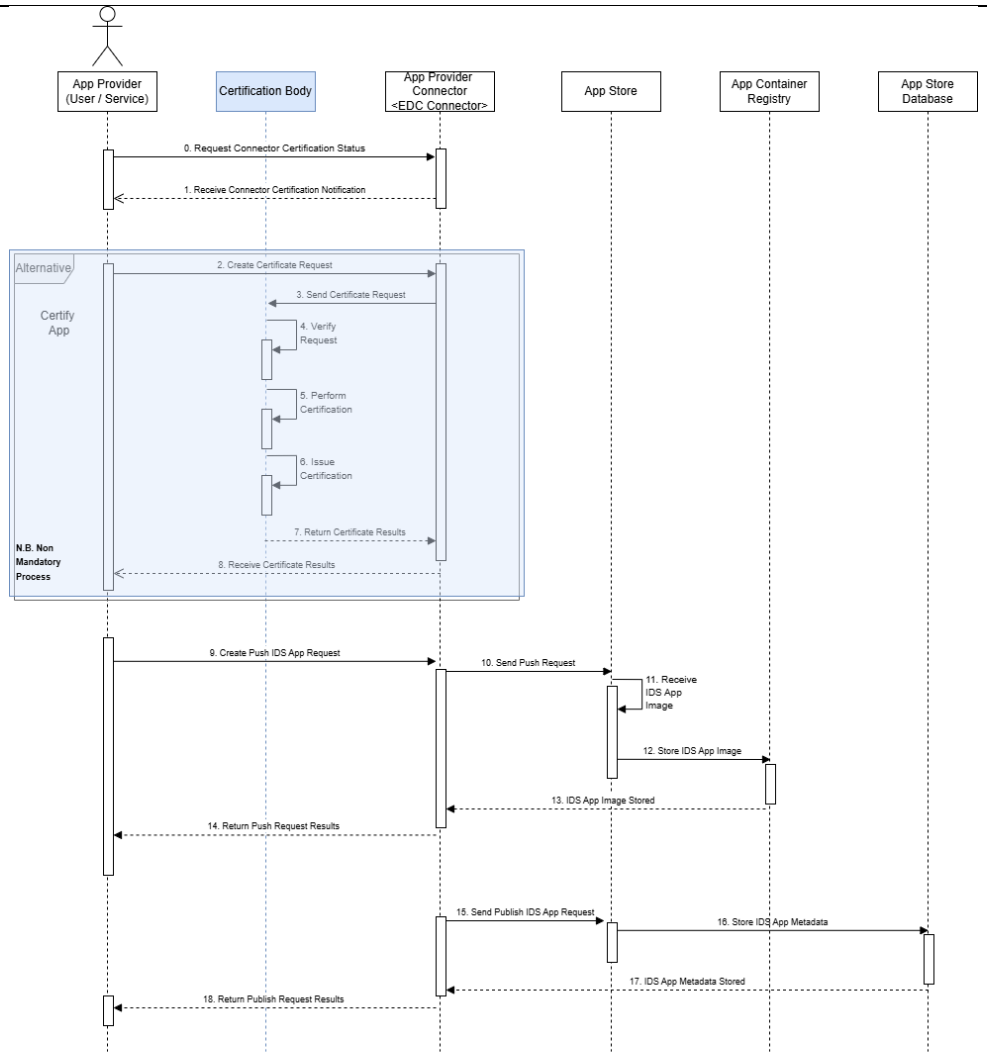

Figure 3: Scenario 1 - 'Publish an App Process'' UML Sequence Diagram

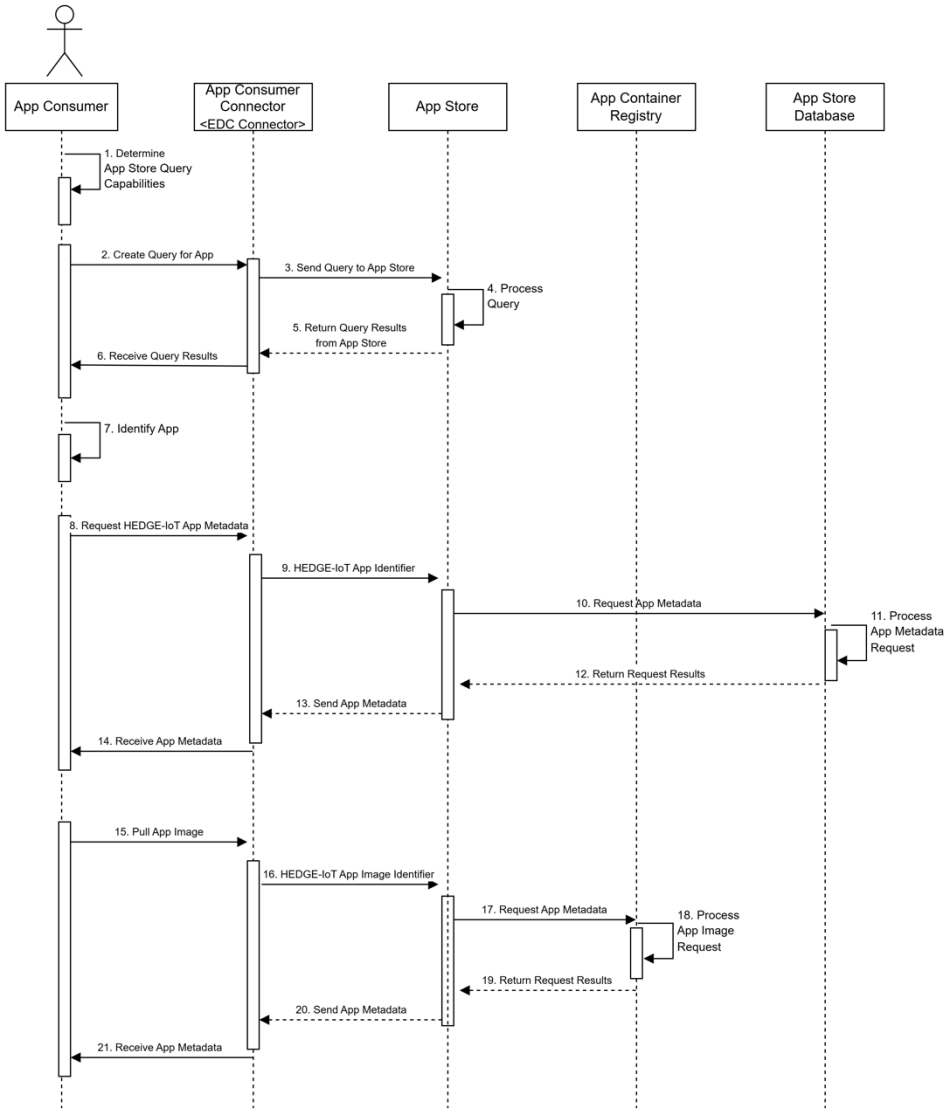

Figure 4: Scenario 2 - "Find/Retrieve/Reuse/Access an App Process" UML Sequence Diagram

3 Technical details

3.1 Actors

Actors			
Actor Name	***Actor Type***	***Actor Description***	***Further information specific to this use case***
Data App Provider	Logical Actor	An entity that publishes an IDS App to the App Store. The App Provider is responsible for creating or owning the app and for supplying its metadata and container image.	HEDGE-IoT Data App Developers (Demos / 3rd Parties from Open Calls)
App Store	Logical Actor	The platform (software component) that manages the lifecycle of data apps. It provides interfaces for: App Registration & Publication: accepting new app submissions, including metadata. App Maintenance: updating or removing apps, managing versions. App Search/Query: allowing users to find apps by various criteria. App Provisioning: delivering the app (metadata and container) to the requesting user's connector. Other Optional services	TBD
Data App User / Data App Consumer	Logical Actor	An entity that discovers and uses apps from the App Store. The role involves searching the castep-by-stepping an app, possibly negotiating usage terms and deploying the app in their environment. The App User operates a HEDGE-IoT Connector that can receive and run the app.	TBD
Certification Body	Logical Actor	A Certification Body can issue a Certificate for a connector for compliance with security, safety, or interoperability standards	TBD
Vocabulary Provider	Logical Actor	Semantic Interoperability Technological Enablers Providers	TBD
App Store Container Registry	Logical Actors	Holds and manages software containers representing Apps in the data space or any required components.	TBD
App Store Metadata Registry	Logical Actors	App Store sub-module	TBD
App Provider Connector <EDC Connector>	Logical Actor	Eclipse Dataspace Connector instance deployed by each participant	Manages metadata, policies, negotiation, and transfer on behalf of the provider/consumer. {DST}
