O-CEI

D2.1:

Technical positioning and market analysis report

Revision: v1.0

Co-funded by the European Union

Open CloudEdgeIoT Platform Uptake in Large Scale Cross-Domain Pilots (O-CEI) has received funding by the EU's Horizon Europe program under Grant Agreement number 101189589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

D2.1: TECHNICAL POSITIONING AND MARKET ANALYSIS REPORT

Work package	WP 2
Task	T2.1 & T2.2
Due date	30/06/2025
Submission date	31/07.2025
Deliverable lead	NSCRD
Version	1.0
Editors	Vasilis Pitsilis (P02 NCSRD), Harilaos Koumaras (P02 NCSRD)
Reviewers	Eduardo Garo (P40 PRO) Constantinos Loizou (P41 EMBIO), Errol Ahmet (P41 EMBIO), Theofylaktos Apostolou (P41 EMBIO)
Abstract	This deliverable aims to assess the potential, depth, and challenges of enabling an open, cross-sector CEI continuum platform that fosters trusted and secure collaboration among market stakeholders. It does so by presenting both the collective technical positioning and capabilities of the O-CEI consortium, which integrates recent research achievements, validated components, and cross-domain experience derived from leading EU-funded projects and aligned with the progress of key European initiatives. In parallel, it provides a market analysis across eight core pillars, examining current adoption levels, the technical, behavioural, and regulatory barriers that hinder uptake, and the opportunities to scale and adapt technologies that can support a trusted, innovation-driven collaborative continuum.
Keywords	Cloud-Edge-IoT Continuum, MetaOS, Federated Orchestration, Cognitive Computing, Trust Management, Data Spaces, Interoperability, Data Space Connectors, Federated Marketplace, Zero Trust Security, Distributed Ledger Technologies (DLT), Function-as-a-Service (FaaS), Semantic Resource Modelling, Infrastructure Automation, Edge Intelligence, Digital Sovereignty, Hourglass Model, Digital Transformation, Market Uptake Barriers

DOCUMENT REVISION HISTORY

Version	Date	Description of change	List of contributor(s)
V0.1	08/04/2025	ToC and Initial Content	NCSRD Contributors
V0.2	05/05/2025	Introduction parts, Section 2	NCSRD Contributors
V0.3	12/05/2025	Section 3 initial content. Guiding framework for expected contributions.	NCSRD Contributors
V0.4	19/05/2025	Integrated section 3 contributions	All Contributors
V0.5	23/05/2025	Additional section 3 contributions	All Contributors
V0.6	25/05/2025	Section 4 Integration by BLUSPECS	BLU Contributors
V0.7	29/05/2025	Second round of contributions integration	All contributors
V0.8	02/06/2025	Updates, reviews, Glossary integration	All contributors
V0.9	12/06/2025	Internal review comments integrated	All reviewers
V1.0	29/07/2025	Submission ready version	NCSRD & UPV contributors

LIST OF CONTRIBUTORS

Name	Partner	Email
Alejandro Fornés	P01 UPV	alforlea@upv.es
Ignacio Lacalle	P01 UPV	iglaub@upv.es
Andreu Belsa	P01 UPV	anbelpel@upv.es
Vasilis Pitsilis	P02 NCSRD	vpitsilis@iit.demokritos.gr
Harilaos Koumaras	P02 NCSRD	koumaras@iit.demokritos.gr
Spyridon Georgoulas	P02 NCSRD	spygeorgoulas@iit.demokritos.ge
Andreas Sakellaropoulos	P02 NCSRD	asakellaropoulos@iit.demokritos.gr
Roger Briz	P03 INNO	rbriz@innovalia.org
Tanya Suarez	P4 BLU	tsuarez@bluspecs.com
Bianca Gresele	P4 BLU	bgresele@bluspecs.com
Sergio Cueva	P4 BLU	scueva@bluspecs.com
Mathieu Schumann	P05 EDF	mathieu.schumann@edf.fr
Francisco Blanquer	P06 CMA-CGM	ho.fblanquer@cma-cgm.com
Mihail Gaianu	P07 CONTI	mihail.gaianu@continental-corporation.com
Ernö Kovacs	P12 NEC	Ernoe.kovacs@neclab.eu
Benjamin Hebgen	P12 NEC	Benjamin.Hebgen@neclab.eu
Tobias Jacobs	P12 NEC	Tobias.Jacobs@neclab.eu
Lara López	P15 ATOS	lara.lopez@eviden.com
Panagiotis Tsarchopoulos	P18 NOVA	ocei@partners.nova.gr
Ioannis Markopoulos	P18 NOVA	Ioannis.Markopoulos@novaict.gr
Kaja Swat	P19 PS	kaja.swat@phoenix-rtos.com
Marek Białowąs	P19 PS	marek.bialowas@phoenix-rtos.com
Andrzej Głowiński	P19 PS	andrzej.glowinski@phoenix-rtos.com
Paweł Szmeja	P21 IBSPAN	pawel.szmeja@ibspan.waw.pl
Wiesław Pawłowski	P21 IBSPAN	wieslaw.pawlowski@ug.edu.pl
Ricardo Simón Carbajo	P22 CeADAR	ricardo.simoncarbajo@ucd.ie
Andrés L. Suárez-Cetrulo	P22 CeADAR	andres.suarez-cetrulo@ucd.ie
Miguel Aspis	P22 CeADAR	miguel.aspis1@ucd.ie
Georgios Stavropoulos	P24 CERTH	stavrop@iti.gr
Georgios Spanos	P24 CERTH	gspanos@iti.gr
Vaios Bolgouras	P26 UPRC	vaiosbolgouras@ssl-unipi.gr
Michalis Takaronis	P26 UPRC	mtakaronis@ssl-unipi.gr
Medina Saracevic	P27 THD	medina.saracevic@th-deg.de
Madalin Neag	P28 ONS	mneag@opennebula.io
Antonio Alvarez	P28 ONS	aalvarez@opennebula.io

Antonio Kung	P29 TRIALOG	Antonio.kung@trialog.com
Cécile Rabrait	P29 TRIALOG	Cecile.rabrait@trialog.com
Philippe Krief	P30 ECL	philippe.krief@eclipse-foundation.org
Rosaria Rossini	P30 ECL	rosaria.rossini@eclipse-foundation.org
Olga Kalinowska	P31 CHOCO	olga@chocolate-cloud.cc
Athanasia Sampazioti	P32 IQB	athanasia.sampazioti@inqbit.io
Hrvoje Ratkajec	P33 XLAB	hrvoje.ratkajec@xlab.si
Carolina Viceto	P34 UBW	cviceto@ubiwhere.com
João Viegas	P34 UBW	jviegas@ubiwhere.com
Ivan Paez	P35 ZSCALE	ivan@zettascale.tech
Filisia Melissari	P36 SYNELIXIS	melissari@synelixis.com
Alexandra Lakka	P36 SYNELIXIS	lakka@synelixis.com
Christos Syllogidis	P36 SYNELIXIS	syllogidis@synelixis.com
Antonio D'Averio	P38 ARES2T	antonio.daversio@ares2t.com
Silvia Canale	P38 ARES2T	silvia.canale@ares2t.com
Vincenzo Suraci	P38 ARES2T	vincenzo.suraci@ares2t.com
Eduardo Garro	P40 PRO	egarro@prodevelop.es
Antonio Alcudia	P40 PRO	aalcudia@prodevelop.es
Constantinos Loizou	P41 EMBIO	k.loizou@embiodiagnostics.eu
Errol Ahmet	P41 EMBIO	e.ahmet@embiodiagnostics.eu
Theofylaktos Apostolou	P41 EMBIO	theo.apo@embiodiagnostics.eu
Nikos Kalatzis	P44 GSC	nikosk@greensupplychain.eu
Nikos Marianos	P44 GSC	n.marianos@greensupplychain.eu
Eusebiu Catana	P45 ERT	e.catana@mail.ertico.com
David Campo	P47 FIWARE	david.campo@fiware.org
Peter Bresseleers	P49 FENIX	p.bresseleers@fenix-network.org
John White	P53 HESSO	john.white@hesge.ch

| **Grant Agreement No:** 101189589 | **Topic:** HORIZON-CL4-2024-DATA-01-03
| **Call:** HORIZON-CL4-2024-DATA-01 | **Type of action:** HORIZON Innovation Actions

DISCLAIMER



Open CloudEdgeIoT Platform Uptake in Large Scale Cross-Domain Pilots (O-CEI) has received funding by the EU's Horizon Europe program under Grant Agreement number 101189589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

COPYRIGHT NOTICE

O-CEI © 2025-2028

Project funded by the European Commission in the Horizon Europe Programme		
Nature of the deliverable:	R	
Dissemination Level:		
PU	Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page)	X

* *R: Document, report (excluding the periodic and final reports)*
DEM: Demonstrator, pilot, prototype, plan designs
DEC: Websites, patents filing, press & media actions, videos, etc.
DATA: Data sets, microdata, etc.
DMP: Data management plan
ETHICS: Deliverables related to ethics issues.
SECURITY: Deliverables related to security issues
OTHER: Software, technical diagram, algorithms, models, etc.

TABLE OF CONTENTS

1 ABOUT THIS DOCUMENT 13

1.1 Deliverable context 13

1.2 The rationale behind the structure 13

2 INTRODUCTION TO O-CEI 15

3 TECHNOLOGY OOSITIONING: RESEARCH OUTCOMES AND INNOVATIVE COMPONENTS 17

3.1 Reference architecture and technological pillars for the CEI continuum 17

3.1.1 TF3 reference architecture for the CEI continuum 17

3.1.2 Core pillars of the CEI technology landscape 19

3.1.3 Continuum orchestration, MetaOS 22

3.1.4 Security, trust and digital identity 31

3.1.5 Data spaces, interoperability 39

3.1.6 AI-Driven computing & cognitive cloud 50

3.2 Analysis of recent research projects and outcomes 58

3.3 Mapping utilities to TF3 RA building blocks 69

3.3.1 Orchestration 73

3.3.2 Network 76

3.3.3 Monitoring and Observability 78

3.3.4 Data Management 80

3.3.5 Resource Management 85

3.3.6 Trust and Reputation 86

3.3.7 Security and Privacy 87

3.3.8 AI 88

4 MARKET ANALYSIS AND ADOPTION BARRIERS 93

4.1 Market landscape and technology adoption trends 93

4.1.1 Key industry sectors impacted by CEI computing continuum 93

4.1.2 A firm level view of CEI 96

4.2 Identification of adoption barriers 98

4.2.1 Technical barriers: challenges in integrating emerging technologies 98

4.2.2 Economic barriers: cost implications, ROI concerns 99

4.2.3 Organisational barriers: skills gap, industry readiness 99

4.2.4 Legal and ethical barriers: data-sharing constraints, liability concerns 100

4.3 Regulatory and standardization considerations 100

4.3.1 Complying with regulations and standards 100

4.3.2 Impact of EU regulations (Data Act, AI Act, NIS2, etc.) 101
4.3.3 Standardisation efforts and compliance challenges 102
4.3.4 Combining ecosystem and technology considerations 103
4.4 A scenario-based approach to CEI adoption 104
4.4.1 Gaining an initial understanding of the market context of each O CEI pilot 104
4.4.2 Key takeaways and industry requirements 116
4.5 From pilot to operations at scale: initial recommendations 117
5 CONCLUSIONS 120
REFERENCES 121
APPENDIX A: O-CEI GLOSSARY 123
APPENDIX B: PILOT WORKSHOP METHODOLOGY 127
APPENDIX C: IMPACT OF KEY EU REGULATIONS ON CEI PILOTS 129

EXECUTIVE SUMMARY

This document, **"D2.1 – Technical Positioning and Market Analysis Report"**, defines the strategic technological foundation of the O-CEI project by aligning recent research outcomes and innovation components with real-world Cloud-Edge-IoT (CEI) deployment needs. It provides a structured view of the current technological landscape, focusing on technologies that enable comprehensive CEI usage across sectors. This technological positioning builds on the combined expertise and domain knowledge of project partners, as developed within a wide range of EU-research funded projects. It also incorporates research outcomes and a variety of innovative tools and components that partners bring into the project. The overall goal is to support the development of a flexible, scalable, and trusted CEI ecosystem for Europe.

At the core of the technical positioning is the **EUCloudEdgeIoT (TF3) Reference Architecture**, developed through the analysis of numerous European research and innovation projects in the fields of edge, cloud, and IoT, primarily under Horizon 2020 and Horizon Europe. The TF3 RA identifies **eight essential building blocks**: Orchestration, Network, Monitoring and Observability, Data Management, Resource Management, Trust and Reputation, Security and Privacy, and AI. These building blocks serve as reusable, modular components, each embedding key capabilities required to operate across heterogeneous CEI infrastructures and manage the continuum effectively.

However, these modular blocks alone do not represent complete, deployable systems. To support real-world application and integration, O-CEI introduces **four broader technological pillars** that reflect how CEI technologies converge in practice:

- Continuum Orchestration and MetaOS
- Data Spaces and Interoperability
- Security, Trust, and Digital Identity
- AI-Driven Computing and Cognitive Cloud

Each pillar brings together functionalities of multiple TF3 building blocks, addressing functional, operational, and domain-specific requirements. The first part of this deliverable describes the current technological landscape based on partner contributions around these pillars, highlighting key innovations, and presenting the collective expertise of the O-CEI consortium, this way setting a unified baseline from which subsequent project developments will be derived. The document also includes in this, first, part an analysis of outcomes from several relevant EU projects, emphasizing their contributions to O-CEI's objectives and pillars, as well as a structured mapping of available tools and components to TF3 building blocks — forming the project's baseline technical inventory.

The second part of this deliverable complements the technology positioning with a structured market analysis, focusing on current adoption status analysis, possible applicability, adoption trends, and perceived barriers. This analysis is grounded in input from the project's **eight pilot use cases**, collected through a series of dedicated **workshops and interviews**, all following a common methodology. These sessions gathered insights from industry stakeholders and domain experts, helping to identify both sector-specific requirements and broader CEI adoption challenges. This dual approach to merge technological insight with real-world market feedback lays the foundation for O-CEI's adoption and integration strategy, and the future innovation roadmap.

In summary, this deliverable serves as both a collective technical positioning identity and a strategic reference, aligning CEI research, architecture, and innovation with actionable directions for real-world adoption and impact.

LIST OF FIGURES

Figure 1: TF3 Building Blocks for CEI 17
Figure 2: Synthesis of TF3 Building Blocks into O-CEI Pillars and Innovation Model..... 20
Figure 3: Compiling CEI on TF3 Building Blocks 21
Figure 4: Cloud Edge IoT Continuum 23
Figure 5: Continuum Components layered interactions 24
Figure 6: Cloud-edge-IoT continuum perspective source: EUCloudEdgeIoT 28
Figure 7: IOTA in aerOS continuum 37
Figure 8. Conceptual architecture of data fabric. 48
Figure 9. Data exchange between heterogeneous sources and consumers is achieved through a data fabric built on a knowledge graph. 48
Figure 10: Indicative MCP architecture 56
Figure 11: IPCEI-CIS Reference Architecture, version 1.0 61
Figure 12: Bridging the devices to the cloud-edge continuum through edge applications .. 64
Figure 13: OpenNebula components 73
Figure 14: Grafana example dashboard 78
Figure 15: Zenoh example topology 81
Figure 16: NGSI-LD Data Model with Meta-Model and RDF/RDFS Grounding (source: ETSI/ISG CIM) 83
Figure 17: Context Broker Deployment Mode 83
Figure 18: AI Agent Distribution across the CEI Continuum 84
Figure 19: OpenNebula AIOps module architecture 89
Figure 20: Example MLFLow tracking dashboard 90
Figure 21: Kubeflow example lifecycle 92
Figure 22: CEI Sphere D2.1 – Preliminary market analysis. European CEI Solution Spending Forecast, 2024-2028 (EUR Bn) (Source: IDC, Worldwide Internet of Things Spending Guide, November 2024. 93
Figure 23: Adapted from CEI Sphere D2.1 – Preliminary market analysis. European CEI Solution Spending Forecast, 2024-2028 (EUR Bn) (Source: IDC, Worldwide Internet of Things Spending Guide, November 2024). 94
Figure 24: Proportion of edge workloads best suited to run in each deployment categories (Source: IDC EC Edge Observatory Custom Survey 2023/2024: Edge Computing (Waves 1, 2), n=600) 95
Figure 25: Top reasons to deploy Edge 97
Figure 26: Compliance ecosystem 102
Figure 27: Compliance integration challenge example 102
Figure 28: Hourglass model 103

LIST OF TABLES

Table 1: Overview of Orchestration and Networking Solutions in O-CEI 30
Table 2: Overview of Security, Trust, and Digital Identity Solutions 38
Table 3: Data sharing Infrastructure in O-CEI 43
Table 4: Semantic and Runtime Interoperability Enablers 49
Table 5: AI Solutions Integration in O-CEI 57
Table 6: Overview of relevant research projects informing O-CEI 58
Table 7: Component capabilities connected to TF3 building blocks 70
Table 8: O-CEI Pilot 1 “Electricity Grid performance optimisation upon RES integration” KPIs 105
Table 9: O-CEI Pilot 2 “Transition to Software-Defined Vehicle (SDV) platform” KPIs 107
Table 10: O-CEI Pilot 3 “Energy consumption and emission reductions in postal Service fleet operation via intelligent BEV charging strategies” KPIs 109
Table 11: O-CEI Pilot 4 “Variable demand in challenging maritime terminal landscape” KPIs… 110
Table 12: O-CEI Pilot 5 “Energetically and environmentally sustainable Halloumi cheese production” KPIs 112
Table 13: O-CEI Pilot 6 “Smart re-charging and efficiency of robot tractors in large fruit production fields” KPIs 113
Table 14: O-CEI Pilot 7 “Trustworthy and secure EV charging upon reliable 5G networks” KPIs 115
Table 15: O-CEI Pilot 8 “Heightened social engagement and acceptability of energy flexibility in urban areas” KPIs 116
Table 16: Initial recommendations for pilots 118

ABBREVIATIONS

AIaaS	AI as a Service
API	Application Programming Interface
CAPIF	Common API Framework for Northbound Interfaces
CEI	Cloud Edge IoT
CI/CD	Continuous Integration Continuous Deployment
CL4	Cluster 4
CNC	Centralized Network Configurator
CNN-LSTM	Convolutional Neural Network + Long Short-Term Memory
COTS	Commercial Off-The-Shelf
CSDM	Common Semantic Data Model
CVEs	Common Vulnerabilities and Exposures
DAG	Directed Acyclic Graph
DID	Digital ID
DIDs	Decentralized Identifiers
DLT	Distributed Ledger Technology
DPL	Distributed Persistence Layer
EDC	Eclipse DataSpace Connector
ETSI	European Telecommunications Standards Institute
FAIR	Findable, Accessible, Interoperable, and Reusable
FL	Federated Learning
FaaS	Function as a Service
HPC	High-Performance Computing
IAAM	Identity, Authentication and Authorization Management
ICT	Information and Communication Technology
IDS	International Data Spaces
IDS	Intrusion Detection System

IDSA	International Data Spaces Association
IE	Infrastructure Element
IPSM-AF	Interoperability Profile Semantic Mediation - Architectural Framework
IaC	Infrastructure as Code
JSON	JavaScript Object Notation
LLM	Large Language Model
LSP	Large Scale Pilot
LSTM	Long Short-Term Memory
LXC	Linux Containers
MCP	Model Context Protocol
MLP	Multilayer Perceptron
MetaOS	Meta Operating System
NEF	Network Exposure Function
NGSI-LD	Next Generation Service Interfaces - Linked Data
ODCT	Ontology-Driven Constraint Tester
ODPS	Open Data Product Specification
ODRL	Open Digital Rights Language
OID4VC	OpenID for Verifiable Credentials
PAP	Policy Administration Point
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PIP	Policy Information Point
QoS	Quality of Service
RA	Reference Archtecture
RAG	Retrieval-Augmented Generation
RBAC	Role-Based Access Control
RDF	Resource Description Framework
RML	RDF Mapping Language

RTOS	Real Time Operating System
SBOM	Software Bill of Materials
SHACL	Shapes Constraint Language
SLO	Service Level Object
TF3	Task Force 3
TOSCA	Topology and Orchestration Specification for Cloud Applications
VCs	Verifiable Credentials
VM	Virtual Machine
VPN	Virtual Private Network
XML-RPC	eXtensible Markup Language - Remote Procedure Call

1 ABOUT THIS DOCUMENT

1.1 DELIVERABLE CONTEXT

Item	Description
Objectives	Deliverable D2.1 directly supports project objective O2 by providing a comprehensive analysis of CEI technologies, research outcomes, innovative components and EUCEI architecture patterns with the aim to converge towards shaping a broader, flexible ecosystem capable of adapting to the diverse needs of industry verticals. Additionally, it directly addresses O3 by bringing together the technological outcomes of research and technology providers with the perspectives of industrial stakeholders. This enables understanding of the current level of CEI technology integration, while also identifying key barriers and constraints that should be addressed to enable wider adoption.
Work plan	Deliverable D2.1 sets the baseline for defining the technical approach of the project and is triggering subsequent design and development technical activities. At the same time, it is also closely related to project pilots (WP4) as it provides analysis of the current technology and status of the market of the various verticals.
Milestones	The submission of Deliverable D2.1 contributes to fulfilment of MS1: Identity definition. Web site is online since M1 and now D2.1 provides the establishment of market and technical positioning for O-CEI. As a result, MS1 is now considered complete.
Deliverables	Deliverable D2.1 is the initial deliverable in the project's series, operating independently without relying on input from previous deliverables. It is expected to provide input and serve as a baseline for subsequent technical deliverables D2.2 and D3.1.
Risks	There are no foreseen nor identified risks associated to this deliverable.

1.2 THE RATIONALE BEHIND THE STRUCTURE

The content of this deliverable is organised in six main sections which can efficiently report on T2.1 and T2.2, with the aim to serve both as a comprehensive technology positioning reference and as a basis for understanding market dynamics and adoption challenges

- **Section 2** introduces the reader to O-CEI project, exposing its motivations, objectives and main proposals. This section sets the stage for understanding how O-CEI aims to contribute towards a sustainable and resilient European CEI ecosystem.
- **Section 3** is the first of the two core parts of this report and is dedicated to positioning O-CEI within the broader technological landscape. It does so by leveraging recent research advancements and identifying how these innovations converge into a coherent framework for enabling flexible and federated CEI continuum deployments. This section is divided into three parts, each contributing a different layer of understanding. The first part functions as the project's technical positioning report. It examines the status and relevance of the TF3 Reference Architecture (RA), highlighting its role in shaping interoperable CEI infrastructures. It then

explores the most impactful technological pillars driving CEI evolution, such as orchestration, data spaces, security, AI, and how these domains together support the vision of an open collaboration space for data and service sharing across industries. The second part complements this view by surveying recent and ongoing European research projects, especially those focused on MetaOS, Cognitive and Swarm Computing, and other CEI-relevant topics, and identifying how their matured research outcomes align with and contribute to the goals of O-CEI. The third part provides a structured mapping of research outputs, tools, and components, either operated or contributed by project partners, to the TF3 RA building blocks. This mapping serves as an early reference for subsequent technical tasks and deliverables, plus helps to inform how platform components and services will be assembled in the O-CEI framework.

- **Section 4** presents the second core part of the deliverable, focusing on the market landscape and adoption barriers. It starts by examining current technology adoption trends and the market impact of CEI technologies across different sectors. It also provides an initial analysis of the regulatory and standardisation environment, acknowledging that more in-depth work will be presented in subsequent deliverables. Section 4 further incorporates insights gathered from interviews and workshops held with pilot stakeholders, highlighting technical, administrative, legal, and industrial barriers to adoption. Based on these insights, strategic recommendations are proposed to enhance the scalability and adoption of O-CEI outcomes.
- **Section 5** concludes the deliverable with an overview of the main findings and challenges that O-CEI may encounter as the project evolves.
- **Last section** provides references and appendices, including a glossary to ensure clarity and consistency in the use of key terms throughout the document.

2 INTRODUCTION TO O-CEI

The digital transformation of Europe is progressing, yet it faces significant challenges. A rapidly growing volume of data, increasingly distributed and heterogeneous computing resources, and fragmented technological solutions have created a complex landscape. While data generation explodes, driven by over 43 billion connected devices globally, including more than 15 billion IoT devices today, Europe continues to rely heavily on non-European digital platforms and technologies. This creates dependencies, market imbalances, and limits Europe's ability to secure its data economy and digital sovereignty.

Amid this complexity, the **computing continuum paradigm** has emerged as a key enabler for Europe's digital autonomy and competitiveness, and **Edge computing** is increasingly prioritized by the European Union as a critical layer in the CEI landscape. While businesses recognize edge computing value in enabling low-latency, secure, and resilient processing close to data sources, reducing transmission bottlenecks and improving responsiveness, they also face significant hurdles in implementation, including complexity, interoperability, and rising operational costs.

With the goal of addressing challenges within these technological fields, European actors, over recent years, have generated a wealth of research and innovation within the grounds of numerous EU funded projects, delivering critical advances in areas such as MetaOS, Cognitive Cloud, Swarm Computing, Data Spaces, or Federated Learning. Research and Innovation projects have produced significant outputs and have developed and delivered valuable building blocks and technological breakthroughs that now form a strong foundation for the CEI continuum. In parallel, major initiatives born in Europe, such as **Gaia-X, Catena-X, Omega-X, and IDSA**, have been launched to foster trusted federated data ecosystems and interoperable digital infrastructures across sectors. BDVA, Gaia-X, the FIWARE Foundation, and IDSA joined forces to create the Data Spaces Business Alliance (DSBA), which played a major role in establishing the consortium behind the Data Spaces Support Centre (DSSC), a project funded by the EU under the Digital Europe Programme, which provides recommendations for the deployment of data spaces

While these efforts are brought together by many partners across a multitude of scientific and technological fields, many of their results remain disconnected, and the opportunity to work together under a flexible, unified framework has yet to be broadly explored. As a result, the full potential to collectively support an ecosystem capable of delivering steady or ad-hoc solutions for European market stakeholders remains, to some extent, underexploited. A natural next step, meaning both a challenge and an opportunity, is to bring these diverse outcomes together under a **synergistic and interoperable framework**, integrating the technological outputs of research projects with the ecosystem models promoted by DSBA. Such an effort would also complement Europe's evolving regulatory landscape, including the Data Act, AI Act, and other digital regulations, to ensure that innovation aligns with principles of **trust, sovereignty, and openness**. Together, these advancements would strengthen Europe's position in building a cohesive CEI ecosystem that enables reuse, integration, and cross-sector collaboration, driving digital transformation forward.

O-CEI approach

Building on the foundational work of initiatives like EUCloudEdgeIoT, and aligned with Europe's vision for digital sovereignty, trust, and competitiveness, the O-CEI project is designed to accelerate innovation, foster collaboration, and promote industrial uptake across the CEI continuum. The project builds upon valuable research outcomes and technological results generated by recent and ongoing EU-funded projects, bringing together diverse innovations in fields such as orchestration, trusted data sharing, interoperability, security, and AI-driven intelligence. By synthesizing these outcomes

into a unified, flexible, and sustainable framework, O-CEI empowers stakeholders to develop and deploy secure, trusted, and interoperable digital services. In doing so, it addresses the challenge of fragmentation by establishing a flexible platform where standard-based architectural patterns support the sharing and exploitation of data, services, and computational resources across industry domains for mutual benefit. O-CEI aims not only to advance technological innovation but also to ensure these innovations are implemented in Europe, safeguarding economic security and enabling participation in the most strategic segments of the digital value chain.

In line with key concerns reflected in Europe's strategic priorities, including the vision of a "Connected Collaborative Computing" Network (3C Network) [1], O-CEI integrates technological advances from research with practical approaches to trusted, interoperable, and collaborative processing across the CEI continuum. Focusing on technical areas such as the combination of best-of-breed technologies for continuum orchestration, efficient data management, security, trust, and AI-driven processing across edge and cloud environments, the project develops patterns, blueprints, and architectural frameworks that enable the integration of computational resources, facilitate data sharing, and support the development of innovative applications and services across all sectors. Through this approach, **O-CEI leverages research findings to support broader European initiatives and enable flexible, reproducible deployments that deliver targeted technical solutions to a diverse range of market stakeholders,** aiming to act as a multiplier for the adoption of digital innovations**.** By aligning technological, governance, and regulatory considerations, the project contributes to shaping an ecosystem that promotes openness, trust, interoperability, and sovereignty, bridging research outcomes with industrial needs and supporting Europe's path towards a more competitive and resilient digital economy.

O-CEI will validate its framework through the development and deployment of **eight large-scale pilots addressing a number of real-life challenges across several strategic sectors**: optimizing electricity grids, advancing software-defined vehicles, enhancing smart charging for electric postal fleets, improving energy management in maritime ports, promoting sustainable dairy production, increasing efficiency in e-tractor operations for kiwi crops, deploying private 5G networks for electric vehicle charging, and fostering social acceptance of energy flexibility in urban areas. These diverse pilots will demonstrate the framework's scalability, cross-sector applicability, interoperability, and adaptability to real-world challenges, while also addressing key concerns such as rising energy consumption, operational costs, and environmental impact.

3 TECHNOLOGY OOSITIONING: RESEARCH OUTCOMES AND INNOVATIVE COMPONENTS

3.1 REFERENCE ARCHITECTURE AND TECHNOLOGICAL PILLARS FOR THE CEI CONTINUUM

3.1.1 TF3 reference architecture for the CEI continuum

During 2023 and 2024, more than 25 ongoing research projects collaborate to design a Reference Architecture for the continuum that set the basis for further implementations in order to be able to develop interoperable, and reusable, software. [2][3][4]

As a result of this work, 8 technological building blocks were identified. Each of the building blocks represents and groups functionalities that are mandatory when developing software for the CEI continuum. Thus, these building blocks serve as the **core components** of any architecture implementing technical solutions. This means, that wherever a new architecture is designed to cover new research challenges or to solve specific sectorial issues, it must include, at least, the functionalities identified in each of these building blocks and evolve or include other as needed, based on the requirements previously identified. Though not all building blocks are mandatory, as their further inclusion relies on specific project needs, as previously stated, they should be carefully considered. Those building blocks were represented as follows:

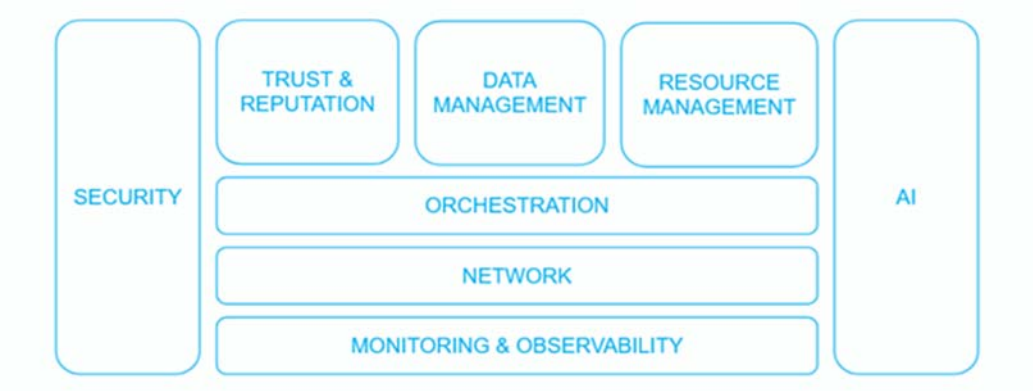


Figure 1: TF3 Building Blocks for CEI

- **Security & Privacy:** Common CEI solutions include a wide variety of stakeholders (including users and providers), physical resources (mainly infrastructure and devices), datasets (data from different sources) and network connections and it is necessary to ensure security at all levels, from simple access to the platform to all communications, even ensuring all agents involved (stakeholders, physical resources, datasets or network connections) comply with the current legislation. Thus, the minimum set of functionalities already identified are as follows:
    - *Authorization and Authentication* mechanisms, including *Attestation* control, managing the system access according to predefined roles to ensure all agents are allowed to be part of it. This also includes:
        - *Registry* management through a catalogue containing information about all agents and their associated role(s).
        - *Identity and access management* mechanisms to ensure only those agents included in the registry catalogue are allowed to access the system.

- Ensure *data integrity and availability*, as well as *System integrity*, through trustable, accessible and coherent data.
- *Secured communications* to ensure the safeness of all communications within the system. Additionally, data *encryption* is also recommended.
- *Isolation* mechanisms to minimize hardware architectures related threats.
- And, optionally, only for economic transactions *accountability and transaction recording* is needed.

- **Trust & Reputation:** Additional to the security layer, system and provider trust and reputation measures are needed to support user informed decisions. Thus, there are different things to be taken into consideration:
  - *Calculation* mechanisms based on the available information about the trust level of a given system, including information from the hardware provider as well as user one based on their actions inside the system.
  - *Reputation* ranking mechanisms based on the available information to provide recommendations to the user for selecting the most suitable provider.
- **Data management:** Data is one of the key pillars of the continuum, as it is consumed internally by the system, but also by the applications deployed on it. In this case, three major functionalities were identified, although the implications of it are quite deep:
  - *Storage* capabilities, as data must be accessible at any time for any action.
  - Registration in a *catalogue,* so specific dataset can be accessed by any component.
  - *Analytics* under demand, to perform informed actions.
- **Resource management:** When performing any action in the continuum, there are a lot of physical resources involved, from cloud/edge infrastructure to IoT devices. Thus, in order to be able to execute an application following the most suitable deployment scheme in an efficient manner, it is necessary to properly manage all of them. The basic identified functionalities are as follows:
  - Basic *management* functionalities for smart decision-making. As well as a *catalogue* containing the list of all the resources allowed to be part of the system.
  - *Allocation and scheduling* functionalities to identify all available resources and set up the needed spot (group of resources needed to perform a given action). This includes *co-allocation* functionalities to create some short of federated spot with resources from different providers.
  - *Discovery* mechanisms to identify not only the available resources, but the ones that can better fulfil the requirements to execute an application.
- **Orchestration:** This is the core pillar of any system, as it allows not only to deploy an application, but also interconnects the network and physical resources. Thus, different capabilities were identified as the basis for designing and implementing any system:
  - Connect with the resource management building block, establishing the needed requirements in order to create a *federation* of cloud/edge infrastructure and IoT devices for any given deployment and/or redeployment.
  - The *service description* itself to include all the requirements at application level.
  - Two types of orchestrator: *low-level* to directly interact and manage network, resources and services as needed, and *high-level* to develop the most suitable deployment scheme. The latter also including *re-orchestration* capabilities in case something is going wrong during the execution.

- **Network:** Although network management is possible in private environments only, network constraints are of paramount importance when interconnecting cloud, edge and IoT environments.
    - Thus, orchestration and slicing functionalities should be applied wherever feasible.
- **Monitoring & Observability:** This building block covers all aspects of QoS and QoE, taking into account that monitoring the system health prevents potential malfunctions. Thus, there are two major functionalities identified:
    - *Monitoring* of physical resources, infrastructure and telemetry, as well as the application performance itself.
    - *Observability* capabilities, collecting the information within the SLOs *catalogue* to perform the system maintenance and raise alerts, sometimes, even before they may occur.
- **Artificial Intelligence:** In this case, artificial intelligence is transversal to all building blocks and applied across all actions. For instance, its supportive role can be seen in its integration with the scheduling process, where it assists the high-level orchestrator in making informed decisions based on available information.

3.1.2 Core pillars of the CEI technology landscape

Building on the analytical foundation established in the previous section through the presentation of the EUCloudEdgeIoT TF3 Reference Architecture, this section shifts focus from the analysis of architectural components toward their synthesis into higher-level technical pillars that reflect key capability areas across the continuum. While TF3 has successfully harmonized concepts and classified technological outcomes into eight key building blocks, it also recognized that establishing a common taxonomy to homogenize concepts across the continuum, from cloud to edge to IoT, was necessary to create a shared language that supports integration and interoperability. This effort not only standardized terminology but implicitly laid the groundwork for combining and aligning these building blocks within higher-level technological pillars capable of addressing the multifaceted challenges of CEI implementations.

In addition to providing a common taxonomy, it is important to recognize that while the eight building blocks offer a clear analytical structure, they remain relatively discrete and segmented by nature, each addressing a specific technical area. Individually they do not fully capture the integrated, cross-cutting technological positioning required to address the interconnected challenges of deploying a functioning CEI continuum. Leveraging the composable nature and foundational capabilities of the building blocks, we provide for broader, cross-functional pillars that integrate multiple components into flexible, architecture layers supporting practical integration pathways and scalable deployments across domains tailored to the demands of CEI implementations.

In accordance with established methodological frameworks, this transition represents a move from analysis, identifying and deconstructing architectural elements, toward synthesis: combining these elements into integrated layers of technological capability. The TF3 Reference Architecture offers a structured inventory of essential building blocks and a shared language to describe them; the challenge now is to contextualize and interrelate these components in ways that can effectively address the operational, technical, and organizational challenges of implementing a CEI continuum.

To this end, we introduce four core technological pillars, designed to integrate and operationalize the concepts and components of the TF3 building blocks. These pillars combine building block-level elements into higher-level functional areas, providing a structured perspective that supports the deployment of interoperable, trusted, and collaborative CEI solutions:

- Continuum Orchestration and MetaOS
- Data Spaces and Interoperability
- Security, Trust, and Digital Identity
- AI-Driven Computing and Cognitive Cloud

It is important to note that these pillars span the entire CEI continuum, leveraging different TF3 building blocks depending on the system layer (e.g., device, edge, cloud) and the specific deployment scenario (e.g., latency-sensitive industrial application vs. cloud-scale analytics). For instance, the “MetaOS and Orchestration” pillar may rely more heavily on the Orchestration and Resource Management building blocks in cloud and edge contexts, while Security and Trust may emphasize Trust & Reputation and Security & Privacy components at the device or IoT level. By leveraging the appropriate tools and functionalities from each building block, the pillars act as integrative frameworks that bridge the functional gaps between IoT, edge, and cloud environments, supporting seamless orchestration, data flow, and service interoperability across the continuum. These pillars reflect the intersection of technological themes recurring across research outcomes, project contributions, and stakeholder needs. They build on TF3’s taxonomy and standardized building blocks by framing them within functional domains that address cross-cutting challenges in CEI implementations. Each pillar integrates multiple building blocks, translating abstract architectural constructs into applied technological capacities tailored to real-world deployment contexts.

Figure 2 below provides a conceptual visualization of the synthesis process described in this section. At its foundation lie the eight TF3 Reference Architecture building blocks, the analytical outcome of EUCloudEdgeIoT efforts to classify, structure, and define the core enablers required to support deployments across the CEI continuum. These modular components are leveraged and composed to form the four O-CEI technological pillars, which collectively represent the project consortium’s technical positioning.

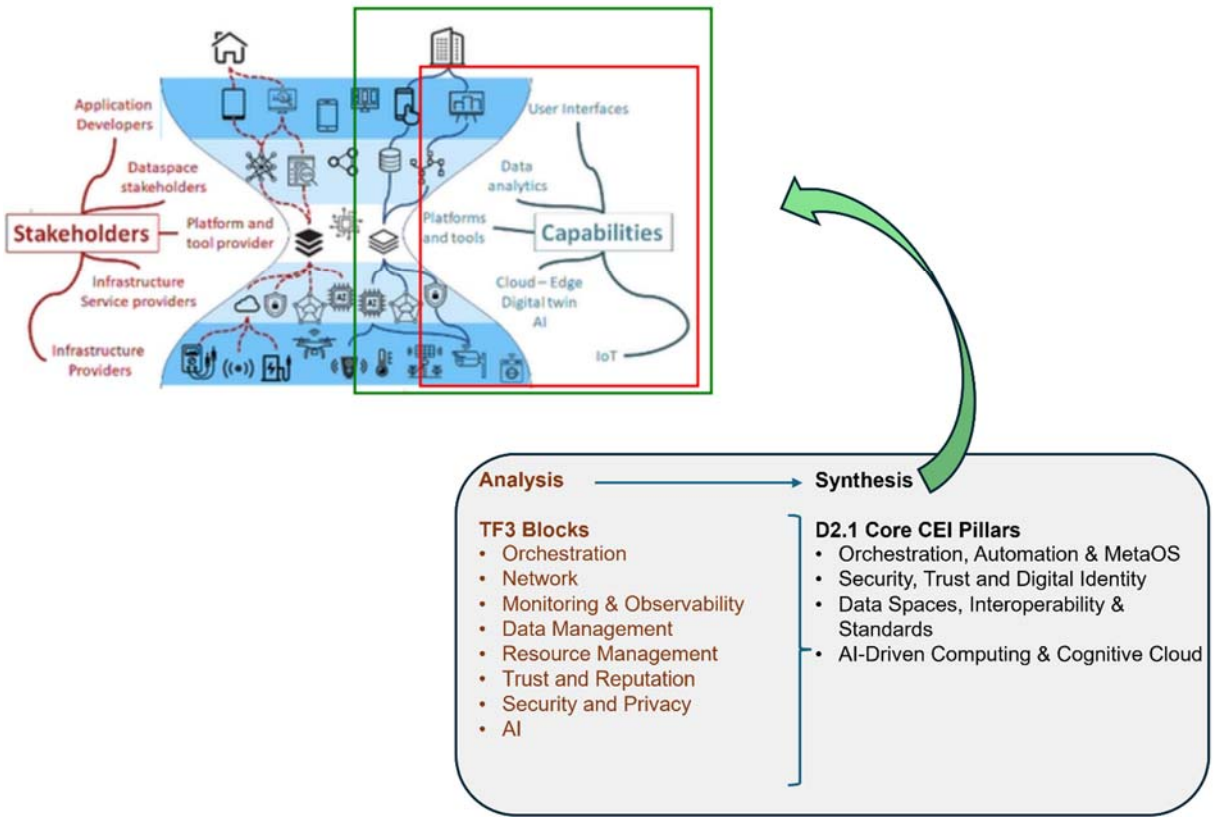


Figure 2: Synthesis of TF3 Building Blocks into O-CEI Pillars and Innovation Model

This synthesis is also mapped onto the hourglass model, discussed in section 4.3.4, which highlights the progression from foundational capabilities to market-facing applications. It reflects the alignment between TF3-derived technical constructs and the innovation-oriented pillars of O-CEI, illustrating how technical assets are translated into broader business capabilities and integration scenarios.

Figure 3, shown below, illustrates how the four O-CEI pillars stack together to ensure CEI functionality:

- At the infrastructure layer, the Continuum Orchestration & MetaOS pillar anchors operations, orchestrating distributed compute, network, and IoT resources.
- Above it, the Data Spaces & Interoperability pillar builds on these services and raw data, transforming them into structured, interoperable, and shareable assets.
- At the top, the AI-Driven Computing & Cognitive Cloud pillar leverages both orchestrated infrastructure and interoperable data to deliver intelligent behavior, automation, and cross-domain collaboration.
- Cutting across and underpinning all layers, the Security, Trust & Digital Identity pillar acts as a transversal enabler, ensuring secure interactions, identity management, and data integrity across the full continuum—from infrastructure orchestration to application execution.

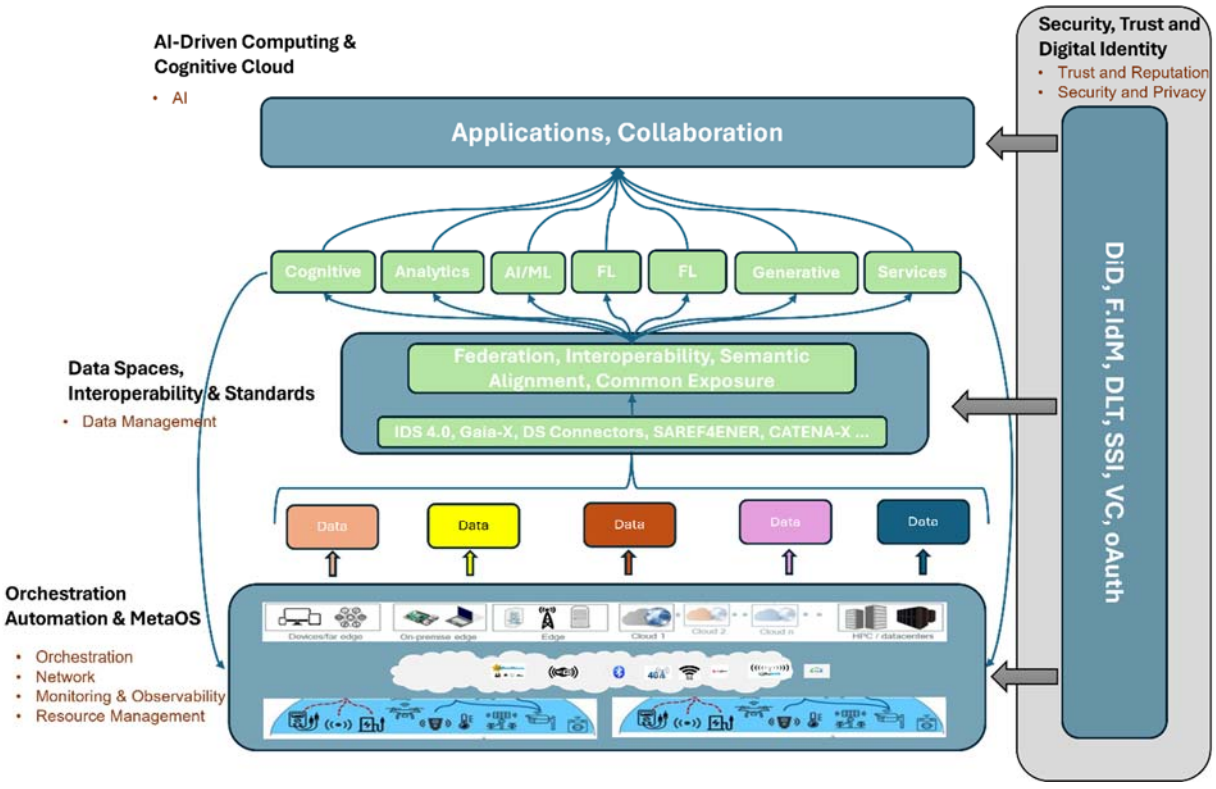


Figure 3: Compiling CEI on TF3 Building Blocks

In the following parts of this section, the descriptions provided within each pillar aim to capture the collective technical positioning of the O-CEI project consortium. Each partner's contribution reflects their insights, expertise, and prior work in relevant research projects, bringing forward mature technological outcomes and innovations that can serve as building blocks for implementation. Together, these descriptions coordinate and merge diverse research knowledge into a cohesive narrative that defines the project's initial technological baseline. This collective effort not only highlights key advances and trends but also lays the groundwork for leveraging these outcomes toward concrete, operational solutions within O-CEI's implementation trajectory.

