Ref. Horizon-CL5-2023-D3-03
Grant Agreement 101160614

EU-DREAM

Effective Uptake of Digital Services to Repower European Consumers and Communities as Active Participants in Energy Transition and Markets

D1.3 Modular Platforms

EU-DREAM is supported by the European Union's Horizon Europe programme under agreement 101160614. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

DOCUMENT INFORMATION

DOCUMENT/ DELIVERABLE ID	D1.3 Modular Platforms
TYPE	Report
DISTRIBUTION LEVEL	Public
DUE DELIVERY DATE	30/09/2025
DATE OF DELIVERY	29/09/2025
VERSION	V1.0
DELIVERABLE RESPONSIBLE	AKKO
AUTHOR (S)	Giovana M. Ciprián (AKKO) Philippe Szczech (AKKO)
OFFICIAL REVIEWER/s	Nicolò Bertozzi (LINKS) Babak Arbab Zavar (AAU)

DOCUMENT HISTORY

VERSION	AUTHORS	DATE	CONTENT AND CHANGES
v0.1	Giovana M. Ciprián Herrera (AKKO)	02/09/2025	Table of Contents
v0.2	Giovana M. Ciprián Herrera (AKKO)	05/09/2025	First Draft
v0.3	Giovana M. Ciprián Herrera (AKKO)	09/09/2025	Consolidated
v0.4	Nicolò Bertozzi (LINKS) Babak Arbab Zavar (AAU)	21/09/2025	Revision
v0.5	Giovana M. Ciprián Herrera (AKKO)	22/09/2025	Release Candidate
v1.0	Giovana M. Ciprián Herrera (AKKO) João Catalão (UPO)	29/09/2025	Final Version

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

ACKNOWLEDGEMENTS

NAME	PARTNER
Joao Catalao	UPORTO
Marialaura Di Somma	ENSIEL/UNINA
Gianfranco Chicco	ENSIEL/POLITO
Philippe Szczech	AKKO
Wail El Bani	AKKO
Maroua Dorsaf Djelouat	AKKO
Nicolò Bertozzi	LINKS
Marco Sacchet	LINKS
Alessandro Mozzato	LINKS
Alfredo Favenza	LINKS
Babak Arbab Zavar	AAU
Amedeo Buonanno	ENEA
Salvatore Fabozzi	ENEA

We would like to sincerely thank all partners involved in this task for their valuable input and for always being available to provide support and guidance whenever needed.

ACRONYMS

ACRONYMS	
AI	Artificial Intelligence
API	Application Programming Interface
CLI	Command Line Interface
DCAT	Data Catalogue Vocabulary
DDIM	Dynamic District Information Management Server
DID	Decentralized Identifier
DLT	Distributed Ledger Technology
DT	Digital Twin
EMS	Energy Management System
EU	European Union
FAIR	Findable, Accessible, Interoperable, Reusable
GA	Grant Agreement
GDPR	General Data Protection Regulation
IDSA	International Data Spaces Association
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
JSON	JavaScript Object Notation
JWT	JSON Web Token
NLP	Natural Language Processing
RBAC	Role-Based Access Control
RDF	Resource Description Framework
SMI	Service Mesh Interface
SPARQL	Query language for RDF
SPIFFE	Secure Production Identity Framework for Everyone
SQL	Structured Query Language
SSO	Single Sign-On
SVID	SPIFFE Verifiable Identity Document
VC	Verifiable Credentials

CONSORTIUM

Participant No.	Participant Organisation Name	Acronym	Country
1 (Coordinator)	Universidade do Porto	UPO	PT
2 (Partner)	Cleanwatts Digital SA	CWD	PT
3 (Partner)	EPRI Europe DAC	EEU	IE
4 (Partner)	SSE Airtricity LTD	SSEA	IE
5 (Partner)	DCSix Technologies Limited	DCS	IE
6 (Partner)	Aalborg Universitet	AAU	DK
7 (Partner)	Teknologian Tutkimuskeskus VTT OY	VTT	FI
8 (Partner)	Akkodis High Tech	AKKO	FR
8.1 (Affiliated)	Akkodis I&S	AKKIS	FR
8.2 (Affiliated)	Akkodis Digital	MODIS	FR
9 (Partner)	Universidad de Castilla - La Mancha	UCLM	ES
10 (Partner)	Agenzia Nazionale per le Nuove Tecnologie, L'energia e lo Sviluppo Economico Sostenibile	ENEA	IT
11 (Partner)	Consorzio Interuniversitario Nazionale per Energia e Sistemi Elettrici	ENSIEL	IT
11.1 (Affiliated)	Politecnico di Torino	POLITO	IT
11.2 (Affiliated)	Università degli Studi di Napoli Federico II	UNINA	IT
11.3 (Affiliated)	Università degli Studi di Salerno	UNISA	IT
12 (Partner)	Fondazione Links - Leading Innovation & Knowledge for Society	LINKS	IT
13 (Partner)	IREN SPA	IREN	IT
13.1 (Affiliated)	IREN Mercato SPA	IME	IT
14 (Partner)	Aristotelio Panepistimio Thessalonikis	AUTH	EL
15 (Partner)	DOMX Idiotiki Kefalaiouchiki Etaireia	DOMX	EL
16 (Partner)	Vlaamse Instelling Voor Technologisch Onderzoek N.V.	VITO	BE

Executive Summary

This deliverable defines the EU-DREAM **Data Platform**, which is designed to be modular, secure, and interoperable. It begins with a set of requirements that explain what the platform must do, covering different areas as data collection, storage, integration, monitoring, access, visualization, security, and deployment. These requirements ensure that the system respects data privacy and sovereignty while staying accessible and easy to use.

The report then explains how these requirements are consolidated into a clear architecture. Each layer of the platform is described with its role, the chosen open-source technologies, and the functions and interfaces it provides. Together, these elements show how the platform can acquire data from many sources, organize it into common formats, provide reliable access, and protect it with strong security measures.

