



**Large-Scale Testing and Experimentation Facility (TEF) for Assessing,
Validating, and Enhancing AI-Powered Next-Generation Energy Solutions**

**D2.2 - EnergyGuard Architecture in alignment with
the European ecosystem**

By ICCS

Athens, August, 31, 2025

Work Package: WP2 – Requirements, Specifications, Synergies and Alignment with
European Ecosystem

Authors: ICCS, LXP, CARTIF

Status: Final

Date: 31/08/2025

Version: 1.0

Classification: Public

Disclaimer:

The EnergyGuard project is co-funded by the Horizon Europe Programme of the European Union under grant agreement No. 101172705. This document reflects only authors' views. The EC is not liable for any use that may be done of the information contained therein.

EnergyGuard Project Profile

Grant Agreement No.: 101172705

Acronym:	EnergyGuard
Title:	Large-Scale Testing and Experimentation Facility (TEF) for Assessing, Validating, and Enhancing AI-Powered Next-Generation Energy Solutions
URL:	http://www.energy-guard.eu
Start Date:	01/01/2025
Duration:	36 months

Partners

	EREVNITIKO PANEPISTIMIAKO INSTITOUTO SYSTIMATON EPIKOINONION KAI YPOLOGISTON (ICCS), Decision Support Systems Laboratory, DSSLab <u>Co-ordinator</u>	Greece
	CENTRO DE INVESTIGACAO EM ENERGIA REN - STATE GRID SA (RDN)	Portugal
	CENTRO DE INVESTIGACIONES ENERGETICAS MEDIOAMBIENTALES Y TECNOLOGICAS (CIEMAT)	Spain
	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA)	France
	FUNDACION CARTIF (CARTIF)	Spain
	RIGA MUNICIPAL AGENCY "RIGA ENERGY AGENCY" (REA)	Latvia
	BLUENERGY REVOLUTION SCRL (BER)	Italy
	METASYMBIOSE SAS (MS)	France
	LUXPROVIDE SA (LXP)	Luxembourg
	FUNDINGBOX ACCELERATOR SP ZOO (FBA)	Spain
	FORA FOREST TECHNOLOGIES SLL (FORA)	Spain
	FYLLA ANASTASIA (AF)	Greece
	ENGREEN SRL (ENGREEN)	Italy
	BUVNIECIBAS VALSTS KONTROLES BIROJS (SCCB)	Latvia
	HAL SERVICE SPA (HAL)	Italy

Document History

Version	Date	Author (Partner)	Remarks
0.1	21/05/2025	Evangelos Karakolis (ICCS)	Initial ToC
0.2	10/07/2025	Evangelos Karakolis (ICCS)	Sections 2 and 3
0.3	14/07/2025	Frederic Pinel (LXP)	Section 4.4
0.4	15/07/2025	José L. Hernández García (CARTIF)	Section 4.1
0.5	20/07/2025	Evangelos Karakolis, Alexandros Tzortzis, Theodosios Pountridis, Panagiota Rempi (ICCS)	Section 5
0.6	23/07/2025	Sotiris Pelekis, Theodosios Pountridis (ICCS)	Sections 4.2 and 4.3
0.7	21/08/2025	Evangelos Karakolis (ICCS)	Sections 1 and 6
0.8	29/08/2025	José L. Hernández García, Francisco Javier Miguel (CARTIF)	Internal review
0.8.1	29/08/2025	Frederic Pinel (LXP)	Internal review
0.9	30/08/2025	Evangelos Karakolis (ICCS)	Addressing review comments
1.00	31/08/2025	Spiros Mouzakis (ICCS)	Final Version

Executive Summary

This deliverable presents the first version of EnergyGuard’s conceptual architecture reporting the work that took place in Task 2.3 entitled “Architectural design of the EnergyGuard TEF”, building on the foundations that have been set up in Tasks 2.1 and 2.2 and the outcomes of the deliverable D2.1 entitled “Landscape review analysis, EnergyGuard User requirement and specifications”. It presents the conceptual architecture of EnergyGuard based on the requirements and specifications gathered from different stakeholders from the Energy sector as well as AI experts and energy service providers. Furthermore, and beyond the overall project architecture this deliverable also presents separate conceptual architectures for the different components of the platform as well as the alignment of the current architecture with existing initiatives and project architectures. To do that existing relevant initiatives are presented in a separate chapter in this document, along with their relevance and alignment.

More specifically, the deliverable at hand presents the most important and relevant existing architecture models and their alignment with EnergyGuard. These models include GAIA-X, IDS-RAM and I-ENERGY among others. As a next step, it presents the overall conceptual architecture along with a detailed explanation of each component. Furthermore, the individual components for Data Management, AI development and testing sandbox, Acceptance environment, and High-Performance Computing are presented in detail along with the requirements they aim to address. Last but not least, the alignment with the European ecosystem is discussed in detail covering also integration aspects with other initiatives and projects including European Digital Innovation Hubs, Data Spaces, AIoD platform, other TEFs, AI Factories and the DOME Marketplace.

The outcomes of this deliverable will be used to develop the first versions of EnergyGuard components under the scope of WP3 and WP4, while the updated and final conceptual architecture of EnergyGuard will be provided in the deliverable D2.4, based on the updated requirements and specifications that might arise as the project moves forward.

Table of Contents

1	Introduction	10
1.1	Purpose and Scope.....	10
1.2	Structure of the document.....	10
2	Landscape analysis and existing architecture models	11
2.1	I-ENERGY	11
2.1.1	Alignment with EnergyGuard	11
2.2	GAIA-X	12
2.2.1	Alignment with EnergyGuard	13
2.3	IDS-RAM	14
2.3.1	Alignment with EnergyGuard	14
2.4	CEEDS.....	16
2.4.1	Alignment with EnergyGuard	16
3	EnergyGuard Conceptual Architecture.....	18
4	EnergyGuard components architectural overview	21
4.1	Data Management.....	21
4.2	AI Development and testing sandbox	24
4.2.1	Energy Analytics applications and DTs.....	25
4.2.1.1	<i>Electricity grid DT.....</i>	<i>25</i>
4.2.1.2	<i>Microgrid with DER DT</i>	<i>25</i>
4.2.1.3	<i>Buildings DT</i>	<i>26</i>
4.2.1.4	<i>Hydrogen testing DT</i>	<i>26</i>
4.2.1.5	<i>REC DT</i>	<i>26</i>
4.2.1.6	<i>DeepTSF</i>	<i>26</i>
4.2.2	AI development toolbox	26
4.2.2.1	<i>AI model development</i>	<i>26</i>
4.2.2.2	<i>Model deployment and inference endpoints</i>	<i>27</i>
4.2.2.3	<i>Experiment monitoring</i>	<i>27</i>
4.2.2.4	<i>AI Workflow orchestration</i>	<i>27</i>
4.2.2.5	<i>AI model repository and benchmarks</i>	<i>28</i>
4.3	Trustworthy AI Acceptance Environment.....	28
4.3.1	Purpose and Position in EnergyGuard Architecture	28
4.3.2	High-level Functionalities	28
4.3.3	Requirements	30
4.3.3.1	<i>Compliance and Trustworthiness Evaluation</i>	<i>30</i>

4.3.3.2	<i>Data Scientists – Continuous Monitoring for Model Integrity</i>	<i>30</i>
4.3.3.3	<i>Solution Developers – Safety, Security, and Compliance Scores ..</i>	<i>30</i>
4.3.3.4	<i>Polymakers – Aggregated and Anonymized Insights.....</i>	<i>31</i>
4.4	HPC	31
4.4.1	Architectural Diagram.....	31
4.4.2	Description	32
4.4.3	Requirements	33
5	Alignment with the European ecosystem.....	35
5.1	AI on Demand	35
5.2	Data spaces.....	35
5.3	European Digital Innovation Hubs (EDIHs).....	36
5.4	DOM Marketpale.....	37
5.5	Other TEFs.....	37
5.6	AI Factories.....	38
6	Conclusion	39
	Annex I: References.....	40