By structuring the technological landscape through these four pillars, this section provides a solution-oriented synthesis that builds on the analytical clarity of the TF3 Reference Architecture while moving toward a more integrative understanding of how these components can work together to enable the O-CEI project's goals, establishing a shared technical baseline from which flexible, operational deployments and cross-domain solutions can be developed and validated.

The following subsections present each of these four pillars in more detail, describing their technological scope, their relationship to the TF3 Reference Architecture, and the insights contributed by O-CEI partners.

3.1.3 Continuum orchestration, MetaOS

A critical enabler of the CEI continuum is the ability to coordinate and manage computational, network, and data resources seamlessly across distributed environments. Continuum Orchestration and MetaOS emerge as foundational pillars to achieve this goal, providing the control plane and abstraction layers needed to integrate heterogeneous infrastructures, from cloud cores to far-edge devices. This pillar reflects the convergence of orchestration solutions, automation frameworks, and MetaOS concepts that collectively drive the operationalization of the continuum.

In O-CEI, this pillar provides both technical positioning and a functional baseline to guide the implementation of continuum orchestration within the project. It aligns directly with WP3 activities such as T3.4 ("Innovative evolution of resource management, network and orchestration utilities") and T3.7 ("Continuous automation of platform & blueprints development, integration and deployment"), ensuring that technical advances in orchestration and automation inform integration, adaptation, and deployment strategies. The following content explores key research outcomes and technological solutions underpinning this pillar, highlighting innovations that inform O-CEI's approach to orchestrating the CEI continuum. To frame the technical contributions and architectural directions explored in this section, we begin by examining the characteristics and requirements of the CEI continuum itself.

The CEI Continuum refers to the integration of computing, networking, and storage capabilities across highly distributed infrastructures, spanning from cloud data centres to resource-constrained IoT devices. This integration enables more efficient and flexible data processing and storage across distributed networks, while also supporting transparent and automated workload management across a heterogeneous set of computing resources—from edge devices to cloud data centres.

The Multi-Provider Cloud-Edge Continuum brings together a diverse mix of infrastructures, including large-scale data centres, co-located facilities, on-premises systems, and on-device edge nodes. This results in a heterogeneous computing landscape that spans COTS servers, telecom infrastructure, and specialized processing units embedded in industrial equipment. As the number of resource-constrained IoT devices continues to grow exponentially, the CEI Continuum emphasizes the need to move computation closer to the data source, while recognizing that the majority of processing should still shift from the device to the edge layer. This model underscores the potential to dynamically utilize compute resources across cloud, edge, and device segments, depending on the specific needs of each application.

This layered and heterogeneous model is visually illustrated in Figure 4, which depicts the conceptual structure of the CEI Continuum, bringing together the three principal domains, Cloud, Edge, and IoT, each of which comprises a diverse set of infrastructures, from hyperscale data centers to near-edge compute nodes and constrained IoT devices such as sensors and gateways. These domains also reflect a range of deployment models and ownership contexts: Cloud spans public, private, hybrid, and multi-cloud setups; Edge includes public edge, near edge, and far edge nodes; and IoT consists of devices such as sensors, actuators, and gateways operating at the extreme edge.

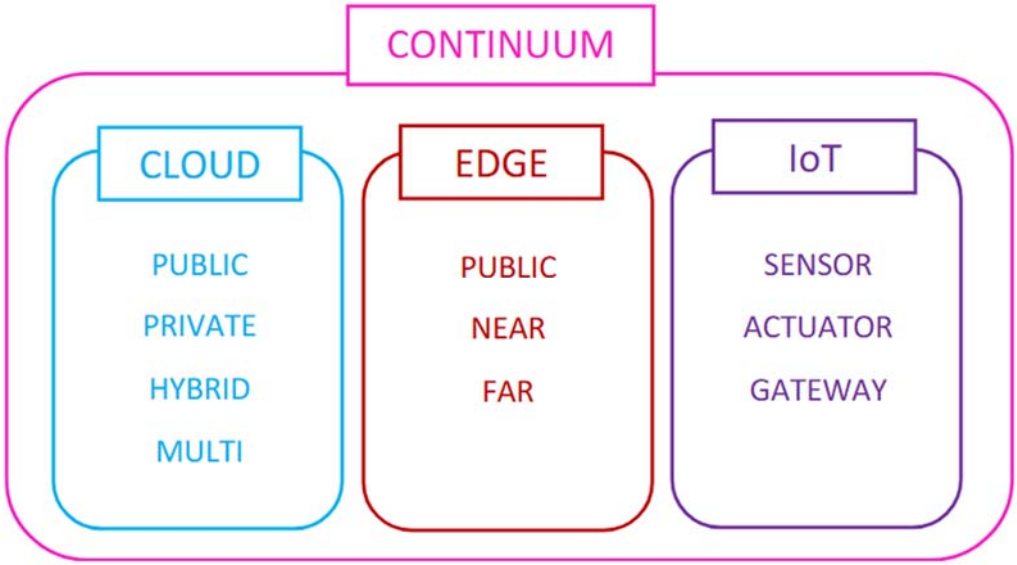


Figure 4: Cloud Edge IoT Continuum[1]

The continuum paradigm emerges as more than the sum of these parts: it allows distributed workflows to span across all layers, dissolving traditional technical silos. Computation and data can move seamlessly, from the farthest cloud to the most constrained device, according to application requirements and available resources. Depending on specific application needs, data and workloads should be dynamically assigned to the most suitable execution environments to meet requirements for latency, privacy, security, and performance. This model facilitates low-latency data processing closer to the source, while leveraging centralized compute for more intensive or less time-sensitive tasks.

This vision demands orchestration systems that go beyond traditional, centralized models. These systems must be intelligent, interoperable, and able to operate across federated, multi-provider environments. To ensure seamless workload mobility and service continuity across heterogeneous execution layers, they must incorporate policy-awareness, runtime adaptation, and automation. A key enabler of this flexibility is the rise of orchestration frameworks and MetaOS architectures, which abstract the complexity of underlying infrastructure. These frameworks support dynamic workload placement, optimized resource allocation, and adaptive service management, fostering an open, federated ecosystem that avoids vendors lock-in and enhances infrastructure resilience.

To realize this vision in practice, the new generation of orchestration and MetaOS solutions embrace existing operating systems, virtualization technologies and virtual infrastructure management providers leveraging their capabilities and hiding their peculiarities. Well established and proven technologies like containerization (e.g., Docker), declarative deployment (Docker Compose, Helm), and orchestration platforms (Kubernetes), provide essential infrastructure abstractions and automation for scalable, reproducible deployments across heterogeneous edge and cloud environments. Their widespread adoption provides a resilient substrate upon which more advanced orchestration logic can be layered.

Figure 5 below presents a logical layering of the various technologies and components discussed throughout this section. It visualizes their indicative placement within the CEI stack based on their roles in infrastructure abstraction, orchestration logic, deployment automation, and service

[1] https://zenodo.org/records/8403593

management. While this representation does not imply direct interoperability among all technologies, it provides a reference architecture to support informed integration choices. Practical adaptations and targeted tailoring—guided by specific use case needs—may enable synergies among these components, facilitating cohesive orchestration and management strategies across heterogeneous deployments. The layering highlights where in the CEI continuum each component adds value, helping identify combinations or substitutions aligned with functional, operational, and contextual requirements.

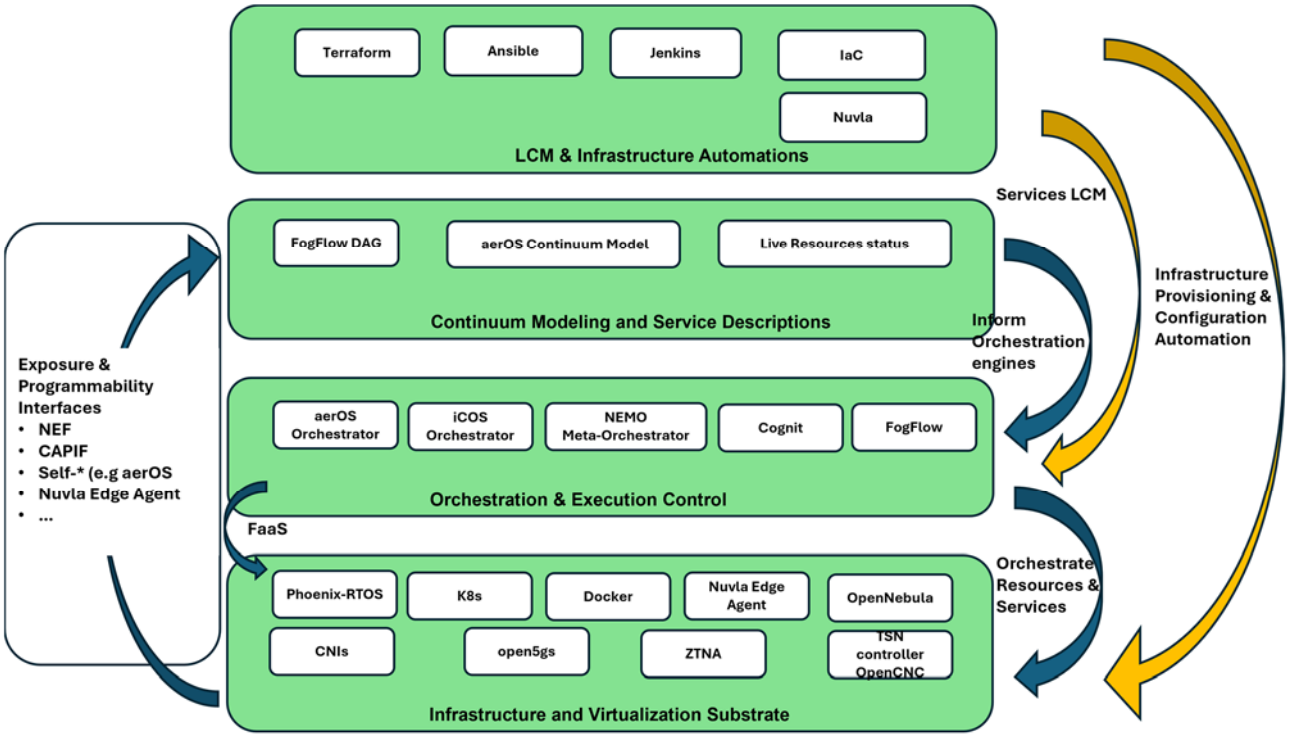


Figure 5: Continuum Components layered interactions

The presented layering structure primarily emphasizes the flow of information and control across the continuum: from raw infrastructure exposure (Layer 1) to orchestration decision-making (Layer 2), to abstract modelling and service descriptions (Layer 3), and finally to infrastructure automation and lifecycle execution (Layer 4). It reflects how orchestration and MetaOS strategies draw on underlying infrastructure capabilities and semantic models to manage dynamic, federated environments. This structure serves as a conceptual scaffold for the more detailed descriptions of individual components that follow, helping to contextualize their roles within the broader orchestration stack.

Thus, a set of concrete innovations has been developed within the framework of numerous EU-funded projects, building on established technologies. These innovations collectively reinforce the Continuum Orchestration and MetaOS pillar and are grounded in the TF3 building blocks—mainly Orchestration, Networking, Resource Management, and Monitoring & Observability. They reflect the collaborative contributions of many project partners, who have provided both technical assets and validated architectural guidance. These outcomes advance the development of O-CEI, enabling scalable, federated, and context-aware orchestration across the CEI landscape.

Among these advancements is the emergence of more sophisticated orchestration paradigms, notably the concept of federated, multi-layer orchestration. Actively explored and demonstrated in recent and ongoing European research projects, this paradigm addresses the complexity of managing distributed resources, diverse administrative domains, and real-time constraints within the CEI continuum. At its core, federated orchestration relies on multi-layered architectures that decouple high-level policy interpretation and intent resolution from low-level resource enforcement.

This architectural separation enhances flexibility, scalability, and responsiveness in dynamic environments. For example, double-layer orchestration—as validated in projects like aerOS by participating partners NCSRD and UPV—aligns intent-driven management with localized execution logic.

Beyond aerOS, more CL4 research efforts have produced orchestration platforms that operationalize this architectural paradigm. **The ICOS orchestrator**, developed within the ICOS project, acts as a Meta-Kernel layer orchestrator that sits on top of lower-level orchestrators like Kubernetes. It provides policy and QoS-aware orchestration across the continuum. ICOS focuses on managing the service lifecycle and making informed deployment and redeployment decisions to meet evolving performance and quality constraints. By abstracting heterogeneous execution environments, it contributes to the modular and policy-driven management of CEI resources.

Similarly, the **NEMO Meta-Orchestrator (MO)** introduces a comprehensive workload orchestration and optimization system across the NEMO cluster network. Serving as a high-level control point for coordinating underlying Kubernetes clusters, the MO handles deployment, migration, and scaling workflows via components such as the Deployment Controller and the Intent-Based Migration Controller (IBMC). The IBMC integrates AI-driven recommendations and service intents, making NEMO a robust and adaptive orchestration layer for CEI workloads.

Together, these approaches—developed within CL4 projects such as aerOS, ICOS, and NEMO—concretely embody the architectural logic envisioned in the EUCEI TF3 framework. This framework decomposes orchestration into key sub-blocks: federation management (for cross-domain resource allocation), service description (including SLA and policy specifications), low-level orchestration (network and compute resource control), and high-level orchestration (overarching deployment strategies). It also underscores the importance of re-orchestration triggers that support dynamic adaptation and system resilience under changing operational conditions.

To efficiently support such orchestration strategies, a detailed and dynamic representation of the continuum is essential. The effectiveness of these strategies across the CEI landscape relies heavily on a deep understanding and accurate modelling of the distributed infrastructure. As orchestration needs have evolved, the importance of robust continuum models has become increasingly evident—driving research efforts to develop frameworks that capture the complexity of heterogeneous environments. In response, CL4 projects have made continuum modelling a core research focus, producing frameworks that represent computing and communication resources not as static infrastructure elements, but as semantically enriched, context-aware entities.

Graph-based resource representations are increasingly used to characterize each node—whether cloud, edge, or IoT—by functional capabilities, contextual attributes (e.g., latency, energy, locality), and compliance metadata (e.g., security, data sovereignty). UPV, drawing on experience from the aerOS project, has contributed to the development of NGSI-LD-based continuum modelling processes, which structure semantic information in a way that supports interoperable, standards-aligned orchestration. These models enable the generation of real-time system state representations, allowing orchestrators to reason dynamically about network topology, node availability, and trust levels. The sharing of this semantic and contextual information across administrative domains is supported by lightweight federation mechanisms, such as Context Brokers further discussed in section 3.1.5, ensuring that orchestration decisions remain informed and coordinated even in distributed environments. The combination of static semantic models with live system metrics, collected through Monitoring and Observability mechanisms, forms the foundation for policy-driven and AI-enhanced orchestration logic. By integrating capabilities from multiple TF3 building blocks, including Trust and Reputation and Artificial Intelligence, this hybrid intelligence enables orchestration systems to make more accurate, trustworthy, and adaptive placement, scaling, and migration decisions across heterogeneous CEI environments.

Another solution that also builds upon a deep understanding of the continuum and supports intent-driven orchestration is **FogFlow**—a distinct implementation that complements existing approaches with an alternative architectural perspective. While sharing core operational goals such as high-level abstraction, model-driven decision-making, dynamic adaptation, and integration with existing infrastructure technologies, FogFlow introduces a unique programming paradigm tailored for the CEI context. Onboarded by NEC within the O-CEI framework, FogFlow offers a high-level programming model based on **data stream processing**, where applications are represented as **directed acyclic graphs (DAGs)**. In this model, each node functions as a processor of incoming data elements, while each link defines the communication flow between processing stages. This structure, referred to as a **service topology**, can be composed within a **no-code environment**, providing an abstraction layer that aligns seamlessly with intent-based orchestration strategies. The service topology is then translated into a concrete deployment plan through a two-step mapping process: first, transforming the DAG into a processing tree by identifying required processing elements and their interconnections; and second, mapping this tree to appropriate execution nodes within the CEI continuum, taking into account quality-of-service constraints and SLA requirements.

While NGSI-LD-based models provide a descriptive understanding of the infrastructure—capturing what entities exist, their capabilities, and how they relate—FogFlow complements this by modelling how data flows and is processed across the continuum. Its use of directed acyclic graphs (DAGs) focuses on the execution logic of applications, while NGSI-LD focuses on contextual awareness and semantic interoperability. Together, these complementary models bridge semantic context and operational behaviour, enabling orchestrators to reason not only about system topology and constraints, but also about dynamic execution paths and data dependencies. This integration supports more intelligent, adaptive, and intent-driven orchestration, particularly in complex, federated CEI environments.

As a natural progression from orchestrating services and resources, a critical concern for MetaOS platforms is the **lifecycle management (LCM) of workloads**, ensuring that applications can be reliably deployed, monitored, updated, and optimized across the continuum. To address this, the **Nuvla platform**, operated in O-CEI by HES-SO, has been integrated as a production-ready solution that strengthens O-CEI's capability for unified edge-to-cloud service management. Through Nuvla, the environment gains the ability to automate the full application lifecycle across heterogeneous infrastructures—simplifying deployment, enabling dynamic updates, and ensuring continuous optimization. Its infrastructure-agnostic design allows seamless management of services across both resource-constrained edge nodes and centralized cloud environments, ensuring operational consistency and portability. By building on standardized containerization technologies, Nuvla contributes a modular and scalable orchestration layer that integrates easily with existing ecosystems. Its decentralized agents improve the resilience and adaptability of service execution, even in partially connected or unstable network conditions. Furthermore, integration with DevOps workflows through programmable interfaces provides the foundation for rapid, policy-driven service delivery, enabling streamlined provisioning and reconfiguration. With its unified management stack embedding orchestration logic, access control, and resource governance, Nuvla helps operationalize key MetaOS principles—supporting secure, flexible, and composable deployments that are critical to realizing the O-CEI vision.

Complementary and supportive to the advanced orchestration and lifecycle management solutions discussed above are the **infrastructure automation pipelines** that underpin them. These pipelines are essential to ensure that CEI environments can be provisioned, configured, and maintained in a repeatable and scalable manner, with minimal manual intervention. They form the operational backbone that allows high-level orchestration strategies to be reliably executed across distributed infrastructures. Concrete, validated practices have been brought into the O-CEI ecosystem through contributions from related initiatives such as the **6G-SANDBOX project**, which has demonstrated how integrated DevOps pipelines can automate the provisioning and configuration of CEI components across both cloud and edge domains. In this approach, continuous integration and

delivery practices are used to manage virtualized infrastructure, coordinate deployment workflows, and ensure precise configuration of distributed systems (e.g., using Jenkins for workflow automation, Terraform for infrastructure provisioning, Ansible for configuration management, and OpenNebula as the virtualization backend). These pipelines combine established tools for workflow orchestration, infrastructure provisioning, and component configuration, integrated within a flexible virtualization environment. Through this toolchain, infrastructure automation becomes tightly coupled with the orchestration logic, enabling consistent, policy-driven deployment of services across the CEI continuum. By embedding this level of automation, O-CEI strengthens its ability to support dynamic and resilient service instantiation, ensuring that orchestration frameworks can operate effectively even under volatile or resource-constrained conditions. This synergy between infrastructure automation and MetaOS orchestration strategies is a critical enabler of secure, scalable, and composable deployments across the edge-to-cloud continuum.

Building on this foundation, and in response to real-life use cases that push orchestration needs even further, O-CEI embraces solutions capable of managing resource-constrained environments such as IoT devices. One such solution is **Phoenix-RTOS**, a modern, microkernel-based real-time operating system designed specifically for low-power, constrained hardware operating at the far edge of the CEI continuum. Strongly aligned with the TF3 “Resource Management” building block, its modular architecture allows the inclusion of only essential components, optimizing resource usage while ensuring a secure and robust foundation for application execution.

Phoenix-RTOS’s potential was demonstrated in the **COGNIT project**, where it was deployed on a smart meter to execute AI-based decision modules using a **Function-as-a-Service (FaaS)** model. This orchestration approach, also developed within COGNIT and tailored for edge environments, enhances system flexibility by enabling the offloading of compute-intensive tasks from IoT devices to more capable cloud or edge nodes. The FaaS framework supports adaptive, serverless execution and facilitates intelligent workload distribution, while maintaining compliance and data sovereignty constraints. Crucially, this approach brings **AI-driven decision-making** into the orchestration process, illustrating how functionalities from different TF3 building blocks—such as *AI* and *Resource Management*—can be combined effectively. The seamless coordination between Phoenix-RTOS at the device level and the FaaS orchestration layer across the continuum exemplifies the kind of cross-layer integration needed for portable, intelligent, and secure service execution—supporting O-CEI’s goals for composable and adaptive CEI infrastructures.

As orchestration expands from centralized cloud systems toward distributed edge and device-level environments, and as automation takes over more decision-making processes, the CEI landscape becomes increasingly open, dynamic, and less directly controllable. This shift, while enabling scalability and adaptability, also introduces new challenges in maintaining trust, ensuring access control, and securing operations across federated domains. In this new reality, traditional perimeter-based security models are no longer sufficient. To address these challenges, O-CEI incorporates **Zero Trust** principles as a complementary layer to enable secure, federated orchestration. Evolving beyond its foundational concepts [5], Zero Trust now enforces continuous verification, least-privilege access, and breach containment strategies, all of which are essential to trust-aware orchestration across the continuum.

In this context, components such as **NEMO’s Access Control** showcase how access policies can be enforced using attribute-based definitions. Semantically enriched continuum models, previously discussed, further support this by allowing orchestration logic to incorporate trust and contextual constraints directly into decision-making. This approach is further reinforced through integration with technologies like Extended Detection and Response (XDR), Single Sign-On (SSO), and Federated Learning for privacy-preserving threat detection. As a result, orchestration decisions are no longer based solely on performance or availability, they are also dynamically informed by trust, context, and risk posture. While these security and trust functionalities are detailed in section 3.1.4 it is important to reference them in conjunction with orchestration of the continuum as the reusable foundation of

TF3 building blocks provision for the integration of capabilities classified within Security and Privacy, Trust and Reputation and showcase how relevant project outcomes are integrated in a modular yet cohesive way to enable secure, trustworthy orchestration operations across the CEI continuum.

While Zero Trust provides one path to enabling secure and trustworthy orchestration through continuous verification, an alternative and complementary approach focuses on ensuring that orchestration decisions can be audited and verified after execution. In this direction, O-CEI supports compliance-aware orchestration by integrating architectures that embed auditability into the orchestration logic. To this end, UPV contributes mechanisms for runtime logging, integrity verification, and traceability, enabling transparent insight into how orchestration actions and data movements occur across the continuum. These mechanisms are further reinforced through the integration of Distributed Ledger Technologies (DLT) such as IOTA, which provide immutable records of orchestration decisions and data access, even within constrained CEI environments. This architecture supports regulatory compliance, reinforces operational accountability, and complements trust models with verifiable, tamper-resistant evidence. The integration of such mechanisms for the continuum orchestration highlights how orchestration-enabling components in O-CEI often synthesize functionalities from multiple TF3 building blocks, such as Monitoring and Observability, Trust and Reputation, and Security and Privacy, demonstrating the modular and cross-cutting nature of the reference architecture.

As decentralization increases and the integration of a growing number of nodes relies on network expansion and interconnection, the management of underlying network resources becomes an intrinsic part of MetaOS design and a major concern for orchestration. The network fabric is no longer a passive enabler—increasingly distributed and adaptive orchestration logic depends on it as a critical component for ensuring system-wide responsiveness, resilience, and coordination. In this context, MetaOS platforms must not only orchestrate compute and services but also integrate and manage advanced networking capabilities—ranging from software-defined networking (SDN) and programmable data planes to telecom-grade B5G infrastructures. Among these, 5G stands out as one of the most impactful ICT innovations in recent years, particularly for industrial and mission-critical use cases. Its native support for low-latency, high-reliability communication, network slicing, and dynamic QoS control makes it a foundational enabler for scalable and context-aware orchestration across the CEI continuum. Figure 6 below illustrates this continuum-wide orchestration vision—showing how service components span from cloud to edge and IoT, with B5G networks playing a key role in connecting resource-constrained devices at the far edge.

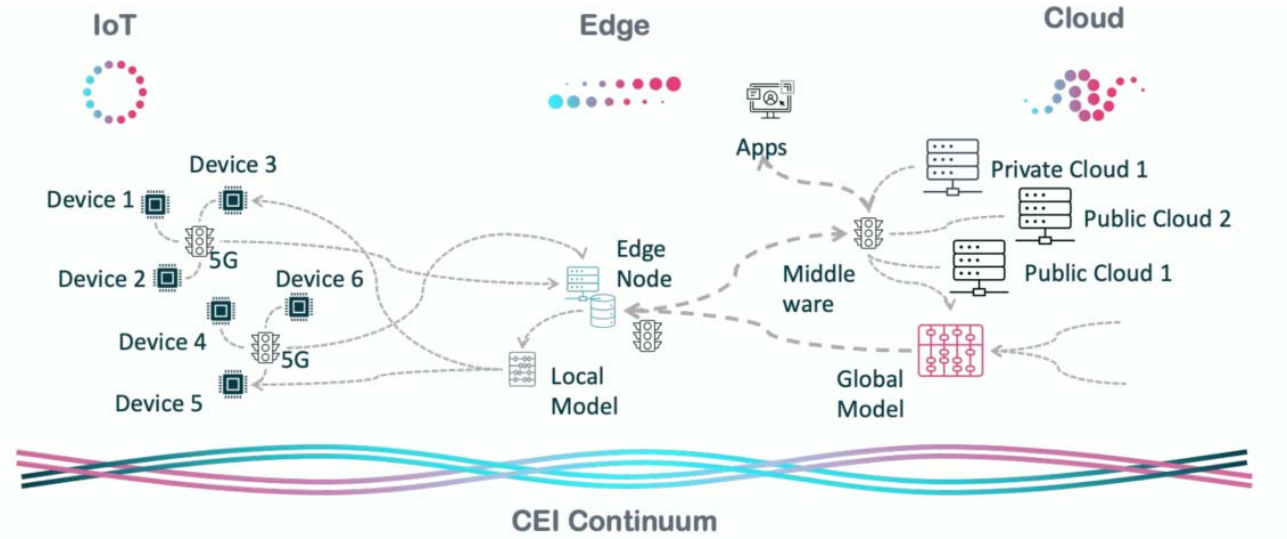


Figure 6: Cloud-edge-IoT continuum perspective source: EUCloudEdgeIoT

To support this vision, open and programmable networking interfaces have emerged as key integration points. The Network Exposure Function (NEF) and the Common API Framework for Northbound Interfaces (CAPIF) enable secure interaction between orchestration platforms and 5G/6G network capabilities. NEF facilitates access to QoS, policy, and session management features, while CAPIF—along with its open-source implementation, OpenCAPIF—standardizes how APIs are managed and discovered across telecom and vertical domains. Within this context, NCSRD brings validated expertise through its contributions to European projects such as EVOLVED-5G, SUNRISE-6G, and ENVELOP. These efforts include the development of programmable 5G core features using open-source tools like Open5GS, where the ease of deployment and configuration plays a pivotal role in enabling rapid prototyping, testing, and integration of advanced orchestration logic. This agility supports the broader vision of bringing network-level programmability closer to CEI orchestrators. By exposing network-level intelligence and programmability to CEI orchestrators, these tools unlock new capabilities such as QoS-aware orchestration, mobility-driven service placement, and real-time network context integration. These networking technologies have also been validated in pre-commercial research infrastructures such as the π-NET network operated by NOVA, where deployment integrates zero-touch orchestration, analytics-driven closed control loops, and open APIs for seamless service exposure and self-service management.

These developments reinforce the importance of network programmability and abstraction as first-class elements in orchestration, aligning with TF3 building blocks including Network, Orchestration, and Resource Management. They also extend O-CEI's ability to bridge vertical domain requirements (e.g., mobility, industry, energy) with the flexibility and reliability of telecom-grade infrastructures, ensuring the continuum is not only compute-aware but also network-intelligent.

Complementing the flexibility of mobile networks, certain domains—particularly in manufacturing, energy, and mobility—require not just low latency but also strict timing guarantees within local networks. To embrace these domain-specific requirements, the project introduces OpenCNC as part of its vision to extend orchestration capabilities across diverse networking environments. OpenCNC, an open-source Centralized Network Configurator for Time-Sensitive Networking (TSN), enables deterministic control over real-time Layer 2 communication, including precise traffic scheduling, stream prioritization, and fault tolerance. Deployed as Docker containers within a Kubernetes environment, OpenCNC integrates seamlessly with orchestration platforms, offering fine-grained control over wired network behavior. When combined with programmable 5G/6G capabilities, this approach allows CEI orchestrators to manage both wireless and wired infrastructures in a unified, policy-driven way. It enables CEI applications to meet stringent latency, reliability, and synchronization requirements, reinforcing the project's commitment to supporting vertical domains through flexible, modular, and standards-aligned orchestration strategies—consistent with the TF3 building blocks for Network, Orchestration, and Resource Management.

However, just as the widespread distribution of computing and services required rethinking trust and access through Zero Trust models, the expansion of programmable and exposed network capabilities introduces parallel security challenges. These networking advancements raise significant security considerations, especially as the number of connected devices and exposed interfaces continues to grow. The Zero Trust model, previously introduced and further discussed in section 3.1.4, becomes increasingly relevant [6], at the network level, enforcing continuous verification and least-privilege access, even within internal or trusted infrastructures.

Building on this foundation, Zero Trust Network Access (ZTNA) is emerging as a modern alternative to traditional VPN-based models [7]. Unlike legacy approaches that grant broad access after initial authentication, ZTNA evaluates multiple dynamic attributes—such as device posture, location, and behavioral context—before granting granular, adaptive access to services and data. These mechanisms are especially critical in orchestrated CEI systems, where network resource access must be context-aware, dynamically adjusted, and policy-driven.

To synthesize the diverse orchestration and networking approaches explored, the table below provides a consolidated overview of the key components, their positioning, and their integration across TF3 building blocks.

Table 1: Overview of Orchestration and Networking Solutions in O-CEI

Component	Type	Purpose	Key Features	Positioning	TF3 Blocks
aerOS	Federated Orchestrator	Multi-layer, intent-driven resource orchestration	Double-layer architecture, semantic modeling, NGSI-LD, trust scoring	Continuum-wide, semantic-aware orchestration (MetaOS aligned)	Orchestration, Networking, Trust & Reputation
ICOS	Meta-Kernel Orchestrator	QoS and policy-aware orchestration over Kubernetes	Service lifecycle management, high-level abstraction	Meta-layer orchestration logic over cloud-edge orchestrators	Orchestration, Resource Management
NEMO MO	Meta-Orchestrator	Intent-based service orchestration and migration control	AI-driven intents, Access Control, IBMC module	Top-layer orchestration over K8s, integrated with access control	Orchestration, AI, Security & Privacy
FogFlow	Model-driven Orchestrator	DAG-based service orchestration using data stream logic	No-code programming, service topology mapping, SLA-aware placement	Semantic-aware orchestrator, aligned with intent-driven approaches	AI, Orchestration
Nuvla	Lifecycle Manager	Unified edge-cloud application lifecycle management	Deployment automation, updates, observability, DevOps integration	Platform-level lifecycle control, edge-cloud consistent deployment	Resource Management, Security & Privacy
Phoenix-RTOS	RTOS + Serverless Orchestration	Lightweight runtime with AI offloading for constrained devices	FaaS model, modular microkernel, smart meter orchestration	Far-edge device-level orchestration	AI, Resource Management
OpenCNC	TSN Network Controller	Real-time deterministic control over wired industrial networks	Traffic scheduling, prioritization, fault tolerance	Local wired infrastructure orchestration integrated into K8s	Network, Resource Management
Open5GS + NEF/CAPIF	Programmable 5G Core + APIs	Expose 5G/6G network programmability to orchestrators	QoS policies, session control, OpenCAPIF integration	Network programmability layer for orchestration platforms	Network, Orchestration, Security & Privacy

These efforts collectively demonstrate how orchestration in O-CEI spans the full stack, from communication infrastructure and network programmability to semantic data models and application execution, guided by open interfaces, trust-aware policies, and scalable automation. Together, these innovations form the backbone of the Continuum Orchestration and MetaOS pillar, building on the TF3 Reference Architecture's core building blocks: Orchestration, Resource Management, Network, and Monitoring & Observability, while also integrating capabilities from additional blocks such as Security and Privacy, Trust and Reputation, and AI. Through this integrated and modular approach—linking infrastructure, data models, execution frameworks, and security policies—the Continuum Orchestration and MetaOS pillar lays a robust foundation for resilient, intelligent, and federated service and resource management across the CEI landscape, aligned with the evolving needs of distributed, multi-domain environments.

3.1.4 Security, trust and digital identity

As distributed, multi-stakeholder ecosystems grow in complexity, ensuring **security, trust, and digital identity** across the CEI continuum becomes indispensable. Security and trust are foundational to the reliable operation of **federated CEI ecosystems**, where distributed actors must collaborate seamlessly and securely across administrative domains and under stringent legal and regulatory requirements.

The **limitations of perimeter-based security models** are becoming increasingly apparent as CEI architectures mature. In response, modern systems are adopting more advanced approaches based on **zero trust principles**, **verifiable identity**, **distributed enforcement**, and **runtime policy control** to ensure both operational integrity and stakeholder accountability.

For O-CEI, this pillar positions the project within ongoing innovation efforts in decentralized security and trust models while also informing technical work under WP3, especially T3.6 (“Adaptation of security, privacy, trust and reputation utilities”) and T3.7 (“Continuous automation of platform & blueprints development, integration and deployment”). It provides the foundation for integrating security and trust mechanisms into both marketplace onboarding and platform-wide operations.

Recognizing the strategic importance of these concerns, this section offers a layered perspective—from identity and trust frameworks to anomaly detection, policy enforcement, and runtime threat mitigation—highlighting O-CEI’s concrete contributions and alignment with the TF3 Reference Architecture building blocks.

Identity and trust Layer

As the O-CEI ecosystem spans a highly distributed and federated landscape of cloud, edge, and IoT environments, the ability to establish secure, verifiable, and privacy-preserving identities is fundamental to building trust across system boundaries. To meet this challenge, O-CEI adopts and integrates a Self-Sovereign Identity (SSI) approach, an emerging paradigm that enables entities such as users, services, platforms, and devices to create and manage their own digital identities without relying on centralized authorities.

SSI architectures leveraged within O-CEI support privacy-by-default through mechanisms such as selective attribute disclosure, which allows only the minimum necessary personal or system information to be revealed during authentication and interaction processes. This ensures both compliance with regulatory frameworks like the GDPR and the ability to support fine-grained access control across multiple domains. Importantly, the SSI layer is closely tied to federated identity services, forming the basis for secure onboarding, authentication, authorization, and the lifecycle management of digital identities for users, devices, and services operating within the ecosystem.

Rather than prescribing a single implementation, O-CEI integrates multiple mature and complementary SSI solutions developed by partners within recent EU projects. For example, IQB provides an implementation, validated through the OASEES and SAFE-6G projects, which enables the secure and automated onboarding of edge and IoT devices via DIDs and VCs linked to device-specific metadata. These credentials provide a cryptographically secure way to assert identity and authorize actions without central intermediaries.

At the core of this model are Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs), as specified by W3C, which form the technological backbone for trusted identity exchange across the continuum. These credentials are issued by authorized parties but are cryptographically verifiable by any actor in the system, without requiring queries to a central registry. This decentralized verification capability eliminates the need for central authorities and enables trusted identity exchange between stakeholders in the cloud-edge continuum, enhancing both the scalability and resilience of identity management in federated infrastructures.

Among the SSI-related implementations within O-CEI, multiple partners contribute mature, production-ready solutions addressing different aspects of identity and trust. CERTH contributes a standards-compliant, SSI wallet architecture developed in the ASSIST-IoT project and further aligned with the **European Blockchain Services Infrastructure (EBSI)**. This wallet supports full identity lifecycle capabilities, including issuance, storage, revocation, and private communication through DID-based channels. It empowers users and organizations to control their credentials and manage consent on data disclosure, thus reinforcing **digital sovereignty** and **data minimization** in cross-sector interactions.

In parallel, the FIWARE Community provides an Identity, Authentication and Authorization Management (IAAM) framework based on DIDs and VCs. Authentication components implement the request of VCs based on OID4VC protocols defined by the OpenID Foundation and verification of VCs relying on trust frameworks compliant with EBSI. VC verifier components are fully compatible with wallets developed under the EU Digital Identity initiative, since OID4VC protocols are supported. Authorization functions implement a standard PEP–PDP–PIP–PAP architecture for the enforcement of policies expressed in ODRL (Open Digital Rights Language), more concretely, ODRL profiles which support rules based on claims linked to VCs.

Together, these implementations provide a baseline for building a **composable and interoperable trust fabric**, capable of addressing the diverse identity requirements of its pilots and stakeholders. They align closely with the **Trust and Reputation**, **Security and Privacy**, and **Resource Management** building blocks of the TF3 Reference Architecture, positioning decentralized identity as a central enabler of secure, collaborative, and regulation-compliant operations across the CEI continuum. Building on this identity foundation, O-CEI introduces trust management mechanisms that extend beyond authentication, assessing and governing the behaviour of interconnected resources in real time.

Trust management and Zero Trust in the O-CEI ecosystem

In the modern cloud ecosystem, multiple computing resources must work together to provide the scalability and flexibility required for dynamic workloads. These resources, ranging from hardware to software stacks, are often integrated into complex environments that require careful management and continuous monitoring. To effectively manage the trustworthiness of these integrated resources, and subsequently control their trust capabilities, it is crucial to employ techniques that assess and act on their security and performance.

The **Trust Management** process focuses on evaluating the trustworthiness of individual integrated resources, which collectively form the backbone of the ecosystem. One effective approach, introduced in prior work and adopted within O-CEI, is to abstract and organize these diverse components as ***Infrastructure Elements*** **(IEs)**, representing fundamental computing or network nodes across the CEI continuum. These IEs interact with one another to create topologies that enable resource pooling and distributed service execution. While some IEs take on more critical roles within these topologies, each contributes to the integrity and trustworthiness of the system.

Given the diversity and complexity of these resources, it is vital to assess their behavior and capabilities to ensure they operate securely and reliably. The Trust Management process collects various performance and security-related data from each IE, including security events, health metrics, service activity, communication patterns, and system status indicators such as CPU and RAM usage. Additionally, the system tracks update statuses and historical reputation data for each IE, aggregating these attributes into a comprehensive, dynamic Trust Score.

This Trust Score has been divided into three sub-scores. The reliability, the security, and the reputation subscores. The Reliability Sub-score evaluates the performance and stability of Infrastructure Elements (IEs) based on key system metrics such as CPU and memory usage. This

sub-score uses the TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) algorithm for the calculation, as it is well-suited for multi-criteria decision-making. TOPSIS ranks IEs by comparing their performance to an ideal solution, ensuring that the evaluation is both systematic and objective. The Security sub-score represents the system's short-term state regarding exposure to cyber threats, specifically the occurrence of attacks. The aerOS Self-Security module is configured to operate by transmitting real-time alerts to a designated endpoint in the Trust Manager. These alerts are triggered whenever Self-security generates a notification, enabling immediate processing and integration into the trust evaluation framework. The Reputation Sub-Score represents an overall measure of confidence in the robustness and reliability of an IE within the system. It evaluates the IE's historical performance by analyzing its exposure to cyber threats and its ability to maintain consistent functionality without frequent failures over a defined finite period. The total trust score (TS) provides a comprehensive evaluation of the trustworthiness of Infrastructure Elements (IEs) by combining multiple sub-scores and applying a penalty for critical events. This score ensures a balanced assessment of reliability, security, and reputation, tailored to the specific priorities of the system. The total trust score is computed using the following formula:

$$TS=(Wrep\times SBrep) + (Wsec\times SBsec) + (Wrel\times SBrel) - Penalty$$

An implementation of this approach, developed by IQB with support from UPV within the aerOS project, contributes a concrete realization of Trust Management within O-CEI. Trust data can also be registered onto the IOTA Tangle, enabling tamper-proof and verifiable trust histories. This architecture supports proactive issue identification and rapid adaptation to changing conditions. Ultimately, Trust Management ensures that only trusted IEs participate in mission-critical operations, reinforcing the ecosystem’s security, resilience, and accountability.

Building upon the dynamic trust assessment mechanisms discussed above, the O-CEI ecosystem further reinforces its security posture through the implementation of a Zero Trust architecture. As a foundational element of the Security, Trust, and Identity Management pillar, Zero Trust plays a central role in enhancing the resilience of the ecosystem. While its principles were also referenced in the MetaOS and Orchestration pillar—particularly in relation to secure orchestration across administrative domains—its core development and enforcement mechanisms are anchored here. This illustrates how Zero Trust not only supports identity and access management but also enables trusted, verifiable interactions across distributed layers, reflecting the intentional interplay between pillars in O-CEI.

Built on the principle of **“never trust, always verify”,** Zero Trust shifts away from traditional perimeter-based models that allow attackers to move laterally once inside. Instead, it enforces **strict, identity-aware access controls** for every request—regardless of location, user role, or network boundary. The core principles of **Zero Trust** include:

- **Explicit Verification**: Continuous authentication and authorization based on all available data points. Zero Trust assumes that attackers could already be present within the network, so user identities, privileges, and device security must be verified at all times. Furthermore, connections should periodically time out to enforce re-verification of users and devices.
- **Least Privilege**: Access should be limited to the lowest possible level based on identity and access management models, such as Just-In-Time (JIT), Just-Enough-Access (JEA), risk-adaptive policies, and data protection principles. Access is granted on a need-to-know basis, minimizing exposure to sensitive information or critical parts of the network.
- **Assume Breach**: Assuming that hostile actors may already exist within the network, this principle drives the implementation of practices like micro-segmentation, verified end-to-end encryption, and the prevention of lateral movement to protect as many resources as possible from compromise

Within O-CEI, Zero Trust complements the Trust Management framework by providing the enforcement layer that secures resource interactions. By verifying every access request and minimizing unnecessary privileges, Zero Trust reduces the risk of breach escalation. This is supported by mechanisms such as **Role-Based Access Control (RBAC)**, proactive threat monitoring, and incident response. **NEMO's Access Control** can enforce fine-grained policies, ensuring that only those with explicit authorization can interact with sensitive components of the ecosystem. Furthermore, **NEMO's Access Control** can be integrated with services that employ two-factor authentication and encryption mechanisms within a service mesh, reinforcing Zero Trust security principles

Moreover, Zero Trust principles also extend into the AI integration efforts within the O-CEI ecosystem. As discussed below, in section 3.1.6, privacy-preserving techniques such as Federated Learning and Homomorphic Encryption, as applied in **NEMO's FREDY** framework, demonstrate how sensitive data can remain secure even during collaborative AI processing. These approaches align with the Zero Trust philosophy by ensuring that data privacy, security, and access control are preserved throughout the entire AI workflow, without relying on implicit trust between participants. Such cross-pillar implementations further reinforce the security posture of the ecosystem and demonstrate the adaptability of Zero Trust beyond identity and access management alone.

Together, **Trust Management** and **Zero Trust** form a comprehensive security strategy that continuously monitors, evaluates, and protects integrated computing resources. By adopting both approaches, the **O-CEI** ecosystem enhances its ability to respond dynamically to changing conditions and emerging threats, ensuring that each resource remains trusted and resilient in the face of potential cyber threats. Building on this trust foundation, next paragraphs discuss how, O-CEI embeds proactive cybersecurity mechanisms directly into runtime environments, providing localized defence and immediate threat response across the continuum.

Cybersecurity and runtime threat mitigation

As part of O-CEI's strategy for resilience across the continuum, cybersecurity must be embedded close to the edge, where services are deployed, and threats may arise in real time. Given the increasing complexity and dynamism of CEI infrastructures, traditional centralized protection mechanisms fall short in terms of latency, precision, and adaptability. O-CEI integrates a suite of monitoring and mitigation tools, enabling localized threat detection, vulnerability analysis, and behavioral anomaly identification across distributed environments. **CERTH** brings forward lightweight **Intrusion Detection Systems (IDS)** and anomaly detection tools deployable at edge nodes, enabling real-time identification of malicious activity such as DoS attacks or data exfiltration. These are enhanced through **AI/ML-based techniques** and **Federated Learning**, improving accuracy while preserving data sovereignty. In parallel, SYNELIXIS contributes automated vulnerability assessment mechanisms that continuously scan deployed services for known weaknesses. These mechanisms leverage widely adopted threat intelligence frameworks such as MITRE, OWASP, and CWE, which serve as the foundation for many vulnerability scanning tools [8]. Examples include Zed Attack Proxy (ZAP) and Burp Scanner, both of which align their reporting with these taxonomies. At the same time, advances in the field are driving the development of scalable, intelligent crawlers that improve both coverage and precision. These newer approaches incorporate techniques such as machine learning [9], hybrid analysis workflows [10], and deeper behavioral inspection to identify complex or previously undetected vulnerabilities [11], [12]. This proactive layer strengthens preventive defense across the stack.

While edge-level threat detection secures the infrastructure, runtime and application-layer defences are critical to ensuring the operational integrity of services as they execute. So, from an application-level perspective, **XLAB** integrates **Dynamic Application Security Testing (DAST)** capabilities using AI and NLP techniques like LogBERT to continuously monitor application logs and detect runtime anomalies. These tools operate autonomously and adaptively, addressing threats that only

manifest during operational phases. Complementing these contributions, **UPV** has also led research on **ransomware mitigation strategies**, such as deploying honeypots within CEI environments to test response behaviors and capture threat signatures. This supports both real-time incident response and long-term security posture. In doing so, they form a strategic security layer that underpins O-CEI's capacity to detect, respond to, and learn from emerging threats across the CEI landscape. These runtime and application-level defenses operate atop a foundational execution environment that must itself be resilient, scalable, and secure by design.

It is worth noting how cybersecurity and threat mitigation components introduced by O-CEI partners natively integrate AI and machine learning techniques. This illustrates how security-focused tools further build on capabilities from more TF3 building blocks, such as AI-Driven Computing and Monitoring & Observability, in order to adapt to the dynamic and distributed nature of CEI environments.

Infrastructure & execution layer

To realize the trust, transparency, and accountability principles introduced above, O-CEI integrates decentralized execution infrastructures as part of its foundational layer. The Infrastructure & Execution Layer of the O-CEI ecosystem plays a crucial role in enabling secure, scalable, and interoperable data-driven services. This layer integrates a variety of advanced technologies and methodologies which form a robust foundation for building a decentralized, trustworthy environment, where secure data exchange, auditing, and reputation management are enabled to ensure the integrity, transparency, and seamless interaction of diverse system components. Key contributions to this layer come from research organizations such as CERTH and UPRC, who have developed and deployed critical components through various research projects.