Finally, the report looks at how the platform fits into the idea of European data spaces. It shows that the EU-DREAM Data Platform includes the main features needed for trusted data sharing: connectors, federated identity and trust, usage policies, catalogues, open interfaces, secure communication, auditability, monitoring, distributed infrastructure, and governed storage.

In summary, the deliverable sets out the requirements, outlines the technical solutions, and demonstrates that the EU-DREAM Data Platform complies with European data space principles.

Table of Contents

ACRONYMS....3
EXECUTIVE SUMMARY....5
1 INTRODUCTION....9
2 INTERDEPENDENCIES BETWEEN T1.3 AND OTHER TASKS....10
3 SYSTEM SCOPE AND BOUNDARIES....11
4 REQUIREMENT COLLECTION....13
4.1 Overall requirements....13
4.2 Data Platform Requirements to Support Integration of Other Components....16
4.2.1 EU-DREAM components integration requirements....16
4.2.2 Data Platform Requirements supporting simulation models execution. 16
4.2.3 Data Platform requirements the AI-based Assistant Tool and NLP intermediator....16
4.2.4 Data integration and Data Interoperability requirements....16
5 TECHNICAL SPECIFICATIONS....17
5.1 EU-DREAM Data Platform Architecture and Component Specifications....17
6 DATA SPACE....22
6.1 Background....22
6.1.1 Definitions....22
6.2 Application of Data Space concept in EU-DREAM Data Platform....23
7 CONCLUSION....24
References....25
ANNEXES....26
Annex A....26
Annex B....31

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

LIST OF FIGURES

Figure 1: Tasks 10
Figure 2: EU-DREAM Data Platform Components in the Overall Architecture (Links credits) 12
Figure 3: EU-DREAM Data Platform 17

LIST OF TABLES

Table 1:Overview of EU-DREAM Living Labs 9
Table 2: EU-DREAM Data Platform Architectural Layers and General Requirements 13
Table 3: EU-DREAM Data Platform Layers and Component Specifications 18
Table 4: Key Definitions for Data Spaces 22
Table 5: Alignment of EU-DREAM Data Platform with Data Space Principle 23

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

1 INTRODUCTION

EU-DREAM – Effective Uptake of Digital services to Repower European consumers and communities as Active participants in energy transition and Markets aims to capture end-users' preferences and expectations in a simplified manner, improving customers' awareness, trust, and confidence in their interactions with the energy market by introducing a smart intermediator – as an energy attorney – for end-users. Real empowerment will come from translating all the energy markets complex rules into everyday concepts such as comfort, temperature, and utility bills, so citizens understand the impact of their preferences and routines associated with energy use, rather than knowing all technical details.

The report defines the requirements and preliminary specifications for a modular, scalable, and secure data platform. The EU-DREAM Data Platform will support tools such as Digital Twin[1], Artificial Intelligence (AI)-based assistant tool and Natural Language Processing (NLP)[2]-based intermediator. It will give users and tools simple access to real-time, historical, and forecasted energy data from many different sources. By removing technical barriers, such as differences in formats, protocols, or storage systems, the data platform will ensure interoperability with third-party systems and enable smooth communication between stakeholders. It will also present data in a standardized, user-friendly way, making participation in the energy market easier for both non-experienced and experienced users.

The requirements will be tested and validated in six Living Labs, which are real-world environments representing different contexts and user needs:

Table 1:Overview of EU-DREAM Living Labs

LL Number	Description	Location
LL1	AI-Powered Renewable Energy Communities	Coimbra, *PORTUGAL*
LL2	Energy Poverty and Vulnerable Consumers Impacts	Genk, *BELGIUM*
LL3	Energy Flexibility at Residential Scale and Interoperability	Northern Italy, *ITALY*
LL4	Consumers Empowerment for Energy Management	Dublin, *IRELAND*
LL5	Smart Energy Services for Multi-Energy Vectors	Thessaloniki, *GREECE*
LL6	DT-enabled Residential IoT Microgrid	Aalborg, *DENMARK*

This deliverable is the result of Task 1.3 activities, which focused on gathering the requirements and specifications for a federated data platform that enables seamless access to services and solutions.

[1] A DT is a virtual replica of a physical object, system, or process that helps us monitor, analyse, and optimize its real-world counterpart

[2] NLP is like teaching computers to understand and talk to us in our everyday language.