Metadata Broker	Logical Actor	Federated catalogue service that stores and exposes only descriptive metadata (no data or binaries). Enables search and discovery of data assets and services across the dataspace.	• IDS-compliant Broker implementation

Identity Provider	Logical Actor	Issues, validates, and revokes identities for all HEDGE-IoT ecosystem participants.	• Manages identities at two levels: (1) Organisation level – each company/platform receives a certified IDS Participant ID; (2) Individual/user level – personal certificates

3.2 References

References						
No.	***Reference Type***	***Reference***	***Status***	***Impact on use case***	***Originator / organisation***	***Link***
1	IDSA Specification	IDSA 3.5.3: App Store and App Ecosystem	Online	Defines app store ops, app types, endpoints	IDSA	IDS Knowledge Base 3.5.3
2	IDSA Processes	IDSA 3.4.5: Publishing & Using Data Apps		Lays out the Publish/Find/Retrieve/Use processes	IDSA	IDS Knowledge Base 3.4.5
3	IDSA Usage Control	IDSA 4.1.6: Usage Control	Online		IDSA	IDS Knowledge Base 4.1.6
4	GitHub Repository	Eclipse Dataspace Connector (EDC) Documentation	Online	Provides docs for EDC	Eclipse Foundation	https://github.com/eclipse-edc1
5	International Data Spaces Information Model	**Revision:** 4.2.0	Online	The Information Model primarily aims at describing, publishing and detecting data products (Data Assets) and reusable data processing software (Data Apps) in the International Data Space. Data Assets and Data Apps are the core resources of the International Data Space and are hereinafter referred to as resources.	IDSA	https://w3id.org/idsa/core
6.	IDS Message Types	Aligned to v4.1.0	Online	Descriptions for each message, information about the payload as well as changes in the properties, i.e., new properties and additional mandatory properties. "Message Class" and "Message Subclass" are just for structural / hierarchial description. The		IDS MessageTypes