CERTH hosts a dedicated Blockchain Team actively engaged in the research and development of various Distributed Ledger Technologies (DLTs), including Ethereum, Hyperledger, and others. This work is conducted through participation in multiple EU-funded research projects, as well as private contracts with SMEs and industrial partners. Additionally, CERTH operates an official node of the European Blockchain Services Infrastructure (EBSI) in Greece. Among its key areas of expertise are Self-Sovereign Identity (SSI) solutions, Blockchain-based wallet implementations, and the design of advanced Blockchain architectures tailored to secure, scalable, and interoperable applications. Within the scope of the Assist-IoT project, CERTH has developed a suite of DLT-based enablers that hold strong potential for further enhancement and integration into the O-CEI project. These enablers are designed to enhance trust, transparency, and accountability across distributed IoT ecosystems. Their functionalities can be effectively leveraged to address the challenges associated with data exchange, integrity, and collaboration among stakeholders in complex environments. Together, the below-mentioned DLT-based enablers form a robust foundation for the development of secure, transparent and scalable data-driven services within the O-CEI project ecosystem.

- The first component, the Logging & Auditing enabler, records critical actions during data exchanges between stakeholders. By ensuring transparency, auditability, non-repudiation and accountability, it enables traceable and verifiable interactions. Additionally, it captures resource access requests and security-related events, facilitating the provision of digital evidence and the resolution of potential disputes.
- The second component, the Data Integrity Verification enabler, introduces mechanisms for verifying the integrity of exchanged data using DLT technologies. Smart contracts (chaincode) hosted on peer nodes implement the verification logic. These contracts store hashed representations of data in immutable structures and compare them against client-submitted hashes to ensure data has not been tampered with, thus strengthening data assurance across the network.

- The third component, the Distributed Broker enabler, facilitates trusted data sharing across heterogeneous IoT devices within the system architecture. It manages metadata associated with data sources and ensures that such metadata remains discoverable, trustworthy and retrievable, thereby enabling secure and seamless data interoperability across domains.
- Lastly, the DLT-based Federated Learning (FL) enabler introduces a decentralized reputation management system within a FL framework. This mechanism safeguards against free-riding and malicious behaviors by evaluating and tracking the trustworthiness of local model contributors. Only participants with sufficient reputation are allowed to influence the global model, thus preserving the integrity and security of collaborative learning.

These DLT-based enablers lay the groundwork for secure, transparent, and scalable data-driven services in the O-CEI ecosystem, directly addressing the challenges of IoT data exchange, integrity, and trust.

Additionally, UPRC introduces work which builds upon the insights gained from the ERATOSTHENES project. It focuses on the design and deployment of decentralized trust and identity infrastructures for IoT ecosystems, exploring Blockchain-based data brokerage mechanisms and inter-ledger trust information sharing. Research within this area has led to the development of smart contract-based trust score propagation mechanisms in distributed IoT contexts. These mechanisms play a critical role in the O-CEI ecosystem, allowing for the decentralized evaluation and propagation of trust scores across different IoT devices and stakeholders. By leveraging Blockchain technologies, trust information can be shared securely and transparently, allowing for the ecosystem to maintain high levels of integrity and reliability. These contributions are particularly relevant for scenarios where devices and stakeholders with varying trust levels need to collaborate in a distributed environment.

Further contributions come from UPV, which has already provided an operational deployment of IOTA in aerOS, where it forms the basis of a secure and auditable data exchange layer across the CEI continuum. By utilizing IOTA's Tangle architecture, nodes (Hornet IEs as seen in Figure 7) within the continuum can exchange critical network state information—such as IP and domain addresses—in a feeless, trustworthy, and traceable manner. UPV's experience includes the implementation of a private IOTA Tangle across multiple infrastructures: Docker-based runtime environments, various Kubernetes clusters and domains, automated via Compose and Helm charts, and enhanced with peering and API services. This experience positions IOTA as a candidate enabler for immutability checking, decentralized trust, traceability, and data integrity as an O-CEI utility framed within the Trust and Privacy block of the TF3 Architecture. In parallel, UPV has also contributed to embedding **auditability mechanisms** into orchestration and data-sharing layers. This includes the implementation of runtime logs, integrity verification tools, and traceability components to ensure that platform decisions and data movements remain transparent and verifiable. These capabilities demonstrate a strong synergy with the TF3 Observability and Monitoring building block, which provides the foundational telemetry and event streams required to feed trustworthy audit trails into DLT systems. Distributed ledger technologies—such as IOTA—have been explored as enablers for immutable tracking of orchestration decisions and secure data usage, even in resource-constrained edge environments. This aligns closely with O-CEI's goals for trust-aware and accountable federated operations.

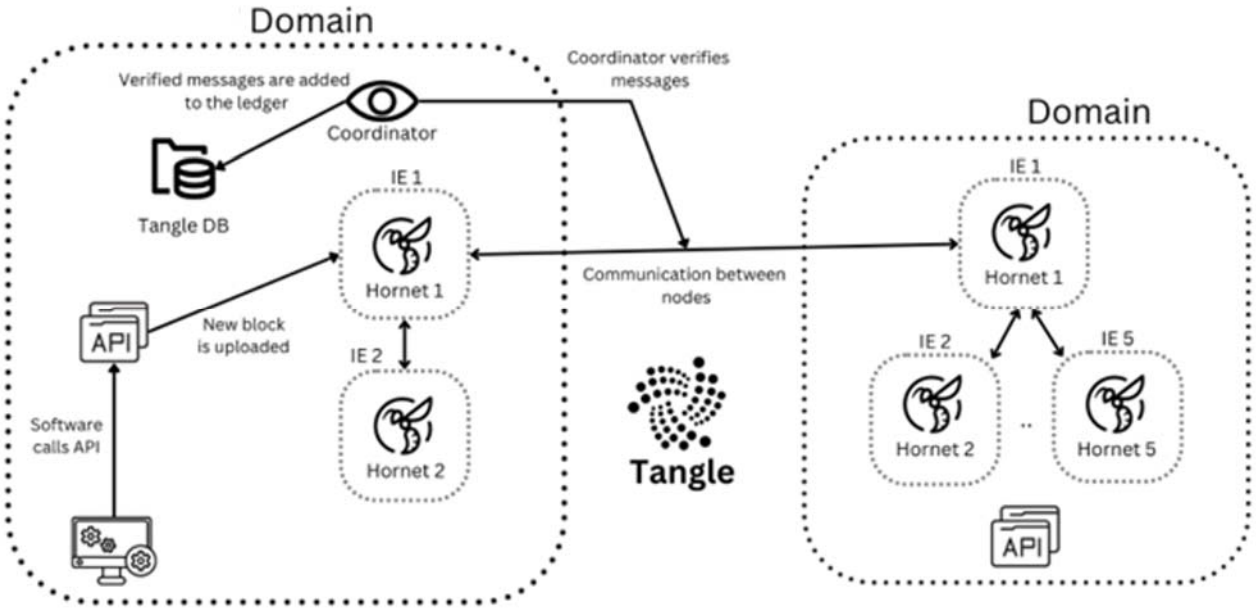


Figure 7: IOTA in aerOS continuum

Secure data brokerage & federated trust enforcement

A core extension of the infrastructure capabilities described above is the **Blockchain Data Brokerage Engine** developed by UPRC. Anchored in the *Security, Trust, and Digital Identity* pillar, this enabler supports **secure, auditable, and policy-driven cross-domain data sharing** across the CEI continuum. It leverages **DLTs** and smart contracts to decentralize the logging of data transactions, enforce customizable access policies, and maintain robust data provenance across federated domains.

Architecturally, the engine aligns with principles of **data sovereignty, GDPR compliance, and decentralization of trust**, offering a scalable foundation for federated data ecosystems. Its hybrid design supports **on-chain governance of inter-domain agreements**, while delegating the actual data exchange to secure, high-performance off-chain channels—thereby balancing **auditability, security, and efficiency**.

Designed with interoperability in mind, the engine demonstrates strong compatibility with key European frameworks such as **Gaia-X, IDS, and EBSI**, enhancing its potential for cross-sector integration and adoption. As such, it acts as a foundational bridge toward supporting Data Spaces, where trustworthy, semantically aligned, and policy-compliant data exchange is required. This connection will be further explored in the following section dedicated to Data Spaces and Interoperability. Moreover, the engine's capabilities, particularly around **traceability, non-repudiation, and policy-based trust enforcement**, are foundational for advanced CEI scenarios involving **federated learning, multi-party analytics, and secure data monetization**.

UPRC's contribution here builds on its prior work in decentralized identity and trust infrastructures and could play a pivotal role in realizing a **trusted data exchange layer** within the O-CEI architecture. As such, this engine not only contributes critical functionality but also aligns tightly with the broader vision of enabling secure, sovereign, and interoperable digital ecosystems across Europe.

To consolidate the comprehensive set of technologies and frameworks presented in this section, the following table provides an overview of the key components, their functional roles, and how they align with the TF3 Reference Architecture.

Table 2: Overview of Security, Trust, and Digital Identity Solutions

Component	Type	Purpose	Key Features	Positioning	TF3 Blocks
SSI (Self-Sovereign Identity)	Decentralized Identity Framework	Decentralized digital identity management across domains	DIDs, VCs, privacy-by-design, selective disclosure	Federated identity infrastructure for CEI participants	Trust & Reputation, Security & Privacy
SSI Wallet (CERTH)	Identity Wallet	Manage credentials lifecycle and secure communications	DID-based, aligned with EBSI and ASSIST-IoT	User-controlled identity lifecycle management	Security & Privacy
IAAM Framework (FIWARE)	Access & Identity Management	Authentication and authorization based on VCs	OID4VC, PEP-PDP, ODRL policy enforcement	Interoperable VC-based access control in CEI	Security & Privacy, Trust & Reputation
Trust Management System (aerOS)	Runtime Trust Assessment	Dynamic trust scoring of infrastructure elements	Trust Agent, Trust Manager, IOTA-based logs	CEI-wide evaluation of node trustworthiness	Trust & Reputation
Zero Trust Security Model	Security Model	Continuous identity-aware verification and access control	Least privilege, assume breach, RBAC, micro-segmentation	Cross-layer enforcement strategy for secure interactions	Security & Privacy
FREDY Framework (NEMO)	Privacy-preserving AI Framework	Secure collaborative learning without exposing raw data	Federated Learning, Homomorphic Encryption, ZT-aligned	Secure AI computation layer integrated in orchestration	Security & Privacy, AI
Runtime Threat Mitigation	Edge Security Toolkit	Local anomaly and vulnerability detection at edge	IDS, DAST, honeypots, ML-enhanced detection (LogBERT)	Edge-deployed runtime protection for services	Security & Privacy, Monitoring & Observability, AI
DLT Enablers (Assist-IoT, IOTA, EBSI)	Blockchain-based Trust Infrastructure	Data integrity, logging, and trust history verification	Audit trails, hashed logs, reputation storage	Infrastructure-level enablers for immutable data assurance	Trust & Reputation, Monitoring & Observability
Blockchain Data Brokerage Engine	DLT-based Data Exchange	Cross-domain data sharing with policy enforcement and auditing	Smart contracts, Gaia-X/IDS/EBSI alignment	Federated data exchange layer for CEI	Trust & Reputation, Security & Privacy

The technologies and methodologies outlined in this section define O-CEI's technical positioning in the domain of Security, Trust, and Digital Identity. By integrating decentralized identity frameworks, dynamic trust assessment, Zero Trust enforcement, and distributed ledger technologies, O-CEI establishes a robust foundation for secure and accountable operations across the CEI continuum. These capabilities, built on mature outcomes from recent EU research, form a modular and composable security fabric that addresses cross-domain authentication, real-time threat detection, verifiable data governance, and regulatory compliance.

Aligned with the TF3 Reference Architecture, this pillar primarily draws from the *Security and Privacy* and *Trust and Reputation* building blocks, while also intersecting with others such as *Monitoring and Observability* and *Orchestration*, enabling the deployment of secure, resilient, and trustworthy CEI infrastructures. Moreover, these security technologies serve as foundational enablers across all other O-CEI technological pillars—supporting secure orchestration and MetaOS runtime operations, enforcing data integrity and access control within Data Spaces and Interoperability layers, and safeguarding AI-driven computing processes and distributed intelligence workflows. As also illustrated in Figure 3, this pillar vertically spans the entire O-CEI architecture, establishing bindings with all other technological pillars. While it may draw on specific facilities from them, it universally delivers security and trust capabilities across the ecosystem. In doing so, they ensure that trust and security are not isolated features but intrinsic properties of the entire CEI architecture.

3.1.5 Data spaces, interoperability

Trusted data sharing and seamless interoperability are fundamental to unlocking the full potential of the Cloud-Edge-IoT continuum. This goes beyond technical integration; it represents a strategic enabler for Europe's digital sovereignty and AI-readiness. The pillar of **Data Spaces and Interoperability** encompasses both technological frameworks and governance mechanisms necessary to enable cross-sector data exchange, federated ecosystems, and interoperable services, ensuring data sovereignty, trust, and semantic alignment. This pillar reflects the ambition to build open, secure, and interoperable data ecosystems across multiple domains.

In O-CEI, this pillar significantly shapes the project's technical positioning and practical integration work, particularly tasks like “T3.3 O-CEI Marketplace and zero-trust onboarding strategy” and “T3.5 Upscaling of data management, observability, and AI-supporting utilities.” By consolidating insights from recent initiatives and research efforts, this pillar provides a concrete baseline for the upcoming implementation of interoperable, secure data-sharing mechanisms within the project. The following sections outline foundational concepts and active standards currently shaping the data-sharing and interoperability landscape within O-CEI.

Data Spaces are foundational ecosystems enabling **secure, sovereign, and interoperable** data sharing. They establish the necessary governance, identity, and contract enforcement mechanisms to ensure that data exchanges between entities are **auditable, policy-compliant, and remain fully under data owners’ control**. O-CEI embraces this model explicitly, integrating it with concrete standards, tools, and frameworks to facilitate cross-domain interoperability and digital sovereignty at scale.

Building upon this foundational vision, O-CEI also anticipates an evolution from conventional data spaces toward more general **Asset Spaces**. In this extended paradigm, not only data but also software components, configuration profiles, workflows, and usage policies become sovereign, composable digital assets. This approach supports advanced collaboration scenarios such as service orchestration, AI deployment, and distributed execution, while preserving sovereignty, interoperability, and trust. This vision enhances the functional richness of data ecosystems and positions O-CEI as a forward-looking enabler of next-generation federated services prepared for advanced AI-driven applications

In the longer term, Trialog participates actively in ongoing discussions about transforming data spaces into broader, more operational asset spaces (including e.g. software, precise configurations, profiles or rules). A new standard on Digital Asset Spaces is currently being developed to support diverse asset exchanges per country, per domain, and across sectors. The European Commission seeks feedback from EU-funded projects on the maturity and replicability of profiles shared by future Digital Asset Spaces. O-CEI aims to contribute concrete use cases (and associated profiles) from large-scale pilots, emphasizing their operational durability and replication potential.

To operationalize these visions, robust technical infrastructures supporting interoperability at multiple levels are essential. O-CEI incorporates implementation patterns and enablers, such as data fabrics, federated catalogs, and knowledge graphs, that transform theoretical concepts into actionable tools. The basis of a Data Space is to ensure products (comprising data and/or services) are easily findable, understandable, accessible, and reusable, following FAIR principles. These mechanisms enable trustworthy, secure, and effective exchanges among participants.

Reference Frameworks

To ensure this vision is well-grounded and aligned with existing best practices, O-CEI draws on established reference frameworks developed by key initiatives shaping data spaces. One of the earliest and most influential of these is the IDS Reference Architecture Model (IDS-RAM), first published by the International Data Spaces Association (IDSA) in 2018. IDS-RAM introduced

fundamental concepts such as the data space connector, detailed later in this document. A particularly significant part of IDS-RAM is the Dataspace Protocol (DSP) specification, covering APIs and protocols connectors must support for:

- Publishing data/service offerings,
- Enabling peer-to-peer negotiation between providers and consumers,
- Controlling interactions after an agreement is reached (with the actual data transfer outside DSP's scope).

Further development of key IDSA specifications, including DSP, has now shifted to the Eclipse Dataspace Working Group (EDWG). IDSA collaborates actively with EDWG on additional protocols, such as the Decentralized Claims Protocol (DCP).

Gaia-X, established in October 2019, pursues a broader scope than data spaces, notably federated cloud infrastructure. However, its contributions—including trust frameworks, standardization of verifiable credentials, and policy enforcement based on ODRL profiles—are highly relevant and supportive of data space requirements.

In September 2021, Gaia-X, IDSA, and FIWARE Foundation established the Data Spaces Business Alliance (DSBA) to align efforts towards interoperable data spaces. This collaboration led to the creation of the EU-funded Data Spaces Support Centre (DSSC), tasked with developing recommendations on building blocks for data spaces. O-CEI closely follows DSSC developments and intends to contribute practical use cases and feedback, helping shape these interoperability building blocks.

Connectors

Connectors represent the foundational building blocks for achieving interoperability and sovereign data sharing within Data Spaces. By providing integrated suites of open-source components, connectors enable the secure publication, discovery, negotiation, and controlled exchange of data and related services. In O-CEI, connectors will be employed as core elements of our interoperability architecture, operationalizing data space concepts such as policy enforcement, semantic alignment, and cross-domain trust. Two connectors, in particular, have emerged as mature, widely adopted solutions: the Eclipse DataSpace Connector (EDC) and FIWARE Data Space Connector (DSC). Both frameworks are grounded in the principles established by International Data Spaces (IDS) and Gaia-X, but each provides distinct capabilities suited to specific scenarios and requirements.

The Eclipse DataSpace Connector (EDC) International Data Spaces (IDS) and Gaia-X principles to facilitate secure, policy-enforced exchange of data assets between organizations without requiring replication of entire datasets. EDC provides essential functionalities such as a catalog API, contract negotiation service, and policy enforcement runtime, encapsulated within a lightweight, extensible microkernel architecture suitable for deployments ranging from Kubernetes clusters to single-board computers. Asset metadata and usage constraints in EDC are expressed as JSON-LD resources conforming to the IDS Information Model. Its transfer layer supports both synchronous (HTTP) and asynchronous protocols such as AS4 [21] and IDS Multipart. Within O-CEI blueprints, an EDC instance can be deployed as a CEI utility within the Data Management block: industrial partners can publish datasets (or real-time topics from Zenoh, MQTT, etc.) and attach ODRL-based licenses. Consumers discover these datasets through the federated marketplace, negotiate smart contracts, and retrieve only authorized data segments. Because EDC integrates seamlessly with external IAM solutions (such as., Keycloak) and supports attribute-based access control, it aligns effectively with O-CEI's zero-trust, cross-domain identity management strategy. EDC has been extensively validated through sector-specific implementations like Omega-X (energy) and Catena-X (automotive/industrial), as described later in this document.

The FIWARE DSC aligns similarly with IDS and Gaia-X principles but extends the range of supported interaction scenarios. Unlike EDC, the FIWARE DSC explicitly supports identity and authentication scenarios involving individual users, software agents, and even devices, all operating under organizational frameworks. Consequently, while EDC primarily utilizes the Eclipse Decentralized Claims Protocol (DCP) to exchange verifiable credentials that authenticate software agents representing organizations, the FIWARE DSC employs the OpenID Foundation's OID4VC protocol, facilitating authentication and credential exchanges involving individual identities linked to organizations.

Additionally, the FIWARE DSC supports the Eclipse Dataspace Protocol (DSP), like EDC, which includes the following sub-protocols:

- Catalog protocol: for discovery of data/services and related offerings.
- Contract Negotiation protocol: for negotiating and reaching agreements between data/service providers and consumers.
- Transfer Process protocol: for initiating, controlling, and terminating data exchanges.

Beyond these similarities, FIWARE DSC provides an alternative binding for the Contract Negotiation protocol based on TM Forum Open APIs. This compatibility simplifies integration with DOME-compatible marketplaces, leveraging TM Forum Open APIs to enhance federation capabilities and streamline contract management workflows.

Together, these connectors provide O-CEI with robust and proven tools for implementing our data-sharing and interoperability vision. The next section will present concrete field-tested experiences from existing sector-specific implementations led by project partners, such as Omega-X and Catena-X, which demonstrate practical deployments and validations of these connectors.

The practical maturity of both the Eclipse and FIWARE connectors is further demonstrated through their deployment in major European data space initiatives. These implementations serve not as abstract proofs of concept, but as operational environments that validate the architectural, semantic, and governance models behind data spaces at scale. In the context of O-CEI, these real-world experiences are especially valuable. Several project partners have played active roles in initiatives such as Omega-X and Catena-X, contributing directly to the development, deployment, and integration of connector technologies. These contributions reflect a high level of technical capacity and provide O-CEI with a strong foundation of proven solutions, architectural patterns, and interoperability strategies that can be adapted and scaled across industrial domains.

Omega-X and Catena-X stand as high-impact European initiatives that have successfully operationalized data space concepts using the connectors described above. Their implementation highlights the technical maturity and readiness of data-sharing infrastructures, while also proving the direct involvement and expertise of O-CEI partners in enabling trusted, cross-sector data exchange.

Omega-X has delivered a reference implementation of a federated, sovereign, and interoperable energy data space, fully aligned with European initiatives such as Gaia-X and IDSA. The architecture builds on foundational frameworks—including EDC, FIWARE, and BDVA/DAIRO—to ensure data sovereignty, trust, and security across diverse energy domains. A key contribution of Omega-X is the Common Semantic Data Model (CSDM), which ensures semantic interoperability across all use cases. This model integrates existing ontologies such as ETSI SAREF, IEC CIM, and SEDMOON and was co-designed through stakeholder engagement spanning flexibility, renewables, local energy communities, and electromobility. The technical architecture incorporates essential building blocks such as connectors, semantic adapters, identity and credential management, and federated service/data marketplaces, establishing a reference for future energy data spaces. Project partners involved in O-CEI played central roles in the design, deployment, and validation of these

components—demonstrating concrete know-how in implementing Gaia-X aligned data spaces. The Omega-X platform has been validated through real-life pilots across Europe, achieving high success rates in data-sharing deployments and interoperability tests.

Catena-X is an EU-driven alliance focused on enabling trusted, federated data exchange in the automotive and industrial sectors. It builds on the architectural foundations of Gaia-X and the IDSA model and uses the EDC as a core enabler. Catena-X provides a blueprint for generating trusted, interoperable, and scalable data spaces in complex, multi-stakeholder environments. The architecture is structured around five key pillars:

- Role-Based Participation
- Service Map
- Operational Procedures
- Data Space Governance
- Lifecycle Management

These experiences ensure trusted operations and interoperability across the data lifecycle—and O-CEI benefits directly from the involvement of partners who contributed to the technical realization and governance frameworks within initiatives like Catena-X. By building on field-tested patterns, reference architectures, and compliance strategies, O-CEI is well positioned to adopt, adapt, and extend validated technologies, accelerating the development of its own trusted and interoperable data space infrastructure.

Marketplaces

To operationalize these validated assets and enable scalable collaboration across domains, O-CEI places the marketplace at the center of its architecture. It serves as a unified environment where all essential resources—datasets, APIs, connectors, semantic models, workflows, and services—are published, discovered, and negotiated. These offerings may include terms and conditions, usage policies, and pricing models, supporting trusted exchange across sectors. Far from being just a transactional tool, the marketplace is a critical enabler for onboarding, cross-domain interoperability, and dynamic service integration, laying the foundation for scalable federation within the O-CEI ecosystem.

In O-CEI, marketplace capabilities are essential to support cross-domain use cases, enabling organizations to publish, discover, and negotiate offerings of digital products, which may include datasets, APIs, services, or composite workflows. These offerings can incorporate terms and conditions, usage policies, and pricing models, aligning with the principles of data sovereignty and trusted exchange.

The FIWARE BAE/Marketplace framework provides a robust, open-source foundation for implementing such marketplaces. It includes a marketplace portal based on TM Forum Open API specifications, enabling the creation of product listings as combinations of data and services. Building on this foundation, the DOME (Distributed Open Marketplace for Europe) initiative, developed under the Digital Europe Program, aims to become the reference marketplace infrastructure for the commercialization of Cloud and Edge services across Europe.

DOME integrates the FIWARE BAE/Marketplace with Blockchain-based mechanisms to support federated marketplace scenarios. This allows providers to register products once and make them discoverable across multiple interconnected marketplaces. Federation respects autonomy: distinct marketplaces operated by different legal entities, such as IaaS providers, platform providers, or domain-specific brokers, can connect to DOME's infrastructure while maintaining their own governance and user experiences.

O-CEI benefits from direct participation in the DOME ecosystem, ensuring alignment with its catalog federation, access control, and contracting mechanisms. These components are implemented using TM Forum APIs, with product offerings represented as Verifiable Presentations (VPs) built from Verifiable Credentials (VCs), in line with Gaia-X specifications and W3C standards. Trusted issuers validate these credentials, enabling product authenticity and regulatory compliance across the federated network.

To complement marketplace capabilities, O-CEI also incorporates lightweight, semantic-enabled data brokers contributed by UPV. These brokers support dynamic discovery, access control, and contract-based usage of data assets, while providing semantic annotation, logging, and policy enforcement. Designed in alignment with IDSA reference models and the principles of the Data Spaces Business Alliance (DSBA), these brokers offer modular, standards-based capabilities that reinforce O-CEI's goals of transparency, interoperability, and federated data control.

Further strengthening its marketplace architecture, CERTH introduces a distributed and trustworthy search mechanism, designed for decentralized discovery across Blockchain networks. In this model, encrypted keywords or search tokens are stored on-chain, while the underlying data remains encrypted and distributed off-chain. This approach supports secure, privacy-preserving search and discovery in both permissioned and permissionless environments. To maintain consistency and integrity across the distributed system, CERTH brings Hyperledger Fabric's consensus algorithm, ensuring synchronized, verifiable state across nodes.

This architecture contributes to O-CEI's vision of scalable, federated data services, where participants can publish and consume digital products through trusted infrastructures. Decentralized identity and access management (IAM) is a key element: DOME incorporates a W3C-compliant IAM framework using OIDC4VP and SIOPv2, ensuring verifiable, decentralized authentication that is compatible with EBSI and the emerging eIDAS2 regulatory framework.

Through these aligned efforts. across marketplaces, semantic brokers, and identity frameworks, O-CEI creates the conditions for trusted participation in a federated European data economy. This infrastructure supports not only data exchange but also onboarding of tools, workflows, and AI services, furthering the project's commitment to digital sovereignty and cross-domain innovation.

The discussion above has outlined the foundational infrastructure components, from connectors to marketplaces, that enable O-CEI vision for sovereign and trusted data sharing. These building blocks formulate the operational core of O-CEI data space architecture, ensuring secure, policy-compliant, and interoperable interactions across domains and stakeholders. The following table consolidates these key enablers and their respective roles within the interoperability infrastructure of O-CEI:

Table 3: Data sharing Infrastructure in O-CEI

Component	Type	Purpose	Key Features	Positioning	TF3 Blocks
Eclipse DataSpace Connector (EDC)	Connector Framework	Secure, policy-enforced data exchange without full dataset replication	Catalog API, Contract Negotiation, Policy Enforcement, Supports JSON-LD and AS4, IAM integration	Data Space Middleware for sovereign data exchange	Data Management
FIWARE Data Space Connector (DSC)	Connector Framework	Interoperable data sharing with individual and device identity support	OID4VC support, DSP compatibility, TM Forum API integration, multi-actor identity	Data Space Middleware for trust and identity-sensitive contexts	Data Management
FIWARE BAE /	Federated Marketplace Platform	Cross-domain publishing, discovery, and	TM Forum APIs, VPs/VCs for products,	Marketplace Layer for cross-	Data Management, Orchestration

DOME Marketplace		negotiation of digital assets	federation across domains	domain interoperability	
UPV Semantic Data Brokers	Semantic Broker	Semantic annotation, dynamic discovery, contract enforcement for data assets	Logging, access control, standards-based broker architecture	Distributed CEI semantic edge integration	Data Management

With these architectural foundations in place, the next layer of interoperability lies in the **semantic alignment** of data — enabling not just connection, but shared understanding across sectors and systems

Semantic Interoperability and Domain-Specific Modelling

While the marketplace provides the unified space for publishing and discovering assets, true interoperability across domains depends on a shared understanding of data — how it is described, structured, and exchanged. This requires more than APIs and connectors: it calls for a semantic foundation that enables consistent meaning across heterogeneous systems, domains, and stakeholders.

Semantic interoperability is a cornerstone of O-CEI's technical architecture. To support scalable and federated collaboration across sectors such as energy, agriculture, mobility, and environment, the project applies a layered strategy: starting with widely adopted semantic standards, extending to domain-specific ontologies, and culminating in tools for mediation, conformance validation, and runtime transformation. These foundations ensure that data assets — whether exchanged through connectors or discovered in the marketplace — can be integrated, understood, and reused in a consistent and machine-actionable way.

To realise this layered strategy in practice, O-CEI integrates semantic technologies such as JSON-LD, RDF, OWL, and JSON Schema to support consistent and interoperable data representation. Widely adopted models like SAREF, FIWARE Smart Data Models, and the W3C SSN ontology serve as shared semantic foundations. These provide the structured vocabularies needed to harmonise data architectures and enable data portability across the Cloud-Edge-IoT continuum.

Recognising that each system or domain may carry its own specificities, O-CEI promotes a design approach where core ontologies are reused and extended, ensuring semantic alignment while preserving the ability to capture detailed, domain-specific requirements. This balance between generalisation and specialisation ensures that data remains both interoperable and expressive. In dynamic, heterogeneous CEI environments, however, interoperability must extend beyond static design. O-CEI supports runtime semantic mediation, enabling data translation and harmonisation on the fly. This is especially important for integrating legacy systems or constrained IoT devices that may not natively support modern semantic models.

Data transport across such systems typically relies on standard APIs (e.g., REST), event streams (e.g., MQTT or Kafka), or structured documents (e.g., XML, JSON-LD). Achieving seamless integration requires the combination of transport protocols with semantic transformation tools, leveraging technologies like SPARQL, RML, or IPSM-AF to map between representations. To operationalise these mechanisms in real-world CEI settings, O-CEI partner UPV contributes a suite of syntactic and semantic mediation tools. These include protocol adapters and data format transformers that facilitate communication across diverse systems — bridging standards like REST, MQTT, RDF, and XML — and enabling interoperability even in industrial and resource-constrained environments.

To ensure that semantic interoperability is implemented in a consistent and replicable way, O-CEI benefits from methodological frameworks developed and validated in large-scale European projects. A key contribution comes from TRIALOG, which has established an interoperability-by-design methodology [13] applied in domains such as energy, health, and IoT.

This framework was developed through work in projects such as GIFT and MAESHA, and is currently being extended in the Horizon RESONANCE project, where it is used to describe and share solutions developed in pilots, facilitate replication, and identify potential standardization gaps at European or international level [14] . Insights gained—particularly through the Int:net project—have contributed to strategic recommendations for an interoperability roadmap for the energy transition as [15], well as guidelines for the testing of Smart Energy Systems[16]. In addition, two reports on the evolution of interoperability standards have been published under Int:net [17], one analyzing trends in semantic standardization and the other offering governance recommendations for coordinating ontologies such as SAREF and SAREF4ENER.

In the context of Int:net, TRIALOG contributed to an interoperability case survey template, designed to identify and document common interoperability challenges across complex systems. This approach supports knowledge transfer and helps highlight standardisation gaps. It is currently applied within RESONANCE to ensure consistency across pilot implementations and to feed recommendations into European roadmaps for the energy transition.

Through this work, TRIALOG has also contributed to the evolution of semantic standards, including coordination and governance recommendations for SAREF and SAREF4ENER, as well as new tooling to support conformance verification (detailed later in this section). These contributions form part of the semantic backbone leveraged within O-CEI.

Building on these frameworks, O-CEI also integrates tested semantic models from real-world data space initiatives. In the Omega-X project, several O-CEI partners contributed to the development of the Common Semantic Data Model (CSDM)—a modular ontology stack tailored for energy flexibility and broader energy data ecosystems. Developed through participatory workshops and technical alignment, CSDM integrates and extends existing standards such as ETSI SAREF, IEC CIM, and SEDMOON. It includes two core ontologies: DataValueOntology, for representing time series and events, and EnergyDataSetOntology, for describing structured energy datasets.

This model supports interoperability across multiple energy-related domains—such as flexibility services, renewables, local energy communities, and electromobility—making it a strong foundation for semantic consistency in O-CEI's energy use cases. Within Omega-X, TRIALOG, in collaboration with **EDF**, led the use case analysis [18] and contributed to the modeling process, ensuring the resulting ontology aligns with both **IDSA** and **Gaia-X** semantic profiles.

Complementing this work, O-CEI also leverages extensive experience in **agriculture and food systems**, where federated data sharing requires similar semantic strategies. Key contributions come from **GSC**, a core partner in projects like **Ploutos**, **Farmtopia**, and **OpenAgri**. In these initiatives, GSC developed a modular and extensible semantic architecture for agrifood data spaces, addressing challenges such as interoperability with legacy systems, stakeholder-specific access control, and reuse of existing standards.

Under Ploutos, GSC coordinated the creation of the Ploutos Common Semantic Model (PCSM), which builds on ETSI SAREF4AGR and reuses several standard ontologies—including ENVO, SSN/SOSA, FoodOn, OM, and SAREF4AGRI—covering environmental, observational, organizational, and measurement aspects. This demonstrates how the model integrates widely accepted semantic assets to support cross-domain interoperability in agriculture. This model supports queries across distributed networks of agrifood stakeholders and has been validated in pilots involving commercial actors in the processed fruit sector [19].

Building on PCSM, the Farmtopia Semantic Model (FSM) introduced refinements to support complex measurement units (QUDT), time periods (OWL Time), spatial alignment (INSPIRE), and semantic reference vocabularies like Agrovoc, EPPO, and OBI. FSM also introduces structures for modeling services and datasets, based on the IDS Information Model and Data Cube Vocabulary, making it especially suitable for integration with the O-CEI data marketplace infrastructure.

These models have been implemented using developer-friendly formats like JSON-LD and integrated with REST APIs, supporting real-world deployment and interoperability in agri-environmental monitoring, remote sensing, and supply chain traceability. In ongoing projects such as Farmtopia and OpenAgri, where GSC plays a leading role, these efforts are being extended and offer clear opportunities for synergy with O-CEI's goals in energy–agriculture integration.

To further support cross-domain semantic alignment, O-CEI incorporates Smart Data Models, an open, harmonized modeling initiative developed by the FIWARE Foundation in collaboration with TM Forum and IUDX. These models provide reusable schemas based on the NGSI-LD standard from ETSI, enabling rich semantic interoperability through linked data principles and context-aware entity relationships.

NGSI-LD extends conventional data modeling with semantic annotations and entity linking, allowing datasets to be queried, subscribed to, and integrated across domains in real time. By leveraging NGSI-LD and Smart Data Models, O-CEI supports interoperability across diverse systems—spanning energy, environment, mobility, logistics, and agriculture—while also simplifying integration workflows for real-time analytics and decision-making.

Relevant Smart Data Model domains for O-CEI include:

- **Smart Energy**: for consumption, production, and grid data.
- **Smart Agrifood**: for field monitoring, AgriApps, and agricultural processes.
- **Smart Cities and Logistics**: particularly port infrastructure and transport nodes.
- **Smart Environment**: for air quality, weather, and IoT-based environmental sensing.
- **Smart Transportation**: including EV charging infrastructure and fleet management.

Each Smart Data Model includes a technical schema, specification, URIs, and example payloads in NGSI v2 and NGSI-LD formats, offering both human readability and machine-actionability. These models are public, open-source, and royalty-free, promoting innovation, portability, and adoption across regions and platforms.

In cases where pilot-specific needs arise, O-CEI partners may define custom extensions or entirely new models, following the same open design principles to ensure continued semantic compatibility and cross-domain reusability.

To make the application of Smart Data Models more scalable and developer-friendly, NEC is piloting the use of **Generative AI and LLM-based agents** for semantic alignment. In this approach, AI agents analyze new data streams to identify the most relevant Smart Data Model, assist in mapping transformations, and generate FogFlow operators for deployment in the CEI infrastructure. While this solution remains exploratory, it showcases a promising strategy for operationalizing semantic interoperability at scale.

To complement these semantic modeling efforts and reinforce trust and interoperability, O-CEI integrates ontology-based validation and conformance testing tools. A key contribution comes from TRIALOG, which has developed the Ontology-Driven Constraint Tester (ODCT)—a tool designed to validate datasets against domain-specific ontologies such as SAREF and SAREF4ENER. Originally

piloted within the *'Flexible Start'* use case of the EU Joint Research Centre's Code of Conduct for Energy Smart Appliances, the ODCT ensures that smart devices (e.g., washing machines, thermostats) operate in accordance with agreed semantic constraints. This increases not only semantic consistency but also automation and trust in CEI-wide data interactions.

As semantic models mature, O-CEI advances toward operationalizing semantics through machine-actionable metadata and lifecycle-aware governance. A critical enabler here is JSON Schema, which defines structural properties of data products but is being extended with linked data semantics using JSON-LD contexts. This decoupling—syntax via JSON Schema, meaning via semantic vocabularies—supports modularity, enhances automation, and aligns with FAIR principles.

In the data marketplace context, this dual approach is reflected in standards such as the Open Open Data Product Specification (ODPS), which encodes licensing, quality, and lifecycle metadata directly into machine-readable formats. Similarly, the Dataspace Protocol (EDP) combines JSON Schema and JSON-LD to define standardized message formats and contracting mechanisms for federated exchanges.

Together, these components define a maturing semantic ecosystem where JSON-compatible formats serve not only validation but also governance, orchestration, and policy enforcement. This evolution underpins O-CEI's strategy to build trusted, scalable, and interoperable data infrastructures, aligning with Gaia-X and EDC principles and positioning the project at the forefront of Europe's emerging data economy. While semantic models ensure that data is intelligible and consistent across domains, their practical deployment at scale requires an infrastructure capable of dynamic integration and delivery. This is where the data fabric paradigm becomes essential.

Data fabric and execution-aware storage

Building upon these interoperability foundations, the **data fabric paradigm** extends the ecosystem's capabilities by introducing mechanisms for semantically aligned data flow across distributed environments. While connectors and protocols ensure interoperable negotiation and access control, the data fabric focuses on the **continuous, metadata-driven management and delivery of data**, including mechanisms for routing, enrichment, and execution-aware exports toward consuming applications and services. It offers a unifying architectural layer that bridges diverse data sources and silos, exposing them through context-aware, federated views tailored to specific operational or analytic needs.

Metadata serves as the foundation for the data fabric, guiding the process of raw data ingestion, connecting diverse data silos, and tailoring the exposure of these integrated data to the consumers' requirements. This new data management paradigm abstracts product consumers – including operational applications, AI/ML models, or visualization tools – from the underlying complexities and location of data and services. In this context, metadata functions as a smart integration layer, facilitating seamless data interoperability between various product providers and consumers, as depicted in Figure 8. To realize the data fabric, the knowledge graph is envisioned as an enabling technology to build an integration layer based on the semantics of the data and the flexibility of graph structures.

To realize this architecture, **knowledge graphs** are envisioned as a key enabler. Knowledge graphs elevate metadata management and integration, facilitating the seamless integration of disparate data and services. Knowledge graphs unlock the creation of a metadata-driven layer, grounding semantic data models such as ontologies and taxonomies, which capture the concepts and the relations that underpin data silos within organizations. In this respect, standards from the Semantic Web such as RDF, OWL, or SPARQL, are leveraged in combination with existing open-source projects such as Chimera .

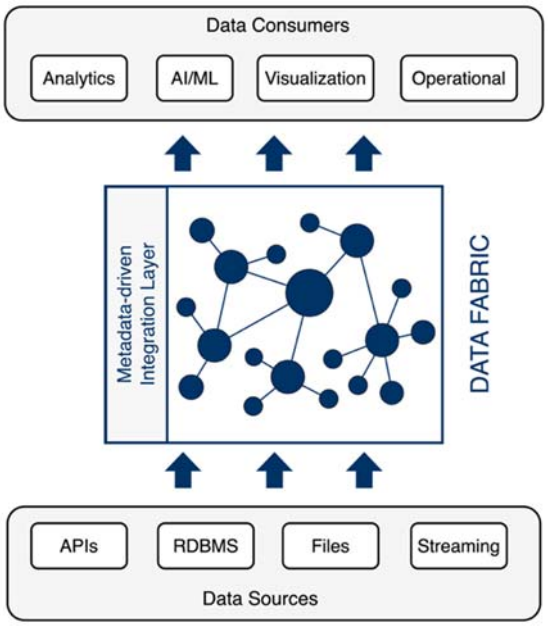


Figure 8. Conceptual architecture of data fabric.

By building upon a knowledge graph, the data fabric enables data sharing by integrating data silos using proper access protocols and schemas and formats and delivering the integrated to the consumers based on their needs (i.e., destination protocol and data schemas and formats). To illustrate this, Figure 9 shows a scenario with consumers that requires combined data from an SQL database and an XML file, with in turn, are transformed and delivered as a JSON file.

In this data exchange flow, the data fabric implements specific connectors to extract data from the different data sources. After this step, other components map these data to domain-specific ontologies, transforming and integrating the data into a knowledge graph. Finally, exposure components convert this graph representation into the requested JSON format and schema and deliver it as a file. This same workflow is followed for other use cases, such as transforming JSON data obtained from a message queue like Apache Kafka into a CSV file that follows the structure indicated by the consumer.

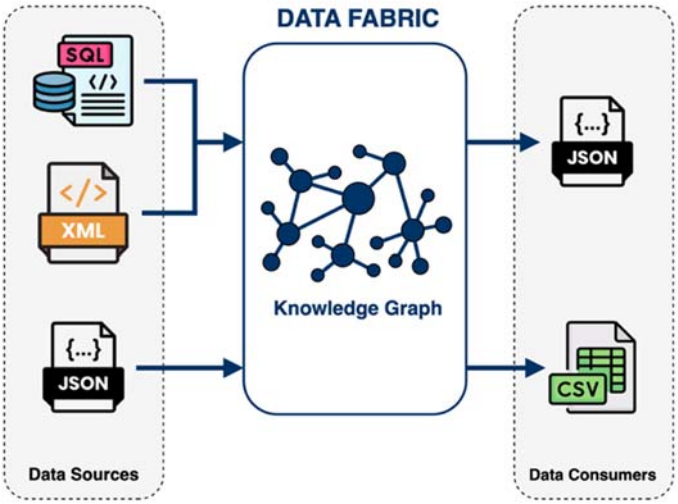


Figure 9. Data exchange between heterogeneous sources and consumers is achieved through a data fabric built on a knowledge graph.

To operationalize such flexible and real-time data integration patterns, the underlying infrastructure must support efficient, scalable, and topology-aware data movement, especially across the heterogeneous and resource-constrained environments of the CEI continuum. Zenoh's, one of the components to be integrated in O-CEI blueprints, architectural model, unifying **data-in-motion**, **data-at-rest**, and **edge computation** into a lean and topology-adaptive protocol, epitomizes the next-generation shift toward federated, real-time data fabrics. Its support for **locality-aware routing**, **low-overhead messaging**, and **versatile topologies** (peer-to-peer, brokered, routed) directly addresses O-CEI's goal of scalable, energy-efficient CEI deployments.

In parallel to transport and routing innovations, O-CEI also explores execution-aware storage paradigms that minimize data movement by enabling "computation-in-place", a crucial step toward efficient, distributed data processing in CEI environments. Emerging object-oriented data stores (e.g. dataClay) support this shift by persisting data in native in-memory formats and executing processing routines directly where the data resides. These models reduce latency and network overhead, especially in simulation or analytics workloads, and offer fine-grained access control, versioning, and data governance, key capabilities for federated CEI infrastructures.

The following table consolidates the semantic foundations and runtime integration mechanisms that make interoperable, context-aware data sharing possible across CEI domains. These components collectively support model-based data governance, semantic mediation, real-time alignment, and intelligent infrastructure.

Table 4: Semantic and Runtime Interoperability Enablers

Component	Type	Purpose	Key Features	Positioning	TF3 Blocks
TRIALOG ODCT	Validation Tool	Semantic conformance testing of datasets against domain ontologies	Ontology validation, smart appliance conformance	Semantic Validation in CEI pipelines	Data Management
Omega-X Common Semantic Data Model (CSDM)	Semantic Ontology Stack	Semantic interoperability for energy use cases	ETSI SAREF, IEC CIM, DataValueOntology, EnergyDataSetOntology	Semantic foundation for energy-related data sharing	Data Management
Farmtopia Semantic Model (FSM)	Semantic Ontology Stack	Interoperability in agriculture and food data sharing	QUDT, OWL Time, INSPIRE, DataCube, JSON-LD mappings	Semantic modeling for agrifood CEI data	Data Management
Smart Data Models (FIWARE, TM Forum)	Semantic Data Model Library	Cross-domain semantic modeling using NGSI-LD	Standard schemas, linked data, domain-specific models	Foundational semantic vocabulary for all CEI sectors	Data Management
TRIALOG Interoperability-by-Design Framework	Methodology	Governance and replication of semantic strategies	Case surveys, roadmap input, energy domain alignment	Semantic governance and methodology	Resource Management
Data Fabric & Knowledge Graph	Execution-Aware Integration Layer	Metadata-driven data routing, enrichment, and delivery	Federated views, semantic metadata, graph-based modeling	Execution layer integration across data providers/ consumers	Data Management
NEC Generative AI Semantic Alignment	AI-based Semantic Assistant	Use AI agents to align new streams with Smart Data Models	FogFlow operator generation, data stream analysis	AI-enabled semantic integration	Data Management
Zenoh	Data Fabric Transport Layer	Unified data-in-motion and data-at-rest transport	Locality-aware routing, peer-to-peer and brokered topologies, low-latency messaging	Real-time, topology-adaptive data routing for CEI fabrics	Data Management, Orchestration

While the data fabric infrastructure routes and semantically aligns data into data spaces for secure and interoperable exchange, emerging protocols are also enabling intelligent data flow outward—toward AI-driven applications and autonomous agents. One such example is the **Model Context Protocol (MCP)**, which provides a standardized interface for AI agents, including large language models, to dynamically and securely interact with federated data sources and external tools. Although not part of the core data space architecture, MCP exemplifies how downstream AI systems can be built on top of CEI data infrastructures through trusted, context-aware interfaces. This emerging layer of interaction forms a bridge between structured data ecosystems and cognitive computing systems—laying the groundwork for the AI-driven orchestration and reasoning capabilities discussed in the following section

3.1.6 AI-Driven computing & cognitive cloud

AI is emerging as one of the most transformative technological shifts in history, with profound implications for innovation, productivity, and the global economy. Especially, generative AI is poised to significantly reshape workflows and accelerate decision-making across a wide range of domains. Recognizing this, O-CEI has established AI-Driven Computing and Cognitive Cloud as one of its key technological positioning pillars, exploring how AI can enhance both infrastructure operations and intelligent service delivery across the continuum. This reporting informs the technical objectives of Task 3.5 (“Upscaling of data management, observability and AI supporting utilities”), highlighting innovations that integrate AI into the O-CEI platform and enable advanced, distributed AI capabilities.