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

2 Interdependencies between T1.3 and other tasks

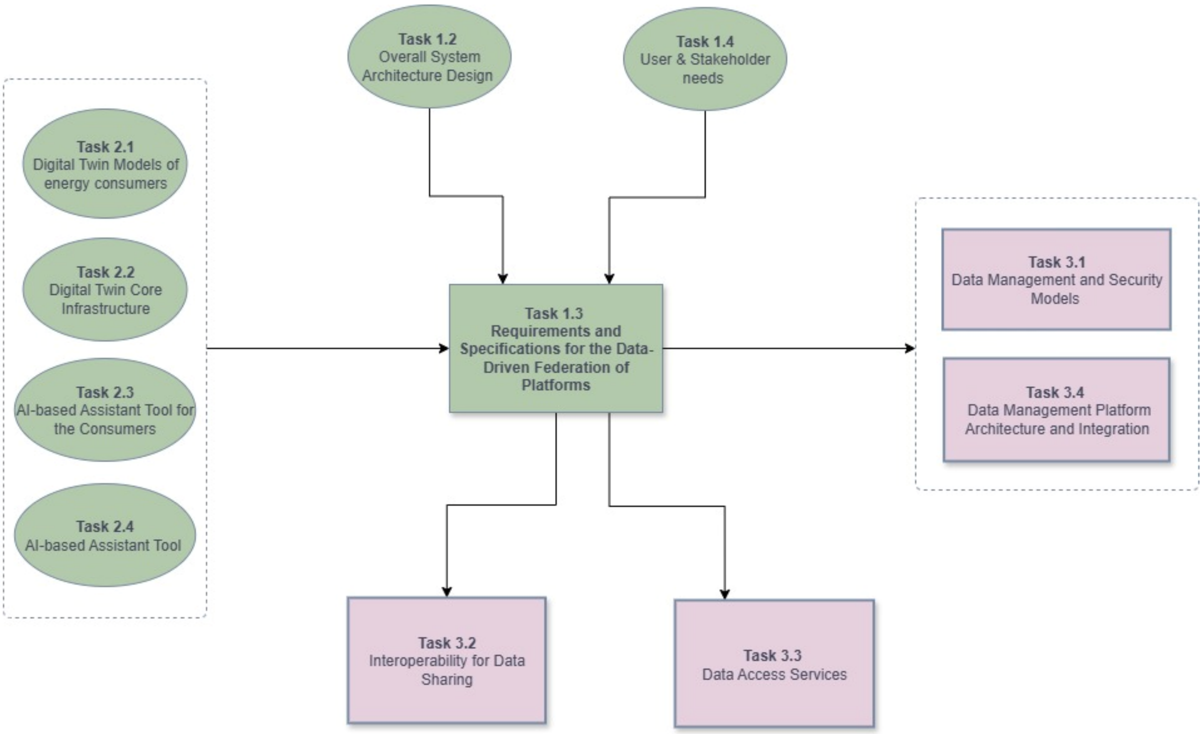

Figure 1: Tasks

To achieve this, it relies on inputs from other tasks and WP's, such as the definition of the overall system architecture in Task 1.2, the identification of user and stakeholder needs in task 1.4, and the description of digital components in WP2, particularly Tasks 2.2 (Digital Twin Core Infrastructure), 2.3 (AI-based services), and 2.4 (NLP-based Intermediator).

Once the requirements are defined, they serve as a reference for other activities involved in the design, development, and deployment of the platform. This includes work related to selecting technologies, developing software modules, and preparing the platform for use in different environments such as the Living Labs.

There is also a strong connection with Task 3.1 and Task 3.4, which address security aspects, and Task 3.2, which deals with data integration, and interoperability. In addition, this task is tightly aligned with Task 3.3, which is in charge of the development of the data platform. The present task sets the conceptual and technical foundations, while Task 3.3 builds upon them to deliver the operational data platform. Technical constraints from those areas are considered to make sure the requirements are realistic and can be implemented using open-source tools and standards.

3 SYSTEM SCOPE AND BOUNDARIES

The purpose of this section is to describe the main role and the main location of the EU-DREAM Data Platform in the overall EU-DREAM system.

The data platform is the technical environment where various digital tools, AI-based services, and energy market interfaces are connected and integrated. Conceptually, it serves as a central component of the overall EU-DREAM architecture for data services.

It enables data to flow seamlessly between components such as sensors, local energy management systems, DT, market tools, and end-user interfaces, while also relying on cross-cutting services such as security, interoperability, and data governance. DT act as virtual replicas of physical energy assets. They allow simulations, forecasts, and analyses to be run safely in the digital space before decisions are applied in the real system. Distributed Ledger Technology (DLT) ensures data decentralization while also guaranteeing that every data transaction linked to these processes is recorded with an immutable timestamp, providing auditability and compliance. NLP-based intermediators simplify interaction by allowing users and systems to issue queries or commands in natural language. Message brokers underpin the platform's communication backbone, routing high volumes of sensor data, commands, and results in real time between all modules.

In other words, the data platform is the backbone that supports the communication and coordination of different modules across the six Living Labs. It ensures that user preferences can be collected, processed, and translated into actionable insights, that service providers can interact with secure and trusted data streams; and that governance mechanisms are respected throughout the system.

The EU-DREAM data-driven platform is designed to operate as a modular, interoperable, and secure digital infrastructure that facilitates the integration and use of energy-related data. Its scope is defined by the need to enable scalable access to services, tools, and datasets, while supporting cross-domain collaboration and end-user empowerment within energy systems.

Practically, the data platform is made of various components that are integrated in the EU-DREAM system. These components are detailed in the specification part of this document, and they have been selected to meet the requirements described in the dedicated paragraph hereafter. The green-circled elements in the diagram below define the EU-DREAM Data Platform and its components.