List of Figures

Figure 1 - I-ENERGY Reference Architecture	12
Figure 2 - GAIA-X Proposed trust framework	14
Figure 3 - IDSA Reference Architecture Model.....	15
Figure 4 - CEEDS Reference Architecture	17
Figure 5 - EnergyGuard Conceptual Architecture	18
Figure 6: Data management components architecture.....	22
Figure 7 - AI Sandbox architectural diagram	25
Figure 8: Conceptual diagram of the EnergyGuard Trustworthy AI Acceptance Environment.....	29
Figure 9 - HPC Architecture (Meluxina)	31

Acronyms and Abbreviation

Terms	Definition
AI	Artificial Intelligence
AIoD	AI on Demand platform
API	Application Programming Interface
DIH	Digital Innovation Hubs
EDIH	European Digital Innovation Hubs
HPC	High Performance Computing
DOME	Distributed Open Marketplace for Europe
DB	Data Base
TEF	Testing and Experimentation Facility
DER	Distributed Energy Resources
REC	Renewable Energy Certificate
IDS	International Data Spaces
IDSA	International Data Spaces Association
IDS-RAM	International Data Spaces Reference Architecture Model
ML	Machine Learning
DL	Deep Learning
EU	European Union
DSO	Distribution System Operator
VM	Virtual Machine
VPN	Virtual Private Network
TSO	Transmission System Operator
DT	Digital Twin
CEEDS	Common European Energy Data Space
GDPR	General Data Protection Regulation

GPU	Graphics Processing Unit
CPU	Central Processing Unit
FPGA	Field-Programmable Gate Array
SSH	Secure Socket Shell
ETL	Extract Transform Load
JDBC	Java Data Base Connectivity
SME	Small Medium Enterprise

1 Introduction

1.1 Purpose and Scope

This deliverable relates to Task 2.3 – Architectural design of EnergyGuard TEF and Task 2.4 – Interoperability and alignment specifications with the European Ecosystem (AIOD, DIHs, Data Spaces, HPC, DOME, AI Factories and other TEFs). The focus of this deliverable is to present the initial version of the EnergyGuard platform’s architecture and the strategic cross-fertilisation and alignment efforts between EnergyGuard project and the European Ecosystem including AIoD (AI on Demand), Digital Innovation Hubs (DIHs), Data Spaces, DOME, AI Factories, other TEF projects and High-Performance Computing (HPC).

The EnergyGuard architecture consists of several groups of services, related to data management, AI development and testing, an acceptance environment in terms of trustworthiness, compliance and security and an HPC component offering the resources for high performance AI model training and serving. All these services and components are protected through an Access Control and Identity management module that grants access only to users with the correct permissions. The provided services and components will be available through a user-friendly dashboard that will provide a smooth user experience to the users to get the most out of EnergyGuard platform based on their business needs.

1.2 Structure of the document

This document is organised in six main sections as described below:

- Section 1 provides the introduction in terms of the purpose, scope, structure, and relation to other tasks from the project.
- Section 2 presents a comprehensive landscape analysis and a number of existing relevant reference architecture models that have been taken into consideration when designing the EnergyGuard Conceptual Architecture
- Section 3 provides a concise description of the initial version of the architecture of the EnergyGuard platform.
- Section 4 presents the individual components of the platform with a focus on the requirements they aim to address.
- Section 5 illustrates the alignment with the European ecosystem.

The document concludes in section 6 with an outlook of the deliverable.

2 Landscape analysis and existing architecture models

In this section, we present several proposed architectures from other projects and initiatives that have been used as a reference for the design of the EnergyGuard conceptual architecture. Specifically, in this chapter we present an architectural overview for I-ENERGY, GAIA-X, IDS-RAM and SEEDS. Each one of the presented architectures poses their own objectives, which are presented in brief, along with the alignment with the EnergyGuard conceptual architecture.

2.1 I-ENERGY

I-ENERGY (Artificial Intelligence for Next Generation Energy) [1], [2] is a Horizon 2020 project launched by the European Commission with the aim of accelerating AI adoption in the energy sector. Running from 2021 to 2024, I-ENERGY focused on building a data-driven innovation ecosystem that connects energy stakeholders with AI solution providers, while fostering trust, transparency, and scalability. The mission of the project was to deliver AI-based energy services, built upon federated, secure, and interoperable data infrastructures aligned with EU digital policies.

At its core, I-ENERGY developed a suite of AI analytics tools and services for areas such as predictive maintenance, energy forecasting, optimization of distributed energy resources (DERs), and smart building management. It also emphasized the creation of standardized, interoperable data frameworks, supported by FAIR principles (Findable, Accessible, Interoperable, Reusable), and used ontologies to ensure semantic consistency across datasets and domains.

2.1.1 Alignment with EnergyGuard

I-ENERGY operated under a multi-actor innovation model, involving technology providers, data owners (e.g., DSOs, building operators), and end users through open calls, pilot projects, and co-creation processes. Importantly, it promoted AI transparency and trust, incorporating tools for explainability, risk assessment, and ethical evaluation—core requirements in AI experimentation and deployment. EnergyGuard aims to provide a trustworthy, high-performance platform for developing and validating digital solutions across the energy sector. I-ENERGY serves as a pivotal precursor, offering reusable components, methodologies, and lessons that can directly enhance EnergyGuard structure and purpose. Specifically I-ENERGY aligns with the following objectives of EnergyGuard:

Reusable AI Toolchains and Workflows

I-ENERGY developed a modular AI toolbox, including pipelines for time series forecasting [3], anomaly detection [4], and DER optimization. These can be integrated into the TEF's AI Sandbox component, reducing development time and enabling users to test and benchmark models under real or simulated conditions.

Data Interoperability and Semantic Data Models

I-ENERGY implemented semantic enrichment mechanisms using ontologies and data harmonization techniques. EnergyGuard's data management and querying layers can benefit from these assets to ensure seamless integration of heterogeneous energy data and improved experiment reproducibility.

User-Centric Ecosystem Design

I-ENERGY's innovation framework focused on multi-stakeholder engagement via open calls and use-case-driven pilots as well as the utilization of AI4EU platform. This approach aligns with TEF's vision to

act as a collaborative, open experimentation environment where users can submit experiments, access shared datasets, and contribute models, as well as the close collaboration and integration with AIoD platform (descendant of AI4EU).

Trustworthiness and Ethics-by-Design

Through its AI trust mechanisms (e.g., fairness metrics, risk assessment), I-ENERGY helps inform the EnergyGuard's acceptance environment, where models are validated for robustness, compliance, and ethical alignment before deployment. This ensures alignment with EU AI Act and energy sector regulations.

In Figure 1 an overview of I-ENERGY reference architecture [5] is presented:




Figure 1 - I-ENERGY Reference Architecture

2.2 GAIA-X

GAIA-X [6] is a European initiative aimed at establishing a sovereign, federated, and trustworthy data infrastructure that enables secure, interoperable, and standards-based data sharing across industries and borders. Founded by European governments, research institutions, and private sector leaders, GAIA-X provides an architectural framework and governance model that promotes data sovereignty, transparency, and trust—core values essential to the EU's digital transformation agenda.

GAIA-X is not a cloud provider itself but rather a federation model that integrates existing cloud and edge providers under a common framework. This includes unified standards for identity management, data usage control, security certification, and service interoperability. Key technical components of GAIA-X include the Federated Catalogue, Self-Descriptions, Identity and Trust Services, Compliance

Frameworks, and Usage Policies—all operating across domains such as energy, mobility, health, finance, and manufacturing.