AI integration into CEI environments typically involves two phases: model training and inference. Each introduces its own architectural, resource, and deployment requirements, which must be addressed when designing for efficient AI enablement in the continuum. Viewed across this lifecycle, there are two main perspectives: on one end of the spectrum, training large foundation models is extremely resource-intensive and currently feasible only for a few well-resourced organizations; on the other end, inference is a high-throughput process already widely adopted across industries. Between these extremes lies a key opportunity: enabling more organizations to participate in the fine-tuning of pre-trained models, either starting with smaller models or scaling up gradually, or to develop lightweight models tailored to their specific domains. By leveraging their own or domain-specific datasets, these organizations can create specialized, cost-effective AI solutions. In some cases, both inference and training may need to occur at the edge to meet data privacy and security requirements. However, this introduces constraints on available resources, making it essential to strike a balance between model size and accuracy to suit edge environments.

“As AI continues to evolve, the cloud-edge continuum will play a critical role, not only in combining processing capacity across distributed HPC and cloud systems but also in addressing the diverse needs of different types of AI systems.”[2]

This brings us to the two complementary approaches involving AI and the cloud-edge continuum: AI for cloud and cloud for AI. While we previously introduced the cloud-edge continuum as a necessary infrastructure to support distributed AI workloads, the reverse perspective is equally important: AI is increasingly essential to managing the complexity of this continuum itself. This idea underpins the concept of the Cognitive Cloud—an AI-enabled cloud-edge continuum capable of intelligently and securely orchestrating the growing scale, heterogeneity, and dynamism of distributed environments. As highlighted in the previously discussed pillars, AI integrates advanced support for a wide range of functionalities across the continuum, from MetaOS and orchestration, to trust and security, and data spaces and interoperability. Its application provides for multiple layers of the architecture, contributing to both infrastructure-level efficiency and application-level intelligence. Bringing these dimensions together, the Cognitive Cloud seamlessly integrates cloud and edge resources while

[2] https://support.opennebula.pro/hc/en-us/article_attachments/22599431725341

dynamically adapting to changing application behaviors and data variability. It enables smart decision support for automatic deployment and service mobility, introducing context-aware decision-making—offering secure, responsive computing for diverse users and use cases.

Extending this vision, O-CEI integrates the **Cognitive Cloud** concept acting as a distributed, intelligent management layer across the cloud-edge continuum which not only enables intent-based orchestration and contextual service management, as previously discussed, but also leverages AI techniques to continuously learn from the system's state, monitoring infrastructure usage, user demands, and application behavior across heterogeneous environments. This allows it to proactively operate to dynamically rebalance workloads, optimize energy consumption, maintain high-throughput, low-latency service delivery and generally optimize many other crucial key performance indexes. For example, AI-driven orchestration can adapt compute placement based on green energy availability, helping reduce environmental impact while meeting performance requirements. Similarly, it supports balanced data traffic routing, real-time resource scaling, and autonomous fault recovery, ensuring a resilient and efficient continuum. To support the vast amounts of data and diverse requirements generated by IoT deployments across sectors, the Cognitive Cloud must operate in a federated and standards-based architecture, one that embraces open platforms, interoperability models, and composable services. O-CEI targets to integrate research results which will guarantee and intelligent orchestration layer which not only empowers developers with fine-grained control over compute, data, and network resources but also ensures that end users experience seamless, secure, and uninterrupted services—regardless of location, context, or device.

In parallel with advancements in platform orchestration and Cognitive Cloud architectures, it is equally important to consider the integration of AI within the rapidly evolving landscape of next-generation networks. AI and ML are increasingly viewed not just as add-ons, but as complementary paradigms to 5G, 6G, Edge Computing, and Open RAN. Their growing presence in various network processes has already demonstrated clear benefits, particularly in automating service and resource management. As these technologies continue to evolve, they are expected to significantly influence the design and operation of future networks by enabling predictive and autonomous capabilities across multiple network layers.

In particular, 6G is introducing the notion of native cognitive capabilities directly into its core network architecture, referred to as the Cognitive Service Architecture [22]. This new model aims to meet the rising demands for service quality and operational flexibility in increasingly complex scenarios. Rather than adding AI post hoc, 6G core systems are being designed to inherently support cognition, enabling networks to sense, learn, and adapt in real time. AI is now widely recognized as a foundational pillar of future 6G systems, prompting intensified research into AI-based enablers. This integration will reshape how 6G systems handle APIs, data pipelines, workload distribution, computational offloading, MLOps practices, and robust privacy and security mechanisms—each aligned with emerging application and user demands.

Given that 6G is expected to support massive parallelism and diverse, dynamic network slices, scalability and sustainability are becoming central concerns. O-CEI integrates, from contributing partners, AI-driven, zero-touch management and orchestration frameworks capable of responding in real time to shifting workloads, mobility patterns, and energy constraints.

6G networks will be AI-native. This next-generation architecture will be defined by several key features:

- *Intelligence everywhere*
- *Zero-touch management*
- *Distributed data infrastructure*

- *AI-as-a-Service (AIaaS)"*[3]

Concluding, intelligence will be seamlessly embedded throughout the network, from centralized nodes through edge to IoT devices, enabling AI/ML workloads to be executed where they are most effective. The synergies with applications classified in more TF3 building blocks, beyond AI, is obvious as this will necessitate a robust distributed data infrastructure to support real-time data availability and processing. The vision is to deploy AIaaS, across the continuum, which will provide access to AI/ML models, datasets, and tools as services. This approach can both empower zero touch management for autonomous network operations and will also empower service providers and users to harness AI capabilities and develop innovative use cases.

Cognitive Computing & Distributed Intelligence

Building on this vision of embedded intelligence, O-CEI explores advanced cognitive architectures that enable distributed, self-organizing decision-making across the CEI landscape. One such approach is based on swarm intelligence and game-theoretic mechanisms, offering a novel model of cognitive behavior across cloud, edge, and IoT devices.

The concept of the **Cognitive Computing Continuum** embodies a paradigm shift in distributed intelligence, where IoT, Edge, and Cloud devices collaboratively self-organize into dynamic, autonomous swarm continuums. These interconnected ecosystems are designed to optimize data processing and enable seamless, real-time service delivery across highly heterogeneous environments. Rather than relying on centralized control, each device operates as a cognitive entity, capable of assessing its own data, computing resources, and contextual environment to participate in distributed decision-making processes.

To enable this level of autonomy and intelligence, a novel framework based on **Asymmetric Multi-Player Competitive Games** is introduced. These games embed self-organizing mechanisms, such as **Pricing**, **Bidding**, and **Auctioning**, into the cognitive logic of each device. Through this embedded game-theoretic intelligence, devices learn to evaluate (or price) their internal resources and data**,** determine (or bid) their optimal level of participation in shared computing tasks**,** and negotiate (or auction) collaborative strategies within broader swarm environments.

This decentralized optimization approach transforms the entire network into a dynamic, self-regulating ecosystem. As devices continuously engage in local decision-making and negotiation, they collectively give rise to **Swarm Continuums**—fluid, purpose-driven computing collectives that adapt in real time to fluctuating workloads, user demands, and contextual conditions. The outcome of this process is a **convergence toward** mathematically sound Computing Equilibriums: optimal operational states where services are fully satisfied, and no computational resource or data asset is wasted.

Beyond performance optimization, these cognitive games are further enriched with capabilities to sense and adapt to additional layers of operational complexity. Devices are equipped to monitor cyber-risk exposure**,** energy consumption**,** hardware faults**,** and other heterogeneity factors that influence system stability. This allows for **context-aware decision-making**, ensuring that swarm behaviors align not only with service performance but also with broader goals of **cyber-resilience**, **digital privacy preservation**, and **energy efficiency**. By granting devices increased autonomy in safeguarding these non-functional properties, the suggested methodology lays the foundation for intelligent infrastructure that is not only self-optimizing, but also self-protecting and environmentally conscious.

[3] https://support.opennebula.pro/hc/en-us/articles/26384602773917-White-Paper-OpenNebula-Telco-Cloud-Model-and-Architecture?utm_content=329630562&utm_medium=social&utm_source=linkedin&hss_channel=lcp-977603

As already discussed in the MetaOS & Orchestration pillar, this vision of intelligent, adaptive orchestration is further exemplified in the COGNIT project, which introduces a **cognitive orchestration** framework centered on an AI-powered, serverless architecture. This provides an AI-powered orchestration layer capable of making autonomous, real-time decisions for task placement, resource allocation, and load balancing—based on live system state, contextual variables, and compliance requirements. It provides AI-driven orchestration logic and cognitive adaptability, strongly aligned with the broader ambition to embed intelligence into operational layers and demonstrates how AI can act not only as a workload optimizer, but also as an autonomous decision-making agent—crucial for enabling adaptive, intent-aware, and self-organizing behaviors across the CEI continuum.

Building upon this vision of distributed autonomy, AI is also being applied to orchestrate and manage system behavior in a **trust-aware, intent-driven** manner**.** A compelling example of applying AI for trust-aware orchestration comes from SAFE-6G project, which introduced the **Cognitive Coordinator**, an advanced AI-driven control layer designed to translate high-level user trust intents into concrete, system-level actions. Rather than relying on static security policies, the coordinator dynamically interprets user expressions of trustworthiness and guides the behavior of specialized network functions, called Trust Functions, to align the system with these expectations.

This AI-driven orchestration loop operates across two coordinated domains: a **regression domain**, which uses deep learning to estimate a preliminary trust score, and a **reasoning domain**, which refines that score using contextual awareness, system telemetry, and historical insights. The result is a calibrated trust target, which is then enforced by deploying containerized functions that adapt the system state accordingly. In essence, this cognitive approach enables **adaptive, intent-aware system reconfiguration**, paving the way for next-generation infrastructures where trust becomes a first-class operational parameter. The concepts introduced here strongly resonate with O-CEI's ambition to embed intelligence into orchestration processes—enabling secure, context-driven, and user-aligned behavior across the CEI continuum.

In parallel with orchestrating trust goals at the system level, AI is also being applied to streamline infrastructure automation and compliance. One example is the evolution of DevSecOps practices through AI-enhanced Infrastructure as Code (IaC) approach for managing and provisioning of IT infrastructure through code instead of through manual processes. With IaC, configuration files are created that contain infrastructure specifications, such as networks, virtual machines, load balancers, and connection topologies. These files can be used to define and deploy infrastructure in a consistent, repeatable, and scalable way.

To address IaC trustworthiness, current SotA is using the DevSecOps approach, integrating security practices into the DevOps pipeline with automated security testing, vulnerability scanning, and compliance checks in the design and runtime phases. The aim is to ensure consistency and uncover vulnerabilities frequently arising from incorrect configurations, parameter values, type errors, and deviations from established best practices and specific standards for specification languages.

Newer approaches shift focus on automating the IaC hardening process with a system for scanning components, libraries and code to identify security vulnerabilities such as CVEs (Common Vulnerabilities and Exposures) and then implementing a resolution plan, such as updating the affected component/library. Additionally, for the secure DevSecOps pipeline, it is also important to focus on automating compliance with requirements and with standards/regulations. This involves checking the entire software supply chain, specifically using an SBOM (Software Bill of Materials), to identify misconfigurations and vulnerabilities. LLMs fine-tuned with expert knowledge, or methods like RAG (Retrieval-Augmented Generation) and Knowledge Graphs, can be employed to enhance this process.

Using AI approaches, especially LLMs (e.g. with RAG and Knowledge Graphs), is shifting the automation towards automating the entire DevSecOps process, e.g. with specialized context-aware AI agents (one for each step – build, test, deploy, monitor…). The end results could be a system that can provide clear, understandable advice about what's happening and what actions developers need to take.

Machine Learning Workflows and Federated Learning for Distributed AI

Having discussed how AI enhances the continuum's capabilities, we proceed to investigate how O-CEI integrates support for AI over the continuum, through scalable machine learning workflows and federated intelligence. Several O-CEI partners, through projects like SWARM, LASAGNE, and Smart Energy District, have developed expertise in setting up machine learning (ML) workflows to predict electricity consumption. The workflows involve the ingestion and preprocessing (cleaning, down-sampling, feature extraction and forming) of raw data from electricity measuring devices. The ML model is trained on the pre-processed data and the results are used to predict energy usage of real-time data. These workflows are applied to electricity consumption measurements from different sources such as electric vehicle chargers, houses or commercial properties. Several machine learning algorithms were tested, including XGBoost, LSTM, CNN-LSTM, CNN-LSTM with attention mechanism, Transformer, Prophet, and DeepAR. Though for the implementation LSTM family algorithms—namely LSTM, CNN-LSTM, and CNN-LSTM with attention mechanism, seem to be better fit. This choice is driven by two key reasons:

- The LSTM family can be used in the case of a Federated Learning (FL) approach.
- These algorithms can be deployed on edge devices (limited CPU and memory resources).

While these workflows already demonstrate the feasibility of deploying ML at the edge, scaling such solutions across heterogeneous environments raises challenges around data privacy, communication overhead, and coordination. To address these, Federated Learning (FL) offers a powerful alternative to centralized training approaches.

FL techniques further support privacy and scalability as they introduce edge-based collaborative training without data centralization allowing thus multiple devices or organizations to collaboratively train a model without sharing their raw data. This approach moves the computation to the data rather than the other way around and enables collaboration without compromising privacy. FL utilizes a client - server logic where each device, the client, trains the model on each own data and sends the training updates to a central aggregator, the server. The server in turn performs averaging to form a global model and redistributes it to clients for further local training. This cycle continues for several rounds until the model converges. FL is a model-agnostic and data-agnostic architecture and is by nature applicable to a setup of IoT devices that work as separate clients. The elimination of data sharing between participants provides enhanced data privacy and security, better compliance with GDPR and makes this approach very well-suited for bandwidth-constrained environments. This innovative approach does not come without challenges like ensuring communication efficiency and securing the training process against malicious participants. In order to tackle those challenges a few approaches are explored including Homomorphic Encryption and Multi Party Computation in weight handling and ensuring client-server TLS/SSL connections.

To operationalize Federated Learning in real-world scenarios, several open-source frameworks have emerged and are brought by O-CEI partners to support deployments. Among the most mature Python-based FL frameworks are Flower, Nvidia Flare, PySyft, FedML, and Google's TensorFlow Federated. Flower is particularly popular in research due to its versatility and ease of use, while Nvidia Flare builds on similar principles with a focus on stricter production-grade deployments. Notably, Flower's upcoming emulation mode is extending its applicability beyond simulation, enabling real-world testing with real devices and network environments.

The potential of FL extends well beyond a single domain, with applications emerging across diverse sectors—from healthcare to robotics—highlighting its flexibility and growing maturity. Recent FL applications range from healthcare (e.g., Google's Kakohealthcare) to robotics, where strategies like FLDDPG—an FL-based deep reinforcement learning method for swarm robotics—have shown promise. Another example is FREDY, a federated learning approach that integrates PATE (Private Aggregation of Teacher Ensembles) with FL-based teacher training to strengthen data privacy and resilience.

As intelligent workflows and federated learning approaches take root across the CEI continuum, the next leap lies in the integration of intelligent agents that can reason, plan, and act across data-rich, semantically structured environments. These capabilities intersect closely with the Data Spaces and Interoperability pillar, as AI agents rely on semantic interoperability and structured data access to identify, retrieve, and integrate relevant information across domains, supporting autonomous, context-aware operations.

AI Agents in the CEI Continuum

Building on this evolution, the emergence of AI agents powered by LLMs marks a critical step toward enabling complex automation and decision-making across distributed systems. These agents operate across the CEI landscape using natural language, reasoning chains, and domain-specific knowledge—advancing the paradigm of cognitive computing introduced earlier.

AI Agents perform automation tasks on behalf of a stakeholder in the system. These agents are equipped with **data space exploration functions**, allowing them to identify available data, services, and peers within the ecosystem. Typically built using **generative capabilities of LLMs**, such agents are increasingly used for **reasoning, text analysis**, and other cognitive automation functions. This capability closely aligns with the principles of Data Spaces and Interoperability, where agents must dynamically discover, interpret, and interact with distributed datasets and services based on shared ontologies and semantic contracts.

A typical implementation architecture for AI Agents uses an **orchestrator/pipeline mechanism**, where the orchestrator analyzes the assigned task, decomposes it into subtasks, and executes the sequence through an intelligent pipeline. A widely adopted design pattern is **RAG**, which enables the agent to retrieve contextually relevant enterprise-specific information (e.g. from documents), augment the LLM prompt with this context, and generate an action plan specialized to the current environment.

These agentic patterns are already being applied across diverse domains, including **agri-commodities, document intelligence, and IoT data processing**, with results deployed in operational and production-grade environments. Such existing tools analyzes **agricultural commodity data** to support informed decision-making. It helps identify risks and opportunities in trade flows, logistics, and market dynamics—contributing to transparency and efficiency in the global agri-commodities sector. From a technical perspective, the tool employs RAG **architecture** to integrate domain-specific knowledge into pre-trained LLMs (e.g., Claude by Anthropic). This enables effective engagement with unstructured sources such as PDF reports and news articles. Retrieved documents are used to condition the model's output, ensuring answers are grounded in verifiable evidence.

In the same context, HESSO has explored the use of LLMs to enhance Named Entity Recognition, focusing on identifying and classifying key elements such as agricultural products, vessel movements, and trade flows without requiring additional training. Our evaluation included Claude, GPT-3.5, and GPT-4 [23]. Building on these capabilities, further developments have focused on **structuring unstructured data** by:

1. Analyzing the content and layout of incoming documents.
2. Identifying document type and semantic structure.
3. Mapping extracted data into predefined datasets or knowledge graphs.

In a separate project, HESSO has contributed to enhancing an **AI assistant for querying IoT-related data**, using a RAG-based system to retrieve documentation and visualize JSON schema. The assistant interprets complex user queries by decomposing them into **multi-step reasoning chains**, thereby improving precision and interactivity. This shift toward **agentic behavior**, where agents plan, reason, and reflect, marks a significant evolution in how LLMs are operationalized in distributed intelligent systems.

To fully realize this evolution, agents require standard mechanisms to interact dynamically and securely with external systems. **The Model Context Protocol (MCP)** provides such a foundation as it enables agents to interact with data and services dynamically and securely. The MCP is an open standard developed in 2024 by Anthropic to enable AI systems, including large language models (LLMs), to securely and efficiently interact with external tools, data sources, and applications. By standardizing these interactions, MCP addresses the complexity and fragmentation associated with custom integrations, allowing AI models to access real-time, structured data seamlessly.

To support such a functionality, each external data source or tool can be abstracted as an **MCP server**, providing a unified interface that AI agents can query or interact with dynamically. This architecture integrates naturally with **AI agents**, as it allows them to autonomously explore available MCP servers, identify relevant sources, and incorporate them into their task execution workflows. For instance, **NGSI-LD Context Brokers** can be exposed as MCP servers, giving AI agents seamless access to rich IoT and contextual data streams. Another strength of MCP is that it can be integrated into the RAG pattern (discussed above). Figure 10 demonstrates how MCP operates on a client-server architecture, where AI applications (clients) communicate with various data sources and tools (servers) through a unified protocol. This design facilitates dynamic context exchange, enabling AI systems to perform tasks such as reading files, executing functions, and handling contextual prompts across diverse environments. The protocol utilizes JSON-RPC 2.0 for message formatting and supports multiple transport mechanisms, including standard input/output (stdio), HTTP with Server-Sent Events (SSE), and WebSockets.

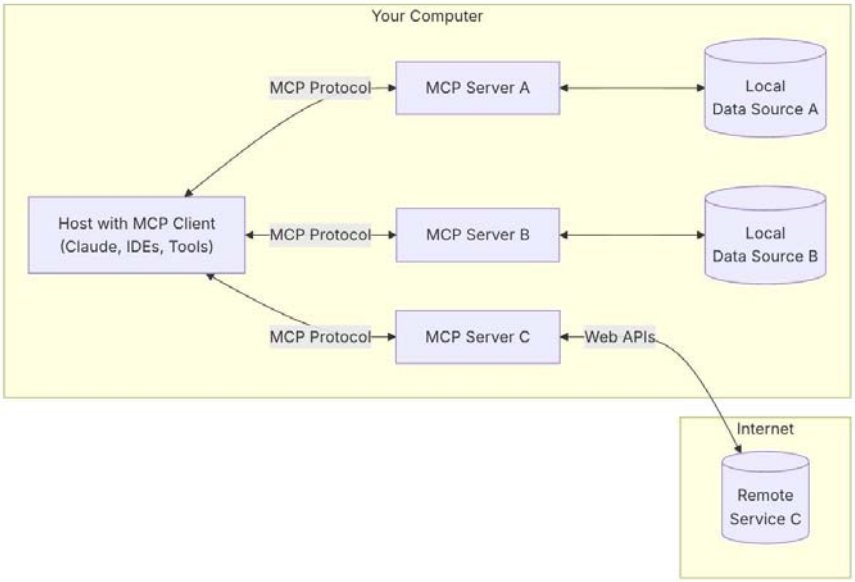


Figure 10: Indicative MCP architecture

To facilitate the integration of MCP-based agentic behaviour within O-CEI environments, we plan to leverage the existing MCP SDKs, which are available in multiple programming languages, including Python, TypeScript, Java, and C#. These SDKs provide a ready-to-use foundation for building both AI agents and their interactions with external systems. In addition, we will explore and evaluate the repository of pre-built MCP servers already available for commonly used platforms such as Google Drive, Slack, GitHub, and PostgreSQL. This will accelerate the development and testing of agentic workflows by offering immediate integration points. Early adopters like Block, Apollo, Replit, and Sourcegraph demonstrate the feasibility and maturity of the protocol in real-world deployments and provide valuable reference implementations for guiding our design and integration strategy.

The preceding sections highlight how O-CEI integrates AI to enable a distributed, intelligent, and secure CEI ecosystem. From federated learning and cognitive orchestration to autonomous AI agents and secure DevSecOps practices, these technologies form the backbone of the project's AI-driven vision. The table below provides an overview of the main AI-related frameworks, architectures, and practices incorporated in O-CEI, summarizing their roles and relevance across technical focus areas.

Table 5: AI Solutions Integration in O-CEI

Component	Type	Purpose	Key Features	Positioning	TF3 Blocks
Cognitive Cloud	Architecture	AI-enabled orchestration and dynamic management of the CEI	Context-aware decision-making, workload optimization, secure service mobility, green computing	Core AI layer for orchestrating services and infrastructure across cloud-edge continuum	AI, Orchestration, Resource management, Security and privacy
COGNIT	Framework	Cognitive orchestration with serverless, AI-powered logic	Autonomous task placement, load balancing, AI-driven orchestration decisions	Showcase of intelligent orchestration using real-time context and compliance	AI, Orchestration
SAFE-6G Cognitive Coordinator	Control Layer	AI-powered trust-aware orchestration in next-gen networks	Trust intent interpretation, dynamic trust scoring, Trust Function management	Adaptive security through AI for Zero-Trust, self-configuring networks	AI, Security and privacy, Trust and reputation
AI-enhanced IaC & DevSecOps	Pipeline / Practice	Secure, automated infrastructure provisioning and compliance using AI	SBOM integration and vulnerability detection, secure IaC scanning and LLM-driven compliance checks	DevOps evolution using AI agents for compliance and security automation	AI, Security and privacy, Resource management
Federated Learning Frameworks (e.g. Flower, NVidia Flare, FedML)	Framework	Distributed ML training across privacy-preserving edge environments	Client-server FL logic, privacy-enhanced training, multi-sector deployment	Foundation for scalable, privacy-preserving distributed AI in CEI	AI, Data management, Security and privacy
Machine Learning Workflows (e.g. LSTM, CNN-LSTM)	Workflow	Energy usage prediction and ML pipeline execution at the edge	Preprocessing, model training, edge deployment, real-time inference	Scalable edge AI with integration in smart energy systems	AI, Data management, Monitoring and observability
AI Agents & RAG Architecture	Agentic System	Intelligent, autonomous decision-making in data-rich, semantic environments	Task orchestration, document reasoning, IoT querying, semantic navigation	Human-aligned intelligent agents in data spaces & smart systems	AI, Data management, Orchestration
Model Context Protocol (MCP)	Protocol	Standardized interface for AI agents to access external tools/data	Secure protocol abstraction, dynamic service access, JSON-RPC, RAG-compatible	Secure, open interface for integrating AI agents with CEI environments	AI, Data management, Security and privacy

Swarm Intelligence and Cognitive Games	Framework / Methodology	Self-organizing distributed intelligence with optimization and adaptation	Bidding, auctioning, pricing strategies, context-aware resilience, energy-aware decision-making	Fully distributed autonomous orchestration inspired by biological and economic models	AI, Orchestration, Resource management, Trust and reputation

This overview illustrates how O-CEI integrates the diverse expertise and experience of its partners to address the multifaceted opportunities and challenges introduced by AI across the CEI landscape. Rather than starting from scratch, O-CEI builds upon proven technologies, research outcomes, and real-world deployments, providing a solid foundation from which to advance the state of the art. Leveraging advances in the field of distributed intelligence, cognitive cloud infrastructure, and responsible AI deployment, O-CEI plans to take on the task of shaping a trustworthy, scalable, and future-proof AI continuum, one that responds to the rising demands and rapid evolution of the field.

3.2 ANALYSIS OF RECENT RESEARCH PROJECTS AND OUTCOMES

O-CEI derives its inputs from recent and ongoing European research and innovation initiatives—particularly Horizon Europe CL4 projects and other large-scale efforts. Rather than providing full project summaries, this section focuses on identifying concrete technological contributions that inform O-CEI's architecture and implementation roadmap. These include research outputs in areas such as MetaOS, cognitive cloud, swarm computing, orchestration, data spaces, and AI lifecycle management—many already validated at TRL 4+ and with proven cross-sector applicability. O-CEI selectively leverages these contributions to develop reusable, context-specific solutions across its pilot implementations. While some technologies, such as marketplace infrastructure or federated trust mechanisms, may serve foundational roles, others are adapted or extended depending on the needs of each deployment. Crucially, O-CEI involves partners who participated in developing many of these technologies, enabling seamless integration, customization, and further innovation.

Table 6 below offers an at-a-glance overview of the most relevant projects, summarizing their key technological outcomes, the positioning pillars they inform, and how O-CEI plans to take them forward. This serves as a reference point for the detailed project descriptions that follow.

Table 6: Overview of relevant research projects informing O-CEI

Project	Relevant Achievements	Associated O-CEI Pillar(s)	Planned Take-Up in O-CEI
IPCEI-CIS	Delivered a modular reference architecture for a sovereign, multi-provider cloud-edge continuum. Includes layers for virtualization, orchestration, AI, data, and management, and domains for federation, sustainability, and compliance.	MetaOS & Orchestration; Data Spaces & Interoperability; Security & Trust; AI Computing	Strategic alignment and adaptation of open architectural components; strengthens foundation for convergence and European cloud-edge collaboration.
aerOS	Developed a MetaOS with federated orchestration, semantic data fabric using NGSI-LD, and secure coordination across heterogeneous domains.	Continuum Orchestration & MetaOS; Data Spaces & Interoperability; Security & Trust	Feeds into orchestrator design, cross-domain workload automation, semantic modeling, and secure federation mechanisms.
ICOS	Introduced AI-assisted orchestrator, AIOps platform (BentoML, MLFlow, dataClay), monitoring tools, and federated learning methods.	Orchestration & MetaOS; AI & Cognitive Cloud; Security & Trust	Supports O-CEI's AI lifecycle, observability and federated AI with transparency and traceability.
NEMO	MetaOS framework for orchestration and adaptive workload deployment across IoT–Edge–Cloud with focus on decentralization, resource awareness, and scalability.	Orchestration & MetaOS; AI & Cognitive Cloud	Provides reusable orchestrator and decentralization logic for continuum resource management.

Project	Relevant Achievements	Associated O-CEI Pillar(s)	Planned Take-Up in O-CEI
COGNIT	AI-driven serverless orchestration and energy-aware platform enabling dynamic, green application deployment across the continuum.	AI & Cognitive Cloud; Orchestration & MetaOS	Powers intent-based serverless orchestration and sustainability-aware deployment logic.
DOME	Developed federated cloud/edge marketplace for trusted service discovery and on-demand composition.	Data Spaces & Interoperability	Serves as foundation for the O-CEI Marketplace architecture, exposing reusable service bundles per pilot.
SAFE-6G	Defined trust-aware orchestration with a Cognitive Coordinator using AI models and verifiable credentials (VCs), decentralized identities, and Blockchain audit trails.	Security & Trust; AI & Cognitive Cloud	Inspires user-intent trust enforcement, decentralized identity, and trust-aware orchestration strategies.
ENACT	Introduced cognitive orchestration with GNNs, DRL agents, and regulation-aware runtime policies. Focus on self-adaptation and trust.	Orchestration & MetaOS; Security & Trust	Informs design of O-CEI dynamic orchestrators with embedded compliance and prediction.
RE4DY	Built federated Industry 4.0 data spaces with IDS-compliant connectors, data usage control, and cognitive integration toolkit.	Data Spaces & Interoperability; Security & Trust	Enables sovereign data exchange and federated industrial integration patterns.
ASSIST-IoT	Delivered microservice-based architecture with modular enablers for semantic interoperability, identity, monitoring, and federated learning.	Orchestration & MetaOS; Data Spaces & Interoperability; Security & Trust	Reusable for modular orchestration, semantic alignment, and trust frameworks—especially for edge/IoT pilots.
DataPorts	Created federated semantic data platform with NGSI-LD models, DLT-based usage control, and policy enforcement mechanisms.	Data Spaces & Interoperability; Security & Trust	Key enablers for secure, policy-aware data exchange; inform pilot-level interoperability and governance.
MANOLO	Delivered energy-efficient AI tools (model compression, NAS), explainability, provenance tracking, and adaptive AI orchestration.	AI & Cognitive Cloud; Security & Trust; Orchestration & MetaOS	Reusable AI lifecycle and observability tools for robust, low-energy intelligent workloads.
CoGNETs	Middleware for decentralized orchestration via cognitive agents, swarm management, and brokerless communication.	AI & Cognitive Cloud; Orchestration & MetaOS; Security & Trust	Supports swarm intelligence and cognitive orchestration patterns for resilient, self-adaptive systems.
IoT-NGIN	Delivered patterns-based IoT meta-architecture with MLaaS and vulnerability scanner for threat detection.	Security & Trust; AI & Cognitive Cloud	Provides tools for proactive threat management and machine learning integration in edge/IoT scenarios.
6G-SANDBOX	Programmable experimentation infrastructure with IaC-based deployment, observability, and reusability.	Orchestration & MetaOS; Monitoring & Observability; Resource Management	Informs automated orchestration, experimentation-as-a-service, and infrastructure validation workflows.
5G-Solutions	Demonstrated distributed orchestration and demand-side energy optimization using 5G-based agents and real-time KPIs.	AI & Cognitive Cloud; Orchestration & MetaOS; Security & Trust (Energy)	Supplies tested smart grid orchestration scenarios and intelligent control patterns for energy pilots.

The following paragraphs provide focused descriptions of each project listed above, giving brief insights on their possible technological contributions for O-CEI plans across its pillars and pilot implementations.

IPCEI-CIS

On 5 December 2023, the European Commission approved The Important Project of Common European Interest on Next Generation Cloud Infrastructure and Services (IPCEI-CIS). Based on the fact that it is the first IPCEI targeting cloud and edge domains aiming to strengthen Europe's digital and technological sovereignty and has developed *"the first interoperable and openly accessible*

European data processing ecosystem, the multi-provider cloud to edge continuum"[4], O-CEI is expected to further investigate strategic collaboration and technical alignment possibilities. Software has already been developed in common addressed technical areas such as data processing capabilities and data sharing tools that enable sustainable, federated, resilient and trustworthy distributed data processing technologies and related real-time services across cloud-edge continuum.

At the end of March 2025, IPCEI-CIS presented its first functional Reference Architecture[5] containing the most important technological components that have been developed by IPCEI-CIS partners across its four workstreams and the transversal project topics (sustainability, federation, security & compliance and management). This open architecture will be investigated for possible adaptations and adoptions within O-CEI.

Looking at Figure 11 below, we can notice 11 building blocks that compose this reference architecture:

- **Physical Layer** - includes all the physical hardware resources needed for the creation of the cloud-edge continuum (compute, storage and networking) and the physical network infrastructure that creates the communication between those resources.
- **Virtualization Layer** - provides the abstraction layer over the physical hardware resources, enabling their dynamic allocation and sharing, hiding the complexity of their heterogeneity.
- **Cloud-Edge Platform Layer** - enables the allocation and lifecycle management of resources across the virtualized Cloud-Edge Continuum, supporting the deployment and orchestration of application components to deliver services within a defined geographical area. This layer includes components designed to handle the complexity of computing environments that span multiple cloud technologies and providers, ensuring interoperability, compatibility, and portability across a diverse, multi-provider Cloud-Edge ecosystem.
- **Service Orchestration Layer** - enables unified and automated integration and management of cloud and edge services across diverse multi-provider environments. In multi-cloud and edge ecosystems, service orchestration is key to efficiently deploying and managing applications at scale across distributed infrastructures.
- **AI Layer** - its objective is to develop core advanced functionalities that promote the adoption of Artificial Intelligence (AI) across distributed data and processing resources, moving away from centralized models. These capabilities are applicable for both end-user applications and internal Edge-Cloud optimization.
- **Data Layer** - provides a new perspective over the data collection, processing and exchange simplifying it, making it easier to scale, to deploy over a distributed cloud edge continuum and having sustainability and trustworthiness embedded. Similarly with the AI layer, this layer can be used both in end-user applications (including by third parties) but also internally to optimize the cloud-edge continuum stack at different levels.
- **Application layer** - offers the tools required to design, develop and update applications and functions and to expose them for use. This layer tackles the end-user applications but also the function catalogues.
- **Sustainability domain** - Improving energy efficiency in cloud and edge infrastructures is essential and involves optimizing hardware, cooling systems, and integrating renewable energy sources. Strategic edge node placement enhances access to green energy and efficient cooling.

[4] https://competition-policy.ec.europa.eu/state-aid/ipcei/approved-ipceis/cloud_en retrieved May 2025

[5] https://eucloudedgeiot.eu/ipcei-cis-reference-architecture-announced-at-the-8ra-general-assembly/

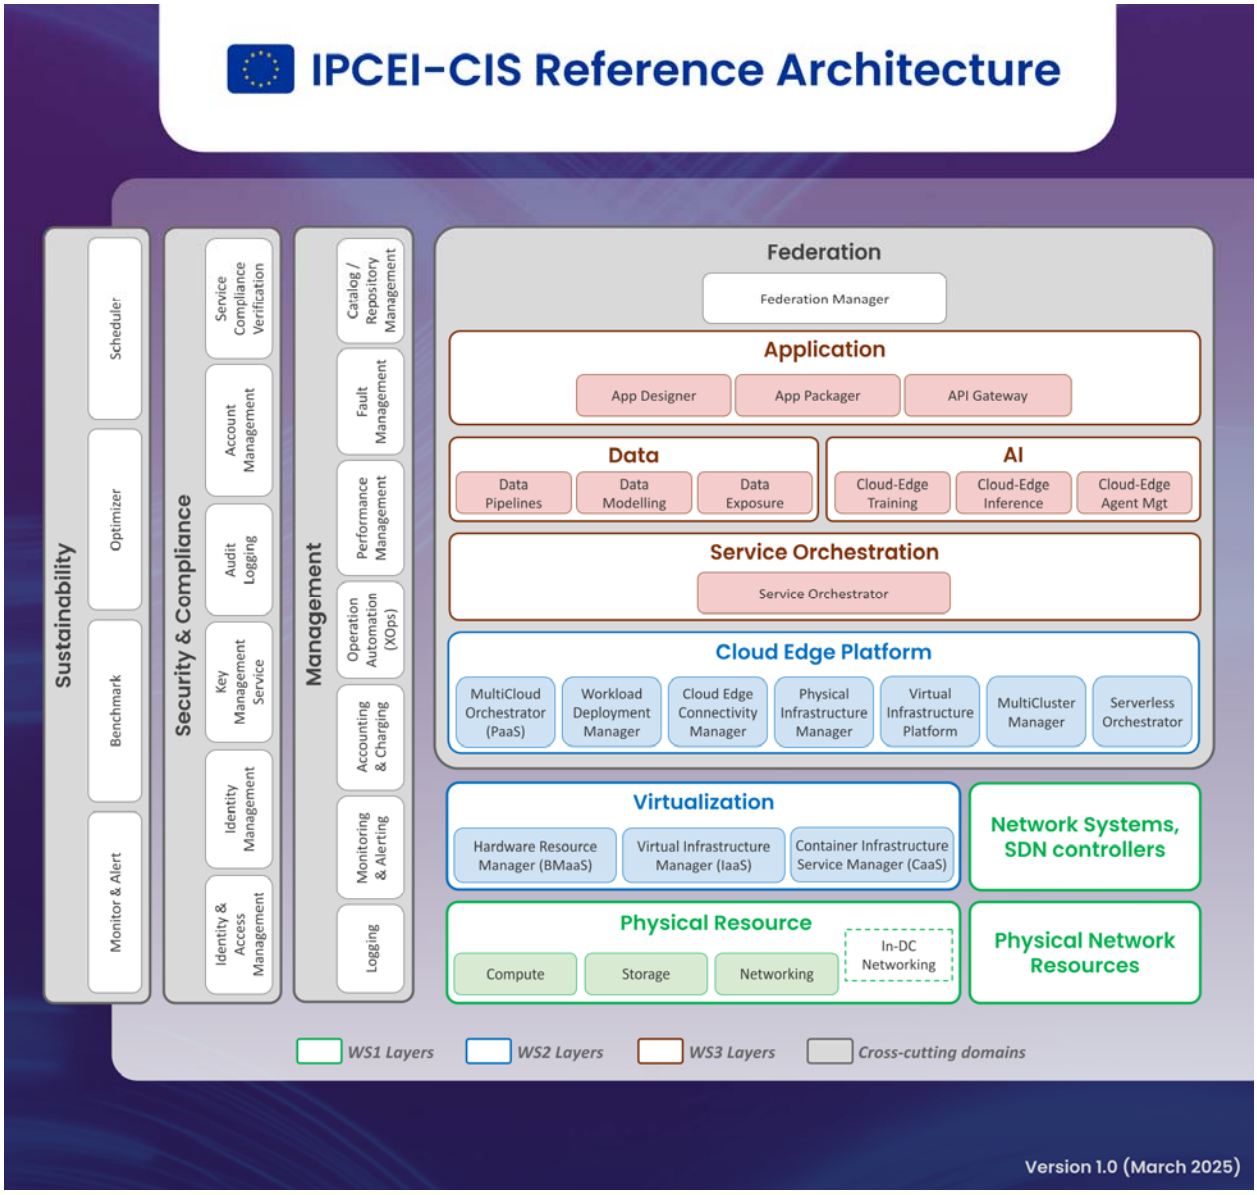


Figure 11: IPCEI-CIS Reference Architecture, version 1.0[6]

Additionally, energy-efficient technologies are being developed for virtualization, application deployment, and lifecycle management to reduce the energy consumption of distributed systems. Sustainability efforts also focus on minimizing redundant data transfers, using AI for energy optimization, and monitoring environmental goals. Efficient data formats and autonomous systems for energy and environmental management further enhance sustainability across the multi-provider cloud-edge continuum.

- **Security and Compliance domain** - ensures data integrity, controlling access, and maintaining compliance with industry standards and regulations. This domain underpins secure cloud infrastructure by protecting sensitive data and enforcing access policies. Through advanced security protocols, real-time monitoring, and compliance frameworks, it mitigates risks, enables rapid threat response, and ensures ongoing regulatory compliance - building trust and reliability in services based on cloud-edge continuum.

[6] https://eucloudedgeiot.eu/ipcei-cis-reference-architecture-announced-at-the-8ra-general-assembly/

- **Management domain -** Efficient cloud infrastructure management requires a comprehensive set of capabilities to maintain quality and performance across all layers. Core functions like logging, monitoring, and resource inventory support real-time anomaly detection, fault management, and performance optimization. Metering, accounting, and charging systems ensure transparent resource usage and accurate billing, while SLA management guarantees service reliability and customer satisfaction. Fault management enhances resilience by enabling rapid issue resolution. Automation across integration, testing, and delivery processes further boosts operational efficiency. Together, these capabilities create a scalable, reliable framework that underpins high-performing cloud-edge environments.
- **Federation domain** - contains the needed functionalities and mechanisms for inter-connecting different providers across cloud-edge continuum on one side, and customers having access to the services and resources from all the federated providers through a single interface on the other side. This concept is applicable at the level of infrastructure, platform, services, data, AI and/or applications levels.

Additionally, in close connection with IPCEI-CIS, the 8ra initiative "is driving the development of a sovereign, interoperable and secure Multi-Provider Cloud-Edge Continuum fostering cross-collaborations to create a resilient and scalable digital infrastructure tailored to Europe's needs"[7]. Given the shared goals of establishing a trusted, standardized, and collaborative cloud-edge continuum, free from vendor lock-in and applicable to a multitude of domains like Metaverse, AI, data spaces or telecommunications, a strategic opportunity, for O-CEI emerges to explore alignment emerges. The paramount goal is to take advantage of the work done at the level of IPCEI-CIS and align to create a cross-collaboration around a sovereign cloud stack.

aerOS (Horizon Europe CL4 – DATA-01-05)

aerOS proposes a MetaOS designed to support intelligent and distributed management of heterogeneous IoT, edge, and cloud resources. Its architecture focuses on abstracting and federating Infrastructure Elements (IEs) across domains. By abstracting infrastructure heterogeneity resources as IEs, aerOS exposes a computing continuum supported by a federated orchestration framework. This enables transparent discovery, orchestration, and deployment of containerized workloads over heterogeneous execution environments.

The orchestration approach in aerOS is based on a double-layer orchestration model, which separates global orchestration intelligence from localized enforcement mechanisms to enable scalable and flexible federation. Thus, there is a combination of AI-powered intent handling by the High-Level Orchestrator (HLO) with localized workload placement executed by Low-Level Orchestrators (LLOs). These orchestrators rely on continuously updated system state implemented by aerOS's continuum model (a semantic knowledge graph), which provides a coherent representation of resource capabilities and service requirements—allowing intelligent placement, abstraction, and negotiation across distributed infrastructures. This architecture contributes directly to the **Continuum Orchestration and MetaOS pillar**, particularly in addressing OS heterogeneity, federated control, and runtime policy-based service adaptation.

The Data Fabric in aerOS facilitates semantic interoperability and distributed data management through NGSI-LD-based knowledge graphs. Each domain contributes to a federated knowledge base, enabling semantic discovery and policy-compliant data exchange. These capabilities align with the **Data Spaces and Interoperability pillar** by supporting cross-domain data integration, standardized exposure, and dynamic data product composition. Additionally, the system embeds mechanisms for verifiable service execution, auditability, and secure coordination across domains, reinforcing the **Security, Trust and Digital Identity pillar.**

[7] https://www.8ra.com/

These results, validated in cross-sector pilots and reaching TRL 4+, have been further consolidated into a working MVP integrating all aerOS components, successfully demonstrated in a representative end-to-end scenario. The outcomes are highly relevant to O-CEI's efforts in orchestrator design, continuum automation, and trusted workload deployment—mapping to TF3 building blocks including orchestration, networking, data management, management, monitoring, trust and security.

ICOS (Horizon Europe 2022-2025)

The ICOS project has yielded a suite of technological components and frameworks that align closely with O-CEI's objectives for orchestrated, intelligent, and secure continuum computing. Chief among these is the ICOS orchestrator, leveraging AI-assisted matchmaking and aligning with O-CEI's goal of enabling scalable service orchestration across edge-cloud continuums, forming this also part of the diverse orchestration toolkit tested in O-CEI blueprints.

A relevant innovation from ICOS is its AIOps platform, which provides AI as a service and supports advanced monitoring and automation by integrating BentoML, MLFlow, and dataClay for model training offloading across the edge to cloud continuum while exploiting data locality [ref]. This asset enhances O-CEI's AI capabilities within its MLOps stack, offering the means to track experiments, automate deployment, and monitor AI services across the continuum. Integration of this platform strengthens O-CEI's capacity to deliver intelligent automation while ensuring transparency and traceability in AI model lifecycles. Moreover, ICOS contributes to the broader AI block in O-CEI through support for Federated Learning, which is key for enabling privacy-aware data processing across distributed environments. The ICOS methodology, applied in conjunction with projects like LASAGNE and NEMO, advances privacy-preserving mechanisms by implementing multi-tiered schemas and secure aggregation techniques. These contribute to O-CEI's vision of trust-centric data sharing and AI computation, especially where data sovereignty and compliance are critical.

In terms of observability, ICOS technologies reinforce the O-CEI monitoring and validation layer through contributions such as log analytics and performance metrics. These insights support dynamic service adaptation and improve reliability in edge-cloud deployments. ICOS's emphasis on observability aligns with O-CEI's blueprint design approach, which prioritizes operational transparency and performance assurance under varied workload and environmental conditions.

NEMO (Horizon Europe 2022-2025)

The NEMO project vision foresees three broad functional layers, which abstract metaOS components' role into functional groups, as well as three vertical (cross-cutting) functions, which apply at all metaOS levels: NEMO Service Management, NEMO Kernel modules and the underlying technology. NEMO developed a metaOS, which enables multi-cluster and multi-network orchestration of containerized workloads across the IoT, edge and cloud continuum. As a metaOS, NEMO is user-centric, facilitating users to develop and deploy on top of the NEMO framework. The NEMO metaOS framework aims to support optimal operation of hyper-distributed applications implemented as microservices in a highly distributed and diverse environment of cloud-edge and IoT technologies.

In NEMO, the meta-Orchestrator is responsible for managing computational resources across the IoT, edge and cloud continuum. The NEMO meta-Orchestrator is driven by the need for efficient decentralization in the IoT to Edge to Cloud Continuum. It simplifies the complexity of distributed computing, making it more manageable. Intelligence guides its decision-making, optimizing workflows for efficiency. Resource efficiency is a priority, especially in resource-constrained edge environments. Adaptability ensures scalability and reliability in a dynamic landscape. Interoperability is key, enabling seamless integration with diverse components and systems. Overall, the meta-Orchestrator addresses fundamental challenges and opportunities in this complex ecosystem and is introducing orchestration and resource management capabilities aligned with TF3 RA building blocks concepts.