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

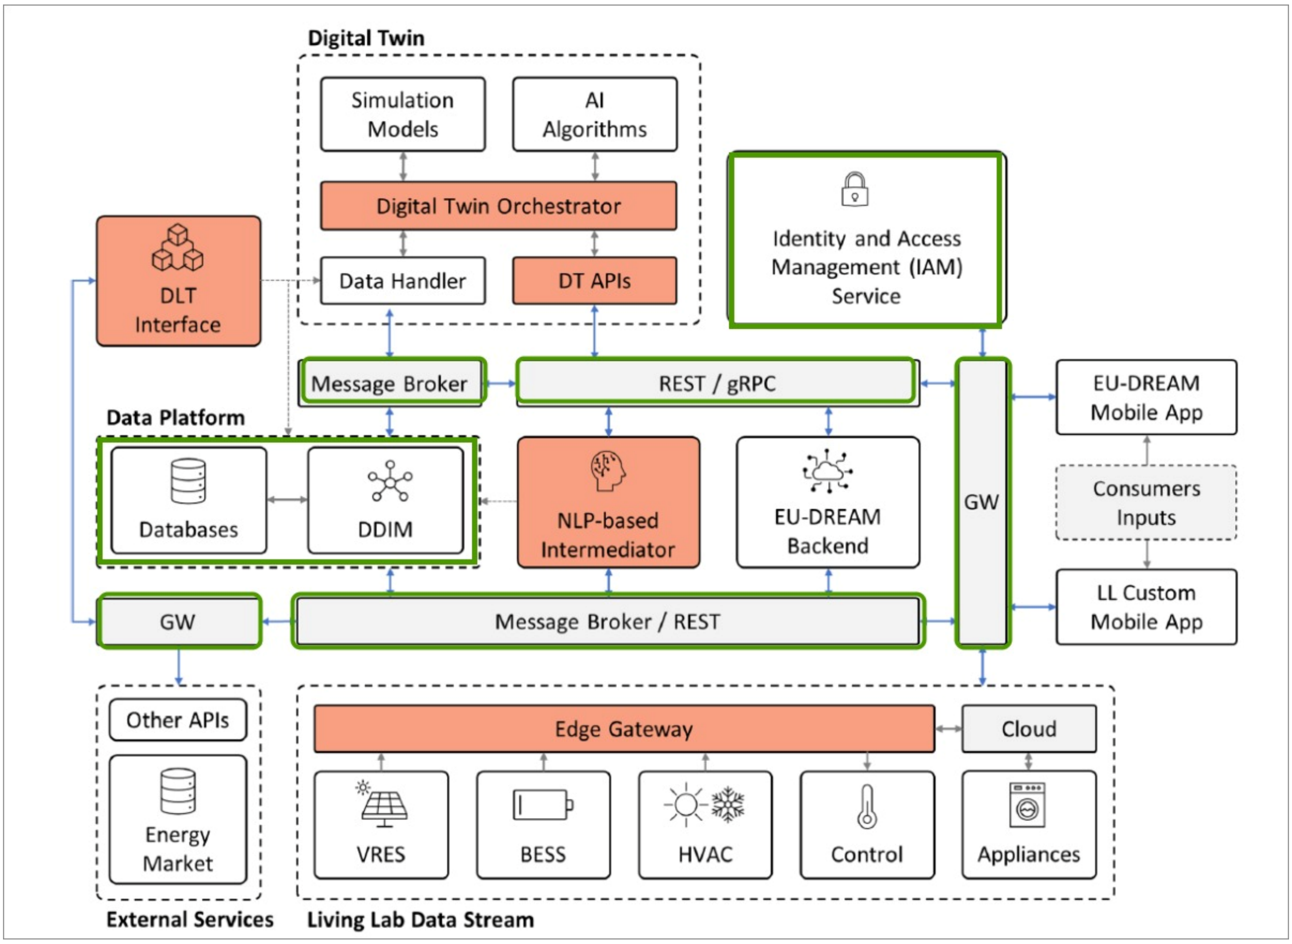

Figure 2: EU-DREAM Data Platform Components in the Overall Architecture (Links credits)

4 Requirement collection

The outcome of the requirements collection activity is a set of concise sentences that describe what the Data Platform shall do and how well it shall do it. Each requirement is filed under a clear category —data Ingestion, storage, integration & processing, access, visualization & Monitoring, security and Infrastructure—so developers, integrators and auditors can all turn to the same reference during design, implementation, and validation. By pinning these requirements down at the start, the EU-DREAM project keeps the platform compliant with EU regulations, interoperable across different energy systems, and robust enough to run DT, DLT, NLP-Intermediator, and future extensions that could be implemented.

4.1 Overall requirements

The following section presents the overall requirements of the EU-DREAM Data Platform. The requirements have been structured into architectural layers which represents a logical grouping of functionalities that address specific needs of the data platform. By structuring the architecture into layers, the system ensures modularity (i.e., each part of the platform can be developed, updated, or replaced independently), as well as clarity of responsibilities, and easier integration between components and Living Labs.

Table 2 summarizes the overall requirements and functions defined for each layer of the EU-DREAM Data Platform architecture. It reflects that the data platform aimed to ensure data integrity, modularity, and full compliance with [3] data principles. These high-level capabilities enable secure, interoperable, and scalable integration of components across modules and Living Labs.

Table 2: EU-DREAM Data Platform Architectural Layers and General Requirements

Layer	Description
Data Ingestion layer	Data ingestion within the EU-DREAM data platform shall have a set of capabilities that capture, validate, and register various types of data produced by the DT, NLP-Intermediator, Market place and the devices in each LL. It shall be able to receive every kind of input from sensor readings, control commands, query results, to large batch files and attached to each of one a "fingerprint" (a cryptographic hash) that shall prove where the data came from and whether it has been altered. Each ingested payload shall be passed through to the DLT, giving it an immutable timestamp and audit trail. Because of this process, any analytics, optimization, or reporting module that later consumes the data could trust that the information was both authentic and current. To summarize, data ingestion shall function as a carefully managed gateway that funnels clean, verifiable, real-time data into the platform's storage, processing, and decision-support layers.

[3] FAIR stands for *Findable, Accessible, Interoperable, and Reusable*.

Layer	Description
Data Storage layer	The system shall store all required data in a structured, and interoperable manner, ensuring compatibility across EU-DREAM components. The data platform also tracks user information—such as comfort limits and other personal preferences. In short, the storage layer provides a single, trustworthy foundation that stays organized, remains compatible with all EU-DREAM services, and scales effortlessly as the project expands.
Data Integration & processing Layer	The Integration and Processing Layer shall federate datasets from distributed sources and harmonize them into a shared semantic model, ensuring interoperability across domains. It shall convert heterogeneous formats into a standard and uniform schema, enrich them with metadata and contextual relationships, and publish them in a FAIR-compliant catalogue to make all assets discoverable and reusable. The layer shall provide standardized query Interfaces for accessing and combining datasets, while performing pre-processing tasks such as validation, cleaning, aggregation, and transformation to guarantee consistency and readiness for analytics and optimization engines. It shall preserve provenance and traceability, maintain data under the control of its owners, and comply with privacy and sovereignty principles. Finally, it shall be adaptable to new data sources, and user-friendly in managing imports, exports, and contextual frameworks, thereby fostering trust and transparency in EU-DREAM services.
Data Access layer	EU-DREAM's data-access layer shall act as the data platform's open gateway, exposing well-documented, standards-based APIs let partner applications invoke any service without vendor lock-in or other proprietary obstacles. Thanks to federated querying, the layer shall be able to provide data from multiple locations—cloud, edge, or local servers—and present them as a single, seamless source. Whether a query targets real-time streams, historical records, or a mix of both, the access layer handles it in the same straightforward way, making EU-DREAM's data easy to find, query, and reuse.
Visualization & Monitoring layer	The Data Visualization and Monitoring Layer shall deliver real-time, interactive dashboards. It shall rely on streaming interfaces to push soft real-time updates support rich visualizations. These dashboards shall include clear, traceable metadata for transparency and within the platform's standard, secure data-access services. This layer shall act as a live performance monitor, displaying metrics like response times and resource consumption, and shall maintain a comprehensive, timestamped audit log of all data actions. By combining real-time observations with detailed historical logs, the system shall enable operators to detect bottlenecks, ensure compliance, and trace subsystems' activities reliably.