				"Message Name" contains the actual Message which is used. New properties, which are mandatory are marked with an (*). Existing properties, which are mandatory for a specific message are listed in the corresponding column.		
7.	Dataspace Protocol		Online	The Dataspace Protocol is a set of specifications designed to facilitate interoperable data sharing between entities governed by usage control and based on Web technologies. These specifications define the schemas and protocols required for entities to publish data, negotiate Agreements, and access data as part of a federation of technical systems termed a Dataspace.		Dataspace Protocol 2025-1-RC1
8.	3.5.4. IDS Metadata Broker					IDS Knowledge Base

4 Step by step analysis of use case

4.1 Overview of scenarios

Scenario conditions						
No.	**Scenario name**	**Scenario description**	**Primary actor**	**Triggering event**	**Pre-condition**	**Post-condition**
Sc.1	Publish a Service/ Sub-service	A service provider develops or containerizes a new IoT/energy service and registers it in the HEDGE-IoT App Store, providing metadata (inputs, outputs, license) and container image. After automated checks (testing/certification) the connector is compliant.	App Provider	The provider deploys a new service or updates an existing one.	1) The service is containerized and meets basic compliance requirements (semantic metadata, security, etc.). 2) App Store is operational and accepting submissions. 3) (Optional) Connector Certification is completed if the store/policies require it.	The new service appears in the App Store catalogue, ready for others to discover and deploy (if certification is active, only after passing this process).
Sc. 2	Find/Retrieve/Reuse/Access a	A developer or system component searches the App	App Consumers	A need arises for a specific functional	1) The user is onboarded with the App Store (valid IDS Connector).	The service is successfully deployed on the user's

	service/sub-service in the App Store	Store for suitable services. After selecting and accepting licensing/usage terms, the app container is retrieved and deployed at the user's edge/cloud. The service can then process data under HEDGE-IoT policies.		ity that could be fulfilled by a published service.	2) The user has the required permissions or license to deploy services. 3) At least one suitable service is published in the catalogue.	node(s) and becomes operational, processing data according to data usage policies and providing the intended functionality.
Sc.3	Interchangeable Common Services/ Sub-services	Multiple providers publish functionally equivalent services using the same data schema/APIs. Users can seamlessly swap one for another without re-engineering data flows—promoting plug-and-play upgrades and vendor-neutral deployments.	App Consumers	The App consumer needs to replace or upgrade an existing service with a new one of similar functionality.	1) Common interface/data schemas are defined (semantic interoperability). 2) Multiple equivalent services are available in the App Store. 3) User's existing workflows and connectors are already configured to consume these standardized services.	The user replaces one service with another having the same interface, maintaining data flows and workflows with minimal reconfiguration. The new service takes over under the same usage policies.

4.2 Steps – Scenarios

Scenario								
Scenario name:		**Sc. 1 – Publishing an App Process**						
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
St.0	Determine if Connector requires Certification	0. Request Connector Certification Status	Data App Provider verifies if its connector is certified to publish a Data App on the HEDGE-IoT App Store.	CREATE	App Provider	App Provider Connector	Inf.00	SEC.2, SEC.4, O.X
St.1	Receive Connector Certification Status	1.Recieve Connector Notification Information	Data App Provider Receives the connector notification message with the up to date the Connector Certification status.	GET	App Provider Connector	App Provider	Inf. 01	SEC.2, SEC.4, D.2,, O.X
St.2	Determination that Certification is Required	2.Create Certification Requests	As certification is required, the provider instructs its local EDC Connector to assemble a certification request.	CREATE	App Provider	App Provider Connector	Inf. 02	SEC.2, SEC.4,, O.X
St.3	Connector Receives Instruction from App Provider	3. Send Certificate Request	Connector creates and sends a new request to the Certification Body.	CREATE	App Provider Connector	Certification Body	Inf. 03	SEC.2, SEC.4,, O.X
St.4	Certification Body receives certification request	4.Verify Request	The Certification Body checks the request and verifies the completeness and schedules the connector for security & interoperability testing.	EXECUTE	Certification Body	Certification Body	Inf. 04	SEC.2, SEC.4, D.2, O.X
St.5	Verification Request Approved	5.Perform Certification	Compliance tests are executed by the Certification Body in an isolated testbed; results are recorded	EXECUTE	Certification Body	Certification Body	Inf. 05	SEC.2, SEC.4, D.2, O.X
St.6	Certification Completed	6.Issue Certification	Upon successful tests, the Certification Body creates a signed digital certificate that attests the Connector's	CREATE	Certification Body	Certification Body	Inf. 06	SEC.2, SEC.4, D.2, O.X