COGNIT (Horizon Europe 2022-2025)

Entitled *"A Cognitive Serverless Framework for Cloud-Edge Continuum"* the COGNIT Project (2023-2025) has in its sights the development of a new distributed Function-as-a-Service (FaaS) paradigm for the smart management and orchestration of edge applications, helping to improve the way in which the applications and services are deployed and used across the cloud-edge continuum.

Answering several cross-cutting requirements related to sovereignty, sustainability, interoperability and security, the COGNIT platform has contributed to the EU digital policies and strategic priorities related to the Cognitive Cloud. The two important policies targeted are European Green Deal and European for the digital age, the project having had an impact on "green cloud" concept dealing with aspects such as optimization of the energy consumption, use of renewable resources and AI-enabled orchestration of resources across cloud-edge continuum. It also has a potential to enable smaller industry players to provide more efficient green services across cloud-edge continuum.

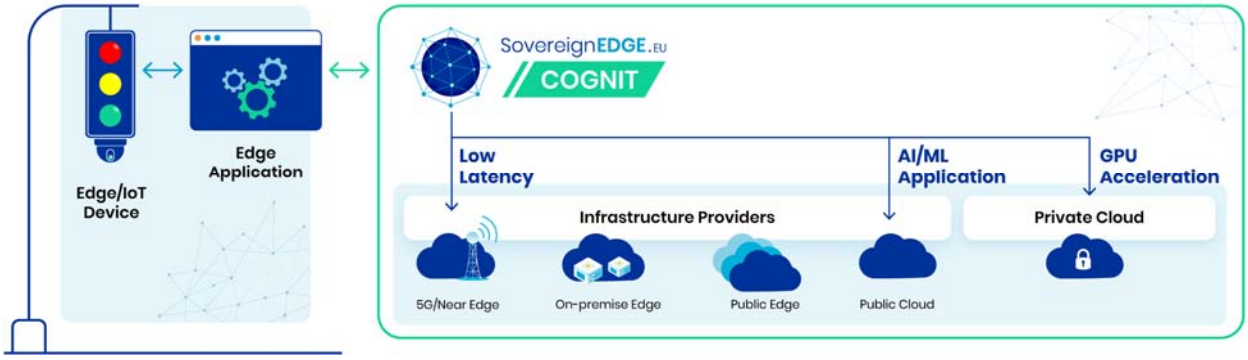


Figure 12: Bridging the devices to the cloud-edge continuum through edge applications[8]

The main objectives of the COGNIT Project are:

- To optimize data processing based on the updates in application demands and behaviour and energy efficiency observations.
- To adopt a new innovative serverless approach for the application management across cloud-edge continuum.
- To support the secure, on-demand deployment of large-scale, highly distributed, and self-adaptive serverless environments by leveraging existing data processing resources from cloud and edge infrastructure providers—including local data centers, cloud platforms, and 5G/telecom operators.

Creating this kind of cognitive cloud-edge continuum platform has had at its core AI-enabled, automation and portability algorithms designed to intelligently manage the increasing complexity and data volumes of distributed computing environments. It enables seamless, secure integration and automatic adaptation to changes in application behaviour and data variability, supporting dynamic service deployment and mobility. By leveraging AI for real-time monitoring and resource management, it ensures efficient load balancing, optimized energy use, and reliable data throughput. The framework also adapts processing capacity based on green energy availability to reduce environmental impact. O-CEI will leverage these capabilities to offers developers enhanced control over infrastructure, while end-users benefit from uninterrupted, adaptive service delivery across the cloud-edge continuum.

[8] https://cognit.sovereignedge.eu/

DOME (Distributed Open Marketplace for Europe Cloud and Edge Services)

DOME is a project initiated under the Digital Europe Programme to address the fragmented landscape of Cloud and Edge services in Europe. Its objective is to deploy and operate a trusted online marketplace for cloud and edge services for EU users from both the private and public sectors, serving as a single point of access. DOME is designed to boost cloud uptake in Europe and facilitate trusted interactions between customers and service providers and brings to O-CEI core marketplace concepts and architecture.

SAFE-6G (Horizon Europe, 2023–2026)

SAFE-6G defines a user-centric trustworthy framework to enhance the user trust of vertical applications in distributed CEI domains of the telco ecosystem. It targets five trust dimensions (namely Safety, Security, Privacy, Resilience, and Reliability) managed by a Cognitive Coordinator that translates user intents into adaptive orchestration decisions. Particularly, the outcomes of these functions, supported by dedicated AI models, can modify configuration parameters of external systems such as the network core, the computing continuum and the vertical services themselves. Trust enforcement relies on verifiable credentials, decentralized identity, and Blockchain-based audit trails, therefore supporting primarily the **Security, Trust and Digital Identity pillar**.

ENACT (Horizon Europe, 2023–2026)

ENACT develops a cognitive cloud-to-edge continuum framework focused on self-adaptation, trust, and regulation-aware orchestration of distributed applications. Its architecture introduces a dynamic graph-based model of resources, a cognitive orchestrator that combines AI-enabled forecasting mechanisms together with a DRL agent and GNN models, and an adaptive Application Programming Model (APM) for runtime optimization. Key components include a Security Risk Modeller and GDPR/AI Act compliance checkers for regulatory assurance. The system supports autonomous orchestration, predictive deployment, and resource-aware adaptation. These contributions support the **Continuum Orchestration and MetaOS pillar** and the **Security, Trust and Digital Identity pillar** through AI-based orchestration and runtime policy enforcement.

RE4DY (Horizon Europe, 2022–2025)

RE4DY aims to establish a federated data space for Industry 4.0 ecosystems, structured around the concept of Data-as-a-Product. The project delivers an integrated suite of frameworks, tools, and protocols that support trusted, sovereign, and interoperable data exchange among stakeholders in the manufacturing domain. Its core components include certified IDS-compliant connectors, distributed ledger mechanisms for auditability, data usage control mechanisms, and a Digital Continuum 4.0 Open Toolkit that enables integration of edge nodes and cognitive services into federated manufacturing environments. These elements are directly aligned with **Data Spaces and Interoperability pillar** and contribute to multiple building blocks, including Data Management, Marketplace, and Trust and Identity.

ASSIST-IoT (H2020, 2020–2023)

ASSIST-IoT developed a modular, enabler-based reference architecture to enhance interoperability, autonomy, and trust across IoT–Edge–Cloud systems. It introduced Core and Transversal Enablers as microservices covering self-configuration, federated learning, semantic interoperability and identity management. The architecture supports containerized deployment with reusable descriptors, APIs, and monitoring tools. These outcomes support the **Data Spaces and Interoperability pillar** (semantic alignment) and the **Security, Trust and Digital Identity pillar** (decentralized identity, auditability). Its orchestration and modular composition patterns align with the **Continuum Orchestration and MetaOS pillar**.

DataPorts (H2020, 2020–2023)

DataPorts developed a modular data platform facilitating secure and interoperable data exchange among stakeholders in port ecosystems. It integrates semantic data modelling, data harmonization, real-time exchange, context-based data brokering, data abstraction and virtualization, and policy enforcement mechanisms utilizing distributed ledger technologies. In addition, its data models were published as standardized schemas (specifically NGSI-LD schemas) to enable cross-vendor interoperability. The platform enabled organizations to share operational data through standardized interfaces while maintaining control over access policies and ensuring traceability. The platform's core mechanisms, including federated data access, ontology-driven interoperability, and smart-contract enforcement of usage rules, are directly applicable to **Data Spaces and Interoperability pillar** (semantic interoperability, standardized APIs, IDS connectors), particularly in supporting cross-domain data federation and policy-aware data sharing across cloud-edge infrastructures. DataPorts shows how a federated data broker combined with IDS principles can enable trusted, cross-domain data exchange for a federated environment. Its Blockchain governance and secure data exchange also support the **Security pillar** (consent management and auditability) of O-CEI.

MANOLO

The MANOLO EU project delivers foundational advances relevant to O-CEI's objectives in energy-efficient AI deployment, trust, and dynamic orchestration across the cloud-edge continuum. Key contributions include hardware-aware model compression, pruning, and neural architecture search, directly supporting O-CEI's aim to deploy lightweight, high-performance AI across heterogeneous infrastructures.

With built-in provenance tracking, quality enhancement, and data distillation, MANOLO's suite aligns with O-CEI's vision for interoperable and trustworthy data flows across federated environments. These capabilities reinforce traceability and robustness in AI pipelines, supporting regulatory compliance (e.g., GDPR, AI Act) across the continuum.

Its benchmarking and explainability tools enhance O-CEI's observability and validation layers, enabling systematic evaluation of model robustness, performance, and trustworthiness under varied hardware and workload conditions. Additionally, MANOLO's work on adaptive resource allocation as well as federated learning in the edge-cloud stack contributes directly to O-CEI's goals for energy-aware, policy-driven orchestration strategies.

CoGNETs

CoGNETs redefines Intelligent Infrastructure Management through a next-generation Middleware Framework designed for scalable, autonomous IoT-to-Cloud integration. By harnessing the collective strength of a vibrant EU-Japan industrial-academic alliance, the project enables devices to seamlessly self-organize into agile swarm continuums—unlocking new levels of efficiency in data processing and service orchestration. At its core, CoGNETs introduces a pioneering fusion of decentralized broker architectures, cognitively adaptive agents inspired by game intelligence, and robust end-to-end security mechanisms. This unique approach significantly improves cognitive resource allocation across diverse, heterogeneous environments. The framework being demonstrated in real-world settings across key sectors—manufacturing, mobility, and healthcare—brings into O-CEI a paradigm of enhanced operational resilience and reinforced digital sovereignty towards industrial competitiveness.

IoT-NGIN (H2020)

IoT-NGIN project started in October 2020 and concluded in September 2023. The project defined a patterns-based meta-architecture, with the aim to encompass different types of IoT architectures and

create an ecosystem that can drive forward and expand their capabilities. IoT-NGIN's platform and tools were tested in multiple use cases, like Smart farming, Smart Energy Grid monitoring etc. The results have yielded progress in multiple aspects of IoT infrastructures, such as networking, MLaaS, data sovereignty etc. Among its contributions, one of the most relevant to the O-CEI project is the developed mechanisms for the identification of vulnerabilities in IoT devices, specifically the IoT Vulnerability Crawler.

6G-SANDBOX SNS HE

The **6G-SANDBOX** project is positioned to become a cornerstone of the emerging European 6G experimentation ecosystem, aligned with the strategic directions set by the SNS JU. It aims to support the research and technology validation processes necessary for the evolution towards 6G over the next decade. Central to this vision is the development of a flexible, modular, and automated experimentation facility that will enable the rapid deployment and evolution of trial networks, incorporating both virtual and physical resources. 6G-SANDBOX contributes a highly relevant foundation to O-CEI's work by advancing methodologies for automated, programmable infrastructure management and experimentation. The project introduces a modular, API-driven facility that supports the dynamic deployment of trial networks using IaC principles. Tools such as Terraform, Ansible, Jenkins, and OpenNebula are employed to manage virtual and physical infrastructure seamlessly across domains, enabling experimentation-as-a-service. These practices align closely with O-CEI's objectives to realize automated, scalable orchestration pipelines, validated through practical deployment workflows. Moreover, 6G-Sandbox emphasized on KPI/KVI analytics and reusable experimentation outputs providing, thus, valuable insights into the design of observability and monitoring mechanisms for CEI platforms. Through its focus on multi-domain resource control, experimentation workflows, and reusability, 6G-SANDBOX directly informs the technical positioning of O-CEI, especially in the areas of orchestration, resource management, and dynamic infrastructure adaptation.

5G-Solutions (H2020)

"5G-Solutions for European Citizens" (website) started in July 2019 with the aim of validating more than 140 KPIs for 20 innovative and heterogeneous use cases in 5 reference Living Labs. One of them was the Smart Energy Living Lab. Among the use cases with the highest commercialization potential, the use cases developed and validated in the Smart Energy Living Lab located in Italy were:

- Industrial Demand Side Management (DSM) with a specific focus on the optimal scheduling of energy loads during normal power plant operation, as well as the computation and actuation of flexibilities offered on the dispatching market and the control actions needed to keep the peak power consumption limited.
- Electrical Vehicle (EV) Smart Charging dealing with the problem of computing real-time the charging power setpoints for active Electric Vehicle charging sessions taking place in one load area, in order to enable the control of the aggregated load and, at the same time, to ensure that the user preferences on final state-of-charge (SoC) and on dwelling time are matched.
- Electricity network frequency stability in a smart charging scenario in presence of power modulation controlled as in the previous use case.

In the first two use cases the Italian partner Ares2t proposed and developed a decentralized scheme to exploit the advantages coming from the 5G technology to distribute the computational effort among communicating session agents and to suitably control power flow for limiting peak power, grid stability and energy efficiency.

European research projects considered, in the O-CEI, for energy flexibility applications:

In addition to the technology-driven projects informing O-CEI architectural foundations, a number of domain-focused initiatives have been studied to guide pilot implementation—particularly in energy flexibility and smart grid services. Power flow control and optimization has been investigated from the methodological, functional and operation point of view in many research and industrial projects. In the following paragraphs research and industrial projects are briefly summarized in the aspects of interest in O-CEI.

ENEL group smart charging (reference partner **Ares2t**) – The co-innovation project developed by the Italian SME Ares2t and the Enel Group started in 2019. The project aimed at integrating charging services offered by the Ares2t's Smart Charging platform, namely Charge Advisor, with the Enel Group's EMSP. The PaaS Charge Advisor enables smart grid players and electric vehicle drivers to access smart services for electric vehicle navigation, charging management and charging control.

Environmental Valorization of Virtual Energy storage (EVVE) launched in 2021 under the leadership of the EDF Group with the aim to deploy 800 vehicle-to-grid (V2G) charging stations in Europe and to demonstrate that by storing low-carbon electricity, electric vehicles can help to reduce CO2 emissions. Based on the "Combo CCS1" European charging standard and supported by the European Union's Innovation Fund, the project develops bidirectional vehicle-to-grid (V2G) technology and services, considering both operational (electricity balance, grid stability, peak shaving) and business goals (competitive price per kWh). Important industrial partners joined the project since the beginning, ranging from EV manufactures to financial operators, with the common goal of reaching a total deployment of 800 bidirectional charging points and offering the first bidirectional charging demonstrator.

Across Continents Electric Vehicle Services (Aces) project started in April 2017 and concluded on March 2020. Technical and economic possible benefits and impacts by large scale EVs integration have been studied on the basis of real charging patterns, grid data and field test to guarantee across continents replicability. A hierarchical approach has been proposed where large scale is basically possible by replicability, while the small scale relied on 50 EVs and 50 V2G DC charging stations.

Intelligent integration of EVs into the power grid for the provision of system services (INEES) project started in 2012 and concluded on 2015. The project explored the real-world technical feasibility of V2G for grid stability and peak shaving. This idea was tested in a fixed one-year fleet trial involving 20 V2G DC charging stations and commercial vehicles. A mobile phone app was developed as a user interface where participants could see how their personal driving behaviour and the requirements of the electricity market were related each to the other.

The above-mentioned research projects and outcomes aimed at developing the ICT infrastructure to allow secure operation and reliable communication of data to guarantee bidirectional flow in the EV charging infrastructure. None of them adopted:

1. Energy optimization methods to exploit bidirectional power flows via closed loop control strategy to guarantee dynamic cost saving and grid stability in the virtual power plant in presence of RES (e.g., photovoltaic plant);
2. Advanced data driven services for energy resource prediction to estimate power availability and consumption in the load area, due to the lack of suitable amount of data and ICT platform enabled them to collect and analyze huge amounts of data.

On top of previous experience and consolidated cloud-based solution, the O-CEI project aims at extending the communication layer to a CEI framework to allow any EV charger included in the 5G network to take autonomously time and context dependent decision on the power setpoint (positive

or negative according to charging or discharging the EV). Distributed and capillary, local micro services support the necessary optimization framework to control the charging operation in the whole load area as well as data management for EV charging pattern recognition and profiling, and power availability/consumption forecasting.

"digitaL frAmework for SmArt Grid and reNewable Energies" (LASAGNE) involved four stakeholders: System Integrator, Independent Software Vendors, Edge Equipment Vendors and Need-Owners. It developed collaborative learning algorithms for local and global consumption/production at medium/low voltage levels. These algorithms rely on a coordination model based on intelligent digital twins, incorporating social attributes by design. A marketplace platform provides a federating backbone where the four above-mentioned stakeholders can easily interact, deploy their technical solutions, implement their legal agreements and financial flows.

More reference projects

Below is a list of additional projects involving consortium partners. While these do not provide direct technical contributions to O-CEI, they may still offer valuable insights or inspiration that can be taken into account during design and implementation.

Project acronym	Access Link
MAESHA	https://maesha.eu/the-project/
ENERGICA	https://energica-h2020.eu/fr/
Int:NET	https://intnet.eu/
HedgeIoT	https://hedgeiot.eu/
RESONANCE	https://www.resonance-project.eu/
ECLIPSE DIGITAL	https://eclipse-digital.eu/
SPADE	https://spade-horizon.eu/
LICORICE	https://www.licorice-horizon.eu/
SUNRISE	https://sunrise-europe.eu/

3.3 MAPPING UTILITIES TO TF3 RA BUILDING BLOCKS

This section presents an initial structured mapping of tools, components, and research outcomes contributed by O-CEI partners to the **eight building blocks of the EUCloudEdgeIoT (TF3) Reference Architecture**. These contributions represent the technical assets developed, owned, or leveraged by the consortium, that will serve as the foundational elements for implementing and integrating O-CEI solutions across the continuum.

The objective here is twofold: First, to provide a **functional inventory** of the technologies available to the project, categorized according to the TF3 building blocks (e.g., Orchestration, Monitoring, Security, AI, etc.), offering a practical overview of the capabilities they bring to CEI applications. Second, this mapping lays the groundwork for the design **O-CEI blueprints** and the implementation strategies that will follow in WP3. It helps establish a shared understanding between technology providers, blueprint designers, and pilot leads regarding what tools and building blocks are at their disposal.

To provide an immediate overview of the current landscape, Table 7 below summarizes the tools and components described in the subsections that follow. Each row identifies a key technological asset and indicates its capabilities according to their classification under the relevant TF3 building blocks. While certain blocks may represent the primary function of a given component, additional functionalities often extend into other TF3 domains as well, reflecting the multifaceted nature of tools used across the CEI continuum. This overview provides a quick orientation within the technological landscape but also serves as a practical reference for identifying potentially reusable components which could address the functional needs of use cases and pilot scenarios and thus support the blueprint compilation process.

Table 7: Component capabilities connected to TF3 building blocks

Component	Orchestration	Network	Monitoring & Observability	Data mgmt.	Resource mgmt.	Trust & Reputation	Security & Privacy	AI
Orchestration Frameworks								
OpenNebula and OpenNebula AIOps	✓				✓			✓
aerOS	✓	✓			✓			
ICOS	✓	✓			✓			
NEMO	✓	✓			✓			
COGNIT	✓	✓			✓			✓
FogFlow	✓	✓			✓			
Continuum Models								
FogFlow DAG		✓	✓	✓				
aerOS Continuum Model			✓	✓	✓			
Declarative Provisioning								
Terraform,	✓				✓			
Ansible,	✓				✓			
K8s Cluster API	✓				✓			
Container Tools								
Kubernetes	✓	✓			✓			
Docker	✓				✓			
CNIs		✓						
OSM	✓	✓			✓			
Edge Node Tools								
Nuvla Edge Orchestrator	✓		✓		✓			
Operating System								
Phoenix-RTOS	✓				✓			
Telecom Enablers								
Open5gs		✓						
NEF		✓	✓				✓	
OpenCNC		✓						
CAPIF/Open Interfaces								
openCAPIF							✓	
Self-Sovereign Identity								
OASIS SSI Framework (DID, VC)						✓	✓	
Assist-IoT SSI Wallet						✓	✓	

Component	Orchestration	Network	Monitoring & Observability	Data mgmt.	Resource mgmt.	Trust & Reputation	Security & Privacy	AI
Blockchain Tools								
Blockchain Data Brokerage Engine				✓		✓	✓	
Hyperledger				✓		✓	✓	
Ethereum				✓		✓	✓	
IOTA Tangle				✓		✓	✓	
Access Control								
FIWARE IAAM						✓	✓	
NEMO Access Control							✓	
Keycloak							✓	
NEMO FREDY						✓	✓	✓
DLT-based Enabler								
Logging & Auditing enabler			✓			✓	✓	
Data Integrity verification enabler						✓		
Distributed broker enabler				✓		✓		
DLT-based FL enabler						✓	✓	✓
IaC Security								
IaC Scan Runner							✓	✓
Vulnerability Detection								
IoT Vulnerability Crawler							✓	
ZAP							✓	
Burp Scanner							✓	
Intrusion Detection								
AI-based IDS							✓	✓
Anomaly Detection							✓	✓
Runtime Visibility								
LOMOS			✓				✓	✓
Thanos			✓					
Prometheus			✓					
ELK Stack			✓				✓	
Grafana and Loki			✓					
Self-* Frameworks								
aerOS self-* suite			✓					
Marketplace								
DOME				✓		✓		
FIWARE BAE				✓		✓		
Data Interoperability								
EDC				✓		✓	✓	
FIWARE IDS Connector				✓		✓	✓	
Orion-LD (MIM1)				✓				

Component	Orchestration	Network	Monitoring & Observability	Data mgmt.	Resource mgmt.	Trust & Reputation	Security & Privacy	AI
Scorpio (MIM1)				✓				
aerOS Semantic stack incl. Semantic Annotator & Translator				✓				
(FIWARE) Smart Models				✓				
Data Infrastructure								
Zenoh	✓			✓				
Skyflok				✓				
dataClay				✓				
Camel	✓			✓				
Chimera				✓				
GraphDB				✓				
AI Driven Computing								
AI Agents							✓	✓
MLOps • MLFLow • ZenML • ClearML • KubeFLow								✓
ML Models Server • BentoML								✓
FL Frameworks • Flower • NEMO FREDY							✓	✓
ML SDKs • TensorFlow • PyTorch • NvidiaFlare • PySyft								✓
AI Algorithms • LSTM • CNN-LSTM • LLMs								✓
MCP				✓			✓	✓
Cognitive AI for Intent-Driven Orchestration								
Intent Based Operations • COGNIT cognitive Orchestrator • SAFE-6G Cognitive Coordinator	✓					✓	✓	✓

This mapping is not exhaustive or static; it reflects the current state and will evolve as the project progresses. The following paragraphs provide more detailed descriptions of these components, highlighting their core functionalities, maturity, and relevance to the O-CEI architectural framework. This elaboration supports their identification as candidate solutions for addressing specific technical needs across implementation use cases and pilot scenarios.

3.3.1 Orchestration

OpenNebula

OpenNebula is a robust and user-friendly open-source platform for building and managing enterprise clouds and virtualized data centers. It integrates established virtualization technologies with advanced capabilities such as multi-tenancy, automated provisioning, and elasticity across private, hybrid, and edge environments. By centralizing the management of IT infrastructure and applications, OpenNebula helps avoid vendor lock-in while simplifying operations, reducing resource usage, and lowering operational costs. A usual OpenNebula Cloud Architecture has two main components:

- The Cloud Management Cluster (including the Front-Ends).
- The Cloud Infrastructure that includes multiple workload Clusters with their hypervisor and storage nodes.

OpenNebula is a modular system that can be adapted to a multitude of cloud architectures across various types of cloud; it is a very flexible solution that can be easily adapted to any type of infrastructure but also can be extended by integrating additional components. Several key components of OpenNebula can be observed in Figure 13 below. For the complex task of orchestration, OpenNebula has developed the next modules:

- OpenNebula Core (oned) - which is the orchestration core, ensuring the management of all cloud resources and handling their lifecycle. It is also the component that connects to the underlying layers via drivers but also, through an XML-RPC API endpoint, to the other components facilitating higher level interactions.
- Scheduler - it is the module responsible for the orchestration of the allocation of VMs using a complex matchmaking algorithm but which, through a set of predefined policies, provides a high degree of flexibility for the VM placement.
- OneFlow - which provides a unified way for managing the lifecycle of the multi-VM applications called services.

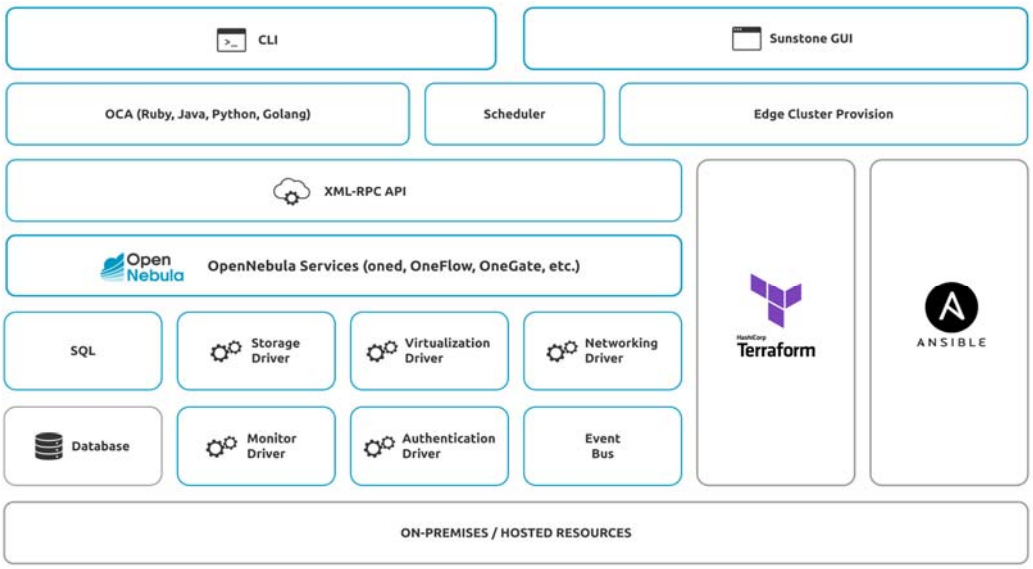


Figure 13: OpenNebula components[9]

[9] https://docs.opennebula.io/6.10/overview/opennebula_concepts/opennebula_overview.html#id1

To operate a distributed cloud-edge infrastructure, there is a need for automation of the deployment and each segment's processes. Once the cloud-edge resources are provisioned, OpenNebula, through its orchestration module, can manage them together with the virtual cloud-edge applications. This module will be responsible for the orchestration of the physical and virtual resources and for their dynamic scaling in accordance with different policies related to the latency, performance, costs or energy efficiency.

aerOS orchestrator

The **aerOS orchestrator** is a composable and interoperable component designed to support dynamic and intelligent orchestration across the CEI continuum. It provides core functionality for enabling distributed workload deployment and lifecycle management, driven by real-time information about the availability and capabilities of computing resources. Leveraging a standardized continuum model and operating within a federated environment, it enables coordinated decisions for resource allocation across heterogeneous and distributed infrastructures.

Its architecture is designed to decouple orchestration logic into distinct layers: one responsible for high-level decision-making based on user-defined requirements and runtime resource conditions, and another focused on enforcing those decisions at infrastructure level. This structural separation allows the orchestrator to adapt easily to new deployment scenarios and integrate with a wide range of platforms and tools already present in CEI ecosystems.

By supporting a model-driven approach to workload orchestration, the aerOS orchestrator can be embedded within broader CEI blueprints—particularly in modules concerned with automation, infrastructure management, and intelligent workload placement. It can be used to support O-CEI ecosystem by delivering orchestration capabilities that respond to application-level demands, execution topology, performance requirements, and other operational objectives. Its flexibility allows for seamless integration into existing automation pipelines or orchestration stacks, offering a means to bridge cloud and edge environments with consistent and responsive control. This makes it well-suited to support use cases that demand workload distribution across the continuum—such as industrial monitoring, smart cities, real-time analytics, or AI deployment at the edge—while abstracting away infrastructure complexity.

The aerOS orchestrator promotes extensibility, enabling enhancements such as AI-based decision logic or custom execution policies to be easily introduced and aligned with evolving needs of CEI applications. Its modular nature ensures that it can grow alongside emerging orchestration strategies, becoming a central enabler for resilient, scalable, and autonomous CEI operations.

ICOS orchestrator

ICOS orchestrator is part of the Meta-Kernel layer, a core component of the whole ICOS framework. Built on top of other low-level orchestrators, such as Kubernetes, it is in charge of managing the service lifecycle across the continuum. It has two main functionalities: i) orchestrating workloads across the continuum, so applications can be easily deployed in IoT devices and/or edge/cloud infrastructures; and ii) take informed decisions for developing the most suitable deployment and redeployment schemes taking into account initial requirements to obtain the maximum levels of QoS and QoE.

NEMO orchestrator

The NEMO meta-Orchestrator (MO) is an open-source software system designed to manage and optimize the distribution of computing workloads across the NEMO cluster network. This core component serves as the NEMO platform's central component, allowing efficient and smooth coordination between multiple NEMO components to allocate intelligently logical resources.

The MO is a high-level controller over container orchestration clusters such as Kubernetes, coordinating resources from the IoT, edge to the cloud continuum, ensuring workloads are deployed without issues. The NEMO MO comprises key subcomponents and services such as MO API, Deployment Controller (DC) and the IBMC (Intent-Based Migration Controller), which effectively orchestrate computing workflows.

The NEMO MO manages workload deployment, migration and scaling, in case there are relevant recommendations from the AI models, while considering monitored performance as per defined intents.

Cognit orchestrator

COGNIT orchestration framework builds around a distributed FaaS paradigm, designed specifically for edge applications. It leverages AI-powered automation and adaptive orchestration to manage dynamic workloads across heterogeneous cloud-edge infrastructures. By enabling real-time decisions for task placement, resource allocation, and load balancing, the orchestrator supports efficient service deployment and operational resilience. Embracing the FaaS model, it facilitates the offloading of compute-intensive tasks from constrained IoT devices, ensuring optimal distribution across the continuum while incorporating security and data sovereignty mechanisms.

FogFlow

FogFlow provides a high-level orchestration programming model for the CEI continuum, enabling the dynamic composition and placement of services through dataflow logic. Applications are described as **service topologies** using a DAG-based model, where each node represents a data processor and each edge represents communication. These abstract workflows are defined in a no-code environment and mapped into a concrete deployment plan across cloud, edge, and IoT resources. The orchestration pipeline follows a two-step mapping: from high-level topology to a process tree, and from process tree to deployment plans optimized based on **QoS and SLA constraints**. During execution, **FogFlow dynamically adapts deployments** to changes in resource availability or operational context (e.g., newly available sensors or shifting monitoring scopes), supporting context-aware and resilient orchestration.

As part of the O-CEI architecture, FogFlow contributes a flexible orchestration layer that bridges user intent with runtime infrastructure behavior, complementing other MetaOS capabilities and deployment frameworks.

Nuvla

Nuvla is a production-ready (TRL9) orchestration and automation platform designed to manage distributed applications and services across the CEI continuum. It provides unified, secure, and flexible application lifecycle management regardless of the underlying hardware or cloud environment. Through its lightweight edge agent (NuvlaEdge), the platform supports the onboarding and management of diverse edge nodes, even in semi-connected or constrained settings. Applications are deployed as standardized containerized workloads (e.g. Docker or Kubernetes), and Nuvla facilitates real-time application supervision, updates, and scaling actions through programmable interfaces, enabling integration with DevOps workflows and CI/CD pipelines.

By combining orchestration with capabilities such as device management, data flow control, and lifecycle governance, Nuvla can bring to O-CEI capabilities to automate and scale services in a modular, resilient, and context-aware manner. It serves as a robust control plane for managing CEI workloads with full visibility, configurability, and adaptability from cloud to edge.

Istio

Istio is a powerful, feature-rich service mesh that provides advanced traffic management, security, and observability for microservices running on Kubernetes. It operates by injecting sidecar proxies (Envoy) alongside application pods to handle service-to-service communication. With capabilities like traffic routing, fault injection, mTLS encryption, and detailed telemetry, Istio helps developers manage complex microservices environments with consistency and security. However, it can be complex to set up and manage due to its broad feature set.

Cilium

Cilium is a modern service mesh and networking solution built on eBPF (extended Berkeley Packet Filter) technology, enabling high-performance, kernel-level processing for cloud-native applications. It primarily serves as a CNI (Container Network Interface) providing advanced networking, load balancing, and network security in Kubernetes clusters. In addition to its CNI role, Cilium can also function as a service mesh, offering features like identity-aware security policies, observability, and traffic management—without the overhead of sidecar proxies. This eBPF-based architecture delivers low latency and efficient resource usage, making it ideal for performance-sensitive environments.

Open Service Mesh (OSM)

Open Service Mesh (OSM) is a lightweight and extensible service mesh developed by Microsoft and built to run natively on Kubernetes. It implements the Service Mesh Interface (SMI) standard, providing a simple and standardized way to manage service-to-service communication, traffic policies, and observability. OSM focuses on ease of use and integrates well with existing Kubernetes orchestration tools, making it a good choice for teams seeking a straightforward service mesh solution without the complexity of larger meshes like Istio. Its design prioritizes quick setup, modularity, and minimal operational overhead.

Ansible

Ansible is an agentless automation tool that supports end-to-end configuration and service orchestration. It maps most strongly to the Orchestration block, where it handles the configuration, deployment, and coordination of services across diverse environments, from cloud to constrained edge but additionally Ansible also supports aspects of Resource Management through post-provisioning setup (e.g., installing software on virtual machines).

3.3.2 Network

OpenCNC for TSN

OpenCNC provides a control plane for **Times-Sensitive Networks**. OpenCNC is designed as a system of microservices: Main Service, TSN Service, Config-Service, and Monitor Service, deployed as Docker containers and orchestrated by Kubernetes. Using an API, the topology is uploaded as a JSON file which is then saved in the corresponding key-value etcd store. The topology file is defined by nodes (endnodes and bridges), and links. Each node is defined by its type, name, model, and ports. The model is to be uploaded via an API as a zip file containing all yang files for that specific model. This makes it possible to have switches from several suppliers, making the tool vendor independent.

- **Main service** initializes OpenCNC by creating the key-value stores and managing the initial network configuration. It manages communication with the user through user-network interface (UNI) to receive requests and send responses.
- **TSN Service** is responsible for engineering the forwarding plane and managing the networking resources through establishing the network configuration. The main functionalities of this service

are scheduling, QoS handling, Stream priority management, forwarding process management, and coordinating the operations between the network optimizer, configuration and reconfiguration entities.

- **Config-Service** compiles the YANG data models, applies configuration to the devices, and checks the configuration against the data models to find any irregularities.
- **Monitor Service** ensures that the forwarding plane setup is operating as intended. QoS guarantee is one of the primary features. The service compiles data and calculates metrics.

Open5gs

Open5GS is an advanced open-source project that provides a comprehensive solution for building and managing mobile networks, supporting both **5G (NR)** and **LTE (evolved)** deployments. It offers a fully featured and modular architecture, based on 3GPP 5G specifications and implementing all core network functions such as AMF, SMF, UPF, and PCF. This flexibility allows operators to customize and scale their network deployments efficiently, supporting a wide range of use cases from small-scale private networks to large public network infrastructures. Open5GS also integrates seamlessly with various RAN technologies, making it a versatile choice for modern mobile network environments. It has been extensively used, from several O-CEI partners, within various research projects.

NEF

NCSRD's 3GPP Network Exposure Function (**NEF**) implementation on top of the Open5GS framework incorporates two core functionalities:

- **Retrieval of device location information** in accordance with **3GPP Monitoring Event (**v1.2.2, Release 17) specification. This functionality enables advanced monitoring capabilities for device location by providing a set of well-defined APIs for subscriptions, notifications, and event-driven interactions. It allows for the efficient tracking of device movements and location updates, enabling real-time notifications for subscribed applications.
- **Device session configuration with appropriate Quality of Service (QoS) priorities**, based on the **3GPP AS-Session-with-QoS (**v1.2.3, Release 17) specification. This ensures optimized network resource allocation by assigning the correct QoS to devices based on service requirements and application needs, improving both user experience and network performance.

By integrating these key specifications, the NEF solution on Open5GS enables secure and efficient interaction with 5G core network services, allowing third-party applications and network operators to leverage critical network functionality.

OpenCAPIF

OpenCAPIF is a general API invocation framework based on the **3GPP Common API Framework (CAPIF)**, designed to facilitate seamless interaction between different network services and applications in 5G and beyond. OpenCAPIF provides a set of standardized interfaces to manage and invoke APIs across various network functions. Key features include:

- **Authentication, Authorization, and Accounting (AAA)**: Ensures secure and controlled access to APIs, tracking the usage of network resources and services.
- **Service API Publishing and Discovery**: Allows for the publication and discovery of available service APIs, enabling dynamic and efficient interactions between network elements.
- **Logging of API Invocations**: Provides detailed logging of API invocations for monitoring, debugging, and auditing purposes.

- **Management of Service APIs and API Invokers**: Facilitates the administration of both service APIs and the invokers, ensuring smooth operation and scalability.
- **Topology Hiding**: Enhances security and simplifies integration by hiding the underlying network topology from external invokers, reducing the complexity of interactions.

OpenCAPIF simplifies the development and deployment of network services by providing a standardized and flexible API invocation framework that supports the evolving needs of modern 5G networks.

3.3.3 Monitoring and Observability

aerOS self-* suite

aerOS self-awareness component enables each integrated computing resource (IE in aerOS) *in the continuum to autonomously monitor and report its current operational status. It continuously* collects detailed telemetry related to the node's computing, memory, disk, network, and power usage, as well as its capacity to handle real-time workloads. The monitored data includes system-level metrics (e.g., CPU load, RAM/disk usage, network throughput, power consumption) and structural identifiers (e.g., hostname, IP, OS type). The collected status is formatted in a semantically modeled way and submitted to Data Fabric. It provides valuable real time information to support context-aware orchestration, fault tolerance, and adaptive workload management across the CEI infrastructure.

Grafana

Grafana is a leading open-source real-time data monitoring, analysis, and visualization platform. At its core, it acts as a unified front end that can pull in metrics and logs from a wide variety of back ends—everything from time-series databases like Prometheus and InfluxDB, to relational stores such as MySQL and PostgreSQL, to search engines like Elasticsearch, and even cloud services like AWS CloudWatch.



Figure 14: Grafana example dashboard

Once connected, Grafana lets you visualise raw datasets as interactive and customisable dashboards. Its flexible panel editor supports charts, tables, heatmaps, and more, all of which can be combined on a single canvas. Users can apply queries and transformations on the fly, zoom into time ranges, and drill down into individual data points for deeper investigation.

Beyond visualization, Grafana includes a robust alerting engine. You define thresholds and conditions directly within your panels, and Grafana will evaluate them at regular intervals. When an alert fires, notifications can be routed via a number of supported channels, ensuring your team stays informed.

Security and governance are built through role-based access controls and fine-grained user management. Administrators can define who can view or edit specific dashboards, folders, or data sources, and audit logs track changes over time.

One of Grafana's greatest strengths is its plug-in architecture. A public library hosts dozens of community-contributed and official plug-ins, supporting additional data sources, custom visualizations, and new panel types.

Thanos + Prometheus

Prometheus is a popular system for collecting metrics from applications and infrastructure, and it is known for its simplicity and powerful query language. It works well for many setups but was originally designed for single-node deployments. This becomes a limitation when trying to monitor large, distributed systems.

Thanos solves this problem by building on Prometheus and adding features like global querying, high-availability storage, and long-term data retention. It does this without changing how Prometheus works locally, making it easy to adopt incrementally. The Thanos Sidecar uploads data to object storage, while components like Querier and Store enable querying across all data sources as if they were one.

ELK Stack

The ELK Stack, comprising **Elasticsearch**, **Logstash**, and **Kibana**, is a powerful open-source suite used for searching, analyzing, and visualizing log and event data in real time. **Elasticsearch** serves as the scalable search and analytics engine, **Logstash** handles data ingestion and transformation from various sources, and **Kibana** provides an interactive web interface for querying and visualizing data. Widely used in observability and security use cases, ELK helps organizations gain deep insights into system behaviour, application performance, and security events from a centralized logging platform.

Kibana

Kibana is a visualization and analytics tool that works with Elasticsearch and OpenSearch. It provides an interface for exploring logs, building dashboards, setting up alerts, and investigating events in a visual way. It's especially useful in situations where understanding system behaviour requires digging through complex logs and data streams.

With features like Lens and TSVB (Time Series Visual Builder), users can create custom views of their data and detect patterns or anomalies over time. It also supports advanced filtering and search capabilities, making finding specific events or errors in large datasets easier.

Loki

Loki is a log aggregation system developed by Grafana Labs that is purpose-built to work seamlessly alongside Prometheus. Unlike traditional logging systems that index the full content of each log line,

Loki takes a more efficient approach by indexing only the metadata—such as labels and timestamps. This makes it more resource-friendly, especially in limited compute and storage environments.

Loki's lightweight nature makes it a strong candidate for edge deployments where full Elasticsearch or OpenSearch stacks may be resource-consuming. Logs are sent via the Promtail agent or other compatible sources and can be queried using LogQL, a powerful query language inspired by PromQL. This enables users to correlate logs with metrics, improving the speed and accuracy of issue resolution.

Loki integrates directly with Grafana dashboards, making it easy to pivot between metrics and logs within a single interface. It supports horizontal scaling, long-term storage in object stores like S3, and retention policies that suit both high-frequency log streams and bursty telemetry from edge workloads. This makes it well-suited for observability in distributed systems where visibility is critical, but infrastructure resources may be constrained.

3.3.4 Data Management

DOME

DOME takes the form of a shared digital catalogue of cloud and edge services. This catalogue is made available through a global DOME portal and through federated marketplaces. The core of the DOME technical architecture is the Distributed Persistent Layer (DPL). The DPL manages the storage and access to information related to:

- The Shared Catalogue of Product Specifications and Product Offerings defined by service providers.
- Product Orders and Product instances throughout their life cycle.
- Information about the actual Usage of Products.

FogFlow

In addition to orchestrating service topologies, **FogFlow** supports automatic data flow management and policy-driven control over data spaces. Its processing elements, typically Docker containers, are tightly integrated with CEI data management frameworks and respect data usage and access control rules throughout execution. The system provides built-in mechanisms for dynamic data routing, integration of new data sources (e.g., sensors), and real-time policy enforcement, all without requiring manual reconfiguration of pipelines.

This capability supports O-CEI's objectives around automated data integration, usage governance, and the reuse of data products within evolving service compositions — especially in dynamic and federated CEI environments.

dataClay

dataClay is a distributed object-oriented datastore that maintains application data in its native format in memory, eliminating the need for serialization. By persisting objects exactly as they are defined in code dataClay provides full support for any user-defined data structure.

Central to dataClay's design is the principle of "computation-in-place." Developers can directly deploy custom data processing routines onto storage nodes instead of transferring large volumes of raw data over the network. These routines run against local object instances and return only the results to the caller, reducing network traffic and accelerating analytics and simulation workloads across both co-located and distributed environments.

dataClay optimizes throughput and latency for data-intensive applications by tightly coupling storage and compute. Its modular back-end allows for horizontal scaling, while built-in replication and fault handling mechanisms ensure resilience under heavy load. Versioning and access control policies are enforced at the object level, offering fine-grained governance without compromising performance.

SkyFlok

Chocolate Cloud's **SkyFlok** Platform provides a solution for secure, privacy-focused data management. It offers two complementary solutions:

- **SkyFlok SaaS Application** – A secure file management and collaboration platform enabling organizations to store, organize, and share sensitive data across multiple clouds. It ensures compliance with GDPR, HIPAA, and other data protection standards through advanced encryption, geographic-aware redundancy, and user-defined data control policies.
- **SkyFlok Smart Storage** – An S3-compatible object storage service which allows customers to spread data across a personalized combination of cloud providers and geographic regions. Users can change where data in their buckets is distributed either manually or automatically based on the current demand.

Both offerings utilize encryption, erasure coding with a distributed-first architecture, and user-controlled settings to ensure reliable access, flexible scaling, and robust data security.

Eclipse Zenoh

Eclipse **Zenoh** is a high-performance, fault-tolerant data distribution protocol inspired by the Named Data Networking paradigm. It treats every piece of information—live streams, stored records, or on-demand computations—as a named resource, and exposes a single, straightforward API for publishers, subscribers, and queryables. It unifies pub/sub messaging with geo-distributed storage, query execution, and in-network computation and delivers bi-directional data flows efficiently without layering multiple middleware stacks.

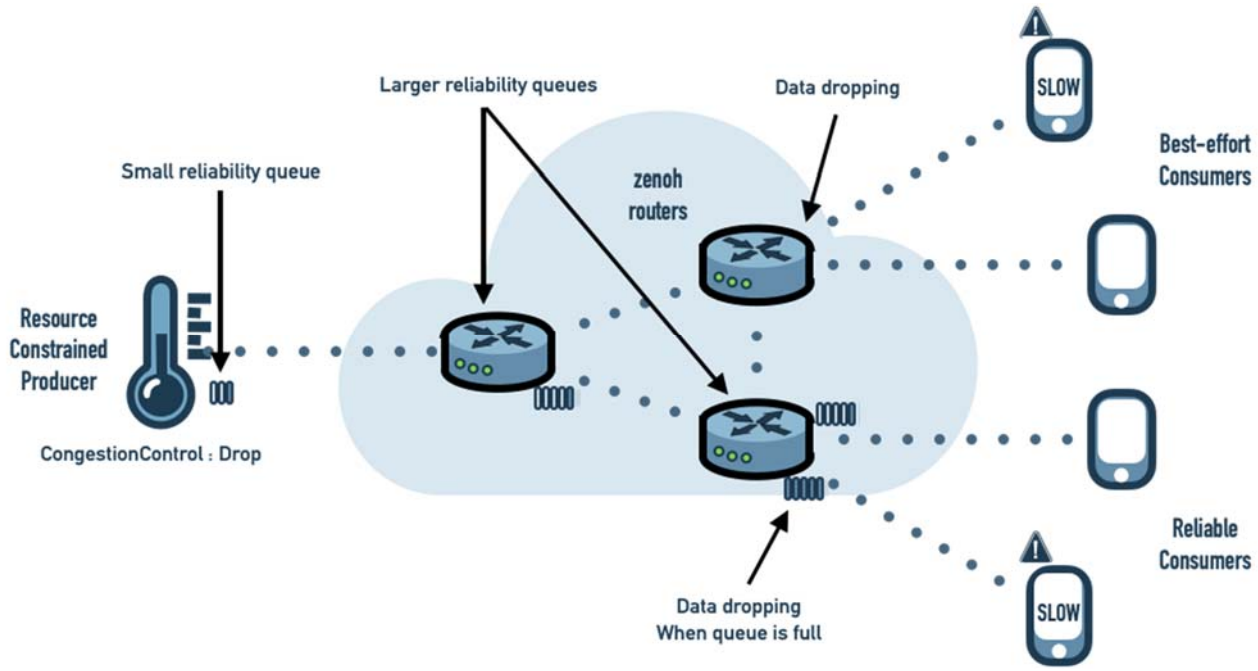


Figure 15: Zenoh example topology

Zenoh scales from microcontrollers on constrained, low-power networks to cloud-scale servers and data centers. Its wire-level batching, dynamic fragmentation, and configurable reliability modes drive consistently higher throughput and lower latency than mainstream alternatives. With only 4–6 bytes of wire-overhead per message—and support for locality-aware routing and in-network processing—Zenoh also minimizes bandwidth use and energy consumption, making it ideal for large-scale IoT and edge deployments.