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

Layer	Description
Data Security - Privacy and Sovereignty layer	The EU-DREAM Data Platform and its security-governance layer shall act as a steadfast guardian of information. It should deliver trust, resilience, and compliance by default, guaranteeing confidentiality, integrity, and availability at every stage. Every dataset should be protected end-to-end with strong encryption, traced by immutable provenance records on a permissioned DLT, and subject to continuous access auditing. Information should remain encrypted both at rest and in transit, while personal data should be automatically anonymized or pseudonymized to satisfy GDPR. All traffic and storage shall follow security controls describe in ISO/IEC 27002 (Strong encryption, access-control, security, and vulnerability management) and the energy-sector cybersecurity standard IEC 62351 (as specific safeguards: TLS or SSH tunnels). Finally, fine-grained authorization should ensure that only verified users or services can view or act on a given asset, and any attempt to modify a record without proper credentials must be blocked.
Data Infrastructure layer	The EU-DREAM Data Platform architecture must be modular and repeatable, so that any new location or project can deploy the same components in just a few minutes. Thanks to Gaia-X and other EU projects mechanisms —technical, governance tools and rules—, the data platform can link multiple clouds or edge nodes as one virtual space, growing smoothly as data volumes or participant numbers rise. In short, the infrastructure must remain flexible and scalable—and easy to replicate so that future growth in users or data never turns into a bottleneck.

A comprehensive list of general requirements is provided in Annex A. Which details each requirement with its identifier, description, and source. The requirements have been collected from the Grant Agreement (GA), the Data Management Plan (DMP), and deliverable D1.2, ensuring consistency with the project objectives.

4.2 Data Platform Requirements to Support Integration of Other Components

The requirements derived from other tasks specify what the EU-DREAM Data Platform must provide in order to support the seamless integration of all project components and services. These requirements make sure that the platform is reliable, secure, and interoperable, while respecting the ownership of data.

4.2.1 EU-DREAM components integration requirements

From Task 2.2, "DT Core Infrastructure", the EU-DREAM Data Platform shall ensure that every architectural component can interface with it through well-defined, event-driven integration contracts. From the registration of a new component to the continuous ingestion of data, the platform must guarantee [4]￼, validation, provenance, and bidirectional control.

4.2.2 Data Platform Requirements supporting simulation models execution.

From the deliverable D2.1, "DT models", the EU-DREAM Data Platform shall provide high-performance storage and[5]to support simulation models and Digital Twins. All simulation data and results must be traceable and immutable, ensuring seamless interoperability across the system.

4.2.3 Data Platform requirements the AI-based Assistant Tool and NLP intermediator

From Task 2.3, "AI-based Assistant Tool" and Task 2.4, "NLP-based Intermediator", the EU-DREAM Data Platform shall provide a unified API gateway for secure access to datasets and services, ensuring ensure low-latency responses[6] that support natural conversational interaction while maintaining compliance and scalability. For Task 2.3, this enables the AI-based Assistant to apply optimisation algorithms that generate personalised consumption plans and recommendations based on data, forecasts, and user preferences. For Task 2.4, it ensures the NLP Intermediator can deliver smooth and user-friendly interactions.

4.2.4 Data integration and Data Interoperability requirements

The EU-DREAM Data Platform shall guarantee that the data from different sources can be understood and used together. To achieve this, the platform shall integrate [7] server as a core architectural component, in line with Task 3.2.

For a complete view, the detailed list of requirements derived from these tasks is presented in the Annex B.

---

[4] In this context, *low-latency transfer* refers generally to transfers fast enough to support real-time monitoring and responsive control.

[5] In this context, *low-latency communication* refers to the fast data exchange needed to enable timely execution of simulations and real-time synchronization of DT.

[6] In this context, *low-latency responses* refer to system replies fast enough to sustain natural conversational flow when interacting with the NLP-Intermediator

[7] The Dynamic District Information Management (DDIM) server

5 Technical Specifications

The present technical specifications will translate the requirements presented in Chapter 6 into concrete, stable, directives. In order to ensure the implementation of a functional and interoperable data platform. The specifications also include the selection of key open-source components aligned with the requirements.

5.1 EU-DREAM Data Platform Architecture and Component Specifications

The EU-DREAM Data Platform follows a layered logical architecture. Each layer is responsible for a specific set of functions, from data ingestion to storage, access, and visualization. Cross-cutting concerns such as security, privacy, sovereignty, and monitoring are integrated across all layers. Below, each layer is explained and accompanied by its Component Specification.

Figure provides a high-level overview of this architecture, showing how the different layers and their components interact to support the complete data lifecycle, from ingestion to visualization.