			compliance level and returns it to the requester.					
St.7	Certificate Issued from Certification Body	7. Return Certificate Results	The Certification Body transmits the signed certificate back to the EDC Connector using IDS Response messages.	GET	Certification Body	App Provider Connector	Inf. 07	SEC.2, SEC.4, D.2, O.X
St.8	Certification Results Issued and Sent to Data App Producer	8. Received Certificate Results	The EDC Connector forwards the received certificate and report to the Data-App Provider	GET	App Provider Connector	App Provider	Inf. 08	SEC.2, SEC.4, D.2, O.X
St.9	Certification Results Successful	9.Create Push IDS App Image Request	The provider now commands the EDC Connector to push the container image to the designated App-Store registry, referencing the certificate ID in the request header.	CREATE	App Provider	App Provider Connector	Inf. 09	SEC.2, SEC.4, O.X
St.10	EDC Received Push Request from Daya App Producer	10.Send Push Request	The EDC Connector creates a Push-Request message that announces the forthcoming upload.	CREATE	App Provider Connector	App Store	Inf.10	SEC.2, SEC.4, O.X
St.11	EDC Pushes the App Image to App store	11.Receive IDS App Image	The App Store receives the container image stream from the connector, performs integrity checks.	CREATE	App Store	App Store	Inf.11	SEC.2, SEC.4, D.2, O.X
St.12	App Store Reives Valid App Image	12.Store IDS App Image	After validation, the App Store creates a new registry image in the App Container Registry.	CREATE	App Store	App Container Registry	Inf.12	SEC.2, SEC.4, O.X
St.13	App image received and stored	13. IDS App Image Stored	The registry reports successful storage by sending an acknowledgement.	REPORT	App Container Registry	App Provider Connector	Inf.13	SEC.2, SEC.4, O.X
St.14	Push Request Results Available and Routed via EDC Connector	14.Return Push Request Message	The EDC Connector relays the positive push result back to the Data-App Provider, completing the binary-upload phase.	GET	App Provider Connector	App Producer	Inf .14	SEC.2, SEC.4, D.2, O.X

St.15	Push Request Successful	15.Publish IDS App Request	The provider instructs the connector to publish the app's metadata to the App-Store catalogue.	CREATE	App Provider	App Provider Connector	Inf.15	SEC.2, SEC.4, D.5, O.X
St.16	EDC Received Publish Request from Daya App Producer	16.Send Publish Request	The EDC Connector sends a Publish-Request message with the metadata to the App Store.	CREATE	App Provider Connector	App Store	Inf.16	SEC.2, SEC.4, O.X
St.17	App Store Reives Valid App Metadata	17.Store IDS App Metadata	The App Store validates semantic fields, checks certificate references, then creates the metadata record into the App Store Database.	CREATE	App Store	App Store Database	Inf.17	SEC.2, SEC.4, O.X
St.18	App Metadata stored	18.IDS App Metadata Stored	The database returns a success message and the new App-ID; the App Store registers this ID in its self-description and notifies the connector that the entry is active.	REPORT	App Store Database	App Provider Connector	Inf.18	SEC.2, SEC.4, O.X
St.19	Publish Request Results Available and Routed via EDC Connector	19. Return Publish Request Results	The EDC Connector forwards the publish-result to the Data-App Provider. The app is now visible in the catalogue and ready for discovery by consumers.	GET	App Provider Connector	App Provider	Inf .19	SEC.2, SEC.4, D.2, O.X

Scenario								
Scenario name:		**Sc. 2 – Find/Retrieve/Reuse/Access a service/sub-service in the App Store**						
Step No.	**Event**	**Name of process/ activity**	**Description of process/ activity**	**Service**	**Information producer (actor)**	**Information receiver (actor)**	**Information Exchanged (IDs)**	**Requirement, R-IDs**
St.1	Data App needed by App Consumer	Determine App Store Query Capabilities	App Consumer identifies how to query the App Store	EXECUTE	App Consumer	App Consumer	Inf.20	SEC.2, SEC.4, O.X
St.2	Query capabilities determined	Create Query for App	App Consumer creates a query request to find a specific App in the App Store	CREATE	App Consumer	EDC Connector	Inf.21	SEC.2, SEC.4, O.X
St.3	App query received by EDC	Send Query to App Store	App Provider Connector creates a new query-message and pushes it to the App Store	CREATE	App Provider Connector	App Store	Inf.22	SEC.2, SEC.4, D.2, O.X
St.4	App query received by App Store	Process Query	App Store processes the received app query	EXECUTE	App Store	App Store	Inf.23	SEC.2, SEC.4, O.X
St.5	App Store processed query	Return Query Results	App Store returns query results back to the App Provider Connector	GET	App Store	App Provider Connector	Inf.24	SEC.2, SEC.4, D.2, O.X
St.6	App Provider Connector received query results	Receive Query Results	EDC forwards the query results to the App Consumer	GET	App Provider Connector	App Consumer	Inf.25	SEC.2, SEC.4, D.2, O.X
St.7	App Consumer received results	Identify App	App Consumer identifies desired app from query results	EXECUTE	App Consumer	App Consumer	Inf.26	SEC.2, SEC.4, O.X
St.8	Desired app identified	Request HEDGE-IoT App Metadata	App Consumer requests detailed metadata for the identified app from EDC	CREATE	App Consumer	App Provider Connector	Inf.27	SEC.2, SEC.4, O.X
St.9	Metadata request received by EDC	HEDGE-IoT App Identifier	EDC creates an App-metadata request and sends it to the App Store.	CREATE	App Provider Connector	App Store	Inf.28	SEC.2, SEC.4, O.X