Zenoh adapts to any deployment topology:

- **Peer-to-Peer:** nodes discover and communicate directly in resilient meshes or cliques.
- **Brokered:** lightweight endpoints offload complexity to full-featured brokers.
- **Routed:** dedicated routers forward Zenoh traffic across wide-area networks.

By consolidating data-in-motion, data-at-rest, and edge computation into one lean protocol, Zenoh slots effortlessly into the “Data Management” block of the EUCEI reference architecture, while also touching on connectivity and in-network compute layers. Its compact API and versatile topology support make it an excellent choice for building interoperable, energy-efficient systems that span from constrained devices to cloud infrastructures

Apache Camel

Apache Camel is an open-source integration framework that implements Enterprise Integration Patterns (EIPs) to facilitate the integration of data sources and processing of data. Camel comes with a wide library of connectors for protocols, such as HTTP, MQTT, or Kafka, enabling seamless connectivity across heterogeneous systems while abstracting underlying protocol complexities.

Camel provides a rule-based routing and mediation engine, allowing developers to define complex message flows and transformations using a variety of domain-specific languages (DSLs) such as Java, XML, Groovy, and YAML. These capabilities make Apache Camel a flexible, modular framework that can facilitate the design, deployment and orchestration of data pipelines within the data fabric.

Chimera

Chimera is a framework built on top of Apache Camel that enables the creation of composable semantic transformation pipelines for schema and data conversion based on Semantic Web technologies. Chimera leverages the RDF4J framework for RDF handling and provides a set of Apache Camel components that enable the following steps in the pipeline:

- Semantic lifting: transforming structured data into RDF graphs.
- Graph management: SHACL validation, SPARQL operations.
- Semantic lowering: extracting structured data from RDF graphs.

GraphDB

Ontotext GraphDB has been identified as a candidate technology for storing the knowledge graph of data Fabric. GraphDB is a highly efficient, scalable, and robust enterprise-ready graph database designed for storing and managing knowledge graphs using RDF and the SPARQL query language.

GraphDB supports advanced reasoning, allowing users to derive new semantic facts from existing data in real time, and offers integration with external search platforms like Elasticsearch. Built on the RDF4J framework, GraphDB provides features such as cluster support, compatibility with industry standards, and both community and commercial support.

Context management using MIM-1 (Context Broker)

ETSI ISG CIM, and in the future ETSI TC Data, are defining the “Next Generation Service Interface – Linked Data (NGSI-LD)” standard. NGSI-LD defines the interfaces of a context broker, a data management architecture that enables the brokering of meta-data rich data items (context). Brokering here means two interaction patterns: query (request-reply, synchronous and asynchronous retrieval of information based on various filter criteria ranging from type-based access to spatial-temporal queries or graph-based selections) and subscribe/notify (asynchronous requesting notifications upon data changes using the same filter criteria as with query). NGSI-LD is one of the MIM-1 approved standards.

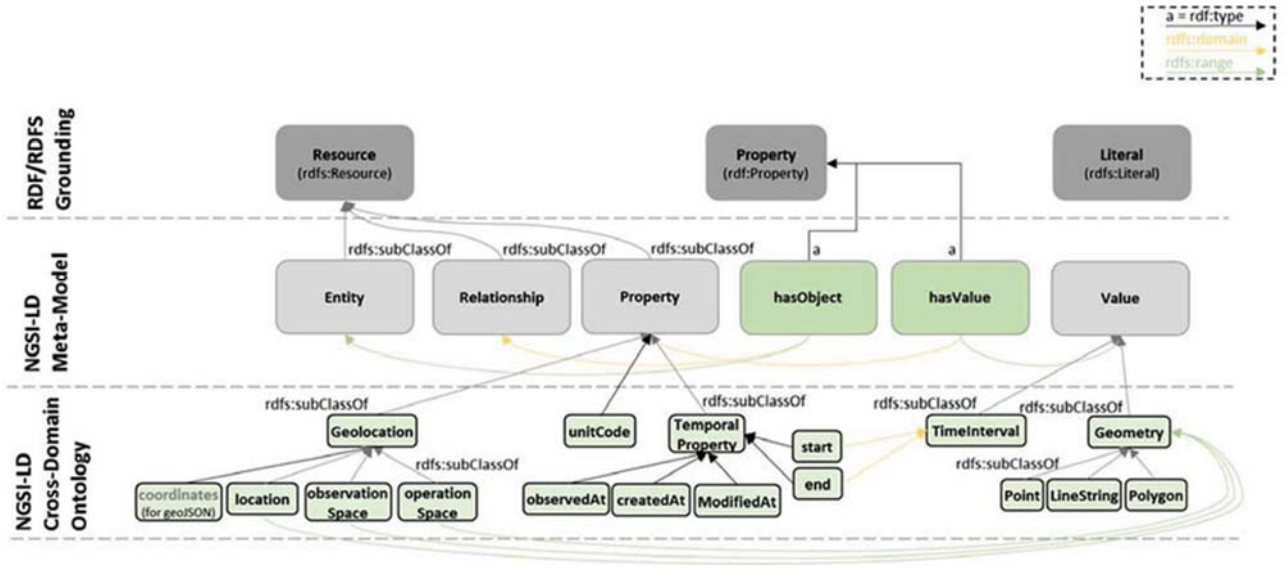


Figure 16: NGSI-LD Data Model with Meta-Model and RDF/RDFS Grounding (source: ETSI/ISG CIM)

Context broker following the NGSI-LD standards can be deployed in decentralized (one owner, many instances) and federated (many owners, many instances) deployments. It has been tested over various data spaces. A lightweight version on the context broker is used in FogFlow as broker in the Edge and IoT part of the continuum.

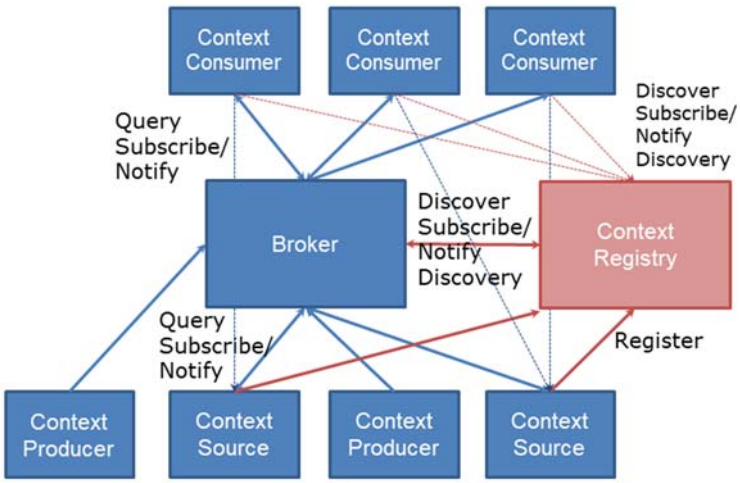


Figure 17: Context Broker Deployment Mode

AI agents using vector databases

NEC provides **AI Agents** which can be distributed across the CEI continuum and sense/actuate from the different parts of the system. Using artificial intelligence techniques to perceive the environment, they make decisions and take actions to achieve specific goals. Overall, AI agents' operation lifecycle includes:

- Perception of its environment through sensors (e.g., data inputs, APIs, user interactions),
- Processing of that information using algorithms (e.g., machine learning, logic, rules),
- Acting upon the environment through actuators or outputs (e.g., sending messages, making decisions, controlling devices),
- Learning or adapting over time to improve its performance or decision-making.

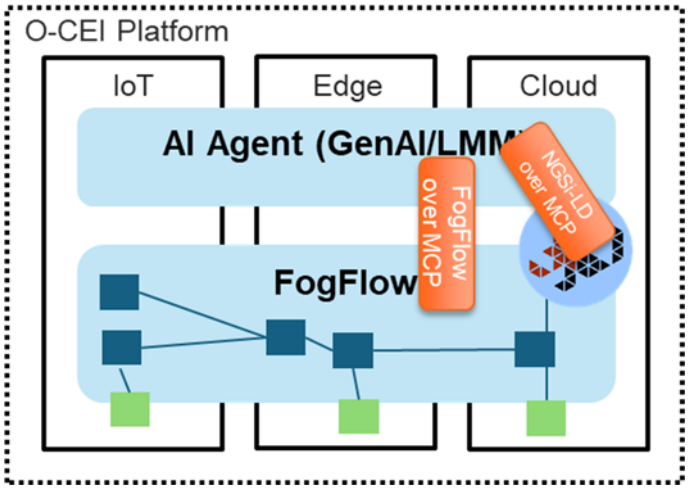


Figure 18: AI Agent Distribution across the CEI Continuum

A special class of AI agents introduced in the project leverages Generative AI and Large Language Models (LLMs) to provide intelligent services. These agents use vector databases to access relevant data objects and rely on the Model Connectivity Protocol (MCP) to consume external tools and resources, extending their capabilities beyond internal knowledge to include dynamic, domain-specific information.

IDS Connectors

IDS connectors are core components within the IDS architecture, enabling secure, sovereign, and standardized data exchange between organizations. They are foundational to the creation and operation of data spaces, federated environments where participants can share data while maintaining full control over its usage. Functionally, an IDS Connector acts as a controlled gateway that allows data providers and consumers to define, enforce, and audit usage policies for shared data. It facilitates trusted communication by ensuring data provenance, consent-based access, and contractual compliance through built-in mechanisms like data usage control, logging, and identity verification. **EDC and FIWARE IDS Connector** open-source implementations are provided and considered for integration within O-CEI blueprints.

aerOS semantic stack

The **aerOS semantic stack** provides a high-level framework for enriching, translating, and managing data across diverse sources and systems within the aerOS infrastructure. At its core, it enables **semantic interoperability** by turning raw data into meaningful, standardized information that can be easily processed, shared, and understood across different platforms. It consists of two main components:

- **Semantic Annotator** – This component enhances incoming raw data (like JSON, XML, or CSV) with semantic meaning, making it easier to interpret and process later. It can handle both individual messages and continuous data streams.
- **Semantic Translator** – After data is semantically enriched, this component ensures it can be understood by different systems by translating it between various semantic formats. It aligns data from producers and consumers using a shared central understanding.

Together, these components allow aerOS to seamlessly integrate and process data from multiple sources, ensuring consistency, clarity, and interoperability across distributed environments.

3.3.5 Resource Management

OpenNebula

OpenNebula, through its involvement in various projects, has developed a mix of AI techniques and Zero-Touch resource management approaches to improve the deployment and management of distributed and heterogeneous cloud-edge environments. OpenNebula provides a unified view of underlying hardware resources belonging to multiple types of cloud, enabling interoperability and multi-cloud portability. OpenNebula enables efficient resource management, reduces energy consumption and optimizes the infrastructure.

OpenNebula infrastructures can be deployed on-premises, in the cloud, at the edge, or across hybrid and multi-cloud environments. It uses the open-source KVM hypervisor for virtualization, with optional support for LXC containers. Cloud resources are managed by one or more OpenNebula Front ends, which coordinate core components—including daemons, services, and interfaces—to handle deployment, orchestration, monitoring, and overall resource management.

OpenNebula, as a Resource Manager, will handle both types of resources, physical infrastructure (e.g. servers, storage, networking elements) but also virtual infrastructure (e.g. virtual machines) distributed across multiple geographical locations. Through a set of interfaces and APIs, OpenNebula eases the provisioning, management and automation of the infrastructure mechanisms. In OpenNebula, the basic management unit is considered the VM.

Phoenix-RTOS

Phoenix-RTOS is a modern, microkernel-based real-time operating system designed for resource-constrained devices, including those at the very edge of the network (IoT devices). Its modular architecture allows for the inclusion of only the necessary components, optimizing resource usage while providing a robust and secure foundation for various applications.

Designed with scalability, modularity, and real-time performance in mind, Phoenix-RTOS supports multiple hardware architectures and integrates essential features such as POSIX compatibility, message-passing IPC, and user-space device drivers. These capabilities enable robust orchestration, monitoring, and integration with cloud-edge management frameworks, making Phoenix-RTOS a viable foundation for deploying AI-driven services and enabling dynamic participation of constrained devices in the CEI continuum.

Terraform

Terraform is a widely used infrastructure-as-code (IaC) tool designed to provision and manage infrastructure resources declaratively across cloud, edge, and on-prem environments. Although, Terraform best fits under Resource Management, as it automates the provisioning of compute, storage, and networking resources required to host services across the CEI continuum, it can also touch **Orchestration** pipelines, where it defines the underlying infrastructure on which services are deployed.

K8s Cluster API

The **Kubernetes Cluster API** extends Kubernetes functionality by enabling declarative provisioning, scaling, and management of Kubernetes clusters themselves. Although this maps directly to Resource Management, where infrastructure becomes programmable and composable Cluster API also integrates with orchestration layers to enable dynamic workload placement and multi-cluster coordination, supporting the Orchestration block.

3.3.6 Trust and Reputation

Suite of DLT-based enablers

Provided by CERTH are designed to enhance trust, transparency, and accountability across distributed IoT ecosystems. These enablers address challenges regarding data exchange, integrity, and collaboration. They can support secure, transparent, and scalable data-driven services. The suite includes **Logging & Auditing enabler** to record critical actions during data exchanges between stakeholders, the **Data Integrity Verification enabler** introducing mechanisms for verifying the integrity of exchanged data, **Distributed Broker enabler** which facilitates trusted data sharing across heterogeneous IoT devices and **DLT-based FL** enabler which introduces a decentralized reputation management system within a FL framework.

Blockchain Data Brokerage Engine

Developed by UPRC enables secure, privacy-preserving, and auditable cross-domain data sharing by leveraging Distributed Ledger Technologies and smart contracts. It provides decentralized logging, customizable access control, and policy-based trust enforcement, supporting traceability and compliance across federated CEI environments. Designed for integration with zero-trust architectures and federated marketplaces, it facilitates on-chain governance while delegating data exchange to secure off-chain channels.

IOTA Tangle

A feeless DLT framework enabling decentralized, traceable data exchange across CEI nodes. It supports secure communication of critical network state information via Hornet nodes, ensuring data integrity and trust in dynamic, federated environments. Automated deployment and APIs facilitate integration with continuum infrastructures.

aerOS Trust Management Component

The **aerOS Trust Management System,** provided by IQB**,** is composed of two core components: the **Trust Agent** and the **Trust Manager**. Together, they enable dynamic, event-driven trust assessment of computing resources (aerOS IEs) across the CEI continuum.

The **Trust Agent** acts as a decentralized tracer responsible for collecting relevant trust-related attributes from IEs via NGSI-LD or MQTT protocols. These attributes include security events, health diagnostics, software and firmware update status, behavioral indicators (e.g., communication and command exchanges), runtime service counts, system status (CPU, RAM) and reputation feedback from peers or users. The **Trust Manager** receives and processes this input to compute a **Total Trust Score (TS)** for each IE. This score is derived using a weighted algorithm that emphasizes the relative importance of each trust-contributing attribute. To ensure scalability and efficiency, only the TS is stored and exposed through Orion-LD and IOTA, while detailed sub-scores and raw data are retained locally in JSON format. The Trust Manager employs both periodic and event-triggered recalculations to maintain timely, accurate evaluations, adapting to environmental changes in real-time. This trust evaluation mechanism ensures scalable, decentralized, and context-aware trust management aligned with TF3's **Trust and Reputation** building block, but also offering capabilities that fall under **Monitoring and Observability**.

XLAB: IaC Scan Runner

The open source **IaC Scan Runner** (enterprise edition is Steampunk Spotter) is a container based service to initiate various open source or proprietary checks over the IaC code. The tool covers multiple IaC languages from TOSCA, Terraform, Ansible, Docker, Bash, Python, Yaml and others. Key features of IaC Scan Runner (Steampunk Spotter) include thorough IaC code analysis, identifying hard-to-catch errors, for improved quality, reliability, and security of IaC code as well as checking IaC code compatibility, identifying issues, and providing advice for smooth migrations. Additionally, it ensures compliance with security standards (e.g. CIS benchmarks for RHEL and AWS), preventing misconfigurations and security vulnerabilities, and allows for the inclusion of custom policies to meet specific security requirements.

3.3.7 Security and Privacy

Cybersecurity tool for anomaly detection

Anomaly detection tool developed by CERTH provides an integrated security system for strengthening the defense of enterprises and organizations against continuously evolving cyber threats. The system adopts a multi-layered approach that includes advanced artificial intelligence algorithms and mechanisms such as data mining, live deep packet inspection, and reinforcement learning. Its goal is to identify anomalies in both network logs (network-based detection) and system/infrastructure logs (host-based detection). By leveraging threat intelligence from external sources as well as historical data, it performs analysis and provides personalized alerts and data visualizations, while also offering recommendations for addressing the detected threats.

Cybersecurity tool for AI-based IDS

This tool, developed by CERTH, leverages deep learning methodologies to detect network intrusions, utilizing PCAP files as the primary input source. The workflow initiates with the extraction and pre-processing of network flows from the captured traffic. Each flow is represented through a transformation of its binary content into integer values, which serve as the basis for the extracted features. In parallel, a set of statistical and temporal metrics—such as packet rates, average packet counts and inter-arrival times—is computed to enrich the representation of network behaviour. The IDS integrates two distinct AI architectures, a one-dimensional Convolutional Neural Network (1D-CNN) in conjunction with Auto-Encoders, and a fully connected Multilayer Perceptron (MLP). These models are independently trained on complementary feature sets to enhance detection across a broad range of attack scenarios. For operational deployment, the system supports real-time monitoring by continuously ingesting new PCAP data from live network streams, analyzing it, and issuing alerts upon identifying anomalous or potentially malicious activity—thereby serving as a proactive cybersecurity defence mechanism.

SSI Wallet

CERTH's SSI Wallet is developed in alignment with W3C standards for **DIDs** and **VCs**, enabling secure, privacy-preserving identity management and credential exchange within decentralized environments.

Self-Sovereign Identity Framework

IQB's **SSI framework** delivers a decentralized, privacy-preserving infrastructure for managing identity and access control across the O-CEI ecosystem. At its core, the system handles the creation, resolution, and management of DIDs for users/devices, ensuring globally unique and verifiable identities without relying on centralized authorities. Leveraging VCs, trusted entities can issue cryptographically secure claims about an entity's attributes (e.g., roles, compliance status), which can then be stored and managed through SSI wallets and presented by the holder to other actors within the ecosystem. These credentials support selective disclosure, enabling entities to reveal only

the minimum required information, thus reinforcing compliance with privacy regulations such as the GDPR. IQB's SSI also facilitates secure onboarding and authentication of all stakeholders using verified credentials, establishing trust at the edge, in the cloud, and across connected domains. Through integration with a federated identity broker—leveraging systems like Keycloak—it bridges SSI with traditional identity infrastructures, allowing seamless interoperability and backward compatibility. Authorization decisions are enforced via smart contracts, which use credential-based access policies to grant or deny access dynamically. Additionally, the system supports credential lifecycle management, including revocation and status checking, ensuring that changes in authorization or trustworthiness are promptly reflected in the ecosystem.

LOMOS

Application running normally has a couple of states it is in – they are typically obtained through its log. It thus follows that observing the application's log is significant for understanding the running and the state of the application. The LOg MOnitoring System (**LOMOS**) is based on an AI engine harnessing Natural Language Processing architecture (LogBERT) to model log streams and characterize their operational states. Unlike traditional log monitoring solutions typically relying on rule-based manual analysis of time series data, this approach operates autonomously without the constraints of manually defined rules or human intervention. LOMOS employs a behavioral model to automatically identify anomalies indicative of abnormal conditions, including potential security threats.

IoT Vulnerabilities Crawler

This tool, developed by SYNELIXIS, aims to offer proactive security to the O-CEI continuum. More specifically, the crawler performs different types of scans to any web service that runs on the continuum and reports back on its findings. The crawler uses different types of vulnerability checks (bundled in the form of plugins) to scan applications or devices. The use of plugins helps to keep a continuously expanding database (e.g. MITRE, OWASP etc.), that guarantees an adequate and up-to-date vulnerability coverage. The crawler is able to start parallel scanning jobs for the different scanning targets. This job produces reports on the vulnerabilities found and makes them available for further exploitation.

Keycloak

Keycloak is an open-source Identity and Access Management (IAM) solution designed to secure applications and services with minimal effort. It provides features such as single sign-on (SSO), user federation, identity brokering, and fine-grained access control through role- and attribute-based policies. Keycloak supports modern authentication protocols including OAuth 2.0, OpenID Connect, and SAML, making it compatible with a wide range of applications and infrastructure. Its administrative console and REST APIs enable centralized management of users, roles, and permissions, making Keycloak a popular choice for enterprises building secure, scalable, and standards-compliant authentication systems.

3.3.8 AI

OpenNebula's AIOps

OpenNebula's AIOps is an AI-powered component designed to enhance the intelligent operation of cloud environments managed by OpenNebula. Its core aim is to optimize the management, monitoring, automation, and IT operations of both physical and virtual infrastructure, including virtualized applications and Kubernetes clusters, by leveraging AI and machine learning technologies.

When integrated into an OpenNebula deployment, this module provides advanced observability, predictive analytics, and anomaly detection—enabling smarter, data-driven management of cloud

resources. By applying sophisticated machine learning models trained on extensive historical data, including metrics such as resource consumption, network activity, application performance, and custom user-defined indicators, it can more accurately forecast workload demands. This enables proactive, efficient scheduling and scaling, ultimately improving overall cloud performance and resilience.

A key feature of OpenNebula AIOps is the intelligent orchestration of Virtual Machines and Kubernetes clusters, driven by predictive insights into application workloads and the utilization of infrastructure and virtual resources. This intelligent orchestration, using the actual state of the system and its AI/ML-driven predictions, can add new hardware resources across cloud-edge continuum, can allocate and re-schedule services and their underlying layers (VMs or containers).

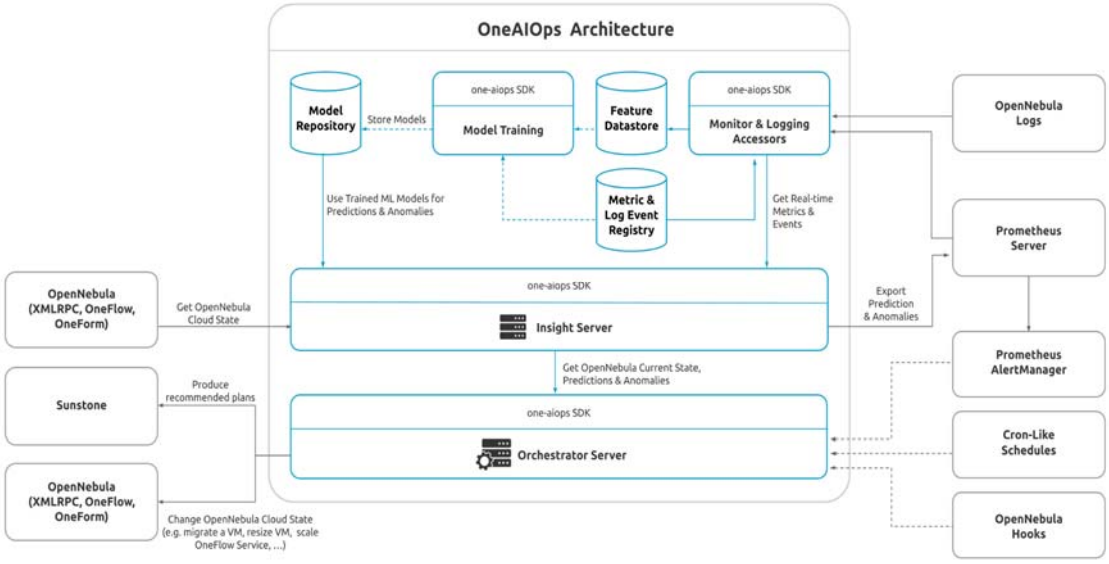


Figure 19: OpenNebula AIOps module architecture

MLFlow

MLFlow is an open-source MLOps framework that streamlines every stage of the machine learning lifecycle—from experiment tracking and artifact management to model versioning and production deployment. Its experiment tracking component automatically captures parameters, metrics, and outputs, while the Projects abstraction packages code and dependencies for reproducible runs. A centralized model registry lets teams catalogue, stage, and annotate models, ensuring transparent governance and traceability across multiple iterations.

MLFlow also boasts seamless interoperability with major cloud and big data platforms—Apache Spark for large-scale data processing, Azure ML and AWS SageMaker for managed training pipelines—and works out of the box with popular Python and R libraries, including scikit-learn, TensorFlow, PyTorch, and caret.

One of MLflow's standout capabilities is its built-in model monitoring, which streams real-time performance metrics and data drift indicators back to a centralized dashboard. This continuous feedback loop alerts you to anomalies or degradation in model accuracy, enabling rapid retraining or rollback. Together, these features provide a cohesive, end-to-end solution for managing complex data science workflows in both research and production environments.

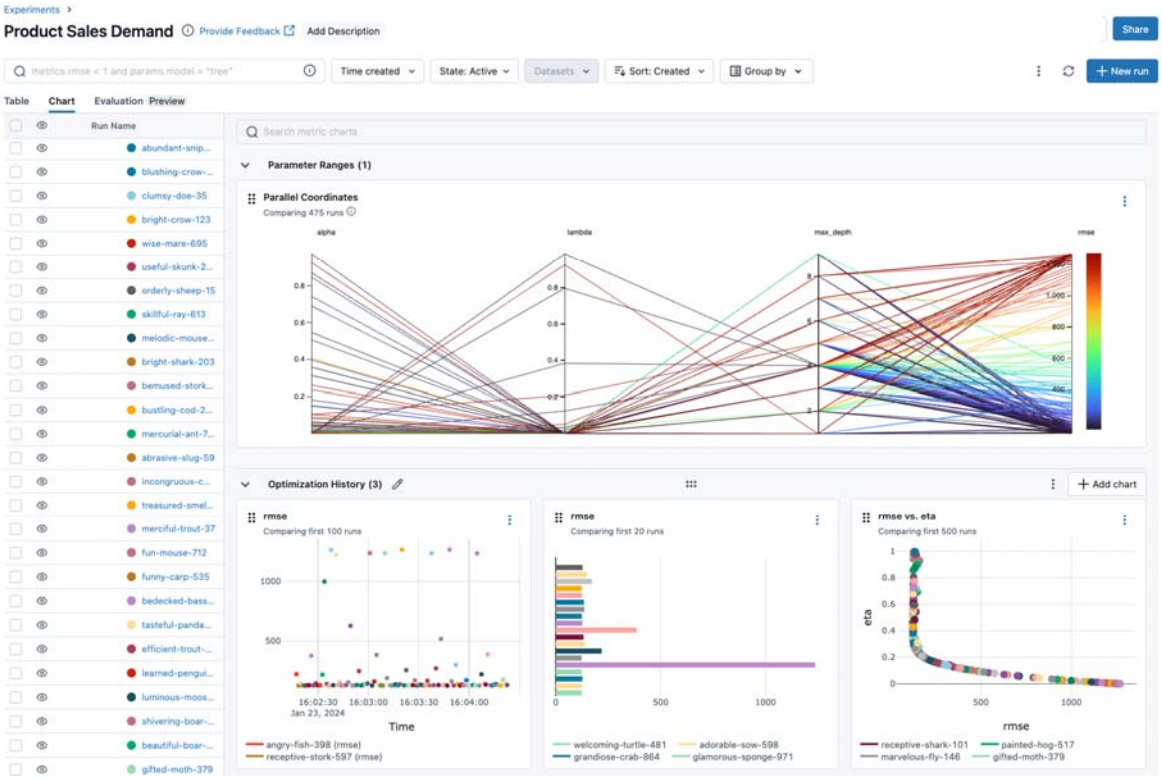


Figure 20: Example MLFLow tracking dashboard

MLFlow can complement an AI execution fabric, as envisioned in O-CEI, by offering robust model governance, performance monitoring, and seamless integration with cloud and edge-based data processing platforms. It supports traceability and operational resilience in CEI-aware AI workflows.

ZenML

ZenML is an open-source MLOps framework that simplifies the construction and management of machine learning pipelines by treating each step—data ingestion, transformation, training, and deployment—as native Python functions. Its metadata service automatically records run configurations, parameters, metrics, artifacts, and code versions, ensuring that every experiment is reproducible, and that lineage can be traced back through the entire workflow.

What distinguishes ZenML is its plugin-driven architecture: you select the orchestrator (Airflow, Kubeflow, or a simple local runner), artifact store (S3, GCS, Azure Blob), secret manager, feature store, and deployment target (KServe, SageMaker) that best fit your needs. In addition, built-in model registry and monitoring extensions provide real-time drift detection and performance alerts, enabling teams to respond quickly to deviations and maintain model quality throughout production.

ZenML provides O-CEI with a flexible and modular foundation for managing the full AI pipeline lifecycle. Its plugin-based architecture allows tailored deployment across edge-cloud environments, aligning with the heterogeneous and federated nature of the CEI continuum.

ClearML

ClearML is an open-source MLOps framework that unifies experiment tracking, data and artifact management, and pipeline orchestration under a single Python-native API. Every script you run is automatically instrumented: hyperparameters, system metrics, source code snapshots, and output artifacts are captured and versioned in a centralized metadata store. This ensures full reproducibility and auditability of all experiments, allowing teams to compare runs, reproduce results, and collaborate seamlessly.

What sets ClearML apart is its integrated pipeline management and automatic resource orchestration. Its DAG-based Workflow API and no-code pipeline builder let you assemble complex ML workflows with conditional logic and parallel steps. Automatic smart-queueing distributes tasks across available agents, while built-in hyperparameter optimization and data profiling modules accelerate experimentation. Together with a comprehensive web UI for monitoring logs, metrics, and resource utilization, ClearML delivers an end-to-end MLOps solution that scales from individual researchers to enterprise teams with minimal custom infrastructure.

Any of the MLOps frameworks described—MLflow, ZenML, or ClearML—can be effectively aligned with the EUCEI reference architecture by supporting key building blocks such as Data Management, Experimentation and Training, Model Deployment, and Monitoring. These frameworks provide modular, interoperable components that facilitate reproducibility, automation, and traceability, all of which are essential for O-CEI.

ClearML offers a powerful, centralized MLOps framework that aligns with O-CEI's goal for traceable, reproducible, and automated AI pipeline execution. Its ability to support parallel experimentation, distributed orchestration, and model lifecycle governance makes it a valuable utility across the CEI continuum.

BentoML

BentoML is an open-source tool for serving machine learning models. It focuses on making it easy to package trained models into production-ready services. This includes the model itself, and all the code needed to prepare inputs, run inference, and format the outputs.

These packaged services, called Bentos, can be containerized and deployed anywhere, from local machines and edge devices to cloud environments. Each Bento includes a lightweight API server and supports multiple runtime backends, including CPU, GPU, and ONNX.

BentoML also integrates with monitoring and deployment tools, allowing users to track performance metrics and automate updates. It's designed to support real-world workflows, where models need to be deployed quickly and reliably, with minimal effort.

Enables scalable, portable AI model deployment across the CEI continuum, supporting low-latency inference and flexible infrastructure integration, key requirements for distributed intelligence in O-CEI pilots.

KubeFlow

Kubeflow is an end-to-end platform for managing machine learning workflows on Kubernetes. It brings together several tools that cover the full ML lifecycle, from data preparation and training to deployment and monitoring. Users can define complex workflows using Python code, which are then executed as pipelines on Kubernetes.

At its core, Kubeflow includes tools for hyperparameter tuning (Katib), distributed training (using frameworks like TensorFlow or PyTorch), model serving (KFServing), and workflow orchestration (Pipelines). It also supports multi-user environments with secure authentication and role-based access.

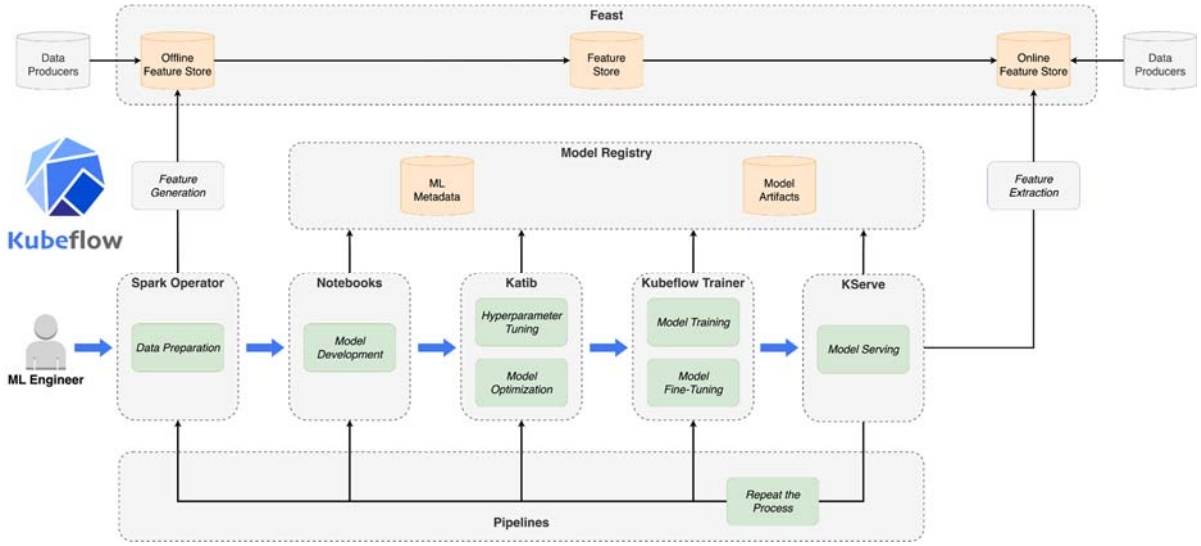


Figure 21: Kubeflow example lifecycle

Kubeflow supports scalable and secure machine learning workflows in containerized CEI environments. Its orchestration capabilities and model serving tools enable automated AI lifecycle management aligned with O-CEI distributed deployments.

Flower

Flower is an open-source framework built specifically for federated learning, enabling the training of machine learning models across distributed devices or silos without requiring centralized access to raw data. Its architecture separates coordination logic from model implementation, giving developers full control over managing training rounds, aggregation, evaluation, and communication. This flexibility makes Flower suitable for a wide range of federated learning scenarios, from mobile and edge computing to multi-institution collaborations.

One of Flower's main strengths is its framework-agnostic design. It supports seamless integration with popular ML libraries like TensorFlow, PyTorch, JAX, Hugging Face Transformers, and even lightweight environments like TFLite or pure NumPy. Developers can plug in existing models and training code with minimal modification, allowing teams to adopt federated learning without reengineering their ML stack.

Customization is central to Flower's design. Users can define their own strategies for aggregation, client selection, model initialization, and evaluation, adapting the learning process to each use case's privacy constraints, hardware limitations, and data characteristics. Server-side and client-side behaviours are configurable independently, which enables highly adaptive and domain-specific workflows.

Flower facilitates federated learning across edge and cloud nodes without sharing raw data, making it well-suited to O-CEI's privacy-aware, decentralized AI goals. Its modular architecture allows custom aggregation and coordination strategies tailored to energy, mobility, or industrial CEI use cases.

4 MARKET ANALYSIS AND ADOPTION BARRIERS

4.1 MARKET LANDSCAPE AND TECHNOLOGY ADOPTION TRENDS

4.1.1 Key industry sectors impacted by CEI computing continuum

The market for Cloud to Edge IoT solutions

The pace of technological development requires O CEI partners to keep a keen eye on market developments and trends that can affect the value of innovation in each of the pilots and scenarios. This section of the report will provide an overview of the market landscape and some of the high-level trends that are likely to continue to affect the development and adoption of CEI. The key challenge is to accelerate the development and adoption of CEI solutions as a way to reduce the productivity and competitiveness gap between Europe and its global competitors [24].

Recent research shows an increase in the forecasted spending on CEI solutions across Europe from the 2024 EUR 158 billion in 2024 to EUR 242 billion in 2028, a CAGR 11.3%. However, this is low margin, high volume market, with end-point hardware accounting for 46.8% of the 2024 spend[10]. It should also be noted that many of these end-point devices are likely to be manufactured in China, despite existing notable specialist or high-end, such as Continental's 5G V2X modules which are manufactured in the Czech Republic.

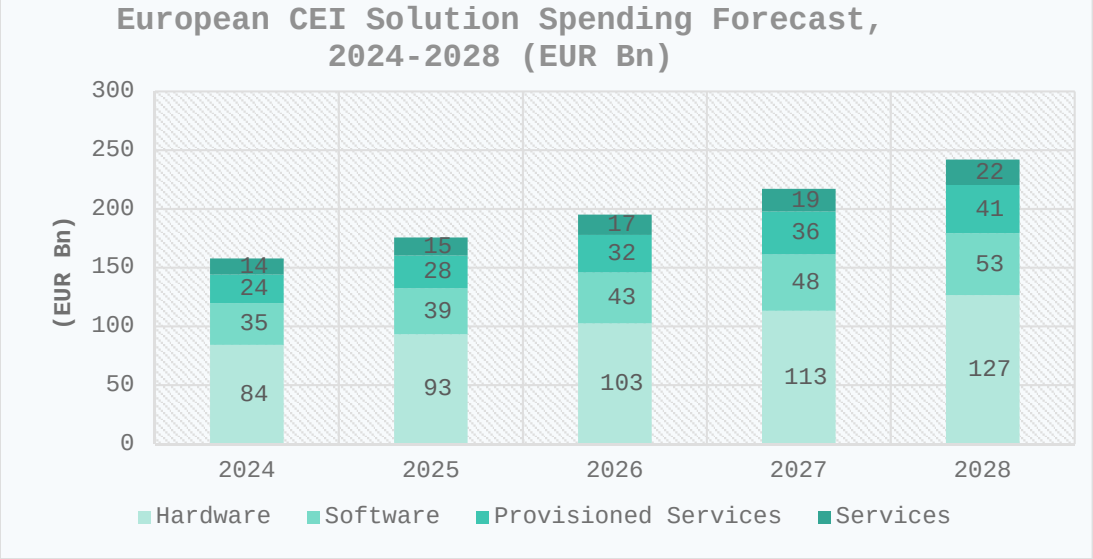


Figure 22: CEI Sphere D2.1 – Preliminary market analysis. European CEI Solution Spending Forecast, 2024-2028 (EUR Bn) (Source: IDC, Worldwide Internet of Things Spending Guide, November 2024.

According to CEI Sphere research, the market is still underdeveloped with the greatest rate of adoption seen in those use cases that are able to batch process data in central locations (e.g., track-and-trace for fleet management, or condition monitoring systems). Real-time control systems, such as those that will be developed in some of the 27 pilot scenarios, are still under used as they require a distribution of compute resources to the edge. The scalability of these systems is closely correlated to their cost. Real-time systems require distributed compute infrastructure, which involves:

[10] IDC, Worldwide Internet of Things Spending Guide (2024)

- Purchasing edge servers, industrial PCs, or gateways for each deployment site.
- Investment in robust, low-latency connectivity (e.g. private 5G, edge networking).
- Edge AI accelerators or GPUs are more expensive per unit of compute than central cloud servers.

On the other hand, deploying real-time control in the CEI continuum shifts costs from cloud-based OPEX to distributed CAPEX and maintenance. Cloud systems offer lower upfront costs and easy scalability, but they also incur ongoing fees, often including wasted spend and can suffer from latency issues in time-sensitive operations.

In contrast, edge architectures require more hardware and local expertise, but they significantly reduce data transfer and energy consumption and can unlock operational savings. With both edge and cloud markets continuing to grow, hybrid strategies are increasingly common—leveraging cloud for routine tasks and edge for critical, low-latency operations. The most cost-effective architecture ultimately depends on workload characteristics and industry needs, with CEI solutions gaining favor for mission-critical real-time applications despite higher initial investments.

Edge as an enabler of AI

The integration of Artificial Intelligence (AI), particularly Generative AI, is emerging as the most transformative force within CEI environments. These technologies enable automated decision-making and support complex, real-time analytics critical for modern applications across energy, mobility and manufacturing.

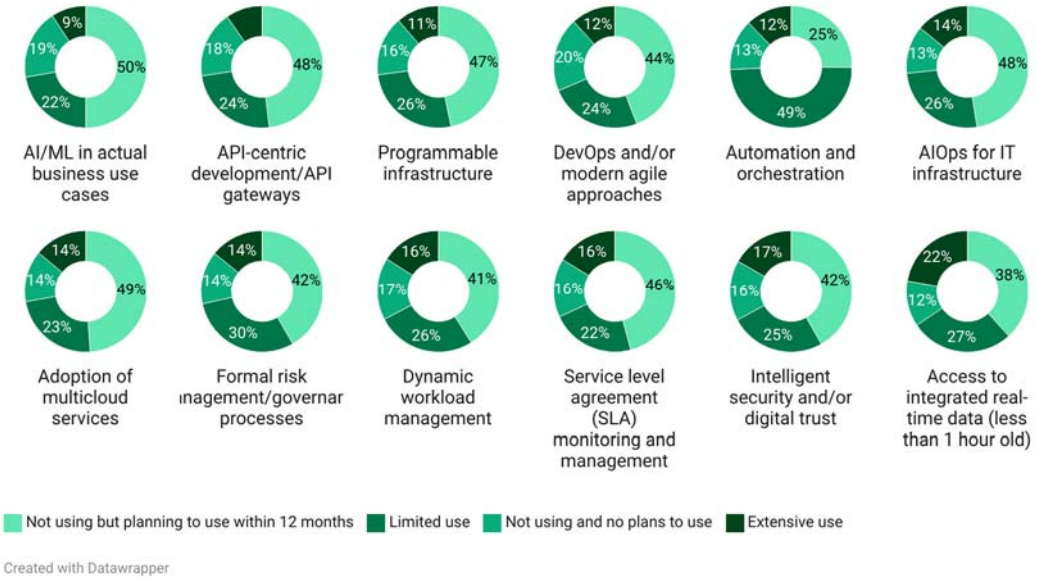


Figure 23: Adapted from CEI Sphere D2.1 – Preliminary market analysis. European CEI Solution Spending Forecast, 2024-2028 (EUR Bn) (Source: IDC, Worldwide Internet of Things Spending Guide, November 2024).

According to the IDC Worldwide AI Spending Guide (2024), AI spending in Europe is projected to nearly triple, from €46.4 billion in 2024 to €133.6 billion by 2028, with Generative AI expected to represent one-third of that total. Despite this momentum, the market remains in the early stages of determining how to effectively scale and manage edge-AI deployments, particularly in latency-sensitive and operational environments such as the mobility sector.

- Only 9% of European companies report extensive use of edge computing for AI/ML in real-world business use cases.
- However, 22% are already using edge compute to access integrated data less than 1 hour old, signalling a growing need for near real-time analytics.

- From a workload suitability perspective, only 12% of edge workloads are natively suited to run on edge devices, whereas 34% are better suited for the cloud. This mismatch underscores the infrastructure and architectural limitations facing edge-AI integration today.

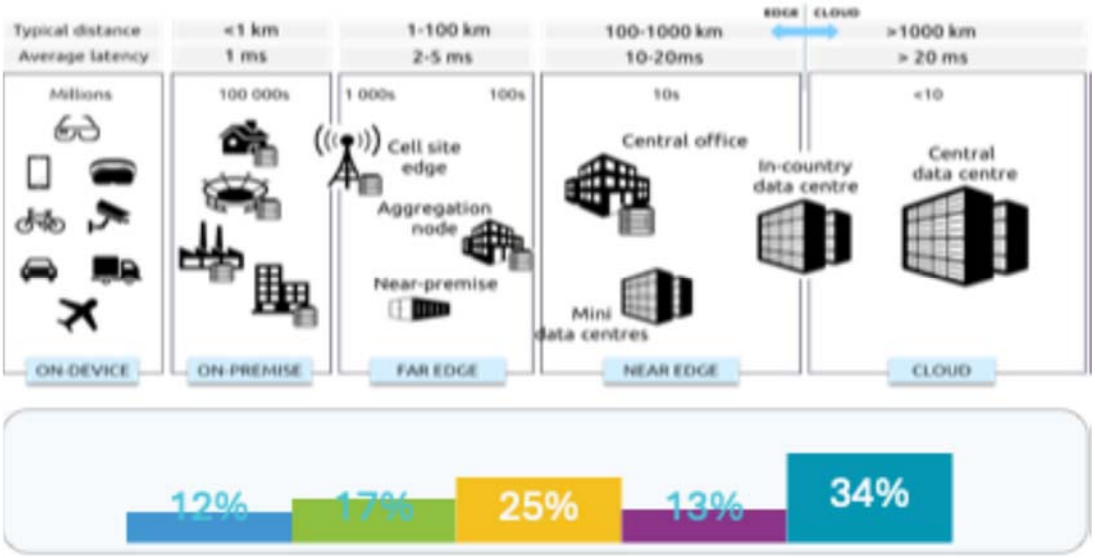


Figure 24: Proportion of edge workloads best suited to run in each deployment categories (Source: IDC EC Edge Observatory Custom Survey 2023/2024: Edge Computing (Waves 1, 2), n=600)

The role of CEI in different sectors

Energy

The European energy sector faces significant challenges, including high costs and the urgent need for decarbonisation. To address these, it is increasingly adopting CEI technologies to enhance grid stability and enable data-driven asset management. Key priorities include reducing energy costs, improving competitiveness in clean energy, and integrating distributed energy resources—such as rooftop solar, EV charging stations, and battery storage—into the grid. Smart grids equipped with IoT sensors collect real-time data to optimise electricity flow, reduce transmission losses, and support demand-response strategies. Cloud platforms enable large-scale AI analytics for forecasting and managing renewable energy. The sector's CEI spending in 2024 reached an estimated EUR 17.6 billion[11], primarily driven by electrical utilities. However, the distributed nature of CEI introduces cybersecurity challenges that must be carefully managed.

Mobility

The mobility sector, covering both passenger and freight transport, is undergoing a digital transformation aimed at enhancing efficiency and sustainability in line with the European Green Deal. IoT adoption is widespread, supporting real-time goods tracking, end-to-end logistics visibility, and optimised route planning. In passenger transport, CEI enables smart traffic systems and Mobility-as-a-Service (MaaS) platforms. Key challenges include seamless connectivity, integrating legacy systems, fragmented operations, a shortage of data analytics expertise, and the need for robust data privacy and security frameworks.