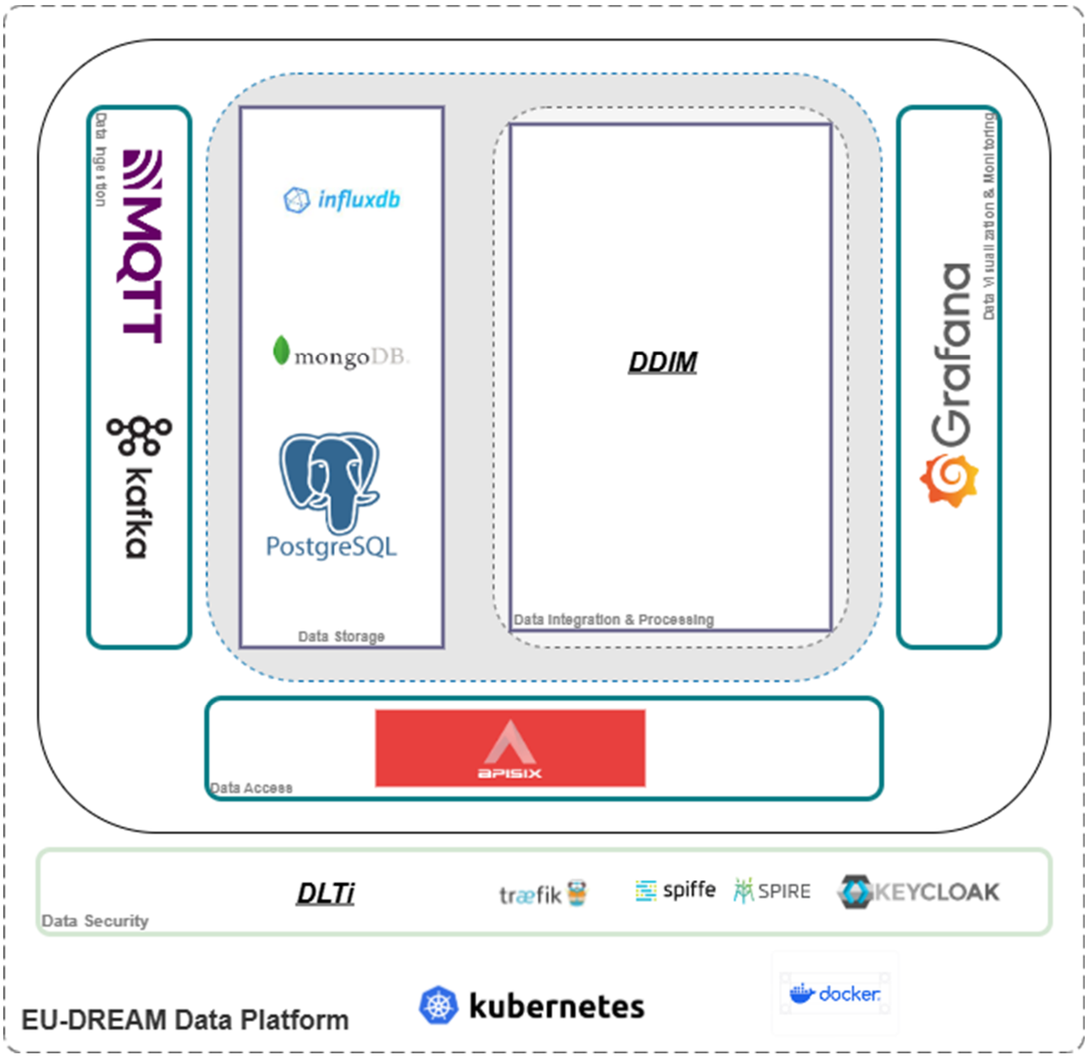

Figure 3: EU-DREAM Data Platform

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

The following table (Table 3) describes each platform layer and presents its Component Specifications, including selected components, main functionalities, and logical interfaces.

Table 3: EU-DREAM Data Platform Layers and Component Specifications

Layer	Description		
Data Ingestion layer	This layer is the front door of the platform. It captures every type of input from the Living Labs and external services. Each batch of data is given a unique ID, timestamp, and cryptographic fingerprint to prove where it came from and that it hasn't been altered. By going through brokers like Kafka (optimized for high-throughput data streaming) or MQTT (lightweight and suited for IoT device communication), the flow is fast and reliable. In short, this layer makes sure all incoming data is authentic, secure, and ready for use by other parts of the system		
	Selected components	**Main functionalities**	**Main logical interfaces**
	Apache Kafka	• Publish/subscribe to real-time streams. • Stores events durably in order, enabling replay. • Scales across producers and consumers. • Handles bursts and ensures no data loss.	• Kafka topic interface: Partitioned streams with offset control and replay.
	MQTT Broker	• Lightweight publish/ subscribe protocol for IoT. • Very low bandwidth and power usage. • Reliable delivery even with unstable connections. • Supports retained messages and will messages.	• MQTT broker interface: Publish/subscribe with QoS and retain options
Data Storage layer	The Storage Layer follows a lake house approach, combining the flexibility of a data lake with the organization of databases. This ensures that both real-time and historical data can be stored efficiently and accessed when needed for control, monitoring, analytics, and compliance. A unified catalogue records schema, quality, retention, and access rights to keep data FAIR.		
	Selected components	**Main functionalities**	**Main logical interfaces**
	MongoDB	• Document database storing flexible JSON-like data. • Supports indexing, aggregation pipelines, search. • Scales horizontally across clusters.	• Document API: REST/GraphQL façade for document CRUD, sharing, change streams

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

EU-DREAM

	InfluxDB	• Specialized for time-series data. • Handles millions of writes per second. • Retention and down sampling policies. • Fast queries for real-time metrics and monitoring.	• Time-series API: Write/query endpoints.
	PostgreSQL	• Advanced relational database. • Powerful SQL queries, joins, and analytics. • Reliable for governance and regulatory records.	• SQL API: ANSI SQL queries and stored procedures.
Data Integration & Processing layer	Once data is ingested, this layer transforms it into a common, understandable language for the whole stakeholders. It applies shared vocabularies and links data into semantic graphs, recording provenance. Information is federated and enriched so that services and partners can query across sources as if they were one. This semantic approach ensures that data is interoperable, searchable, and can be reused fairly under EU standards		
	Selected components	**Main functionalities**	**Main logical interfaces**
	DDIM Server	• Converts diverse formats into RDF linked data. • Records provenance and ownership. • Provides SPARQL endpoints for semantic queries. • Publishes metadata catalogs (DCAT) for discovery.	• SPARQL endpoint: Semantic queries. • DCAT catalog: Metadata registry.
Data Access layer	This layer is the gateway through that each partner and apps will interact with EU-DREAM Data Platform. It offers secure, well-documented APIs that allow queries and subscriptions to live streams. It gives a single-entry point for accessing real-time and historical data across edge, cloud, and local servers. An API gateway controls who can see what, protects against misuse, and logs all activity for compliance		
	Selected components	**Main functionalities**	**Main logical interfaces**
	Apache APISIX	• Cloud-native API gateway. • Routing, load balancing, traffic splitting. • Authentication (OAuth2, JWT, key auth). • Rate limiting, observability, plugin system	• REST/GraphQL gateway: Queries and subscriptions. • Admin API: Manage routes and plugins.
	Operators and users see the state of the platform real time. Dashboards and alerts highlight performance, reliability, and energy footprint, while logs and		