2.2.1 Alignment with EnergyGuard

EnergyGuard requires a robust, compliant, and interoperable architecture to support experimentation at scale. GAIA-X offers foundational building blocks that directly enhance EnergyGuard’s mission in the following ways:

Data Sovereignty and Usage Control

Among EnergyGuard’s core goals is to enable stakeholders to share and access energy-related data—e.g., from grid operators, smart buildings, DER systems, and digital twins—while retaining control over how that data is used. GAIA-X enforces policy-based data access, ensuring that all data flows within EnergyGuard are transparent, auditable, and permitted, in alignment with GDPR and sector-specific regulations.

Federated and Interoperable Architecture

EnergyGuard is a distributed experimentation platform, potentially spanning multiple sites, countries, and infrastructure providers. GAIA-X’s federated service model allows EnergyGuard to integrate cloud, edge, and on-premise components under a unified policy and identity framework. This supports seamless orchestration of workloads (e.g., AI training, simulations) while maintaining trust between different nodes.

Trusted Data Sharing and Identity Management

EnergyGuard’s architecture includes data lakes, access-controlled APIs, and user dashboards that must interoperate securely. GAIA-X provides federated identity and trust mechanisms (via Self-Descriptions and Trust Anchors) that can be integrated into the TEF’s identity and access control layer (e.g., via Keycloak), ensuring only verified actors can exchange data and run experiments.

Ecosystem Alignment and EU Strategic Autonomy

By aligning with GAIA-X, EnergyGuard becomes compatible with broader EU efforts to build sovereign digital ecosystems—especially in sectors like energy where reliance on external platforms poses risks. It also ensures synergy with projects like Enershare, and TwinEU, all of which implement GAIA-X standards and components. In the following figure, the GAIA-X trust framework is illustrated:




Figure 2 - GAIA-X Proposed trust framework

2.3 IDS-RAM

The International Data Spaces Reference Architecture Model (IDS-RAM) [7] is a standardized framework developed by the International Data Spaces Association (IDSA) [8]. It provides the conceptual and technical foundation for building trusted data ecosystems, where organizations can share and exchange data while retaining full sovereignty over it. IDS-RAM is designed to meet the stringent requirements of data privacy, usage control, interoperability, and trust—all of which are critical in energy-related testing and experimentation environments.

At its core, IDS-RAM addresses a central challenge in the digital economy: how to facilitate secure and policy-compliant data sharing across organizational boundaries without compromising ownership, control, or trust. It does so through a layered architecture that defines roles (e.g., data provider, data consumer, broker, clearing house), interfaces, identity mechanisms, and governance frameworks.

2.3.1 Alignment with EnergyGuard

EnergyGuard is envisioned as a federated, multi-party platform where public and private stakeholders—utilities, regulators, tech providers, and researchers—can collaborate on energy innovations using shared data, digital twins, and AI-driven simulations. IDS-RAM is highly relevant to EnergyGuard for several reasons:

Data Sovereignty and Usage Control

IDS-RAM enforces policy-based data sharing, meaning data owners can define how, when, and by whom their data is used—even after it leaves their system. For EnergyGuard, this is vital: grid operators, building managers, or microgrid owners need assurances that their datasets used in experiments will not be misused, redistributed, or reverse engineered.

Interoperable and Secure Connectors

A central component of IDS-RAM is the IDS Connector—a secure interface for exchanging data with embedded policy enforcement mechanisms. These connectors can be deployed within EnergyGuard to enable automated, secure, and standardized data transactions between data spaces, experiment nodes, and AI model sandboxes.

Identity and Trust Management

IDS-RAM introduces a federated identity model, including mechanisms for participant certification, trust anchors, and metadata catalogs. EnergyGuard can use these tools to verify the identity, credentials, and trustworthiness of users, services, and data providers before any interaction occurs—aligning perfectly with its need for security, auditability, and regulatory compliance.

Metadata and Data Catalogs

In IDS, all participants and data assets are self-described using rich metadata formats. This enhances discoverability and semantic interoperability within EnergyGuard. A researcher, for example, can query for building energy data from a specific climate zone with known privacy constraints, all enabled through standardized metadata and catalogs.

Cross-Sector Alignment and EU Compatibility

IDS-RAM is compatible with other major European frameworks, notably GAIA-X and Eclipse Dataspace Components. As EnergyGuard aims to integrate components like Enershare data space, Trino, and AI pipelines, IDS-RAM ensures architectural cohesion and long-term sustainability across sectors and initiatives.




Figure 3 - IDSA Reference Architecture Model

2.4 CEEDS

The Common European Energy Data Space (CEEDS) [9] is a major initiative by the European Commission aimed at creating a shared, secure, and sovereign digital infrastructure for energy data exchange across Europe. CEEDS is part of the broader European Data Strategy and supports the Digital Europe Programme, providing the foundation for new digital services, market transparency, smart grid coordination, and decarbonization efforts. Its central goal is to unlock the value of energy data across the value chain—from production and transmission to distribution, consumption, and flexibility services—while ensuring compliance with EU regulations, including GDPR, the Electricity Directive, and the Data Governance Act.

CEEDS is envisioned as a federated and interoperable data ecosystem, not a centralized platform. It will connect diverse stakeholders—utilities, grid operators (TSOs/DSOs), energy service providers, regulators, prosumers, and technology developers—via common standards, governance models, and secure data exchange mechanisms. These include semantic interoperability, data usage control, consent management, certification, and identity trust frameworks.

CEEDS takes into account the outcomes of several EU funded data space projects including ENERSHARE, OMEGA-X, SYNERGIES among others.

2.4.1 Alignment with EnergyGuard

CEEDS is compliant with GAIA-X and IDS-RAM architectures, while it poses similar objectives to the ones of EnergyGuard. In particular, both CEEDS and EnergyGuard aim at:

- Providing Trusted Access to Real-World Energy Data
- Providing Interoperability through standardized semantic data models
- Supporting Federated architectures
- Aligning with the EU energy policy and digital regulation

EnergyGuard will comply with CEEDS architecture for trusted and sovereign data sharing in the energy sector.




Figure 4 - CEEDS Reference Architecture

Of course, EnergyGuard will also pay attention to the outcomes of other energy related initiatives and projects such as BRIDGE and OpenDEI. However, in this section we provided an overview of the most important projects for EnergyGuard's objectives.

3 EnergyGuard Conceptual Architecture

In the following figure, the EnergyGuard conceptual architecture is illustrated.




Figure 5 - EnergyGuard Conceptual Architecture

EnergyGuard platform conceptual architecture is a comprehensive, modular software architecture designed to facilitate secure, interoperable, and intelligent testing and experimentation in the energy domain. It supports a wide range of data types, AI model development, and experimental deployments, enabling researchers, developers, and stakeholders to ingest, manage, analyze, and experiment with energy-related data, build their own AI models and services and evaluate them in terms of compliance, trustworthiness and security. With its robust access control, data governance, and AI-driven analytical capabilities, the platform integrates real-world energy data sources with advanced high-performance computing (HPC), AI toolchains, and regulatory compliance modules. The ecosystem supports end-to-end experimentation, from data ingestion to model evaluation and deployment, underpinned by transparency, trust, and user collaboration. It consists of four backend components namely the Data management, the AI sandbox, the HPC and the Acceptance environment, alongside the user dashboard, which is the user interface that acts as the access point for all functionalities of the EnergyGuard platform.

The **User dashboard** offers a centralized interface for interacting with the platform. It allows users to submit experiments, manage applications, upload and share data, and access tools such as cost calculators, emission calculators, and collaboration hubs. It also integrates secure payment processing and compliance evaluation to streamline experiment funding and trust evaluation.