St.10	Metadata request received by App Store	Request App Metadata	The App Store creates a metadata-lookup request for its Database	CREATE	App Store	App Store Database	Inf.29	SEC.2, SEC.4, O.X
St.11	Metadata request received by Database	Process App Metadata Request	Database processes and retrieves app metadata	EXECUTE	App Store Database	App Store Database	Inf.29	SEC.2, SEC.4, D.2, O.X
St.12	Database completed processing	Return Request Results	Database returns app metadata to App Store	GET	App Store Database	App Store	Inf.31	SEC.2, SEC.4, O.X
St.13	App metadata available at App Store	Send App Metadata	App Store sends retrieved app metadata to App Provider Connector	GET	App Store	App Provider Connector	Inf.27	SEC.2, SEC.4, O.X
St.14	EDC received metadata	Receive App Metadata	EDC sends app metadata to App Consumer	GET	App Provider Connector	App Consumer	Inf.27	SEC.2, SEC.4, D.2, O.X
St.15	App Consumer received metadata	Pull App Image	App Consumer requests app image downloads from App Provider Connector	CREATE	App Consumer	App Provider Connector	Inf.34	SEC.2, SEC.4, O.X
St.16	Image request received by EDC	HEDGE-IoT App Image Identifier	EDC creates and forwards an image-download request to the App Store	CREATE	App Provider Connector	App Store	Inf.29	SEC.2, SEC.4, O.X
St.17	App Store received app image request	Request App Metadata	The App Store creates a fetch-request message for the Container Registry	CREATE	App Store	App Container Registry	Inf.33	SEC.2, SEC.4, D.2, O.X
St.18	Image request received by Registry	Process App Image Request	App Container Registry processes request and prepares app image	EXECUTE	App Container Registry	App Container Registry	Inf.37	SEC.2, SEC.4, O.X
St.19	Registry processed app image request	Return Request Results	App Container Registry returns app image details to App Store	GET	App Container Registry	App Store	Inf.38	SEC.2, SEC.4, D.2, O.X
St.20	App Store received app image details	Send App Metadata	App Store sends app image details to App Provider Connector	GET	App Store	App Provider Connector	Inf.39	SEC.2, SEC.4, O.X

St.21	EDC received app image details	Receive App Metadata	EDC sends app image details to App Consumer	GET	App Provider Connector	App Consumer	Inf. 40	SEC.2, SEC.4, D.2, O.X

5 Information exchanged

Information exchanged			
Information exchanged (ID)	***Name of information***	***Description of information exchanged***	***Requirement, IDs***
Inf.00	Request Connector Certification Status	[HTTP] Request to query the status of a connector specified by its ID	O.X
Inf. 01	Receive Connector Notification Information	[HTTP] Status of a connector certification identified by its ID	O.X
Inf. 02	Create Certification Requests	[HTTP] Issue a certification request targeting the one instance of a connector	O.X
Inf. 03	Send Certificate Request	[HTTP] Summary of the certification request, including reference to the app with the data space identifier and app store registry internal identifier.	O.X
Inf. 04	Verify Request	[HTTP] Process Trigger	O.X
Inf. 05	Perform Certification	[HTTP] Certification process trigger	O.X
Inf. 06	Issue Certification	[HTTP] Trigger to issue a new certificate	O.X
Inf. 07	Return Certificate Results	[HTTP] Digital certificate, including the signature hashcode.	O.X
Inf. 08	Received Certificate Results	[HTTP] Result status of the certificate issue process. (can include the certificate as detailed in Inf.08)	O.X
Inf. 09	Create Push IDS App Image Request	[HTTP] Request to create a new App image, including the App internal identification, including name and version.	O.X
Inf.10	Send Push Request	[HTTP] Push OIC compliant App image	O.X
Inf.11	Receive IDS App Image	[HTTP] App image details and processing repository details.	O.X
Inf.12	Store IDS App Image	[HTTP] App binary content is received, processed, stored.	O.X
Inf.13	IDS App Image Stored	[HTTP] Confirmation and digest for OIC image stored.	O.X
Inf .14	Return Push Request Message	[HTTP] Confirmation of success of failure of process.	O.X
Inf.15	Publish IDS App Request	[HTTP] Request for visibility change and publication of the App.	O.X
Inf.16	Send Publish Request	[HTTP] Request for visibility change and publication of the App.	O.X
Inf.17	Store IDS App Metadata	[HTTP] App details and linked software images and version details.	O.X
Inf.18	IDS App Metadata Stored	[HTTP] Confirmation that App details and linked software images and version details where stored.	O.X
Inf .19	Return Publish Request Results	[HTTP] Confirmation that App details and linked software images and version details where stored.	O.X
Inf.20	Determine App Store Query Capabilities	[HTTP] Poll query mechanism for App store users.	O.X
Inf.21	Create Query for App	[HTTP] Create an app catalogue request targeting the App Store instance in the dataspace.	O.X
Inf.22	Send Query to App Store	[HTTP] Query of Apps in the dataspace catalogue, targeting the App store instance.	O.X

Inf.23	Process Query	[HTTP] List with App metadata available in the catalogue.	O.X
Inf.24	Return Query Results	[HTTP] List with App metadata available in the catalogue.	O.X
Inf.25	Receive Query Results	[HTTP] List with App metadata available in the catalogue.	O.X
Inf.27	App Metadata	App metadata available in the catalogue	O.X
Inf.28	HEDGE-IoT App Identifier	App identifier in the catalogue.	O.X
Inf.29	Request App Metadata	Issue request for App metadata.	O.X
Inf.31	Return Request Results	App metadata details	O.X
Inf.34	Pull App Image	Binary data of the App image together with the digest confirmation	O.X

ids:Message	Core ids:Message class with it's properties, which are equal for all messages. For communication in the IDS, the specific message types (second table, below) are used.		
Property	**Always mandatory property**	**Can have multiple values at the same time**	**Description**
modelVersion	✓		Information Model version, against which the Message should be interpreted
issued	✓		Date of issuing the message
correlationMessage			Correlated message. Usually needed, if a messages responds to a previous message. A Connector may, e.g., send a MessageProcessedNotification as a response to an incoming message and therefore needs this property to refer to the incoming message.
issuerConnector	✓		Origin Connector of the message
recipientConnector		✓	Target Connector
senderAgent	✓		Agent, which initiated the message
recipientAgent		✓	Agent, for which the message is intended
securityToken	✓		Token representing a claim, that the sender supports a certain security profile
authorizationToken			Authorization token
transferContract			Contract which is (or will be) the legal basis of the data transfer
contentVersion			Version of the content in the payload