<table>
<tr><td rowspan="3">Visualization & Monitoring layer</td><td colspan="3">traces help investigate issues. Makes complex data visualizations into clear trends, for experts, they enable deep diagnostics and compliance reporting.</td></tr>
<tr><td>Selected components</td><td>Main functionalities</td><td>Main logical interfaces</td></tr>
<tr><td>Grafana</td><td>Interactive dashboards for metrics, logs, traces.
Alerts and notifications.
Works with Prometheus, Loki, Tempo.
Extensible with plugins.</td><td>Grafana HTTP UI: Dashboard/alert
REST API: Programmatic dashboard, data source, and alert management</td></tr>
<tr><td rowspan="6">Security, Privacy & Data Sovereignty layer</td><td colspan="3">Guarantees confidentiality, integrity, and availability in line with IEC 62351, the key indicators for smart grid cybersecurity. Every dataset is encrypted, verified, and auditable. It ensures GDPR compliance and allows organizations to keep full control of their data</td></tr>
<tr><td>Selected components</td><td>Main functionalities</td><td>Main logical interfaces</td></tr>
<tr><td>Keycloak</td><td>Identity and access management.
Single Sign-On (SSO).
OAuth2/OIDC, SAML support.
User federation & social login.
Fine-grained roles and permissions.</td><td>OAuth2/OpenID Connect: Authentication flows.
Keycloak REST API: Manage users and roles.</td></tr>
<tr><td>SPIFFE/SPIRE</td><td>Provides secure identities (SVIDs) for workloads.
Attests nodes and workloads automatically.
Issues and rotates short-lived certificates.
Enables zero-trust communication between services.</td><td>SPIFFE Workload API</td></tr>
<tr><td>DLT API</td><td>Anchors cryptographic hashes of data in a distributed ledger.
Provides tamper-evident audit trails.
Builds trust across organizations without a central authority.</td><td>Hash Verification API: Submit/verify hashes; query immutable audit logs</td></tr>
<tr><td>Traefik Mesh</td><td>Lightweight Kubernetes service mesh.</td><td>Service Mesh proxy (SMI interface)</td></tr>
</table>

		• Encrypts service-to-service traffic (mTLS). • Provides load balancing, retries, circuit breaking, rate limiting. • Observability via metrics and tracing.	
Infrastructure layer	The foundation of the data platform runs on edge, local servers, or cloud. Uses containers and orchestration so services run the same everywhere, making deployments consistent across environments and simplifying updates, recovery, and scaling. Ensures repeatable, secure deployments across Living Labs		
	Selected components	**Main functionalities**	**Main logical interfaces**
	Docker	• Packages apps and dependencies into portable containers. • Ensures consistent runtime across environments. • Provides isolation for stability and security. • Simplifies build, share, and update of software.	• Docker CLI: Command-line tool to build, run, and manage containers. • Docker REST API: Lets other programs automate and control the Docker engine.
	Kubernetes	• Orchestrates containers across machines. • Monitors health and restarts failed apps. • Autoscaling of workloads. • Manages deployments, replicas, and service discovery.	• Kubernetes API server (REST): Central interface for creating, reading, updating, and watching all cluster resources.

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

6 Data Space

Data spaces are an important part of Europe's plan for a fair and trusted digital economy. They are not single platforms or big databases, but open ecosystems where different actors can share and use data safely. Each participant keeps control of their own data, while common rules, open standards, and trust services make sure the exchange is secure and fair.

European initiatives like the Data Spaces Support Centre highlight three main ideas:

- Data sovereignty meaning keeping control of your own data.
- Interoperability, where systems can work together across organizations.
- Openness, no lock-in, equal access and equal opportunities for participants.

In EU-DREAM, the project aims to apply these principles as much as possible within the capacities of the data platform. This ensures that the DP remains aligned with the vision for trusted data sharing.

6.1 Background

With the progress of digitalization, the need for trusted and interoperable environments for data exchange has become essential. Data spaces provide such environments by connecting existing systems through federated infrastructures and by relying on open standards, decentralized identity, usage policies, and interoperability services. The EU-DREAM Data Platform aims to integrate several building blocks inspired by data space concepts, making it data-space-ready for possible future integration.

6.1.1 Definitions

Table 44: Key Definitions for Data Spaces

Term	Definition
Data Space	A federated data ecosystem where participants (data providers, consumers, and intermediaries) share data based on mutual trust, common governance, and technical interoperability, ensuring data sovereignty and compliance with shared policies and legal requirements.
Data Sovereignty	The principle that data owners retain control over who accesses their data, for what purpose, and under which conditions.
Connector	A software component used by each participant of a data space to publish, discover, and exchange data. It ensures secure data transfer and policy enforcement.
Federated Catalogue	A distributed directory of available data products and services, enabling metadata-based search and discovery across the data space.
Identity and Trust Framework	Mechanisms that verify and authenticate participants and services using decentralized identifiers (DID) and verifiable credentials (VC).
Usage Policy	A machine-readable set of rules defining how data may be used once it is accessed or transferred, often using standards such as ODRL (Open Digital Rights Language).
Compliance Service	Validates the conformity of participants and services with technical and legal standards required by the data space or federation.

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

6.2 Application of the Data Space concept in the EU-DREAM Data Platform

A typical data space architecture brings together components such as connectors, federated catalogues, identity and trust services, and compliance and logging mechanisms to guarantee data sovereignty, interoperability, and secure exchange. Within EU-DREAM, the Data Platform integrates analogous functions through its selected technologies, allowing it to reflect and operationalize these principles of sovereignty, interoperability, and openness in line with the technical design.

The correspondence between these architectural principles and the EU-DREAM Data Platform can be observed in the way its components have been selected and integrated. The table below summarizes how the platform's technologies and services align with the typical building blocks of data spaces, demonstrating the extent to which the project incorporates sovereignty, trust, interoperability, and openness into its design.