The **Data management component** handles the secure ingestion, storage, and processing of data coming from the different data sources of the EnergyGuard platform (e.g., electricity grids data, microgrids with distributed energy resources (DER), buildings, hydrogen infrastructure, renewable energy certificates (REC), and weather data). Tools like Kafka and Pentaho manage ingestion and harmonization, while Timescale and MongoDB act as secure data lakes. Ontology and semantic enrichment ensure data interoperability and contextualization using technologies such as Apache Jena and Neo4J. In this context, several common data models will be defined and developed in accordance with the pilots' needs leveraging and extending existing ontologies and data models such as SAREF, OEO, etc.. Data anonymization will be a standalone tool that will be used to anonymise data before their ingestion to EnergyGuard platform. On top of the data lake, an efficient query engine will facilitate a wide variety of queries across diverse databases and technologies, leveraging technologies like Trino. Last but not least, this component will facilitate trusted data sharing utilising modern data space technologies like the Eclipse Dataspace as well as the Enershare energy data space to ensure GDPR compliance and seamless cross-domain data access. Trusted data sharing will provide information about the available datasets, their owner, and the terms of use through its metadata catalogue, it will facilitate peer-to-peer energy data sharing through data space connectors, while the data broker will ensure that users access only the data they are authorised to access.

The **HPC and experiments infrastructure** is responsible for the execution of AI experiments with very high performance, powered by components like SLURM for workload management. It includes GPU/FPGA nodes for AI training acceleration, large memory modules, and CPU resources for service deployment. MELUXINA HPC supports easy execution of AI experiments using Jupyter notebooks, APIs and custom framework deployment as well as SSH connection for more customized experiment executions. Besides experiments execution in a high performance setting, this component is responsible for providing proper cloud resources (storage, memory, CPUs) to host all EnergyGuard services.

The **AI Sandbox** is the core experimentation and development environment within the TEF platform, providing the tools and infrastructure necessary to build, test, evaluate, and deploy AI and ML models for the energy sector. It abstracts the complexities of high-performance computing, data management, and orchestration, allowing researchers and developers to focus on domain-specific innovation. This environment supports both the lifecycle of AI development and the deployment of smart energy services through integrated toolchains and scalable runtime support. Regarding AI development, this component offers AI model development and training capabilities utilizing libraries such as PyTorch and Tensorflow, AI workflow orchestration using Dagster, experiment monitoring and evaluation using MLFlow, model deployment and configuration of inference endpoints as well as an AI model repository that will provide available AI models and benchmarks to be used for experimentation and comparison when developing new AI models. On the other hand, Applications and Digital Twins (DTs) subcomponent provides the domain-specific context and simulation capabilities that make the EnergyGuard platform highly relevant to the energy sector. It includes a variety of digital twins such as Electricity Grids DTs, Microgrid with DER DTs, Building DTs, Hydrogen Testing DTs and REC DTs along with other energy analytics applications such as DeepTSF[10] for energy time series forecasting.

The **Acceptance environment** is responsible for applying rigorous evaluation to the developed AI models and services. The evaluation includes vulnerability scanning (leveraging technologies like Wazuh), benchmarking, code quality checks (through SonarQube or other similar tools), and stress testing (via Locust). The environment also incorporates ethical considerations via AI Fairness 360, risk databases, trust marks, and human-in-the-loop verification, supporting both technical validation and regulatory compliance. It is worth mentioning that this component can be used for constant AI model evaluation through AI agents in order to ensure that AI services remain compliant, trustworthy and secure during their entire lifetime. In cases that security, compliance or trustworthiness requirements are compromised, the people responsible for such models are notified in order to proceed with corrective actions.

Last but not least, the **Access Control and Identity management** service is responsible for managing centrally user identities and role-based access control through Keycloak, supporting authentication, authorization, and user provisioning. It safeguards access to the platform's services and data, ensuring that only verified users can interact with sensitive datasets and computational resources.

4 EnergyGuard components architectural overview

In this section, we present a wide insight of the different components of the EnergyGuard platform providing more details regarding their technical implementation.

4.1 Data Management

The depicted data management architecture (see Figure 6) presents modular framework designed to manage the end-to-end lifecycle of data. It aims at the management of data from raw ingestion to high-level interoperable data sharing. This architecture embraces interoperability, semantics, privacy, and governance by design and each architectural layer performs a distinct set of responsibilities:

- Data ingestion and harmonisation enable structured and unstructured data integration from multiple sources, which is harmonised in the common data model.
- A secure data lake stores and manages both time-series and contextual data, enforcing data quality and cataloguing.
- An intelligent querying layer allows semantically enriched data in order to provide data access to end-users via SQL and business intelligence tools.
- Finally, interoperable data sharing ensures that harmonised information can be securely exchanged via APIs, linked data, or federated systems.

The architecture also provides mechanisms such as access control, authorisation, anonymisation, encryption, and auditing throughout the security and privacy transversal layer. On the other hand, the ontology provides alignment between entities.




Figure 6: Data management components architecture

The explanation of the architecture layer-by-layer:

1. Data sources layer: This is the raw data level, which relies on the TEFs and the data the real environments provide. It should be highlighted this is the foundational layer that includes heterogeneous raw data sources. From each TEF, different protocols, sampling frequency and data access are envisaged (which are documented in D3.1).
2. Data ingestion and harmonisation layer: Its main objective is to interface the data sources from the TEFs (data sources layer). This layer collects, integrates and formats raw data into a unified structure. It ensures data harmonisation before it reaches the secure data lake. For that end, it contains a set of ETL (extract, transform and load) modules (namely ETL1, ETL2, ... ETLn) tailored for each data source. These ETLs are periodically triggered by a Scheduler. Moreover, there is a Data bus formed by a Data Integrator, a Data Model Formatter and a JDBC. This component aims to the integration of the data samples from diverse ETLs, formatting in the common data model to be passed to the JDBC component to connect and access to databases.

3. Secure data lake layer, which is responsible for secure storage, data quality, and metadata management. It segregates temporal and static datasets while ensuring consistency between them. Apart from ensuring a proper storage of data, it also provides data curation mechanisms to assure a minimum quality of data to be used by the digital twins. Additionally, the data catalogue provides a structured repository of metadata, enabling efficient discovery, indexing, and semantic search of available datasets across the platform.

This layer is composed by the following components:

- Timeseries database to store dynamic datasets (the ones that changes over time and are sequential). The envisaged technologies are Postgres + TimescaleDB extension.
 - Static data database that holds contextual / static data using graph-oriented systems (e.g. GraphDB or Neo4J). It should be highlighted the sync mechanism, which ensures alignment between dynamic and static databases.
 - Data quality service to ensure data curation mechanisms. Two procedures are included within this layer. First of all, the checker to validate data integrity and consistency with indicators such as completeness or accuracy. Secondly, the corrector, which applies rules to repair or enhance data quality.
 - Data catalogue service to support the metadata in order to provide efficient data discovery mechanisms about the available datasets in the data management structure. Its core is the Data Catalog DB that stores catalogue metadata. The manager maintains metadata and data organisation, while the Search Engine allows retrieval via semantic or attribute-based queries.
4. Intelligent querying layer acts as the core access point for extracting harmonised data stored in the system. Its goal is to reduce the complexity in the data sharing mechanism by enabling intuitive and enriched data access. Through the application of business intelligence, it allows the calculation of pre-analytics and, then, sharing enriched and pre-processed data. In short, it supports enriched querying and provides user-friendly access to structured, processed data.

The key components are:

- SQL query engine, in charge of executing structured queries on the data lake (both dynamic and static repositories) using standard SQL syntax. It abstracts the complexity of underlying databases to query both databases.
 - Business Intelligence (BI) tools to process data obtained by the SQL query engine, combine datasets and generate the data marts. It is in charge of the pre-processing by combining or aggregating data to generate pre-analytics. It supports advanced analytics by providing contextual or computed insights that go beyond simple retrieval.
 - Data enricher, which enhances raw or processed data with additional context or semantic data to provide contextual information to the timeseries (which only includes timestamps and values, without additional enrichment).
 - Data Marts in form of domain-specific subsets of the data optimised for querying and performance. They store curated subsets of the data lake and are oriented in a specialised topic (e.g., thermal energy use).
5. Interoperable data sharing layer to facilitate secure, standardised, and semantically interoperable data exchange across entities. From the data marts, it provides federated data to be formatted in a common data representation, as it is JSON-LD (JavaScript Object Notation

for Linked Data), and, thus, sharing via standard API (Application Programming Interface) to make data available and accessible, at the same time then understandable.