IDS Message Types	The table contains descriptions for each message, information about the payload as well as changes in the properties, i.e., new properties and additional mandatory properties. "Message Class" and "Message Subclass" are just for structural / hierarchical description. The "Message Name" contains the actual Message which is used. New properties, which are mandatory are marked with an (*). Existing properties, which are mandatory for a specific message are listed in the corresponding column. **All mandatory property declarations of the core ids:Message above still hold**.

Message Class	Message Subclass (Abstract)	Message Name	Description
Request Messages		**RequestMessage**	Client-generated message initiating a communication, motivated by a certain reason and with an answer expected. May be used for messages, which are not covered by the core IDS messages.
		CommandMessage	Command messages are usually sent when a response is expected by the sender. Changes state on the recipient side. Therefore, commands are not 'safe' in the sense of REST.
		InvokeOperationMessage	Message requesting the recipient to invoke a specific operation.
		ContractRequestMessage	Message containing a suggested content contract (as offered by the data consumer to the data provider) in the associated payload (which is an instance of ids:ContractRequest).
		ArtifactRequestMessage	Message asking for retrieving the specified Artifact as the payload of an ArtifactResponse message.
		AccessTokenRequestMessage	Message requesting an access token. This is intended for point-to-point communication with, e.g., Brokers.
		QueryMessage	Query message intended to be consumed by specific components.
		DescriptionRequestMessage	Message requesting metadata. If no URI is supplied via the ids:requestedElement field, this message is treated like a self-description request and the recipient should return its self-description via an ids:DescriptionResponseMessage. However, if a URI is supplied, the Connector should either return metadata about the requested element via an ids:DescriptionResponseMessage, or send an ids:RejectionMessage, e.g. because the element was not found
		ParticipantRequestMessage	This class is deprecated. Use ids:DescriptionRequestMessage instead. Message asking for retrieving the specified Participants information as the payload of an ids:ParticipantResponse message.
		UploadMessage	Message used to upload a data to a recipient. Payload contains data

		AppRegistrationRequestMessage	Message that asks for registration or update of a data app to the App Store. Payload contains app-related metadata (instance of class ids:AppResource). Message header may contain an app identifier parameter of a prior registered data app. If the app identifier is supplied, the message should be interpreted as a registration for an app update. Otherwise, this message is used to register a new app.
		AppUploadMessage	Message that usually follows a AppRegistrationResponseMessage and is used to upload a data app to the app store. Payload contains data app. Note that the message must refer to the prior sent, corresponding AppResource instance. The IRI of the ids:appArtifactReference must must match the IRI of the artifact which is the value for the ids:instance property. The ids:instance is specific for each representation. Therefore, if someone wants to upload multiple representations for an app, he has to state them using multiple ids:instance properties inside the AppRepresentation (and therefore inside the AppResource). Otherwise, no mapping between payload and app metadata can be achieved.
ResponseMessage		**ResponseMessage**	Response messages hold information about the reaction of a recipient to a formerly sent command or event. They must be correlated to this message. May be used for messages, which are not covered by the core IDS messages.
		ArtifactResponseMessage	Message that follows up a ArtifactRequestMessage and contains the Artifact's data in the payload section.
		AccessTokenResponseMessage	Response to an access token request, intended for point-to-point communication.
		ContractAgreementMessage	Message containing a contract, as an instance of ids:ContractAgreement, with resource access modalities on which two parties have agreed in the payload.
		ContractResponseMessage	Message containing a response to a contract request (of a data consumer) in form of a counter-proposal of a contract in the associated payload (which is an instance of ContractOffcer).
		ResultMessage	Result messages are intended to annotate the results of a query command.
		RejectionMessage	Rejection messages are specialized response messages that notify the sender of a message that processing of this message has failed.
		OperationResultMessage	Message indicating that the result of a former InvokeOperation message is available. May transfer the result data in its associated payload section
		ParticipantResponseMessage	This class is deprecated. Use ids:DescriptionResponseMessage instead.

			ParticipantResponseMessage follows up a ParticipantRequestMessage and contains the Participant's information in the payload section.
		ContractRejectionMessage	Message indicating rejection of a contract.
		DescriptionResponseMessage	Message containing the metadata, which a Connector previously requested via the ids:DescriptionRequestMessage, in its payload.
		UploadResponseMessage	Message that follows up a UploadMessage and contains the upload confirmation.
		AppUploadResponseMessage	Message that follows up an AppUploadMessage and contains the app upload confimation.
		AppRegistrationResponseMessage	Message that follows up an AppRegistrationRequestMessage and contains the app registration confimation.
NotificationMessage	ConnectorNotificationMessage	**NotificationMessage**	Notification messages are informative, and no response is expected by the sender. May be used for scenarios, which are not covered by the core IDS messages.
		LogMessage	Log Message which can be used to transfer logs e.g. to the clearing house.
		ContractOfferMessage	Message containing a offered content contract (as offered by a data provider to the data consumer) in the associated payload (which is an instance of ContractOffer). In contrast to the ids:ContractResponseMessage, the ids:ContractOfferMessage is not related to a previous contract request.
		ContractSupplementMessage	Message containing supplemental information to access resources of a contract.
		MessageProcessedNotificationMessage	Notification that a message has been successfully processed (i.e., not ignored or rejected).
		OperationResultMessage	Message indicating that the result of a former InvokeOperation message is available. May transfer the result data in its associated payload section.
		RequestInProcessMessage	Notification that a request has been accepted and is being processed.
		ConnectorInactiveMessage	Event notifying the recipient(s) that a connector will be unavailable. The same connector may be available again in the future
		ConnectorUpdateMessage	Event notifying the recipient(s) about the availability and current configuration of a connector. The payload of the message must