CEI is also a key enabler of the Software Defined Vehicle which is at the centre of the Industrial Action Plan for the European automotive sector[12]. The plan predicts that autonomous driving

[11] IDC's Worldwide Internet of Things Spending Guide (2024)

[12] https://transport.ec.europa.eu/document/download/89b3143e-09b6-4ae6-a826-932b90ed0816_en?filename=Communication%20-%20Action%20Plan.pdf

technology will generate EUR 400 billion by 2035. Scalable edge computing is a strategic enabler. Real-time decision-making, autonomous control systems, and over-the-air updates are all edge-dependent. However, the gap between potential and current capability means that:

- Investments must prioritise the preparation of edge sites for AI/ML workloads, particularly in sectors like transport and logistics where real-time control is mission-critical.
- Cloud-edge orchestration tools, AI-friendly hardware, and common open standards are essential to reduce deployment costs and overcome fragmentation.

Manufacturing

To counter declining competitiveness, the manufacturing sector is adopting CEI to improve automation, agility, and sustainability. European manufacturers show strong IoT adoption, particularly in edge-to-cloud integration, which enables real-time monitoring and optimised asset performance. Edge-to-cloud solutions are essential for achieving operational KPIs and enabling low-latency data processing. However, the deployment of more advanced use cases remains limited due to long asset life cycles and the evolving maturity of these technologies. Main barriers include IT/OT integration challenges, data security risks, lack of standardisation, and difficulties in incorporating legacy systems.

4.1.2 A firm level view of CEI

CEI readiness

Work completed under the previous Meta OS CSA led to the development of a **CEI readiness framework – the CEI-ling matrix** [25] **-** that assessed an organisation's ability to adopt cloud to edge technologies through five dimensions- Risk & Compliance, Finance, Industry & Market, Technology, and Organisation- assessed distinct levels of maturity across sectors and reveals critical factors influencing organisations' ability to adopt and scale CEI solutions. The analysis of the responses provided some interesting insights that can help shape the adoption of the technology

- **Infrastructure investment acts as a productivity catalyst**: firms with advanced IT infrastructure and high uptime are not only more CEI-ready but also more efficient, as reduced downtime directly supports productivity and business continuity.
- **Strategic R&D aligns closely with organisational readiness**: organisations investing in research and development are better equipped for CEI adoption, benefiting from continuous innovation, experimentation with emerging technologies like edge computing, and enhanced capabilities for rapid deployment.
- **Legacy system integration is both a challenge and an opportunity**: organisations that successfully integrate legacy hardware and software maximise asset value, reduce transition costs, and maintain smooth operations. Poor integration, however, delays adoption and limits data insights.
- **Flexibility and scalability drive revenue potential**: scalable technology infrastructures enable organisations to respond to market shifts and monetise CEI through offerings like data-as-a-service and predictive maintenance, generating recurring revenues with lower marginal costs.
- **Compliance and certification enhance CEI credibility and trust**: high regulatory compliance correlates with greater CEI readiness, reflecting strong governance, risk management, and stakeholder confidence essential for effective implementation and partnerships.

Overall, CEI readiness depends on a balanced combination of technology, financial planning, compliance maturity, and organisational agility. While manufacturing is well positioned to scale CEI strategies, domains such as research and innovation still require significant investment in foundational systems to fully engage in CEI-driven transformation. Targeted improvements in

infrastructure, compliance integration and financial modelling across all sectors are critical to unlocking CEI's full potential, as are the ability to understand where AI, and particularly LLM and agentic AI, can increased productivity in the technology stack.

What are companies looking for in CEI solutions?

Organisations are increasingly adopting CEI solutions to unlock a wide range of business and technical benefits. On the business front, increased performance (61%), improved cybersecurity (53%), and reduced operational costs (45%) are top motivators. From a technical perspective, companies prioritise enhanced security capabilities (60%), reliability and serviceability (52%), and integrated management (33%) as key drivers for edge adoption.

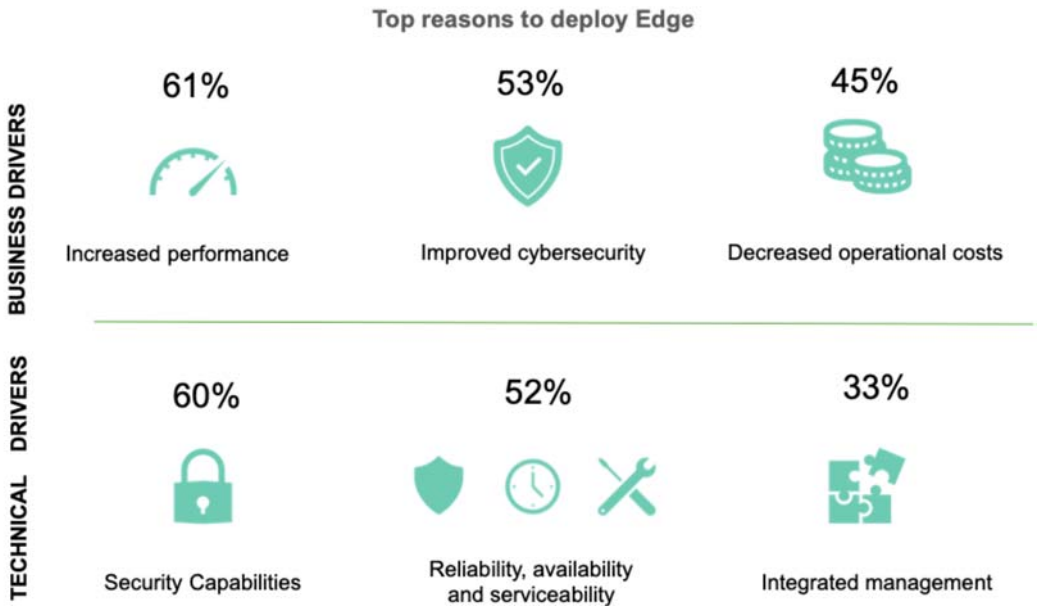


Figure 25: Top reasons to deploy Edge

Despite these benefits, organisations remain concerned about data privacy and security (64%), connectivity and power costs (64%), and the complexity of manageability and automation (52%). Other CEI deployment challenges include:[13]

- Managing data (24%)
- Integrating IoT with existing enterprise systems (21.2%), and
- Selecting suitable IoT platforms (19.2%)

It is also worth noting that many organisations struggle to transition small-scale proofs of concept (PoCs) into full-scale production deployments. This challenge must be proactively addressed during the pilot phase, with a thorough analysis of the following scalability parameters:

- **Technical architecture**: systems must be designed with modular, cloud-edge-compatible infrastructure that can handle increasing volumes of data and device endpoints without compromising performance.
- **Data interoperability**: pilots should use standardised data models and APIs to ensure integration across systems, platforms, and legacy environments.

[13] CEI Sphere D2.1 Challenges Inhibiting CEI Projects (Source: European 5G IoT Survey, August 2024, n = 950 (Europe, all respondents)

- **Security and compliance**: security mechanisms (encryption, authentication, access control) must be extensible and comply with GDPR and sector-specific regulations at scale.
- **Lifecycle management**: clear frameworks must be established for managing assets, updates, and policies consistently across distributed deployments.
- **Operational processes**: manual processes in pilots (e.g. data labelling, device provisioning, exception handling) must be replaced with automation and orchestration tools for scale.
- **Cost-efficiency**: economic viability must be tested under scaled conditions, including TCO analysis (hardware, connectivity, compute, maintenance), and aligned with sustainable revenue or service models.

4.2 IDENTIFICATION OF ADOPTION BARRIERS

4.2.1 Technical barriers: challenges in integrating emerging technologies

The complexity of project development and deployment poses a significant hurdle for CEI projects, particularly for end-to-end solutions involving edge computing. Integrating diverse technologies and systems is an inherent challenge in CEI solutions, as is addressing interoperability issues across various CEI components, platforms, devices, and security products. These issues are exacerbated by the lack of standardisation and open standards or APIs, particularly in areas such as IoT, edge orchestration, AI, and CEI security [24].

Integrating legacy systems and OEM systems with modern CEI solutions presents a recurring challenge across sectors such as energy, mobility, and manufacturing. This includes problems with legacy production assets that are not instrumented or connected. The fragmentation of the IoT device market also adds complexity and slows down deployments, often requiring custom hardware development. Furthermore, the absence of standardised data formats and communication protocols for IoT devices hinders seamless integration.

Limitations in edge computing infrastructure and the need for its further development are also notable. Challenges include ensuring power availability, managing and automating edge nodes, and supporting the high-power demands of AI and GenAI workloads. Edge environments are inherently complex due to their diverse hardware and software configurations, as well as dynamic network conditions.

Managing data across CEI systems introduces additional challenges, particularly in collecting, cleaning, storing, securing, and disposing of large volumes of data generated by numerous devices. The lack of adequate data infrastructure and the poor quality of available data also hamper effective AI implementation. Network connectivity and reliability issues persist, including ensuring seamless coverage and handover, especially in remote or mobile environments. These issues are compounded by diverse connectivity standards, immature 5G capabilities, limited spectrum availability, and the challenge of managing a patchwork of networks. Meeting the high data demands of CEI applications necessitates high-bandwidth, low-latency networks supported by specific technologies.

The development and deployment of AI/ML models at the edge also present a series of technical barriers. These include the need for model compression and optimisation on resource-constrained devices, distributed orchestration across numerous nodes, and maintaining reliable performance in dynamic environments. Managing the increasing complexity and volume of edge data, ensuring data quality and availability for AI training, and addressing issues of AI explainability and trustworthiness further complicate deployment.

Finally, orchestration and management of distributed CEI solutions remain challenging. This includes managing distributed data, applications, and AI models—tasks for which traditional tools are ill-equipped. The absence of open standards, as well as a lack of consensus on terminology, taxonomies, and architectures for edge orchestration, makes coordination more difficult. Furthermore, the ongoing evolution of component technologies such as AI and edge computing continues to limit the robustness and reliability of CEI applications in the short term.

4.2.2 Economic barriers: cost implications, ROI concerns

High costs are a common concern in the adoption of CEI solutions, particularly due to the complexity involved in project development and deployment. The financial burden of deploying and maintaining CEI infrastructure, especially at the edge, represents a significant challenge for many organisations.

The cost of specialised AI and GenAI hardware is another major barrier to adoption, adding to the overall expense. Connectivity costs further contribute to these challenges, as the deployment of technologies such as 5G and fiber can be prohibitively expensive, especially in large-scale or remote edge environments.

Additionally, the need for custom device development and the large number of devices required for IoT solutions create substantial cost barriers that hinder scalability. Demonstrating a clear return on investment (ROI) for complex CEI implementations is also a critical hurdle. This is particularly evident in sectors like mobility, where justifying expenditures can be difficult. Identifying strong AI use cases that clearly demonstrate ROI remains a significant challenge for stakeholders aiming to adopt CEI technologies.

4.2.3 Organisational barriers: skills gap, industry readiness

A significant organisational barrier to CEI adoption is the widespread skills gap and shortage of skilled professionals. This issue spans several key areas, including general CEI deployment, AI and GenAI infrastructure, edge orchestration, edge analytics (encompassing data science, IoT, and edge computing), as well as specialised security and mobility analytics expertise. The brain drain of skilled personnel to other regions further compounds the difficulty of advancing AI adoption in particular.

The complexity involved in project development and deployment does not only present technical challenges but also organisational ones. This complexity can make it difficult for companies to scale projects beyond the proof-of-concept stage and into full production environments. Moreover, organisational resistance, a lack of leadership, and the absence of strategic planning for the adoption and development of CEI solutions are key hurdles—especially in sectors like energy. In addition, the difficulty in identifying clear and compelling AI use cases continues to hinder broader adoption efforts.

Another challenge lies in data management and governance. Many organisations are not yet adequately prepared for these processes, and the lack of clear data governance models can impede effective orchestration. Integrating IT and OT (Operational Technology) domains also presents a major barrier. This stems not only from technical difficulties but also from conflicting organisational perspectives, particularly around issues such as security. These differences can result in mistrust and resistance, and poorly defined governance models or collaboration structures between IT and OT teams can limit the full utilisation of OT data.

Finally, achieving end-to-end visibility in highly outsourced and fragmented operational environments, such as those common in the mobility sector, remains a significant challenge. This lack of visibility can restrict effective coordination and decision-making, ultimately impeding the successful deployment and scaling of CEI initiatives.

4.2.4 Legal and ethical barriers: data-sharing constraints, liability concerns

Security risks and concerns are paramount throughout the CEI continuum due to the distributed nature of deployments, the growing complexity of cyber threats, and increasingly diverse attack surfaces. Ensuring end-to-end security across CEI solutions is particularly challenging, especially when managing security for large numbers of devices. Specific to edge computing, security concerns include the physical protection of edge nodes and securing communication between the edge and the cloud. The task of safeguarding distributed environments and data, as well as managing access control and authentication across multiple sites, adds significant complexity. Additionally, protecting sensitive data processed at the edge and securing AI and ML models is critical.

Data privacy and security concerns are amplified by the growing volume of data and its distribution across the CEI continuum. In collaborative ecosystems, ensuring data security and privacy is a complex and persistent challenge. Data-sharing arrangements raise concerns around data sovereignty and security, especially within ecosystems that process and store sensitive or personal data at the edge. Such edge processing can trigger specific regulatory and compliance issues, making it essential to ensure that edge deployments uphold stringent privacy and security standards to attract investment and promote adoption.

Regulatory complexity and compliance obligations contribute further to operational burdens. Navigating diverse regulatory frameworks—such as the GDPR, EU AI Act, NIS2, and the Cyber Resilience Act—can delay or complicate CEI deployment, particularly for solutions spanning multiple jurisdictions. These legal complexities are accompanied by challenges in data governance and compliance, especially in the absence of well-defined governance models.

Digital sovereignty is another strategic concern, particularly in relation to dependencies on non-European technology providers and the market dominance of large, non-European tech companies. These dependencies represent potential vulnerabilities that Europe is actively seeking to address, particularly in the context of edge computing. Data sovereignty, especially in manufacturing data-sharing and edge deployments, remains a central issue.

Finally, ensuring the explainability and trustworthiness of AI models represents a significant ethical challenge. This is particularly relevant in scenarios involving automated decision-making, where transparency and accountability are essential to building trust and ensuring responsible AI deployment.

4.3 REGULATORY AND STANDARDIZATION CONSIDERATIONS

4.3.1 Complying with regulations and standards

Despite the innovative potential of the O-CEI pilots, several regulatory and standardisation aspects present constraints to development, deployment, and scalability. These challenges affect everything from technical integration to data sharing, compliance, and market participation. Key issues highlighted by the pilots include:

- **Data privacy and cybersecurity.** Data protection regulations, especially GDPR, introduce strict requirements for managing personal and operational energy data. The pilots must ensure secure handling of real-time energy data, including from smart meters, EV chargers, and consumer interfaces. Maintaining data lineage, sovereignty, and implementing secure, decentralised architectures are key concerns. Pilots also identified the need for robust cybersecurity protocols, especially in contexts like ports and island communities, where infrastructure may be more vulnerable.

- **Regulatory uncertainty.** Shifting legal frameworks across the EU impact energy flexibility use cases. Pilot 3[14] highlighted the challenge of designing around uncertain future legislation for renewable energy communities and bi-directional EV charging. Regulatory ambiguity increases the risk of non-compliance and discourages early adoption of innovative solutions.
- **Sector-specific compliance obligations.** Compliance with industry-specific and EU-level regulations is an ongoing concern for pilots in highly regulated domains. Pilot 2 must navigate international automotive standards (e.g., Automotive SPICE, AUTOSAR, ISO 26262) and regulatory schemes such as Euro NCAP, particularly for vehicle-to-grid (V2G) solutions. Pilot 5 addresses the Corporate Sustainability Reporting Directive, EU labelling laws, and Lifecycle Assessment (LCA) reporting in the context of agrifood production, particularly for halloumi cheese.
- **Lack of standardisation and interoperability.** Many pilots face challenges due to fragmented standards, which hinder replicability and integration. Pilot 1 works to align with open protocols like Zigbee, Matter, and IEC 61850 to enable compatibility between legacy and new systems (e.g., boilers, meters, EV chargers). Pilot 2 noted the need for a standardised approach to EV roaming, with current centralised market models creating barriers to fair access. Pilot 5 focuses on creating interoperable semantics and data models for agrifood data exchange. Pilot 6 and Pilot 7 emphasise replicable and scalable solutions, requiring formalised standards to support wide deployment.

These regulatory and standardisation considerations are not just compliance challenges; they are central to the success and scalability of energy flexibility solutions. Addressing them early through coordinated stakeholder engagement, regulatory alignment and investment in interoperable systems will be critical for scalability.

4.3.2 Impact of EU regulations (Data Act, AI Act, NIS2, etc.)[15]

O-CEI stakeholders will have to comply with an increasing number of EU regulations that in turn will have an impact on standardisation. European Standardisation Organisations (CEN, CENELEC or ETSI) have developed harmonised standards, following standardisation requests from the European Commission emanating from Union legislation. These standards are then used by manufacturers, other economic operators or conformity assessment bodies to demonstrate that products, services or processes have a presumption of conformity with this legislation. Examples are:

- The AI act (AIA) which is associated with a standardisation request in support of safe and trustworthy artificial intelligence, currently being addressed by CEN-CENELEC JTC21
- The Cyber Resilience act (CRA) which is associated with a standardisation request in support of union policy on cybersecurity requirements for products with digital elements: CEN-CENELEC JTC13.
- The Data Governance act (DGA) which is associated with a standardisation request as regard European trusted data frameworks: CEN-CENELEC JTC25.

  In the domain of electric vehicles, the regulation on batteries and waste batteries which is associated with a standardisation request on digital product passports: CEN-CENELEC JTC24.

Figure 26 shows the initial compliance ecosystem that O-CEI has identified. Compliance relies on the following elements:

- Regulations such as ESPR, RBMC, NIS2, EUCC, FFDR.

[14] Information derived from the O-CEI Large Scale Pilots (LSPs) presentations at the PIC Paris meeting.
[15] For further information refer to Appendix A “Impact of Key EU Regulations on Cloud-Edge-IoT Pilots”

- Certification schemes, such as LEED, BREEAM.
- Standards needing alignment such as ISO 16739 (Industry Foundation Classes (IFC) for data sharing in the construction and facility management industries), ISO 19650 (Information management using building information modelling) or ISO/IEC 30188 (Digital twin reference architecture).
- European standard organisation committees such as JTC24 DPP (digital product passport) or JTC25 DDCE (Data management, Dataspaces, Cloud and Edge).

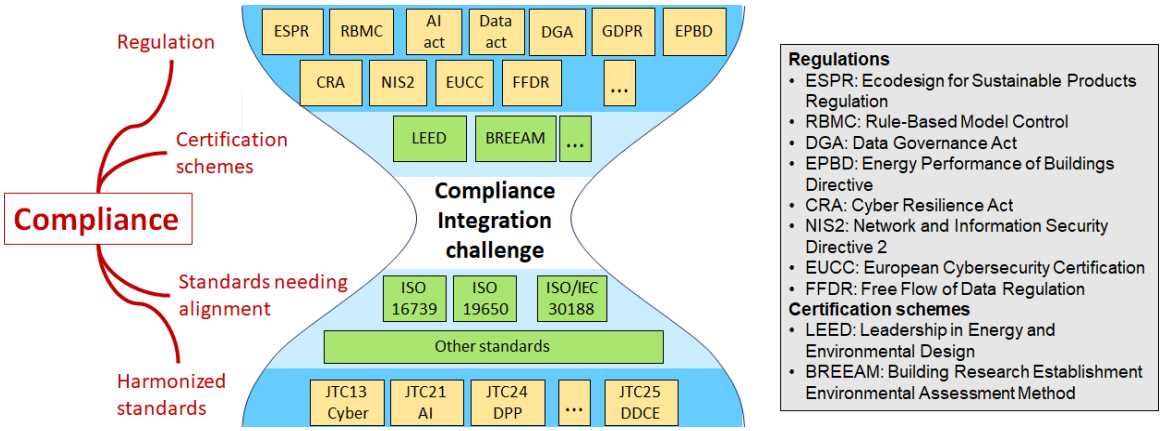


Figure 26: Compliance ecosystem

4.3.3 Standardisation efforts and compliance challenges

The O-CEI platform will require a holistic understanding of the standardisation effort and compliance challenges as shown in Figure 27. It shows that:

- The 8 Regulations indicated in the Figure 27 are associated with 4 standardisation committees who are working on harmonised standards.
- Each standardisation committee (e.g. CEN-CENELEC JTC 21 AI) is responsible for the development of harmonised standards.
- In the case of the AI act, the standards under development include:
  - Cybersecurity for AI systems,
  - AI risk management,
  - AI trustworthiness framework.

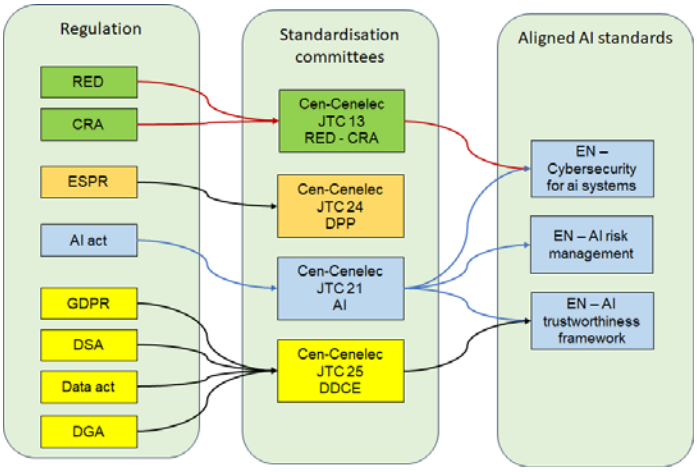


Figure 27: Compliance integration challenge example

The figure shows the need for alignment:

- The development of the AI standard “Cybersecurity for AI systems” must be aligned with the development of standards on the CRA in JTC13.
- Likewise, the AI standard “AI trustworthiness framework” must be aligned with the development of standards on the Data act in JTC 25.

4.3.4 Combining ecosystem and technology considerations

Approach

O-CEI will use the hourglass model developed by CEI-Sphere (Figure 28) which consists of two layers.

- An application layer (upper part) which focuses on user interfaces and data analytics.
- An infrastructure layer (lower part) which focuses on infrastructure, include IoT, digital twins, AI, cloud-edge.

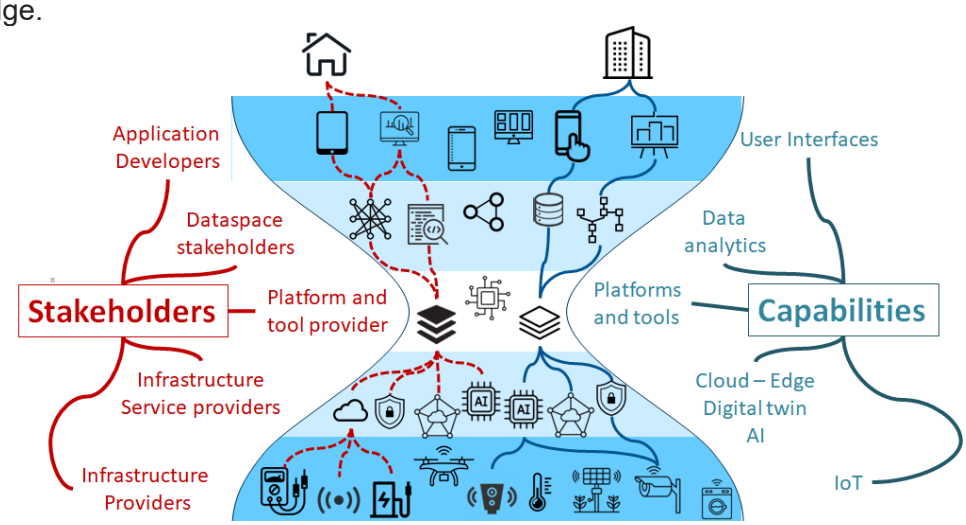


Figure 28: Hourglass model

Objectives and resulting toolkit

This visual representation helps understand:

- The construction the ecosystem that includes the demand side (stakeholders’ business objectives) and the supply side (capabilities),
- The reusability of building blocks (the icons in the hourglass model) for multiple markets,
- The avoidance of fragmentation (e.g. similar building blocks but incompatible in different domains).

CEI-Sphere has consequently proposed a toolkit (CEI Tech backbone toolkit) that will define and help apply to both Large Scale Pilots. The toolkit will help carry out an analysis and devise standardisation plan in the specific area of CEI (Cloud Edge IoT). CEI-Sphere has identified the following methods:

- Method to identify architecture patterns, which will help ensure interoperability through shared foundations, enable flexibility of architecture blueprints, and enable provision of increasing enriched capabilities.
- Method to foster interoperability through information models, architecture patterns and MIM (Minimum Interoperability Mechanisms).

- Method to build, use accessible/open-source implementation.
- Method to anticipate conformity requirements and meet harmonized standard requirements.

The outcome of using the CEI Tech backbone toolkit by O-CEI will be to facilitate the provision of the following assets:

- Information models (for interoperability),
- Architecture patterns (for building blocks reuse in blueprints),
- Minimum interoperability mechanisms, and
- Open-source based implementation.

Challenges of the toolkit

The challenges of developing the toolkit so that it meets the needs of heterogeneous O-CEI pilots and other CEI environments include:

- Defining a holistic assessment plan: Information model assurance plan, architecture pattern assurance plan, minimum interoperability mechanism assurance plan, and open-source based implementation assurance plan.
- Performing relevant analysis (information model analysis, architecture pattern analysis, minimum interoperability analysis, open-source implementation analysis. To this end the VCIO analysis approach will be used (Value – criteria – indicators – observables), used successfully in the case of AI trustworthiness (see AI Trust Label initiative).
- Evaluating and exploiting the model, architecture pattern, minimum interoperability, and open-source based implementation.
- Defining and gaining recognition for a trust label that provides confidence on compliance and define processes and means for obtaining the label.

4.4 A SCENARIO-BASED APPROACH TO CEI ADOPTION

4.4.1 Gaining an initial understanding of the market context for each O CEI scenario

As part of Task 2.2 "Market Analysis and Adoption Barriers Identification", a structured approach was adopted to gather insights from stakeholders across the eight pilots. The objective was to develop an initial understanding of the different pilot scenarios, the market positioning of each use case, the relevance of the solutions proposed, and the key barriers that could affect their adoption. To achieve this, two complementary methods were used: desk research and targeted stakeholder workshops. See Annex B, for the methodology.

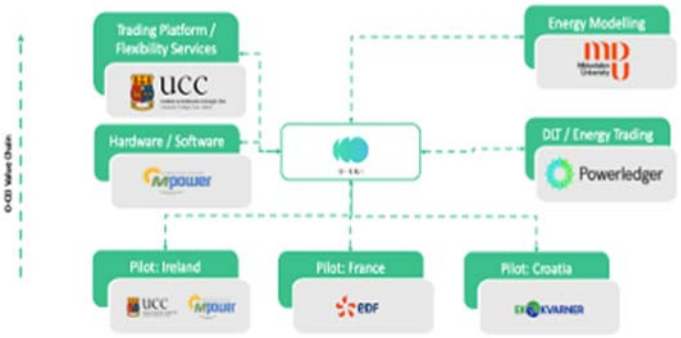
The following constraints were common to all the O CEI pilot scenarios:

- Technical constraints included scalability, security, interoperability.
- Administrative barriers: organizational complexity, funding gaps.
- Legal and ethical challenges; compliance, privacy, liability.
- Industrial mass-market adoption challenges: business model limitations.

Pilot 1: Electricity Grid performance optimisation upon RES integration

Technical constraints:

- Scalability: scaling up is challenging due to the high costs of infrastructure and devices, both hardware and software. AI and cloud-based solutions also require a lot of computing power, which limits wider adoption.
- Interoperability: data is spread across different systems and formats and without shared standards, it is hard to connect everything smoothly.



- Security: bringing AI and cloud tools into energy systems makes it harder to ensure that data is secure and properly protected.

Administrative barriers:

- Organisational complexity: there are many players involved- DSOs, utilities, municipalities, tech providers- with overlapping or unclear roles, which makes coordination difficult.
- Licensing and integration costs: high costs for software licenses and integrating tools make it harder for companies to enter the market or scale up.
- Local funding limitations: while cities are setting aside budgets for the energy transition, they often lack ready-to-use, standardized solutions to implement.

Legal and ethical challenges:

- Data privacy and GDPR: energy consumption data must comply with strict privacy laws, especially when it comes to smart meters.
- Liability: in decentralized and virtual systems, it is not always clear who is legally responsible when something goes wrong.
- AI transparency and ethics: AI-based energy tools need to be explainable, fair, and compliant in order to gain users' and stakeholders' trust.

Challenges to mass market adoption:

- Unclear business models: it is still hard to define clear value propositions and revenue models, which holds back broader market adoption.
- Low consumer engagement: many users do not see the financial benefits of participating as prosumers (e.g. savings or return on investment).
- Fragmented competition: existing solutions, like CHP systems, private microgrids, or local storage, often do not work well together or are not built for large-scale deployment.
- High upfront costs: the initial investment for devices, integration, and data handling is still a major barrier, especially in rural or less developed areas.

TRL: The pilot has a TRL level 6

KPIs:

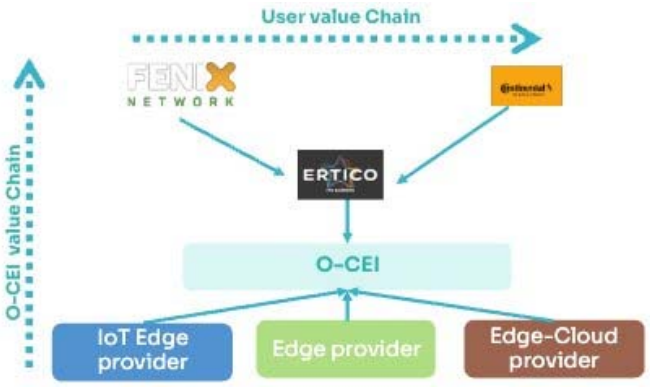
Table 8: O-CEI Pilot 1 "Electricity Grid performance optimisation upon RES integration" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach

KPI-P1-S01	Reduction in energy consumption across pilot sites (residential/commercial) due to control strategies, behavioural engagement, and system efficiency improvements.	Ireland: , ≥ 20%, Croatia: ≥ 15% bill reduction, France: ,≥ 12% (info only), ,≥ 20% (with management)	Pre-intervention energy use (kWh, €/bill) per site	Smart meters and sub-metering at 15-30 min intervals, pre/post comparison
KPI-P1-S02	User satisfaction improvements in terms of comfort, convenience, and control from implemented energy systems and interfaces.	Ireland & Croatia: ≥ 40% improvement	Survey results from pre-pilot phase	Surveys, interviews, feedback via digital interfaces
KPI-P1-S03	Reduction in CO,ÇÇ emissions in residential and commercial buildings due to increased efficiency and cleaner energy sources.	Ireland & Croatia: ≥ 30% (residential), ≥ 35% (commercial)	Estimated emissions based on baseline consumption and carbon factors	CO2 estimates using energy data and emission factors from sensors/meters
KPI-P1-S04	Share of renewable energy in total energy consumption across pilot sites.	Ireland & Croatia: ≥ 35%	Proportion of RE in current consumption	Energy source breakdown using smart meters, inverter data, utility data
KPI-P1-S05	Reduction in grid peak load demand through demand response, smart scheduling, and flexibility measures.	Ireland, Croatia, France: ≥ 20%	Peak load values during comparable pre-intervention periods	Grid load data monitoring, control signal response tracking
KPI-P1-S06	Improved grid voltage regulation and overall grid stability due to distributed energy resources and control systems.	Up to +5% gain in voltage regulation performance	Voltage fluctuation profiles from pilot sites	Voltage sensors and grid monitoring infrastructure

Pilot 2: Transition to Software-Defined Vehicle (SDV) platform

Technical constraints:



- Scalability: the growing complexity of software and the need to process data in real time require a lot of computing power. Edge computing infrastructure like roadside units and edge servers is expensive and not widely available, especially in rural or cross-border areas.
- Interoperability: vehicle systems often do not follow shared standards. Different data formats and market fragmentation make it hard for systems to work together across different vendors or countries.
- Security: protecting data is still a challenge in many cases. Encryption tools are often still being tested, and older systems do not easily connect with newer cloud platforms.
- Latency: processing vehicle data in real time needs extremely low delays, but current networks, especially in less connected areas, cannot always keep up.

Administrative barriers:

- Infrastructure investment needs: a significant amount of funding is required to build edge infrastructure like 5G towers or smart charging points. But it is not always clear when or how that investment will pay off, especially for public-private partnerships.
- Co-ordination complexity: many different players (car makers, telecom companies, insurers, startups, and cities) must work together, both legally and operationally, which makes co-ordination complicated.
- Slow regulatory adaptation: legal and technical standards (like ISO or GDPR) are evolving slowly, which delays progress and makes rollout harder.
- Limited incentives: there are too few national or regional policy-based incentives to support early adopters or people who feed energy back into the grid, which lowers motivation to participate.

Legal and ethical challenges:

- GDPR and data ownership: sharing data between vehicles, insurers, and other actors raises important questions about who owns the data, how consent is given, and who is responsible.
- Liability in automation: when something goes wrong in an automated system, it is unclear who's legally responsible, especially in international scenarios.
- Ethical use of data: it is increasingly important to help users understand how their driving or energy data is used, especially when AI is involved.
- Sustainable software development: developing advanced software often uses a lot of energy, raising questions about its environmental impact.

Industrial mass-market adoption:

- Unclear return on investment (ROI): long timelines for development and deployment make it hard to prove short-term financial benefits, especially for edge-based systems.
- Vendor lock-in and fragmentation: without open standards, companies risk getting stuck in closed systems that limit their flexibility and slow down innovation.
- Low awareness and education: many users do not fully understand the benefits of things like energy-efficient driving, software updates, or rewards for sharing data.
- Regulatory and policy lag: slow updates to laws and regulations delay wider rollout and make it hard to test solutions across borders.
- Competing partial solutions: many tools exist (over-the-air updates or predictive maintenance) but they are often disconnected, hard to combine and not designed to work together.

TRL: The pilot has a TRL level 4

KPIs:

Table 9: O-CEI Pilot 2 "Transition to Software-Defined Vehicle (SDV) platform" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
KPI-P2-01	Edge Computing Performance Latency at the Edge The time taken for data to be processed at the edge (from data collection to decision-making).	<100 mils	500 mils	Average and maximum
KPI-P2-02	Cloud Connectivity and Data Transmission Latency The time taken for data to travel from the edge to the cloud and back	<500 mils	750 mils	Average and maximum

KPI-P2-03	Data Transmission Reliability The percentage of data successfully transmitted to the cloud without errors or losses	> 99%	87%	Ratio of successfully transmitted data
KPI-P2-04	System Performance and Scalability - Uptime	> 99,9%	92%	Total uptime divided by total pilot
KPI-P2-05	System Performance and Scalability - Scalability	>1000 simultaneous vehicles	100 simultaneous vehicles	Maximum number of vehicles supported without performance degradation
KPI-P2-06	System Performance and Scalability – Error rate	< 1%	20%	Ratio of errors to total operations
KPI-P2-07	Data Analysis and Insights - Predictive Analytics Accuracy	< 90%	70%	Accuracy of predictions compared to actual
KPI-P2-08	Data Analysis and Insights - Data Processing Time in the Cloud	< 1 sec	5 sec	Average processing time for cloud-based analytics

Pilot 3: Energy consumption and emission reductions in postal Service fleet operation via intelligent BEV charging strategies

Technical constraints:

- Limited infrastructure for V2G: most grid systems are not ready yet to support two-way energy flows. Even where the technology exists, regulations and technical gaps are holding it back.
- Data fragmentation: real-time data exchange between fleets, storage systems, grid operators and energy markets, is difficult due to poor interoperability and the lack of unified platforms.
- Edge computing limitations: many EVs and chargers do not have the computing power needed for local AI decision-making or real-time optimisation.



- Lack of predictive tools: there's a shortage of tools for forecasting things like battery needs, route planning, or cost optimisation, especially ones that can keep up with how fast things change in fleet and energy operations.

Administrative barriers:

- Unclear system boundaries: AVL and other stakeholders often do not have a full view of how the whole system works, across vehicles, storage, the grid, and markets.
- Permitting and regulatory delays: even when the tech is ready, V2G cannot be used in many places because regulations do not allow it yet.
- Standard compliance issues: car manufacturers are still working on adopting standards like ISO 15118-20, which slows down the rollout of V2G features.
- Cross-sector coordination: Bringing together actors from energy, transport, and infrastructure requires complex partnerships and shared responsibilities—which are hard to manage.

Legal and ethical challenges:

- Guarantee and liability concerns: car makers are cautious about supporting V2G because they are unsure who is responsible if the battery degrades or a problem arises.
- Data privacy and ownership: sharing data about energy use, vehicle location, and user behaviour involves many parties, raising questions about consent and control.
- Market access limits: current regulations often do not allow EVs and batteries to fully participate in energy markets, particularly dynamic ones such as spot trading.
- Lack of standardisation: without clear and shared legal or technical rules for using EVs as energy assets, it is hard to scale adoption.

Industrial mass-market adoption:

- Gaps in the value chain: while the components- EVs, batteries, and smart chargers- exist, complete business solutions that tie them all together are still rare.
- High entry barriers: fleet operators face big investments for energy management systems, infrastructure upgrades and staff training.
- Slow deployment of smart charging: charging systems that respond to grid signals or market prices are not yet widely used, often due to technical or organizational resistance.
- Underused energy assets: EVs and batteries could help balance the grid or reduce energy costs, but these opportunities are mostly untapped.
- Scalability issues: solutions that work in one area need to be adapted for different grid types and infrastructure levels, which slows down large-scale adoption.

TRL: The pilot has a TRL level 4

KPIs:

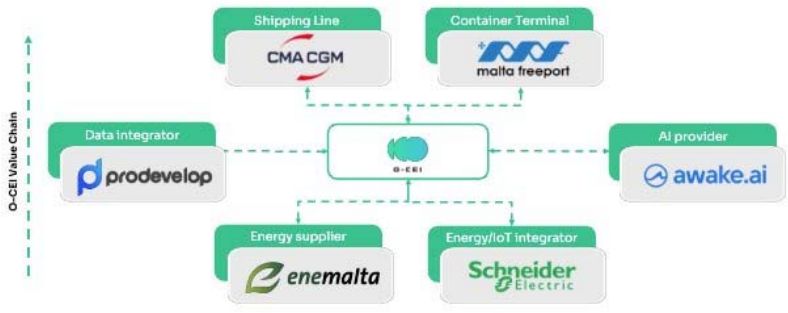
Table 10: O-CEI Pilot 3 "Energy consumption and emission reductions in postal Service fleet operation via intelligent BEV charging strategies" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
P3-KPI1	Reduction of total cost of ownership for fleet operator	10-15%	TBD	Determined for a reduced number of vehicles with simulation
P3-KPI2	Reduce vehicle idling time without productive activity	20-30%	TBD	Determined for a reduced number of vehicles with simulation

Pilot 4: Variable demand in challenging maritime terminal landscape

Technical constraints:

- Scalability: during peak times such as summer or when many port operations happen at once (vessels, reefers, EVs, AC units)—energy demand exceeds what the grid can handle. The current infrastructure is not built to scale up quickly.



- Interoperability: there is a lack of smart energy systems that can work together across residential, industrial, and port users. Systems like OPS, reefers, and cranes are not well connected to each other or to renewable energy sources.
- Data and latency issues: many energy-heavy systems (like reefers) cannot be controlled in real-time, making it hard to optimise energy usage effectively.
- Infrastructure limits: The island still depends mostly on non-renewable energy, and its power plants cannot always increase output unless smarter load-balancing systems are introduced.

Administrative barriers:

- High infrastructure costs: Upgrading the grid or building new power plants is expensive, and the financial return is still unclear.
- Coordination challenges: Energy suppliers, port authorities, shipping lines, local government, and citizens all need to work together—but they have different goals and incentives.
- Regulatory friction: Environmental and energy rules are becoming stricter, but there's no standard across the EU for port energy systems or OPS, which makes things more complex.
- Low renewable investment: Without stronger policies or financial support, investment in local renewable energy is still too low for real long-term change.

Legal and ethical challenges:

- Environmental compliance: Ports must reduce emissions to meet EU rules, or risk fines and a loss in competitiveness.
- Fair access to energy: If industrial energy needs always take priority over residential ones, it could cause public backlash and go against citizen-focused energy policies.
- Accountability gaps: Without smarter systems, it is unclear who's responsible if energy shortages happen—like if reefers shut down or EVs cannot charge.
- Data privacy: Smart systems will collect and use sensitive data, so strong privacy and transparency rules are needed.

Industrial mass-market adoption:

- ROI uncertainty: Smart systems, OPS upgrades, and grid improvements are expensive, and it takes a long time to recover the investment.
- Immature market: While companies like ABB, Siemens, and Schneider offer these technologies, they're rarely used together or at full scale in ports or islands.
- Fragmented vendors: Different providers offer different solutions, but they do not always work well together due to lack of shared standards.
- Misaligned incentives: Residential and industrial users often have energy peaks at the same time, but there are few policies to manage this conflict fairly.
- Weak renewable base: Without more local renewable energy, smart systems can only help with costs—not with real sustainability goals.

TRL: The pilot has a TRL level 4

KPIs:

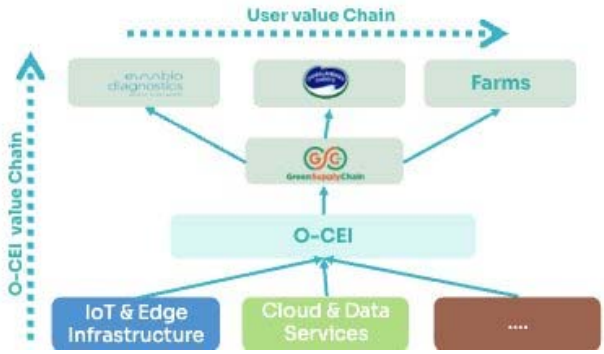
Table 11: O-CEI Pilot 4 "Variable demand in challenging maritime terminal landscape" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach

P4-KPI1	Reduction of energy billing due to vessel peaks (monthly cost)	20%	~57 MWh/day per feeder	Not defined yet
P4-KPI2	Accuracy of energy consumption predictions at MFT and its surroundings by AI/ML models	R2 and MAE: 80%	N/A	Not defined yet

Pilot 5: Energetically and environmentally sustainable Halloumi cheese production

Technical constraints:



- Lack of digital tools: Most farms do not have the basic tech like IoT sensors or sub-metering to monitor energy use during specific stages of production (like per litre of milk or kilo of halloumi).
- Heavy data processing needs: To scale up, farms would need to handle and analyze large amounts of data to adjust energy use based on changing conditions like weather, seasons, and product types.
- No system interoperability: Farm, transport, and factory data are not connected through common standards, making it hard to track energy use across the whole value chain.
- Manual operations: Most livestock and dairy processes are still done by hand, which makes automation or real-time monitoring almost impossible without major upgrades.

Administrative barriers.

- High startup costs: Buying energy-efficient tech and digital tools is expensive, especially for small or traditional farmers.
- Weak coordination: Farms, processors, and regulators do not work closely enough together, so energy savings efforts often stay limited in impact.
- Lack of funding support: There are few accessible grants or incentives to help small farms digitize their operations. Even if they exist the grants are very slow.
- Hesitance to share data: Farmers often worry that sharing data could lead to penalties or loss of control, so they avoid engaging with digital systems.

Legal and ethical challenges:

- Unclear data rights: Farmers do not have guaranteed legal protection when it comes to sharing their data, raising fears it could be misused.
- No consent mechanisms: There are not trusted tools or systems to let data owners decide how their data gets used.
- Complex certifications: Getting certified for sustainability (like CSRD or LCA) requires consistent tracking and reporting across many different actors, which is hard to manage.
- Sensitive data concerns: Some energy data, like HVAC use or animal welfare conditions, may be seen as private, making farmers reluctant to share.

Industrial mass-market adoption:

- Doubtful return on investment: The savings or added value from “green” products might not yet be high enough to justify spending on digital systems.

- Digital skills gap: Many farmers and factory workers do not have the training to use or maintain digital tools.
- Wide variety of farm types: Any solution needs to work for very different kinds of farms, with different sizes and levels of tech readiness.
- Scepticism about sustainability benefits: Without strong consumer demand or policy pressure, producers do not believe that investing in sustainability will pay off.
- Siloed solutions: Tools like OpenLCA, smart HVAC, or herd management systems exist, but they're often expensive, isolated, or do not work well with each other.

TRL: The pilot has a TRL level 4

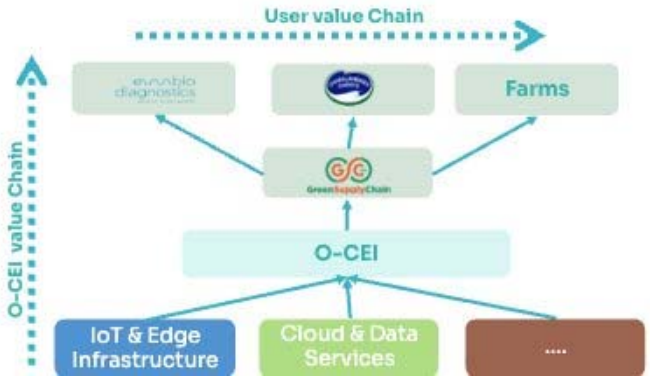
KPIs:

Table 12: O-CEI Pilot 5 "Energetically and environmentally sustainable Halloumi cheese production" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
P5-KPI1	20% reduction of total energy consumption for dairy production	39.38 kWh/1,000 kg).	49.23 kWh to produce 1,000 kg of halloumi cheese	Aggregate energy consumption of all production steps.
P5-KPI2	15% of reduced labour time (cost) for livestock management	1,190 hours	1,400 hours annually	Use Farm/Heard Management Information System (log) to record labor activity.
P5-KPI3	20% reduction of heat stress impact on cattle	8% decrease	10% decrease in milk production during hot months	Recording of calculated ThermalHeatStress Index combined with Milk production.
P5-KPI4	15% reduction of environmental footprint of the final halloumi cheese product	approximately 25.24 kg CO_2e/kg)	29.69 kg CO_2e/kg	Use LifeCycleAssesment mechanisms on data collected in P5-KPI1

Pilot 6: Smart re-charging and efficiency of robot tractors in large fruit production fields

Technical constraints:



- Inconsistent data & AI accuracy: AI models struggle to perform well in farming because data from fields (like terrain, weather, or soil) is often unstructured and varies a lot.
- Weak energy infrastructure: Many remote farms do not have stable energy access or fast-charging options for electric machines. Solar or wind power isn't always viable due to location or cost.
- Immature hardware: Current electric tractors need large batteries and fast charging — both are still expensive and not easy to scale.
- Limited automation tools: Existing tools for planning routes and tasks do not yet adapt well to the real-world variability in farming environments.