As key components:

- Federated data as decentralised data, allowing queries and operations without requiring centralised storage.
 - JSON-LD: Uses linked data format for semantic interoperability.
 - APIs: Enables programmatic access and integration.
 - Data Spaces: Represents virtual environments or platforms where different stakeholders can access shared data under common standards.
6. Privacy and security with the aim of overarching enforcement of data protection regulations and cybersecurity protocols. Basically, it allows:
- Auditing: Tracks access, usage and modifications to maintain traceability.
 - Access Control: Regulates who can query or interact with each component.
 - Encryption: Protects data at rest and in transit.
 - Authorisation: Ensures only permitted users or systems can perform certain actions.
 - Anonymisation: Applies data masking or pseudonymisation to protect personal/sensitive data.
7. Ontology and semantics layer with the goal of integrating semantic meaning and ontological structuring to the data. This allows interoperability and semantic search through well-defined relationships between data points. This layer represents the data model to represent data in standard formats; therefore, the components are the entities conforming the data model itself (Entity 1, Entity 2, ..., Entity n) and their relationships.

4.2 AI Development and testing sandbox

The AI Development and testing sandbox is the platform's core component. All AI experiments and tests are initiated there, and the users have access to DTs, AI models and standalone applications like DeepTSF.




Figure 7 - AI Sandbox architectural diagram

The sandbox comprises of two subcomponents: Energy Analytics applications and DTs, and the AI development toolbox.

4.2.1 Energy Analytics applications and DTs

This subcomponent consists of DTs and energy analytics applications. DTs will be available through a REST API to all components of the AI development toolbox. Also, their data will be stored in the data lake. If necessary, they will also have access to HPC. Specifically, the applications and DTs will be deployed either on premises or on the cloud resources provided by the HPC. Furthermore, they will leverage the HPC to train and deploy AI models and run experiments by submitting these jobs to the scheduler via APIs or through Jupyter notebooks or even direct SSH access to the GPU nodes of the HPC.

4.2.1.1 Electricity grid DT

This is a DT of the Portuguese Transmission System Operation System, provided by RDN. It uses TSO-supplied historical load, topology, and fault data, which are saved in the EnergyGuard data lake, creating a real-time emulation of the grid. Its secure interface lets AI/ML models test forecasting scenarios on the fly (e.g., demand forecasting, fault detection) without risking real-grid stability.

4.2.1.2 Microgrid with DER DT

A 4.2 km medium-voltage microgrid, featuring a 1 MVA substation, combines 161 kW wind, 189 kW PV, 40 kW hydro, Pb-acid / Li-ion storage, and a 7 kW PEM electrolyser. Real-time control and

monitoring supervise laboratory, workshop and EV-charging loads via a dedicated communications and power-analysis backbone.

4.2.1.3 *Buildings DT*

Built on the MATRYCS framework, the Riga residential Digital Twin fuses GEORIGA 3-D models, cadastral files and city-infrastructure layers with real-time solar-potential and weather feeds. Heating-system telemetry from 5000 buildings and SCCB audit records stream into a common data lake, supplying high-resolution historical and live inputs

4.2.1.4 *Hydrogen testing DT*

Hydrogen testing facilities will collect operational data from PEM electrolyser integrated in the CEDER-CIEMAT microgrid, including power input, hydrogen production rate, temperature, pressure, system status, and fault logs. Moreover, data from MOSE electrolyser will be included, such as all BOP measurements, system status and faults. Finally, a DT of SOC electrolysis system MURPHY will be included, with hydrogen production rate, current density applied, temperature and pressure at different locations.

4.2.1.5 *REC DT*

Antrodoco, a small town in central Italy with around 2,800 inhabitants, hosts an energy community that integrates public, residential, and commercial buildings along with two major energy consumers—a sawmill and a wastewater treatment plant. A cloud-hosted Digital Twin will be provided that ingests generation, load, storage and weather streams to mirror the low-voltage community network in real time, exposing live and historical data through open APIs for operators and third-party apps.

4.2.1.6 *DeepTSF*

DeepTSF is an application enabling the development and deployment of time series forecasts using highly accurate, state-of-the-art AI models. Users simply upload a CSV time-series file into the user interface, choose options such as forecast horizon and model type, and a complete pipeline —pre-processing, training, and evaluation— is run automatically. It is based on the Darts library, which offers automated and optimized time series forecasting. DeepTSF is fully integrated with MLflow, which is used for tracking, management, and deployment of AI experiments. There, interactive diagrams, plots, and metrics allow users to evaluate their model's performance. Dagster orchestrates and monitors the workflows, while also offering more advanced functions like hyper-parameter tuning and SHAP-based explainable AI.

This dockerised application will be utilized for forecasting generation and consumption of energy of the Portuguese grid. Then results will be sent to the DT for grid analysis to validate market results. However, it can be used for any application involving forecasting.

4.2.2 **AI development toolbox**

4.2.2.1 *AI model development*

The AI lifecycle comprises of several stages: Data loading, data pre-processing, model training and evaluation. EnergyGuard architecture will facilitate all of them in a seamless and a user-friendly manner. A separate environment will be provided to each user, so they can conduct their experiments. This could be an isolated dockerised container or an environment in a Jupyter notebook or any equivalent

interface. There, they can install all packages and libraries required by their use-case. In addition, these stages will be tracked using MLflow, and the user will be able to choose which ones to run.

4.2.2.1.1 Data loading

Data will be provided from the data lake in the appropriate format desired by the experiment. All data formats produced by the DTs such as csv, json, and xls will be supported. Metadata stored in the ontology will also be available.

4.2.2.1.2 Data preprocessing

The user will be able to preprocess their data in any way (e.g., data imputation, handling missing values, anomaly detection etc.) necessary to conduct their experiment. Pilot data without the need for preprocessing will also be available if there is no need for custom preprocessing.

4.2.2.1.3 Model Training

Then, the model chosen by the user will be trained on the HPC environment. Information about the progress of training will be shared with the user by API, as this can be a time-consuming process. Various model libraries (e.g. PyTorch, TensorFlow) will be supported. The final model will be stored into MLflow.

4.2.2.1.4 Model evaluation

Finally, the trained model will be evaluated, and the results (plots, metrics and other files the user chooses) will also be stored into MLflow. The appropriate metrics (e.g., MAPE, RMSE, accuracy, F1 score) and plots can be used by the user to choose the best model for each use case and register it into MLflow.

4.2.2.2 *Model deployment and inference endpoints*

The final model will be served in a containerised environment for inference. An endpoint with fixed specifications will be made available for use by the creator themselves, or by other users and experiments if the model is made public. This way, collaboration and transparency will be encouraged, making new users able to build on existing solutions.

4.2.2.3 *Experiment monitoring*

It is very important that all experiments conducted within the platform are monitored, and their results are saved. MLflow is a framework made for AI experiment tracking and monitoring, with advanced features such as separate experiments and pipeline stages, a model registry, and the ability to save metrics, plots and other experiment metadata. MLflow saves all its experiments in a Postgres database. This includes all metrics and metadata. Minio artifact store will also be used to store large experiment files such as models, preprocessed datasets and plots.

4.2.2.4 *AI Workflow orchestration*

As mentioned before, an AI pipeline consists of several stages. Therefore, a workflow orchestrator is necessary to run the appropriate scripts at the correct order. EnergyGuard will use Dagster for this purpose, which offers a variety of useful features. Firstly, it can save the logs of scripts and show them live to the user. It also records the time each stage took to complete, which is useful for benchmarking.

Finally, it gives users the ability to insert parameters to its runs in YAML format and supports cron jobs for scheduling and automating runs.

4.2.2.5 *AI model repository and benchmarks*

All models produced by the users will be saved and available in a repository. Versioning will be supported, and their evaluation results such as metrics and plots will also be available there.