			contain the updated connector's self-description
		ConnectorCertificateGrantedMessage	Whenever a Connector has been successfully certified by the Certification Body, the Identity Provider can use this message to notify Infrastructure Components.
		ConnectorCertificateRevokedMessage	Indicates that a (previously certified) Connector is no more certified. This could happen, for instance, if the Certification Body revokes a granted certificate or if the certificate has just expired.
	ResourceNotification Message	**ResourceUnavailableMessage**	Message indicating that a specific resource is unavailable. The same resource may be available again in the future.
		ResourceUpdateMessage	Message indicating the availability and current description of a specific resource. The resource must be present in the payload of this message.
	AppNotification Message	**AppAvailableMessage**	Message indicating that a specific App should be available (again) in the AppStore.
		AppUnavailableMessage	Message indicating that a specific App should be unavailable in the AppStore.
		AppDeleteMessage	Message indicating that an App should be deleted from the AppStore.
	ParticipantNotification Message	**ParticipantUnavailableMessage**	Event notifying the recipient(s) that a participant will be unavailable. The same participant may be available again in the future.
		ParticipantUpdateMessage	Event notifying the recipient(s) about the availability and current description of a participant. The payload of the message must contain the participant's self-description
		ParticipantCertificateGrantedMessage	Whenever a Participant has been successfully certified by the Certification Body, the Identity Provider can use this message to notify Infrastructure Components
		ParticipantCertificateUnavailableMessage	Indicates that a (previously certified) Participant is no more certified. This could happen, for instance, if the Certification Body revokes a granted certificate or if the certificate has just expired.

6 Requirements

Quality of Service Requirements		
Categories ID	***Category name for requirements***	***Category description***
QoS	Quality of Service	Generic properties that service/SUC should provide – quality attributes.
Requirement ID	***Requirement name***	***Requirement description***
QoS.1	Elapsed time response requirements for exchanging data	More than 10 seconds
QoS.2	Availability of information flows	Continuous availability not required but must be available at specific times or under specific conditions
QoS.3	Accuracy of data requirements	Adequate accuracy can be assumed
QoS.4	Frequency of data exchanges	Upon event

Security Requirements		
Categories ID	***Category name for requirements***	***Category description***
Sec	Security	Authentication of user, confidentiality, integrity, prevention of denial of service, non-repudiation or accountability, error management.
Requirement ID	***Requirement name***	***Requirement description***
Sec.2	Eavesdropping	Ensuring confidentiality, avoiding illegitimate use of data, and preventing unauthorized reading of data, is: Quite important
Sec.4	Authentication and Access Control mechanisms commonly used with this data exchange	Public key encryption (e.g. SSL/TLS)

Data Management Requirements		
Categories ID	***Category name for requirements***	***Category description***
D	Data Management	Type of source of data, correctness or validity of data, timeliness or time stamping of data, volume of data, synchronization, or consistency of data across systems, timely access to data, validation of data across organizational boundaries, transaction management, data naming, identification, formats across disparate systems, maintenance of data and databases.
Requirement ID	***Requirement name***	***Requirement description***
D.2	Correctness of source data	Source data is always correct (e.g. by definition)
D.5	Management of data across organizational boundaries	Data exchanges go across organizational boundaries

D.6	Data maintenance effort: human versus automation	Data maintenance is mostly automated but requires occasional intervention

Discovery and Configuration Requirements		
Categories ID	***Category name for requirements***	***Category description***
Conf	Configuration	Locations, distances, communication layout, commonly used communication protocol media, network bandwidth, existing protocols, number of devices, systems, volume of data items, expected growth, etc.
Requirement ID	***Requirement name***	***Requirement description***
Conf.2	Distance between entities	Varies and/or is not relevant
Conf.3	Number of Information Producers	Few to a hundred
Conf.4	Number of Information Receivers	Two to a few
Conf.6	Data exchange methods	Other: REST API (client-server)
Conf.7	Communication access services requirements	Request-response
Conf.8	Commonly used communication protocol	Natively REST, other protocols (e.g., MODBUS, MQTT) supported via semantic adapters.

Other Requirements		
Categories ID	***Category name for requirements***	***Category description***
O	Regulatory obligation related to privacy	2016/679 GDPR (General Data Protection Regulation)
Requirement R-ID	***Requirement name***	***Requirement description***
O.2	Personal data use	Personal data may not be processed unless there is at least one legal basis to do so.
O.3	Right to access, rectify, erasure, restriction	Data retention policy outlines the specific sensitive time period data can be retained, plus how it will be disposed of when the time to do so comes.
O.4	Data transfer consent	The data subject shall have the right to obtain from the controller without undue delay the access/rectification/erasure/restriction of inaccurate personal data concerning him or her.
O.5	Data retention policy	Personal data may not be transferred to a third-party if the data subject did not agree and the third party provide appropriate safeguard.
O.X	All constraints also apply.	All requirements in this category.