Administrative barriers:

- High customization costs: Tailoring electric vehicles and energy systems for different types of farms requires significant investment from both the tech providers and the farmers.
- Low digital confidence: Many farmers are not trained or comfortable using autonomous equipment, even if the tools exist.
- Weak collaboration: EV makers, AI developers, and energy providers are not working together closely enough, which slows down real-world deployment.
- Rural rollout challenges: Bringing charging and digital infrastructure to rural areas is expensive and logistically difficult unless there's strong support from both public and private sectors.

Legal and ethical challenges:

- No common standards: There's no agreed framework or safety certification for self-driving farm machines, which creates risk.
- Unclear data rules: It is not obvious who owns or can use the data collected by autonomous farm equipment.
- Liability issues: If something goes wrong — for example, if a tractor runs out of power — it is unclear who's responsible: the manufacturer, the energy supplier, or the software platform.
- Lack of data transparency: Farmers may be reluctant to share operational data unless they know it won't be used against them.

Industrial mass-market adoption:

- Tech still in early stages: Many required technologies (like fast-charging, AI planning, hybrid systems) are still being tested and not yet ready for wide use.
- Low trust in machines: Many farmers still prefer manual methods, especially for important tasks like harvesting, where timing is critical.
- Difficult system integration: Merging EVs, AI, and energy systems into a simple, reliable product is still a big technical challenge.
- Uncertain benefits: If it is not clear how much time or money will be saved, farmers are unlikely to invest.
- Scattered solutions: There are many small, specialized tools out there, but no leading provider yet offers a complete, all-in-one solution.

TRL: The pilot has a TRL level 5

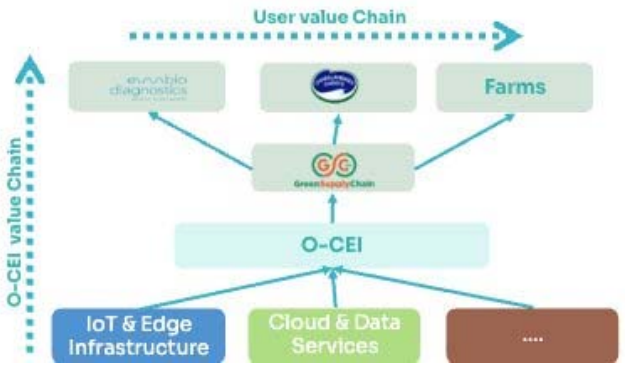
KPIs:

Table 13: O-CEI Pilot 6 "Smart re-charging and efficiency of robot tractors in large fruit production fields" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
P6-KPI1	50% reduction in energy consumption	0.5 kWh/ha (per year)	1 kWh/ha (per year)	-
P6-KP2	Time to collect, analyze and act upon (un)expected event	36 mins	2h	-
P6-KP3	Reduction of greenhouse gas emissions	0.24-0.48kgCO2-eq	0.6 to 1.2kgCO2-eq	-

Pilot 7: Trustworthy and secure EV charging upon reliable 5G networks

Technical constraints:



- Non-scalable infrastructure: current EV charging networks are not designed to adjust in real time based on fleet needs or grid conditions, limiting efficient energy use.
- Weak edge-computing hardware: AI for energy optimisation needs powerful and efficient edge devices, but most of today's equipment cannot handle the data or processing required.
- Immature bidirectional charging and renewables: technologies like V2G and renewable integration are not yet ready for large-scale use, limiting smarter charging strategies.
- Connectivity issues: network delays and unreliability, especially in remote or crowded areas, make it hard to optimise charging in real time.
- Cybersecurity gaps: even though tools like encryption and Blockchain exist, secure communication systems (especially at the edge) are not mature or widely adopted yet.

Administrative barriers:

- Disjointed ecosystem: charging systems are fragmented across operators, fleets, and logistics hubs, leading to inefficiencies and duplicated effort.
- No common standards: without shared protocols for energy and data management, it is hard to scale or integrate across different players and regions.
- Old infrastructure: many telecom and charging operators still use outdated systems that cannot support real-time data, AI tools, or edge computing.
- High upgrade costs: upgrading infrastructure to add renewables or bidirectional charging is expensive and slows progress.

Legal and ethical challenges:

- Privacy and security: managing data securely when using newer technologies like Blockchain or edge computing is still a major concern.
- Regulatory uncertainty: rules differ between regions and do not always cover new tech, which creates confusion and slows scaling.
- Data trust issues: even with Blockchain, ensuring data is trustworthy and secure across distributed systems remains challenging.

Industrial mass-market adoption:

- Lack of usable edge data and AI models: without reliable real-time data and efficient AI that works on smaller edge devices, intelligent local energy systems cannot take off.
- Rigid charging systems: most charging setups do not allow for flexible or dynamic energy use, limiting how well fleets and logistics hubs can scale.
- Unstable networks: Telecom providers struggle to keep networks fast and stable enough for real-time charging decisions.
- High running costs: without better integration of renewables, operators rely heavily on the main grid, driving up energy costs.

TRL: The pilot has a TRL level 7

KPIs:

Table 14: O-CEI Pilot 7 "Trustworthy and secure EV charging upon reliable 5G networks" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
P7-KP1	Cost savings for the end user w.r.t. to the Italian market	>5%	€1-0.90/kWh	Cost saving (€/kWh) percentage w.r.t the baseline
P7-KP2	Daily peaks in power consumption	<0.1%	0 – 10%	Mean daily peaks interval (h) percentage w.r.t 24h
P7-KP3	Up time of smart charging monitor service	>99%	85-95%	Mean monitor service up time (h) percentage w.r.t. 24h
P7-KP4	Reaction time to smart charging requests	<0.01s	3-5 minutes	Elapsed time from charging request and the communication of charging profile

Pilot 8: Heightened social engagement and acceptability of energy flexibility in urban areas

Technical constraints:

- Data gaps: energy forecasting tools (like LLMs or EMS systems) often lack complete, high-quality data, limiting their reliability.
- Fragmented EMS market: EMS vendors rely on complex, proprietary systems with no standard forecasting models, making integration and replication difficult.
- Infrastructure dependency: smart meters (e.g., Zigbee, CLEMAP) and reliable microgrids are critical, but rollout is slow due to high costs and complexity.
- LLMs not yet in use: Large Language Models are not currently applied in energy forecasting because they require validated, high-volume datasets.
- Edge-to-cloud challenges: building a seamless, adaptive edge-cloud system depends on both infrastructure maturity and end-user willingness to engage.

Administrative barriers:

- Weak user feedback loops: there is little direct involvement from end users to test or validate forecasting tools, which slows refinement.
- Low priority for flexibility markets: hotel operators and similar users see little financial incentive to participate in these systems.
- Slow uptake by infrastructure owners: high upfront costs and uncertain returns delay adoption of smart systems like CLEMAP.
- Lack of education: many potential participants (residents, tourists, businesses) are not informed or engaged in energy-related decision-making.

Legal and ethical challenges:

- No clear data governance: there is no standardised approach for sharing data across providers, which is a barrier to LLM development and use.
- Limited social acceptance: local trust, behaviour patterns, and political dynamics, especially in tourism zones, influence adoption but are hard to manage.
- Regulatory uncertainty: legal frameworks around microgrids and market participation are unclear and fragmented, making scaling difficult.

- AI ethics: though not always stated, applying AI/LLMs in residential and tourism energy contexts raises concerns about bias, transparency, and explainability.

Industrial mass-market adoption:

- Unclear business models: energy forecasting and flexibility tools still lack convincing value propositions for key users like hotels or EMS vendors.
- Hard to engage users: getting households, tourists, and citizens involved requires strong onboarding and incentives.
- No shared standards: without common algorithms or technical frameworks, it is hard to scale or coordinate across vendors.
- Tourism adds complexity: seasonal demand shifts and diverse stakeholders make it difficult to replicate solutions in tourist-heavy areas.
- LLMs untested in the field: there is no existing user base or validated models for LLMs in this domain, delaying commercial traction.

TRL: The pilot has a TRL level 4

KPIs:

Table 15: O-CEI Pilot 8 "Heightened social engagement and acceptability of energy flexibility in urban areas" KPIs

Code	Elaborated description	Target value	Baseline	Measurement approach
P8-KP1	Deploy at least two Smart Energy Services (End-user applications) per scenario and make them available to the end-users.	6	0	SES deployed and used.
P8-KP2	A better control of electricity consumption/production through LLM-enhanced flexibility applications: Increase the efficiency (in terms of flexibility) of an edge-to-Cloud solution with LLM by 12% compared with an edge-to-Cloud platform without LLM.	12%	0%	Use the Smart Readiness Indicator (SRI).
P8-KP3	Inventory of technical social salient attributes and a utility functions (part worth) of different prosumers profiles, share of choice (market shares) and behavioral patterns of prosumers.	500	20	Collection of energy prosumers profiles.

4.4.2 Key takeaways and industry requirements

Insights from the eight pilots confirm recurring barriers that slow down the adoption of digital energy and mobility solutions. Technologies like 5G, AI, edge computing, and V2G show strong potential, but are constrained by infrastructure gaps, low interoperability, and fragmented ecosystems.

The insights gathered through the pilot workshop both confirm and provide further nuance to the barriers to deployment described earlier at s4.2 above and, eventually, wide adoption of the CEI solutions. These include:

- Lack of standardisation across energy, mobility, and telecom systems prevents interoperability and limits scalability.
- Outdated infrastructure, especially in EV charging and grid systems, that cannot support edge computing, real-time data exchange, or secure communications.
- Edge AI deployment remains limited due to low compute capacity, insufficient training data, and early-stage federated learning frameworks.

- Cybersecurity and regulatory gaps reduce trust in edge-to-cloud communication and data-driven operations.
- Stakeholder fragmentation leads to isolated pilots that cannot be replicated or scaled without systemic coordination.

All pilots highlight the need for integrated, secure, and scalable solutions. Without targeted action to address these technical and coordination challenges, large-scale adoption will remain limited.

4.5 FROM PILOT TO OPERATIONS AT SCALE: INITIAL RECOMMENDATIONS

The European Commission's Competitiveness Compass (January 2025) sets out a strategic framework for strengthening Europe's position in key enabling technologies, particularly AI, Cloud-Edge-IoT, quantum computing, and green tech as part of its broader push for digital and industrial sovereignty. The report acknowledges the critical role of computing infrastructure in closing the performance gap in AI capabilities and Europe ability to develop sovereign technologies.

Despite these measures, Europe is unlikely to achieve complete technological sovereignty, particularly in semiconductors and hyperscale cloud infrastructure, where global dependencies persist. A large proportion of CEI spend will find its way to low margin, high volume end-point devices manufactured outside Europe.

Similarly, there is no Cloud-to-Edge IoT without AI: AI is indispensable for realising the full value of CEI systems. It enables autonomous decision-making, real-time inference, and adaptive control across distributed devices. Spend on the computing infrastructure required to run more sophisticated AI systems will flow to the makers of the datacenter chips. The value must therefore flow from the system (or tech stack) rather than the component.

While adoption strategies must reflect the systemic interdependence of infrastructure layers, from connectivity (fibre, 5G/6G) and compute (HPC, edge nodes) to orchestration (containerisation, AI pipelines), open-source strategies can help to lower the cost of certain parts of the system and avoid overreliance on certain providers. In any case, effective integration requires coordinated public and private investment across the entire tech stack.

Here the pilots come into their own as they adopt the federated architectures and data abstraction layers developed in the Meta OS. Similarly, the CEI Sphere digital backbone toolkit will support composability and allows modules to be adapted to differing infrastructure levels, urban, rural, maritime, etc., without extensive ex- novo redesign.

Pilot level strategies to drive market take-up

- **Build for heterogeneity (**easier said than done): most pilots are struggling with data fragmentation, interoperability, and context-specific constraints (e.g. local energy systems, rural infrastructure, proprietary EMS, varied farming environments). Open-source approaches can help further develop or adjust the software layers.
- **Scaling roadmap:** a clear, staged roadmap is needed to transition research and innovation (R&I) outputs into industrial-scale deployments, with defined maturity levels, use-case readiness assessments and integration guidance.
- **Business case clarity:** many pilots validate technical feasibility but overlook financial viability. Pilots should include parallel development of commercial models (e.g. service-as-a-product, data monetisation, public-private co-investment) and test them with early users. Future deployments must account for market-oriented parameters such as total cost of ownership (TCO), return on investment (ROI), and operational scalability.

- **Agentic AI:** the rise of goal-oriented, autonomous AI systems ("agentic AI") heralds a new era of increased productivity through automation. However, these systems require robust oversight mechanisms, particularly in cloud-to-edge deployments where decision-making occurs close to the physical world.
- **Compliance**: Unclear liability, data ownership, and weak consent protocols affect all scenarios using AI or distributed data flows. With the support of the CSA, Pilots could develop risk matrices, dataset datasheets to log the inputs to AI and 'grey zone logs' to document regulatory or legal ambiguities encountered during pilots. These artefacts can feed directly into CSA and policy feedback loops.

The table below includes some initial recommendations for each of the pilots. Given the complexity of the pilots, it is not possible to make recommendations to address all of the obstacles to market take up. Instead, we recommend that each pilot sets up a 3-month task force to select a single recommendation to implement by:

- Setting up systems that try to find workarounds and/or
- Presenting the challenge and inviting expert inputs on how to resolve the issue at the CEI Sphere meetings (CSA) and/or
- Inviting solutions to be presented through open calls.

Table 16: Initial recommendations for pilots

Pilot	Recommendations to increase market take up
Pilot 1: Electricity Grid performance optimisation upon RES integration	Scalability and interoperability of heterogeneous sources of energy data are two of the mayor problems. For **scalability**, the pilot leads could perform a comparative cost-benefit analysis across the 3 sites to identify where optimising the grid makes most economic sense. These costs can also be compared to the €400M direct economic cost of the recent blackout in Spain. For **interoperability**, testing data flows based the hourglass model could provide a basic recipe for the
Pilot 2: Transition to Software-Defined Vehicle (SDV) platform	**Cross border connectivity** is likely to be a significant obstacle to market take up. An open call may reveal novel ways of addressing the problem.
Pilot 3: Energy consumption and emission reductions in postal Service fleet operation via intelligent BEV charging strategies	There are **existing market solution**s to some scalability problems that could be explored compressing AI solutions on the market that can reduce model size by 20x. Open calls could focus on solutions to (i) **adapt the pilot tech stack** to different grid types and infrastructure levels or (ii) battery-related forecasts.
Pilot 4: Variable demand in challenging maritime terminal landscape	The Malta terminals are extremely resource constrained. It may be interesting to understand what other terminals

	across the world operate under similar conditions and create a **boarder business case to support exploitation**.
Pilot 5: Energetically and environmentally sustainable Halloumi cheese production	While there are technical challenges, the main obstacle is still likely to be human- moving from manual to automated farming. The pilot could therefore explore take up in other farming environments in which the availability of labour is constrained and the **incentive for automation** is greatest.
Pilot 6: Smart re-charging and efficiency of robot tractors in large fruit production fields	The lack of **stable access to energy or fast charging options** across European farms is a significant hurdle to overcome. An Open call might yield some novel insights to improve/solve this problem.
Pilot 7: Trustworthy and secure EV charging upon reliable 5G networks	As with most other pilots, one of the barriers is the lack of **common standards**. One approach might be to use the CEI Hourglass model to map the interaction points between stakeholders and ensure all charging infrastructure supports open standards (e.g., OCPP, IEC 61850, MQTT) for seamless integration across vendors and regions. This could be part of the final blueprint.
Pilot 8: Heightened social engagement and acceptability of energy flexibility in urban areas	Focusing on the **business case for a lead user** category (e.g., hotels) could help prove the value more effectively than a general case that applies to all end consumers. **Data governance** could also be discussed in the **Energy & Infrastructure Sphere** that is being created under the CEI Sphere CSA.

5 CONCLUSIONS

The work presented in this deliverable provides, on one hand, a technology-facing analysis that consolidates the technical expertise, validated research outcomes, and integrated solutions previously developed and demonstrated by consortium partners in recent EU-funded initiatives. On the other hand, it offers a market-facing perspective by identifying sectoral challenges and barriers to innovation adoption.

From a technical standpoint, it presents a comprehensive recording of the technological assets, experience, and validated tools contributed by the consortium partners, many of which originate from relevant recent or even ongoing EU-funded research. This allows the project to form a clear picture of the technological capacity already available, and how it can be recomposed into integrated CEI solutions. The outcomes of this mapping effort serve as the solid basis for subsequent technical work in WP2 and WP3:

- In WP2, the catalogued utilities and their corresponding capabilities will guide the design of pilot-specific blueprints, ensuring that each pilot is grounded in both sectoral needs and proven components.
- In WP3, the identification of utilities and architectural contributions informs the design of required adaptations, enhancements, and integrations necessary to implement the O-CEI platform and its continuum-aware functionalities.

The technical reporting is structured around four strategic technological pillars — *MetaOS & Orchestration*, *AI-Driven Computing*, *Security & Trust*, and *Data Spaces & Interoperability*. These pillars are built by composing EUCloudEdgeIoT (EUCEI) TF3 building blocks, illustrating how these foundational concepts can be synthesized into higher-level, cross-cutting technical areas. This layered approach ensures both compliance with EUCEI architectural direction and scalability across multiple verticals.

Looking ahead, this deliverable serves as a cornerstone for identity building within O-CEI. It articulates the problem space (sectoral needs, adoption barriers) and matches it with a solution space (technological capabilities, validated components, architectural patterns). In doing so, it creates a shared understanding among all stakeholders and lays a robust foundation for the design, implementation, and evaluation phases that follow.

Moreover, by consolidating partner contributions and surfacing their potential for integration, the document enhances the project's capacity to respond to market needs with coherent, federated, and trust-aware CEI solutions, reinforcing O-CEI's role as a strategic enabler of the next-generation European digital infrastructure.

From a market perspective, the analysis has been informed by pilot-specific workshops and sectoral consultations. These activities helped uncover behavioral, regulatory, and technical barriers that still hinder the wide adoption of CEI-enabling technologies in various market domains. By mapping these barriers against sectoral maturity and specific use case challenges, the deliverable identifies both risks and opportunities for CEI innovation deployment.

REFERENCES

[1] European Commission. How to master Europe's digital infrastructure needs? (2025). https://digital-strategy.ec.europa.eu/en/node/12451/printable/pdf

[2] Open Continuum / EUCloudEdgeIoT Task Force 3. (2024). Compositional View of the Continuum Reference Architecture: Graphical representation of common and potential capabilities. Zenodo. https://doi.org/10.5281/zenodo.11656784

[3] Task Force 3: Architecture. (2023). Developing a Reference Architecture for the Continuum - Concept, Taxonomy and Building Blocks. Zenodo. https://doi.org/10.5281/zenodo.8403593

[4] Open Continuum / EUCloudEdgeIoT Task Force 3. (2024). Functional View of the Continuum Reference Architecture: Minimum set of expected functionalities. Zenodo. https://doi.org/10.5281/zenodo.11656674

[5] Gambo, M. L., & Almulhem, A. (2025). Zero Trust Architecture: A systematic literature review. King Fahd University of Petroleum and Minerals. https://kfupm.edu.sa

[6] Liu, C., Tan, R., Wu, Y., Feng, Y., Jin, Z., Zhang, F., Liu, Y., & Liu, Q. (2024). Dissecting zero trust: research landscape and its implementation in IoT. *Cybersecurity,* 7(1), 20. https://doi.org/10.1186/s42400-024-00212-0.

[7] Mavroudis, V. (2024). Zero-Trust Network Access (ZTNA). Alan Turing Institute. https://arxiv.org/pdf/2410.20611

[8] Jun Ye; Wentao Zhao; Dong Wang. A Tool Design for SQL injection vulnerability detection based on improved crawler. Procedia Computer Science, Volume 247, 2024, Pages 1331-1339, ISSN 1877-0509, https://doi.org/10.1016/j.procs.2024.10.159

[9] Bhandari, G.P., Assres, G., Gavric, N. et al. IoTvulCode: AI-enabled vulnerability detection in software products designed for IoT applications. Int. J. Inf. Secur. 23, 2677–2690 (2024). https://doi.org/10.1007/s10207-024-00848-6

[10] P. Nunes, J. Fonseca and M. Vieira, "Blending Static and Dynamic Analysis for Web Application Vulnerability Detection: Methodology and Case Study," in IEEE Access, vol. 13, pp. 3139-3153, 2025, doi: 10.1109/ACCESS.2024.3522094

[11] Yang, W.; Wang, E.; Gui, Z.; Zhou, Y.; Wang, B.; Xie, W. An MLLM-Assisted Web Crawler Approach for Web Application Fuzzing. Appl. Sci. 2025, 15, 962. https://doi.org/10.3390/app15020962

[12] Moreira, D.; Seara, J.P.; Pavia, J.P.; Serrão, C. Intelligent Platform for Automating Vulnerability Detection in Web Applications. Electronics 2025, 14, 79. https://doi.org/10.3390/electronics14010079

[13] Sebilleau, D., Farre, M., & Genest, O. (2023). How to ensure interoperability in demand response systems: the examples of the European projects H2020 GIFT and MAESHA. 27th International Conference on Electricity Distribution (CIRED 2023), Rome, Italy, 1793-1797. https://doi.org/10.1049/icp.2023.1020

[14] Resonance Project. D2.3 Report on Eco-system Building and Co-creation Activities. https://www.resonance-project.eu/d2-3-report-on-eco-system-building-and-co-creation-activities

[15] Reif, V., Strasser, T. I., Jimeno, J., Farre, M., Genest, O., Gyrard, A., McGranaghan, M., Lipari, G., Schütz, J., Uslar, M., Vogel, S., Bytyqi, A., Dornmair, R., Corusa, A., Roy, G., Ponci, F., Dognini, A., & Monti, A. (2023). Towards an interoperability roadmap for the energy transition. https://arxiv.org/abs/2309.08284v1

[16] Strasser, T. I., Widl, E., Kuchenbuch, R. A., Lázaro-Elorriaga, L., Tellado Laraudogoitia, B., Ginocchi, M., Penthong, T., Ponci, F., Gyrard, A., Kung, A., Mac Gregor, C. A., Garcia Montero, C., & Relano Algaba, E. (2024). Towards interoperability testing of smart energy systems – An overview and discussion of possibilities. https://arxiv.org/pdf/2411.03369

[17] AIOTI. Joint deliverables with Int:net and Stand.ICT. https://aioti.eu/joint-deliverables-with-intnet-and-stand-ict

[18] Lambert, E., Maqueda, E., Valiño, J., Genest, O., Janev, V., Traverson, B., Lefrançois, M., Nachabe, L., Gyrard, A., & Kung, A. (2024). Omega-X: Energy data space for improving efficiency of electric power systems leveraging semantic interoperability and AI. CIGRE. https://www.e-cigre.org/publications/detail/d2-10344-2024-omega-x-energy-data-space-for-improving-efficiency-of-electric-power-systems-leveraging-semantic-interoperability-and-ai.html

[19] Brewster, C., Kalatzis, N., Nouwt, B., Kruiger, H., & Verhoosel, J. (2023). Data sharing in agricultural supply chains: Using semantics to enable sustainable food systems. Sustainability, 7(1), 3287. https://doi.org/10.3233/SW-233287

[20] Hossain Chy, T. M., Lamboro, H. M., Genest, O., Kung, A., Rabrait, C., Sebilleau, D., & Gyrard, A. (2025). Design of an ontology-driven constraint tester (ODCT) and application to SAREF and smart energy appliances. In Advances in Smart Energy Systems (pp. 183-198). Springer. https://doi.org/10.1007/978-3-031-81221-7_13

[21] OASIS. *ebXML Messaging Services Version 3.0: Part 1, Core Features*. OASIS Open. Available: https://groups.oasis-open.org/higherlogic/ws/public/document?document_id=24618

[22] Y. Li, J. Huang, Q. Sun, T. Sun and S. Wang, "Cognitive Service Architecture for 6G Core Network," in IEEE Transactions on Industrial Informatics, vol. 17, no. 10, pp. 7193-7203, Oct. 2021, doi: 10.1109/TII.2021.3063697

[23] Chebbi, A., Kniesel, G., Abdennadher, N., & Dimarzo, G. (2024). *Enhancing named entity recognition for agricultural commodity monitoring with large language models*. In *EuroMLSys '24: Proceedings of the 4th Workshop on Machine Learning and Systems* (pp. 208-213). https://doi.org/10.1145/3642970.3655846

[24] Gole, J. (2025). D2.1 – Preliminary market analysis. Zenodo. https://doi.org/10.5281/zenodo.15259914

[25] BluSpecs. (2024). D4.4 – Market developments learnings. Zenodo. https://doi.org/10.5281/zenodo.14761676

APPENDIX A: O-CEI GLOSSARY

Glossary establishes a shared vocabulary for terms that are frequently used in both technical and pilot contexts, to ensure clarity and alignment across the diverse teams involved in O-CEI. As many of these terms, such as *pilot*, *cross-domain*, *blueprint*, or *marketplace*, can carry different interpretations depending on perspective, this effort aims to prevent misunderstandings and foster a common understanding of key terms used and the meaning they convey. This glossary is the result of a collaborative process involving all partners, beginning with initial definitions shared during the kick-off meeting and iteratively refined through partner contributions. While it is formally included in Deliverable D2.1, the glossary will remain a living document throughout WP2 duration, continuously updated to reflect the evolving consensus within the consortium.

- **Large-Scale Pilot**: Real-world experimentation activities aimed at demonstrating the impact potential of O-CEI technologies. Their large-scale nature is evidenced by factors such as user base, data volume, number of devices or physical elements involved, population affected, grid coverage, or exploitation capacity.
- **Domain**: An area of knowledge or activity characterized by a shared set of concepts and terminology, as understood by practitioners [ISO/IEC 19506:2012]. In the context of O-CEI pilots, all domains are framed within the broader energy vertical and aligned with the SGAM Smart Grid Domains.
- **Vertical Domain**: A domain focused on a specific industry or sector, such as energy, mobility, or health. In O-CEI, all vertical domains build upon the energy system as a baseline—especially grid management, consumption, production, and flexibility.
- **Horizontal (Technology) Domain**: A domain centered on enabling technologies that span across verticals, such as Artificial Intelligence, Digital Twins, Data Spaces, Edge Computing, and Distributed Ledger Technologies (DLTs).
- **Cross-Domain**: Elements or functionalities originating in one domain that are leveraged by others across domains. These can include shared data exchange mechanisms, AI/ML models, software components, methodologies, or strategies. Cross-domain applicability encourages vertical and horizontal domains to reuse and work together (e.g. G. DLT application in data spaces pertaining to energy.
- **Energy Flexibility**: The modulation of electricity production, storage, or consumption at one or more sites, over a given period, in response to an external signal—such as a tariff, market signal, or grid request. Flexibility may involve shedding, shifting, or adjusting demand or supply, either manually or through automated systems. In the context of O-CEI, digital tools enable real-time, cross-sector flexibility in heterogeneous grid environments.
- **Value Chain**: The participation of all required technical operational and digital entities in concert to deliver a particular service or use case. In O-CEI, the value chain spans both the energy domain (e.g., DSO, energy providers) and the digital service domain (e.g., software developers, AI experts, IoT manufacturers, data providers), enabling end-to-end deployment of energy flexibility solutions.
- **Open Platform**: The O-CEI Open Platform is a modular digital infrastructure designed to overcome barriers to cross-domain data sharing, resource utilization, and service provisioning. It facilitates smooth integration across various domains by building upon open standards APIs and interoperable data models. As illustrated in DoA technical diagram, the O-CEI open platform comprises several key building blocks. These include a management portal, Open API for unified external access, and integrated trust and security mechanisms and the marketplace. As open, the platform allows users to share and co-develop resources, and to extend its capabilities with

third-party applications. Such openness is sustainably achieved through governance under an open-source model.

- **Blueprint**: Intentful declarations of an operative target in (one or more) pilots that will have a materialization in terms of O-CEI technological deployment. Blueprints consist of instructions specifically designed to guide a technical rollout of Cloud-Edge-IoT technologies for achieving such target, considering technical elements, strategic conditions, decisions to make, and standards and expert references where they inspire from. Blueprints are delivered in a simple and concise language via a short document aiming at replication across pilots and domains.
- **Marketplace**: A central hub of the O-CEI open platform, which will be designed to introduce and establish collaboration and sharing within ecosystem. Marketplace within O-CEI is the place to host and share catalogues of Blueprints, Utilities and Data assets. Actors of the continuum can upload, discover and utilize blueprints and all accompanying diverse resources, which enable consistent and efficient deployment of the services provided. Trustworthiness and security of all hosted assets are guaranteed through the O-CEI Open Platform framework. O-CEI will align its efforts with the DOME initiative and will base the federated Marketplace implementation on its specifications with the aim hosted resources to be later part of a bigger community and be shared beyond the project's scope.
- **Utility/ies**: Software components, either included in or referenced by blueprints, that are the actual assets deployed to implement domain-specific or agnostic solutions. O-CEI utilities are technically aligned with open specifications and operate by producing and consuming data through standardized ontologies and acting/monitoring as necessary. They adhere to and build upon EUCEI building blocks.
- **Services**: In the context of O-CEI, Services refer to the technical and data-oriented functionalities provided by the Open Platform. These encompass core provisions such as traceability, security, and trust, as well as domain-specific functions tailored to operational needs. Services are built on standards and interoperable data formats in order to benefit from and promote collaboration within the ecosystem. They support both backend operations and end-user applications, enabling cross-domain collaboration and integration across the ecosystem.
- **Resources**: Are the foundational supplies required from services to accomplish their tasks. Resources include essential assets such as data, compute, storage, networking, infrastructure, and energy, all of which are employed by services to fulfill their intended operations.
- **Data Sharing**: Data sharing refers to the process of making data retrievable and accessible across organizations, sectors, or domains, by leveraging open standards, ontologies, and APIs. The process is designed to be interoperable, transparent, and adaptable to diverse requirements. This capability is supported by the data catalogues hosted within the O-CEI platform, alongside the critical role of utilities and services hosted in the open platform. These utilities and services enable the production, consumption, translation, annotation, and transformation of data and might provide connections to wider data spaces services.
- **Data Sovereignty Mechanism**: Data sovereignty refers to the framework and mechanisms that O-CEI will provide to enable data owners to have control when sharing their data. They will be able to define who can access their data, under what conditions, for how long, for what purposes and also ensure confidentiality, integrity.
- **Data Provenance**: Refers to the framework and mechanisms that O-CEI will provide to track the entire lifecycle of data, including its origination, transformations, processing, storage, and flow across domains all over the continuum and through various systems, applications, and processes. Its functionalities will be offered to actors of O-CEI open platform by utilities designed for this purpose. These utilities will implement functionalities which are part of the Monitoring and Observability building block of EUCEI RAM.

- **Trust and Trustworthiness**: Trust is an attitude that a tenant has towards a system. In contrast, trustworthiness is a system property that creates trust to the tenant/user. A user/tenant trusts (or requires a specific level of trust from) a system, because the system is trustworthy. In other words, the trustworthiness of a system contributes to building the trust level of the tenant/user of the specific system. Thus, the more trustworthy the system is, the higher the trust level of the tenant/user will be.
- **Decentralized Intelligence**: Decentralized intelligence, often referred to as Decentralized Artificial Intelligence (DAI), is a subfield of artificial intelligence that focuses on developing distributed solutions for complex problems. It involves multiple autonomous processing nodes, or agents, that work together to solve reasoning, planning, learning, and perception problems. These systems are designed to be robust and capable of functioning even with unreliable resources.
- **CEI Programming System**: Cloud-Edge-IoT programming refers to a programming system that allows to effectively program applications that are combining resources from the cloud-edge-IoT continuum. CEI programs can be distributed in the CEI continuum. Basic abstractions are centered around an abstract programming concept for the CEI continuum (e.g. data streams, workflow orchestrations, process choreography) as well as the mapping to concrete execution systems which are then orchestrated across the CEI continuum based on explicitly formulated application intents.
- **Cloud Computing**: The practice of using a network of remote servers hosted on the internet to store, manage, and process data, rather than a local server or a personal computer for the delivery of computing services over the internet.
- **Edge Computing**: Distributed computing model that brings computation and data storage closer to the sources of data, allowing devices in remote locations to process data at the "edge" of the network, either by the device or a local server.
- **Internet of Things**: The Internet of Things refers to the collective network of connected devices and the technology that facilitates communication between devices and the cloud, as well as between the devices themselves.
- **Software Defined Vehicle**: Vehicle as Software (VaS) is the next level of development according to the data space concept. There, a vehicle becomes an active Infrastructure Element (as per aerOS terminology), with an active role in: (i) computation, (ii) data exchange, (iii) connectivity with other elements to consume and provide context information.
- **Interoperability**: Interoperability refers to the capability of systems, platforms, and devices to seamlessly exchange, interpret, and utilize data & services across different organizations, sectors, and technologies. It relies on standardized protocols, open specifications, and regulatory frameworks (e.g., IEC 61850, CIM) to ensure compatibility between diverse industrial ecosystems.
- **Battery-to-Grid**: The ability of different batteries, networks, and systems to seamlessly communicate and function together, ensuring compatibility across platforms.
- **Vehicle-to-Battery**: A technology that allows electric vehicles (EVs) to transfer energy stored in its batteries to external batteries.
- **Vehicle-to-Grid**: A charging technology that enables batteries to discharge electricity back into the power grid, assisting with energy balancing and grid efficiency.
- **Federated Learning**: A decentralized approach to machine learning where multiple devices or servers work together to train a shared model without exchanging raw data. Each client trains a local model on its own private data and only sends updates to a central server. The server then combines these updates to improve the global model. This way, data stays private,

communication overhead is lower, and models can still learn from diverse datasets while following data protection rules and regulations.

- **Bidirectional Equipment**: Devices and systems capable of transferring energy in both directions, allowing for the flow of power between sources and loads as needed. These are essential in modern energy grids, particularly in renewable energy integration, energy storage, and smart grid applications.
- **Battery Management System**: Proprietary technology that monitors the charge level and state of rechargeable batteries, ensuring safety during charging/discharging and regulating temperature.
- **Peak Loads**: Times of highest electricity demand, often short-term spikes occurring during specific periods due to increased energy usage.
- **Battery Energy Storage System**: A system storing electrical energy in batteries for later use. The BESS helps manage energy supply-demand imbalances, store surplus energy, and provide power during peak periods.
- **Megastorage**: Large-scale energy storage solutions capable of storing substantial amounts of electrical energy. These systems offer significant capacity for storing and delivering power when needed.
- **Low Voltage**: Refers to relatively small electrical potential typically ranging up to 1,000 volts.
- **Medium Voltage**: Spans between 1,000 volts (1 kV) and 72,000 volts (72 kV).
- **Smart Readiness Indicator:** SRI is a Commission initiative under the Energy Performance of Buildings Directive that measures a building's ability to use smart technologies.
- **Vehicle-to-Vehicle Communications:** V2V communications allow Vehicles to be continuously aware of each other.

APPENDIX B: PILOT WORKSHOP METHODOLOGY

As part of Task 2.2 "Market Analysis and Adoption Barriers Identification", a structured approach was adopted to gather insights from stakeholders across the eight pilots. The objective was to develop an initial understanding of the different pilot scenarios, the market positioning of each use case, the relevance of the solutions proposed, and the key barriers that could affect their adoption. To achieve this, two complementary methods were used: desk research and targeted stakeholder workshops.

- Desk Research and Secondary Data Sources: an initial phase of desk research was conducted to collect background information and context for each pilot. This activity included the review of available project documentation and secondary data sources related to the domains of application of each pilot. The aim was to build a preliminary understanding of the existing solutions on the market, key industry challenges, and initial stakeholder groups. The results of this research were used to inform the design and focus areas of the subsequent workshops.
- One-to-one workshops with pilots: following the desk research phase, a series of eight online workshops were organised to collect direct input from the pilot leads and scenario leaders. These workshops were designed to complement the desk research findings with qualitative insights from those directly involved in the pilots' execution. The specific objectives of the workshops were to:
    - Clarify the purpose of each pilot and ensure a shared understanding of the goals.
    - Identify the key challenges, gaps, and stakeholder needs that the pilots aim to address.
    - Map the stakeholders involved and their roles within the ecosystem.
    - Explore existing market solutions and assess how the pilot differentiates itself.
    - Identify potential barriers to adoption and their implications for the business model.

Each workshop lasted approximately 90 minutes and was conducted online using the collaborative platform Miro. The structured agenda included:

- Welcome and Instructions: Introduction to the session objectives and tools.
- Pilot Purpose and Stakeholder Identification: Summarising pilot use cases and identifying all relevant stakeholders.
- Scenario – Problem Definition: Defining the main problems faced by stakeholders and the obstacles to adoption.
- Current Solutions Mapping: Identifying technical and non-technical existing solutions in the market.
- Stakeholders and Actors Mapping: Categorising stakeholders and partners using the hourglass model to visualise relationships and roles.

The outputs of the workshops included:

- Stakeholder maps for each pilot, highlighting roles and interactions.
- An overview of the key industry challenges and innovation gaps.
- An initial analysis of competing solutions and pilot differentiators.
- Early identification of adoption barriers that could impact future business models.

These findings will directly contribute to the preliminary market analysis provided in Deliverable 2.1 and will serve as foundational inputs for the business case development activities foreseen under Task 6.1.

APPENDIX C: IMPACT OF KEY EU REGULATION ON CEI PILOTS

The Cloud-Edge-IoT Large Scale Pilots (LSPs), O-CEI and COP-PILOT, are designed to develop interoperable and sustainable platforms that drive digital transformation across various European sectors. These pilots support the EU's strategy to enhance technological sovereignty and foster data-driven business models. To ensure scalability and legal compliance, the pilots must align with key EU regulations that shape the design, deployment, and operation of CEI systems. The most relevant regulations include the AI Act, Data Act, Data Governance Act, GDPR, Cyber Resilience Act, and NIS2 Directive. This paper provides an overview of these regulations, examining their impact on the LSPs' architecture, operation, and business models.

AI Act

Application dates: Prohibited AI practices February 2, 2025, General obligations including High-Risk AI Systems (Annex III) August 2, 2026, High-risk AI systems (Annex I) August 2, 2027.

The AI Act lays down harmonised rules for the placing on the market, the putting into service, and the use of AI systems in the Union. It aims to mitigate risks associated with AI while fostering innovation. The regulation follows a risk-based approach, categorizing AI systems into four risk levels: unacceptable, high, limited, and minimal. The focus is on high-risk systems, which must meet strict requirements such as risk management, data quality assurance, transparency and human oversight, and conformity assessment. Main actors under the AI Act are providers and deployers. Providers are entities that develop and place AI systems on the market under their own name or brand. They bear the most obligations. Deployers are those who integrate and operate AI systems in real-world environments, being responsible for ensuring proper use and ongoing monitoring. Transparency obligations apply across all risk levels. The role of standards is crucial, as meeting harmonised standards leads to a presumption of conformity, thereby reducing compliance efforts.

In the CEI context, it is essential to determine the roles of involved parties accurately. Depending on the setup, a cloud service provider offering AI-based data analysis tools could be classified as an AI provider, while an industrial operator using predictive maintenance AI at the edge would be considered a deployer. High-risk AI systems are particularly relevant for LSPs since they require compliance with risk management, data governance, robustness, and cybersecurity requirements etc. These systems primarily include those regulated under EU safety regulations (Annex I), such as AI components used in medical devices, machinery, or automotive safety features. Additionally, high-risk systems include applications used in critical infrastructure, such as energy management, as well as education, employment, and public safety scenarios (Annex III). To ensure compliance, LSPs must clearly define the roles of providers and deployers within their CEI ecosystems and identify potential high-risk AI applications early in the project lifecycle.

Data Act

Application dates: General obligations including data access and sharing September 12, 2025, Accessibility by Design September 12, 2026.

The Data Act establish lays down rules for the placing on the market, the putting into service, and the use of connected products (examples) in the Union. It addresses the challenge that much of the data generated by these products remains underutilised due to restricted access rights. The regulation imposes obligations on data holders, typically manufacturers, to provide users with access to data generated by their connected products. On request, the data holder must make the data continuously, in real-time, and free of charge available to the user. Additionally, the user has the right to instruct the data holder to provide the data to a third party, such as a service provider or an aftermarket operator. The data holder must ensure that this transfer takes place without unreasonable delay and under fair, reasonable, and non-discriminatory (FRAND) conditions. The

Data Act sets the contractual framework for such data-sharing agreements, ensuring that the data holder cannot impose unfair conditions or restrict third-party access arbitrarily. The regulation also addresses cloud switching by making it easier to move data and services between providers. It reduces contractual barriers and mandates that switching costs should be minimal, often limited to the actual expenses incurred. To further support flexibility, cloud service providers must ensure that data migration is technically feasible, and that interoperability is maintained.

In the CEI context, it is crucial to identify the roles of involved parties accurately. Typically, the manufacturer of a connected device is the data holder, while the user could be an industrial operator or a consumer owning or leasing the device. Third parties could be service providers, maintenance companies, or other actors that require data access to offer their services. Clear contractual arrangements between these actors are necessary to ensure that data access and sharing comply with the Data Act. High-impact areas for LSPs include data access and sharing obligations, contractual requirements, protection of trade secrets, and cloud switching provisions. For example, an industrial IoT platform that collects data from factory machinery must grant users access to the data and enable sharing with third-party maintenance providers. Compliance also means ensuring that users can switch between cloud providers without significant disruption. To meet these requirements, LSPs should proactively define the roles of data holder, users, and data recipients and establish contractual frameworks that reflect FRAND principles.

Data Governance Act

Application dates: General obligations September 24, 2023. Existing data intermediation services must comply by September 24, 2025.

The Data Governance Act (DGA) establishes harmonised rules for data sharing and governance within the EU. It primarily addresses three areas: access to sensitive public sector data, regulation of data intermediation services, and promotion of data altruism. The DGA mandates that public sector bodies ensure secure access to sensitive data, including personal data (anonymised) or business secrets (aggregated). Exclusive data licensing is prohibited to maintain competitive access. Data made available must be processed in a secure environment, and public entities can charge reasonable fees for access. A key element of the DGA is the regulation of data intermediation services. These services act as neutral intermediaries between data holders and data users, enabling trusted data exchanges. To maintain neutrality, data intermediation services are prohibited from monetising the data they facilitate. They cannot use the data for product development or resale, nor can they offer data-based services directly. The DGA requires that these services be legally and structurally separated from any other data-related business activities, and that they are registered with the competent authorities before operating.

In the CEI context, it is crucial to identify which platforms qualify as data intermediation services. Platforms that merely aggregate data for resale or provide closed data environments are not considered intermediation services. Instead, data marketplaces that facilitate open data exchanges fall under the DGA's regulations. For LSPs, the main challenges involve establishing compliant data intermediation structures and ensuring the neutrality of data sharing mechanisms. For example, an IoT data marketplace offering industrial data to multiple users must clearly separate the data exchange from any analytics services to maintain compliance. Ensuring that data governance aligns with the DGA's neutrality requirements is essential for building trust and avoiding conflicts of interest.

General Data Protection Regulation

Application dates: In force since May 25, 2018.

The GDPR is the primary legal framework governing the processing of personal data within the EU. It aims to protect the fundamental rights and freedoms of individuals by setting strict rules for data collection, processing, storage, and transfer. A key concept of the GDPR is the designation of the

data controller, who determines the purposes and means of data processing, and the data processor, who processes data on behalf of the controller. In Cloud-Edge scenarios, this distinction can be challenging, as multiple actors—such as device manufacturers, cloud service providers, and users—may jointly determine data processing. In such cases, they are considered joint controllers and must establish clear agreements on data protection responsibilities. Cross-border data transfers outside the EU/EEA are subject to stringent requirements under Articles 44-49 GDPR, including the use of Standard Contractual Clauses (SCCs) or adequacy decisions to ensure equivalent data protection standards. The GDPR also mandates data protection by design and by default (Article 25), requiring technical and organizational measures (TOM) to minimize data processing risks. In the CEI context, this includes encryption, pseudonymization, and secure data transmission between edge devices and cloud services to maintain compliance. For LSPs, the main challenge is to accurately identify data controllers and processors in complex, distributed systems and to implement robust data protection strategies that align with GDPR requirements.

Cyber Resilience Act

Application dates: Expected to enter into force in autumn 2027, with a three-year implementation period.

The Cyber Resilience Act (CRA) establishes binding IT security requirements for products with digital elements within the EU. It aims to enhance the cybersecurity of both hardware and software products that rely on remote data processing to function. The regulation applies to manufacturers, importers, distributors, and authorized representatives who place products on the EU market. The CRA covers all connected products, including hardware devices and software applications that communicate through remote data processing solutions, such as APIs, databases, or wireless communication interfaces. Manufacturers bear the primary responsibility for ensuring cybersecurity throughout the product lifecycle. This includes risk assessments, vulnerability management, secure configurations, and regular updates. Depending on the criticality of the product, additional requirements such as conformity assessment and certification may apply. The CRA distinguishes between non-critical, important (Class I and II), and critical products, each with different obligations, including self-certification or third-party assessment. Harmonized standards play a central role under the CRA. Products meeting these standards are presumed to comply with the regulation, reducing the burden on manufacturers. Compliance includes implementing security-by-design principles, maintaining a SBOM, and providing clear information about updates and security measures.

For LSPs operating in CEI contexts, accurately identifying the role of each actor (e.g., manufacturer, operator, service provider) is crucial. Devices with remote data processing capabilities must comply with CRA security requirements, including robustness against cyber threats and the ability to update securely. Compliance strategies should include assessing whether harmonized standards are applicable and ensuring that cybersecurity measures are integrated into product development from the outset.

NIS2 Directive

Application date: October 17, 2024.

The NIS2 Directive aims to enhance cybersecurity in the EU by establishing uniform requirements for network and information system security. It applies to medium and large companies in critical sectors like energy, healthcare, and digital infrastructure, while excluding micro and small enterprises. The directive categorizes entities as essential or important, with essential entities facing stricter requirements. They must both implement measures to manage cybersecurity risks, report incidents within 24 hours and submit a detailed follow-up report within 72 hours. To demonstrate compliance, entities are encouraged to adopt harmonized standards, such as ISO/IEC 27001. For LSPs, the NIS2 Directive is significant as it covers digital infrastructure providers and Managed Service Providers (MSPs), including logistics data management platforms. If an LSP offers services like data storage, processing, or operational technology integration, it may be classified as an

essential or important entity. LSPs must also ensure that cybersecurity responsibilities are clearly assigned within the organization, especially if they manage critical logistics data or supply chain information. Preparing for NIS2 means assessing existing security measures, updating contracts with third-party providers, and training staff to meet new reporting and accountability standards.