4.3 Trustworthy AI Acceptance Environment

As part of the EnergyGuard TEF infrastructure, the Acceptance Environment serves as a dual-layer mechanism for evaluating the technical robustness, legal compliance, and ethical acceptability of AI applications deployed in energy sector scenarios. This environment, led by ICCS, is built upon the requirements and objectives defined in Tasks 5.1–5.4, and is aligned with the EU’s AI Act regulatory vision.

4.3.1 Purpose and Position in EnergyGuard Architecture

Sitting atop the EnergyGuard layered architecture (see Figure 5), the Trustworthy AI (TAI) Acceptance and Regulatory Sandbox is positioned after experimentation and inference, serving as the final compliance validation checkpoint before AI tools are released into production or market settings.

The system allows EnergyGuard to act as a technical enabler for national and EU-level AI regulatory sandboxes, offering standardized services for:

- Pre-deployment acceptance checks
- Policy experimentation
- Risk profiling
- Ethics and legal supervision

This is particularly relevant for high-risk AI applications in the energy domain, as defined by the AI Act, where real-world testing and cross-sectoral collaboration are needed to inform future governance frameworks.

4.3.2 High-level Functionalities

The acceptance environment consists of interconnected modules that span automated assessments, human-in-the-loop evaluations, and legal-ethical review loops. The following figure illustrates this integrated flow.




Figure 8: Conceptual diagram of the EnergyGuard Trustworthy AI Acceptance Environment

The core components include:

- **Automated code quality checks:** Software verification, static analysis, performance benchmarking.
- **Security testing and vulnerability scanning:** Evaluation of AI robustness against adversarial conditions and cyber risks.
- **AI Act compliance scoring:** Evaluation of obligations based on risk level of the system (we anticipate that most systems and services on EnergyGuard will be high risk). In a high risk setting the tool will evaluate to what extent the system addresses the requirements for a risk management system (identification, evaluation and mitigation of risks); relevant representative and bias free data; detailed technical documentation; record keeping for traceability; transparency and explainability; human oversight; and accuracy, robustness and security. The scoring will be calculated using other components like the security testing and vulnerability scanning component. Integration of ALTAI tools to assess alignment with trustworthiness pillars (e.g. fairness, accountability, robustness) will also be examined for the scoring as a secondary option, since ALTAI guidelines are considered non-binding.
- **Testing metrics and experimental reporting:** Includes the generation of standardized experiment logs for reproducibility and benchmarking.
- **Common AI Risks Database:** A shared repository populated through use cases and partner feedback (Task 5.1), providing mappings of known risks and recommended mitigations for each AI asset tested.
- **Supportive Tool For Regulatory Sandboxes:** While EnergyGuard will not operate as a regulatory sandbox in the formal sense defined by the AI Act (i.e., established and managed by competent authorities), it aims to support such initiatives by offering a research-oriented, controlled testing environment where AI applications can be assessed under realistic conditions. The platform will facilitate structured experimentation involving external stakeholders—such as legal experts, ethics committees, and national authorities—who are engaged in supervised validation trials. The Supportive Tool for Regulatory Sandboxes will provide high-level evaluation, early feedback loops and technical insights that can inform internal validation efforts as well as the design and operation of national or EU-level regulatory sandboxes.

This environment will be continuously fed by operational Digital Twins (DTs) and AI assets originating from EnergyGuard pilots, enabling testing under realistic conditions.

4.3.3 Requirements

The Acceptance Environment will be designed to serve the diverse needs of multiple user groups, ranging from startups and data scientists to solution developers and policymakers. Each group interacts with the environment differently, depending on their role in the AI development and validation lifecycle. Below, we outline the core requirements derived from these stakeholder perspectives based on D2.1 interviews and questionnaires.

4.3.3.1 *Compliance and Trustworthiness Evaluation*

User story

Start up intends to evaluate compliance and trustworthiness for their energy analytics AI based solution.

Derived specifications:

- A compliance assessment dashboard showing alignment with key regulations (e.g., GDPR, AI Act, energy sector norms).
- Automated trustworthiness evaluation tools (fairness, explainability, robustness).
- Access to ethical impact checklists and scoring templates.
- Exportable compliance reports for internal or investor use.
- Easily connect existing AI services with the TEF for compliance evaluation (e.g. inference endpoints)

4.3.3.2 *Data Scientists – Continuous Monitoring for Model Integrity*

User story:

Data scientist intends to monitor models continuously for bias, drift and reliability to improve trustworthiness and robustness.

Derived specifications:

- Continuous model monitoring pipeline with bias detection modules (e.g., demographic parity, disparate impact).
- Drift detection on input features and predictions (statistical and concept drift).
- Reliability scoring based on performance under different conditions (e.g., weather, grid load)
- Alert system for detected performance degradation.

4.3.3.3 *Solution Developers – Safety, Security, and Compliance Scores*

User story:

Solution developers need safety, security and compliance scores for their solution

Derived specifications:

- Automated scoring system with configurable weights for safety (fail-safe behavior), security (vulnerabilities), and compliance.
- Modular testing suite based on standards (e.g., ISO/IEC 27001 for security, NIST AI RMF).
- Color-coded report cards (e.g., green/yellow/red flags).
- Exportable certifications and audit logs for external stakeholders.

4.3.3.4 Policymakers – Aggregated and Anonymized Insights

User story:

Policy maker needs anonymised insights into the status of the developed AI services over time to create better public policies and incentives

Derived specifications:

- Anonymized dashboard with aggregated analytics on AI services: use cases, performance metrics, and compliance trends.
- Filtering by time, sector (e.g., buildings, grid, transport), and technology type.
- Policy recommendation engine highlighting gaps and trends.
- Secure data sharing protocols and access control to protect IP and privacy.

4.4 HPC

This section details how LuxProvide's HPC MeluXina underpins EnergyGuard's AI experimentation, training, and serving needs. The LXP HPC will host the AI workloads of the EnergyGuard users. Data generated by TEFs will be processed on LXP infrastructure.

4.4.1 Architectural Diagram




Figure 9 - HPC Architecture (Meluxina)

4.4.2 Description

Meluxina, the HPC hosted by LuxProvide (LXP), is composed of several components. We describe in this section the components that are most relevant to the EnergyGuard platform.

Cloud module

LXP currently offers virtual machines (VM) managed via OpenStack. With a proper account, the user can create and launch VMs of their chosen configuration (for example, 32 vCPUs and 64GB RAM, 1TB storage).

The different flavours available are Linux VMs, but Windows VMs can be arranged (however, LXP does not provide Windows licenses). VMs can also provide GPU cards, Nvidia H200, but in limited numbers. Network access is restricted: a VPN is needed to access a VM from the public internet. The VPN configuration files are provided for network setup.

GPU module

The main computing infrastructure of Meluxina are the GPU and CPU modules. EnergyGuard has access to the GPU module, defined by (in total):

Total GPU nodes	200
CPU	2x AMD EPYC Rome 7452 32c @2.35GHz
GPU	4x NVIDIA Ampere 40 GB HBM 64 Int/FP32 per SM
RAM	512 GB
Local disk	1.92TB temporary, SSD
Network	Infiniband 2x HDR200 SP, dual-rai
Total cores	12'800
Total SMs	86'400
Theretical performance	>15 PFlops, 500 PFlops 'AI'

The 200 GPU-equipped nodes are accessible via the Slurm scheduler. In a nutshell, using the GPU nodes requires that first resources are allocated to the users, and then jobs (processes) can be launched on the nodes allocated. There are restrictions to the number of nodes and duration of the allocation. Slurm scheduling can be requesting from a REST API (behind the VPN), or from a ssh connection to the login nodes.

The VMs from the Cloud module have direct access to the Slurm partition.

S3 - Data transfer to and from the Meluxina GPU nodes is possible via S3.