7 Common Terms and Definitions

Common Terms and Definitions	
Term	**Definition**
App Container Registry	A secure repository that stores and serves container images (e.g., Docker/OCI images) for download by authorised connectors.
App Metadata	The JSON-LD (or equivalent) description of an app: name, version, licence, required inputs/outputs, semantic tags, certification status, etc.
App Store	The marketplace component of the dataspace that lets participants publish, discover, and download certified data-apps and related services.
App Store Database	The internal catalogue that holds all published app-metadata records, making them searchable for users and connectors.
Certification Body	An independent organisation that tests and certifies apps or connectors for security, interoperability and compliance with IDS rules.
Container Image	A packaged, runnable file system (e.g., OCI/Docker image) that contains all binaries and dependencies required to execute an app.
Data Consumer	A participant that retrieves data or services from the dataspace under agreed usage policies.
Data Provider	A participant that offers data assets or services to other parties through an IDS-compliant connector and usage policies.
Dataspace	A federated, governed data-exchange ecosystem where sovereign data sharing occurs via certified IDS connectors and common policies.
EDC (Eclipse Dataspace Connector)	An open-source implementation of an IDS Connector used to connect participants, enforce usage control and host apps.
IDS (International Data Spaces)	A reference architecture, specification and governance model enabling secure, sovereign and interoperable data sharing.
Metadata	Descriptive information about a resource (data asset or app) that enables catalogue search, discovery and correct technical use.
Semantic Interoperability	The ability of systems to exchange data/app interfaces with unambiguous, shared meaning, enabled by common vocabularies and ontologies.
Usage Policy	A machine-readable rule set (e.g., ODRL profile) that defines how, by whom, and under what conditions data or apps may be used.
Vocabulary Provider	A service that curates and exposes controlled vocabularies or ontologies so that all participants use consistent terms.

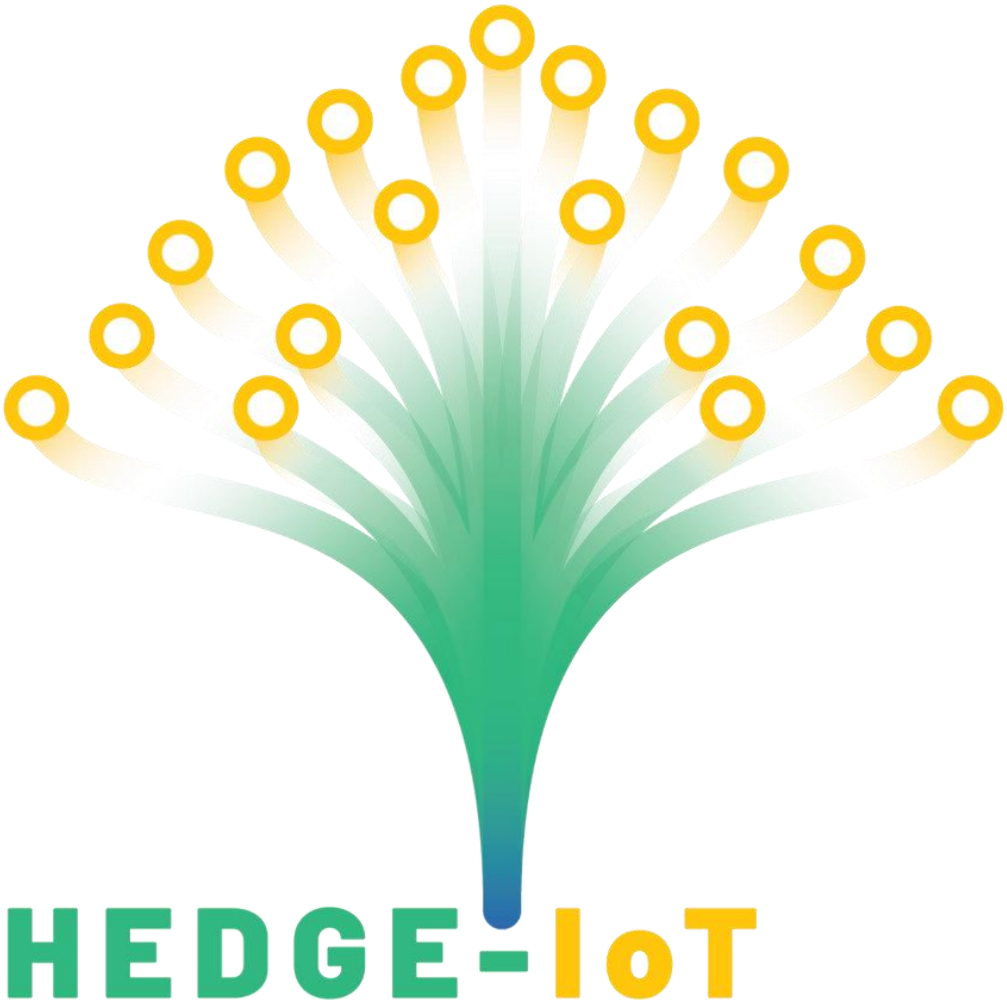

Co-funded by the European Union

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