JupyterLab is also available on the GPU nodes, a user-friendlier interface is available (Open On Demand) to launch a JupyterLab notebook.

Access to the GPU module require an account, and access to the machines is through ssh protocol (public keys). Storage is dedicated to a project, composed of one or more users.

Access to the VM requires a VPN access, and ssh protocol is enabled at the VM. This is in addition to any security mechanism at the application level.

Keycloak is also available.

Software stack on GPU nodes

For security reasons, users of GPU nodes do not have full administrative rights (unlike the allocated VMs). The user can install software in user-space (including those provided by the EasyBuild package manager). Additionally, pre-configured software is available to the user. As the list is constantly evolving, https://docs.lxp.lu/first-steps/software_env/ provides an up-to-date list.

Finally, Docker containers can be executed on Meluxina nodes, with some modifications: Meluxina uses AppTainer instead:

- Docker & OCI compatible
- Non-privileged mode for improved security
- Support for GPU accelerated applications
- Support for creating and running encrypted containers
- Support for trusted containers: PGP signed & verified
- Compressed read-only Singularity Image File (SIF) format, suitable for production (default)
- Writable (ch)root directory called a sandbox, for interactive development (--sandbox option)

Meluxina-AI

LXP is in the process of acquiring an additional AI supercomputer, expected in production in Q4 2026.

This new AI supercomputer is planned to contain the following features:

- A dedicated Kubernetes partition alongside Slurm for elastic serving and pipelines.
- More recent accelerators and options for public ingress to simplify controlled external access. EnergyGuard will provide a single abstraction layer so users can target Slurm now and transparently exploit Kubernetes later without refactoring. Some public IP addresses will be available to users (avoiding the need for a VPN) if needed.

4.4.3 Requirements

From the perspective of EnergyGuard use cases, the LXP HPC should meet the following functional and non-functional requirements:

1. On-demand access to several nodes, each equipped with multiple GPUs, in order to support the fast execution of AI applications.
2. Permanent access to several machines hosting EnergyGuard backend components.
3. Permanent storage space of sufficient capacity.

4. Security mechanisms to separate AI workloads from different users of EnergyGuard applications and data. From the EnergyGuard perspective, the HPC must provide a way for each AI workload to be isolated from another, despite the account management of Meluxina.
5. The allocation model of Slurm, with limited durations (maximum 144 hours per job, or 48 hours by default) must be abstracted, or managed at an appropriate abstraction level, for the EnergyGuard users.
6. Proper accounting of resource used must be available to provide billing information to EnergyGuard.

These requirements ensure that EnergyGuard can offer users a coherent, secure, and high-performance experience - from data ingestion to large-scale training to compliant serving - while keeping operations maintainable and future-proof as the HPC roadmap evolves.

5 Alignment with the European ecosystem

In this section, the alignment with the EU ecosystem is presented in brief. Specifically, we present the alignment with AIoD, Data Spaces, EDIHs, DOME, AI Factories and we consider other TEFs.

5.1 AI on Demand

The AI on Demand platform is a digital hub that fosters European AI research and innovation, championing the region's principles of high quality, trust, and transparency. Various components are available for use by the services offered.

The core ones are the following:

- **MyLibrary** allows the user to browse some types of assets of the AIoD platform (AI models, Datasets, Experiments).
- **RAIL** is a tool that allows AI practitioners to use AI assets directly in platform through a python SDK. RAIL is based on REANA, which is a platform allowing data scientists to run containerised pipelines remotely
- The **AI-on-Demand Metadata Catalogue** contains code and configurations for the platform. It collects AI assets and resources stored across various sources (e.g. Zenodo, Hugging Face and OpenML).
- **AI Builder** is a Graphene based application, which enables users to connect dockerised AI solutions seamlessly, using a graph-based UI.
- The **Deploy AI Marketplace** based on Virto commerce is available where users can purchase AI related products.

EnergyGuard will take the architectural design of these services into account to offer high quality services and a seamless user experience. For example, running experiments and serving models in an isolated environment like RAIL, is something the platform's architecture will support. Furthermore, access to both public and private models and experiments will be provided to each user through the user dashboard, like MyLibrary.

5.2 Data spaces

As outlined in Deliverable D2.1, the EnergyGuard initiative identified data space architecture as a foundational element for enabling secure, scalable, and AI-driven data exchange in the energy sector. This approach supports multi-stakeholder collaboration and ensures compliance with evolving EU data governance frameworks. D2.1 defined key evaluation criteria, including security, data sovereignty, interoperability, transparency, scalability, real-time analytics, innovation, and cost efficiency, reflecting the unique demands of sharing sensitive energy data across diverse systems such as IoT devices, smart homes, and energy marketplaces. It also presented a comparative analysis of open-source data space connectors, assessing their technical features, deployment models, and alignment with standards like GAIA-X, IDS, and NGSI-LD. This evaluation provided a basis for identifying candidate technologies capable of supporting EnergyGuard's secure and policy-compliant data exchange goals.

Building on that foundation, further internal discussions and a clearer understanding of project priorities, the list of candidate connectors has been refined. Specifically, the TRUE Connector, Trusted Connector by Fraunhofer, and the FIWARE Data Space Connector have been excluded from further consideration. Reasons for this include limited ecosystem adoption, lack of alignment with

EnergyGuard's scalability and cross-sector needs, and architectural dependencies that complicate integration with energy-domain use cases.

As a result, the remaining focus is on a smaller set of promising connectors, including the Eclipse Dataspace Connector (EDC), Sovity Open-Source EDC, and TNO Security Gateway, which demonstrate stronger alignment with EnergyGuard's needs in terms of flexibility, compliance, and integration potential. Compatibility and integration with initiatives such as the ENERSHARE dataspace is also being considered, ensuring alignment with emerging standards in the energy sector. With such connectors, the user will be able to securely share and access energy-related data while maintaining full control over how that data is used. This is achieved through a combination of key data space functionalities integrated into each connector.

The user will be able to expose internal datasets to external stakeholders in a controlled way. This can be done through the data interface layer, which connects internal systems (e.g. databases, APIs) to the data space environment.

The user will also be able to define and enforce rules on how shared data can be used. This is made possible by the policy enforcement engine, which interprets usage policies and ensures that consumers comply with terms such as purpose limitation and duration of access.

To ensure data is only accessible to trusted parties, the user will rely on identity and access management features. These authenticate other participants and manage authorizations, leveraging federated identity systems and trust frameworks within the dataspace.

The user will be able to publish and discover datasets using a metadata catalog, which includes information about dataset contents, ownership, licensing terms, and availability. This facilitates transparent and efficient data discovery and negotiation.

All data exchanges initiated by the user will be protected and traceable. This is ensured through secure communication protocols and audit logging mechanisms, which monitor who accessed what, when, and under what policy.

Overall, these connectors allow the user to participate in a trusted and decentralized data sharing environment. They enable compliance with regulations, protect sensitive information, and support EnergyGuard's goal of establishing a federated data ecosystem tailored to the energy sector.

5.3 European Digital Innovation Hubs (EDIHs)

European Digital Innovation Hubs (EDIHs) are regional centers co-funded by the European Commission to support the digital transformation of SMEs and public sector organizations. Their core services include technology testing, training, support with financing, and innovation ecosystem building. Integrating EDIHs with the EnergyGuard platform allows local experimentation, knowledge transfer, and AI adoption to scale across Europe through a federated innovation infrastructure.

EnergyGuard platform, as a high-performance, trustworthy, and AI-driven testing facility, can act as a central backbone that provides advanced AI models, data infrastructure, digital twin simulation environments, and compliance testing capabilities to EDIHs and their clients.

In this direction, EnergyGuard will provide several service interfaces (APIs) along with an API catalog to allow EDIHs easily access AI models and workflows, energy digital twins as well as datasets and dashboards, using different technologies including data space connectors. Moreover, federated identities with EDIHs and AIoD will be thoroughly examined to make user's life easier. With this

approach, SMEs supported by EDIHs, can port their use cases into EnergyGuard environment utilizing several resources, including existing workflows and AI models, HPC capabilities, trustworthiness evaluation tools, Digital Twin capabilities and available datasets.

Furthermore, EnergyGuard's common data models and semantic mapping tools can foster collaboration through data and knowledge exchange. Specifically, SMEs supported by EDIHs will be able to harmonise their local data in accordance with EnergyGuard's common data models enabling cross domain experimentation, and better utilization of EnergyGuard's datasets, reinforcing data sharing through data spaces. Moreover, EDIHs' SMEs can benefit from EnergyGuard through the compliance evaluation tools (bias analysis, GDPR conformance, cybersecurity checks), improving their AI services maturity levels and trust.

Of course, given the early phase of the project and the limited documentation regarding the integration with EDIHs, the potential integration with EDIHs will be further analysed and refined in the next instance of this deliverable.

5.4 DOME Marketplace

The DOME Marketplace, initiated under the EU's Data Spaces Support Centre (DSSC) framework, is a federated platform designed to facilitate cross-sector, cross-border discovery and trading of data assets, services, and applications. It operates under Gaia-X and IDS-RAM principles and supports data providers, consumers, and intermediaries with tools for cataloging, orchestration, consent, and trust.

Integration between EnergyGuard and the DOME Marketplace would transform EnergyGuard from a standalone testbed into a federated node in the European data economy, capable of offering and consuming assets at scale. It fosters trust, transparency, and EU-wide collaboration, while aligning with key frameworks like IDS-RAM, Gaia-X, CEEDS, and Open DEI.

Through the integration with DOME, the EnergyGuard will enable:

- Publishing of EnergyGuard datasets, AI models, APIs, and experimentation services
- Discovery and access of external datasets or services for EnergyGuard experiments
- Cross-data space orchestration of AI pipelines and digital twins
- Participation in cross-domain use cases (e.g., energy + mobility + buildings)

In terms of technical integration, the integration with DOME's federated catalog will be examined in order for EnergyGuard assets (datasets, services, etc.) to be discoverable through DOME, using GAIA-X self-descriptions and IDS connectors, enabling EnergyGuard assets to be visible in the broader EU data and AI landscape.

Moreover, Data spaces compatibility will be examined (along with federated identities) to enable direct connections and peer-to-peer data sharing between DOME and EnergyGuard users.

Of course, given the early phase of the project, the potential integration with the DOME Marketplace will be further analysed and refined in the next instance of this deliverable.

5.5 Other TEFs

EnergyGuard has been funded in the same call with two other TEF projects namely EnerTEF and AI-EFFECT that pose similar objectives as EnergyGuard. Hence, EnergyGuard is in close collaboration with

these two projects in order to avoid standalone siloed solutions including coordination with CoordinaTEF. To this end, significant attention is paid to interoperability and common energy data models, as well as data space interoperability through comprehensive API catalogs and documentation. Of course, given the early phase of all three TEFs, the potential integration among them will be further analysed and refined in the next instance of this deliverable after the release of the first version of each one.

5.6 AI Factories

EU AI Factories are a network of ecosystems that leverage Europe's supercomputing infrastructure to support the development of advanced AI models and applications by providing computing power, data, and expertise to industries, startups, and researchers. They are a strategic initiative to foster European AI innovation and competitiveness, particularly for SMEs, by providing access to AI-optimized supercomputers and related services. The initiative involves significant public and private investment, aiming to become a world leader in AI through collaboration and the creation of a pan-European AI ecosystem.

Although there is some complementarity between the EnergyGuard and AI Factories offerings, the integration between them can undoubtedly unlock significant value as they play a different role in the AI applications development process. Specifically, EnergyGuard can be used for safe pre-market experimentation for new AI-enabled energy analytics services, while the AI factories can be used during the operation and deployment of such application on production environments. To this end, EnergyGuard can be used for testing, benchmarking and validation of new AI services, while the AI factories will focus more on MLOps providing continuous integration, deployment and monitoring of the AI applications. Last, EnergyGuard will provide compliance and ethics evaluation and suggestions for improvement, while the AI factories will provide production scaling, governance and auditability.

In this context, we envision that AI models and services developed in the sandbox environment of EnergyGuard, will be exported and handed off to an AI Factory for production deployment. On the other hand, AI Factories will feed back real-world performance data to EnergyGuard, for revalidation, and testing under new conditions and retraining. Moreover, using common metadata standards, data models and linked ontologies will enable seamless traceability between training (EnergyGuard) and production (AI Factory) as well as semantic interoperability.

Integrating AI Factories with EnergyGuard bridges the gap between responsible innovation and scalable deployment. EnergyGuard validates trustworthiness, performance, and compliance, while the AI Factory ensures industrial-grade AI operations. Together, they enable a mature, EU-aligned, and cross-sector AI ecosystem capable of accelerating innovation in critical areas like energy, climate, and infrastructure. Of course, given the early phase of the project, and the AI Factories, the potential integration between them will be further analysed and refined in the next instance of this deliverable.

6 Conclusion

The goal of this deliverable was to set the initial version of the EnergyGuard architecture and its alignment with the European ecosystem and relevant initiatives including AIoD platform, AI Factories, Data Spaces, Digital Innovation Hubs, the DOME and other TEFs. The open issues related to creating a novel architecture were efficiently addressed by reviewing the current state of the art and other relevant projects and architectures, understanding their key offerings and aligning with the ones we see fit with EnergyGuard, as well as addressing the requirements collected from different stakeholder groups (AI developers, energy stakeholders, researchers etc.) under the context of the deliverable D2.1.

The next version of this deliverable (D2.4) will be published in M24 and will include updates based on the early feedback we plan to gather from our stakeholders as well as the updated requirements and specifications that will be derived in T2.2 and reported in D2.3 which will present the updated requirements and technical specifications.

Annex I: References

- [1] I-ENERGY, "I-ENERGY project," 2021.
 - [2] E. Karakolis *et al.*, "Artificial Intelligence for Next Generation Energy Services across Europe – The I-ENERGY Project," in *20th International Conference e-Society 2022*, 2022.
 - [3] S. Pelekis *et al.*, "DeepTSF: Codeless machine learning operations for time series forecasting," *SoftwareX*, vol. 27, p. 101758, Sep. 2024, doi: 10.1016/J.SOFTX.2024.101758.
 - [4] E. Karakolis, K. Alexakis, P. Kapsalis, S. Mouzakitis, and J. Psarras, "An end-to-end approach for scalable real time Anomaly detection in smart buildings," *13th International Conference on Information, Intelligence, Systems and Applications, IISA 2022*, 2022, doi: 10.1109/IISA56318.2022.9904410.
 - [5] E. Karakolis, S. Pelekis, S. Mouzakitis, G. Kormpakis, V. Michalakopoulos, and J. Psarras, "THE I-ENERGY REFERENCE ARCHITECTURE FOR THE PROVISION OF NEXT GENERATION ENERGY SERVICES THROUGH ARTIFICIAL INTELLIGENCE," in *ES 2023: 21st International Conference e-Society 2023*, 2023.
 - [6] GAIA-X, "GAIA-X project."
 - [7] B. Otto, S. Steinbuß, A. Teuscher, and S. Lohmann, "IDSA Reference Architecture Model," 2019. [Online]. Available: <https://internationaldataspaces.org/wp-content/uploads/IDSRAM-3.0-2019.pdf>
 - [8] IDSA, "International Data Spaces | The future of the data economy is here," 2021.
 - [9] A. Dognini *et al.*, "Blueprint of the Common European Energy Data Space", doi: 10.5281/ZENODO.12635358.
 - [10] S. Pelekis, T. Pountridis, G. Lampropoulos, G. Kormpakis, and E. Karakolis, "DeepTSF | Release 1.0.2," 2023, 1.0.2. doi: 10.5281/ZENODO.12749074.
-