Table 55: Alignment of EU-DREAM Data Platform with Data Space Principle

Principal	Data Platform Technical Specification
Connector Integration	Kafka MQTT and Apisix enable secure, policy-controlled exchange of data streams. Kafka manages ordered event streaming, MQTT supports lightweight publish/subscribe for IoT devices, and APISIX adds routing, authentication, and logging for traceable exchanges.
Federated Identity and Trust Framework	Keycloak (federated login, RBAC) and SPIFFE/SPIRE (workload identities), aligned with Gaia-X and IDS requirements.
Usage Policy Enforcement	APISIX enforces usage rules through OAuth2/JWT issued by Keycloak, ensuring fine-grained access control and traceable interactions.
Federated Catalogue Capabilities	DDIM server provides RDF/DCAT metadata and provenance, enabling semantic discovery and interoperability.
Interoperable Interfaces and APIs	REST, GraphQL, and Kafka interfaces support both queries and event subscriptions, ensuring openness and easy integration with external systems
Data Provenance and Auditability	DLT interface anchors hashes of data to provide immutable audit trails, ensuring tamper-evidence and compliance with security standards.
Monitoring and Observability	Grafana offers dashboards, logs, and alerts for transparency and proactive detection of issues across services.
Scalable and Distributed Infrastructure	Docker and Kubernetes support modular, distributed deployments across cloud, edge, and on-premises environments, ensuring scalability and resilience..

7 CONCLUSION

The specifications and requirements described in this report are intended to support the implementation of a cloud-based, modular, scalable, and secure data platform, making it capable of integrating AI services, DT models, and an NLP-based intermediator, while enabling modular deployment across various operational contexts.

This deliverable has presented the technical specifications, architectural layers, and component interdependencies necessary for the development of the EU-DREAM Data Platform. It consolidates the requirements identified through system analysis and alignment with European standards for interoperability, data sovereignty, and trusted data exchange.

The platform incorporates key principles of data spaces within its design. By embedding elements such as connectors, catalogues, trust and identity management, policy enforcement, provenance, and observability, it is positioned to meet the project's needs while leaving open the possibility of evolving towards broader cross-sector data sharing in the future.

This document represents the current state of development and system design at the time of writing. As the project advances and validation activities take place across the six Living Labs, new insights, constraints, and needs may emerge, leading to adjustments in the data platform design, the technical stack, or integration strategies. For this reason, the document should be seen as a dynamic reference that will be updated as needed to incorporate refinements and updates from implementation and stakeholder, ensuring that the EU-DREAM Data Platform remains aligned with technical requirements and user needs.

References

[1] Gaia-X. [Online]. Available: https://docs.gaia-x.eu/technical-committee/architecture-document/25.05/trust_framework_architecture/#terminology-sources.

[2] OPENDEI, "Design Principles for Data Spaces – Position Paper".

[3] "The International Data Spaces Association on GitHub," [Online]. Available: https://github.com/International-Data-Spaces-Association/idsa.

[4] R. H. e. al., "Interim Report Data Space Design," 2022.

Annexes

The following annexes consolidate the requirements identified within the EU-DREAM project. Annex A presents the general requirements, covering the overall needs of the data platform and Annex B specifies the requirements for the integration of the different EU-DREAM components, ensuring that all modules, services can interconnect through standardized and interoperable interfaces.

Annex A

Req. ID	Requirement Title	Description	Category	Author	Source
GA-001	Integrate WP1 data stream	The system shall integrate the data identified by WP1 and required by the EU-DREAM system.	Data Integration	EU-DREAM Consortium	GA
GA-002	Collect and serve energy data	The system shall collect and serve the energy data expected by the EU-DREAM services.	Data Ingestion	EU-DREAM Consortium	GA
GA-003	Collect real time data	The system shall be able to collect real-time data generated by the physical domain.	Data Ingestion	EU-DREAM Consortium	GA
GA-004	Unified Data Storage Service	The platform shall provide a persistent data-lake that can ingest, store, and serve all raw, integrated and operational datasets required by every EU-DREAM component. The capacity shall scale with demand and guarantee availability to downstream modules	Data Storage	EU-DREAM Consortium	GA
GA-005	Secure data ingestion	The system shall ensure that no real-world data samples are missed during data ingestion.	Data Ingestion	EU-DREAM Consortium	GA
GA-006	Provide a cloud-based solution	The platform shall be designed as a cloud-based solution.	Data Infrastructure	EU-DREAM Consortium	GA
GA-007	Secure data	The platform shall ensure the security of data shared and used by the components of the EU-DREAM system.	Data Security - Privacy and Sovereignty	EU-DREAM Consortium	GA
GA-008	Interoperable Interfaces for Cross-Module Data Exchange	The system shall provide standard, interoperable interfaces and supporting mechanisms that allow all data sources and EU-DREAM	Data Integration	EU-DREAM Consortium	GA

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

		modules to exchange and integrate data seamlessly across the entire platform.			
GA-009	Federated Integration of Local & Remote Data Sources	The platform shall enable a federated data space in which multiple distributed data sources—whether imported and managed locally by the server or hosted remotely—can be accessed and administered as one unified virtual dataset without physically consolidating the data in a single location.	Data Integration	EU-DREAM Consortium	GA
GA-010	Incorporate semantics and the linked data principle	The platform shall incorporate semantics and linked-data principles to facilitate interoperability among different data sources.	Data Integration	EU-DREAM Consortium	GA
GA-011	Open, Standardized APIs with Public Specifications	The system shall expose publicly documented, open-standard APIs—built on widely accepted protocols—and make their full interface specifications available to all EU-DREAM partners for consistent access to every platform service.	Data Access	EU-DREAM Consortium	GA
GA-012	DLT-based Trusted & Immutable Data-Exchange Layer	The system shall integrate a Distributed-Ledger-Technology (DLT) layer that records every data exchange for provenance and guarantees that all stored transactions and data remain immutable unless a formally authorized change is issued.	Data Security - Privacy and Sovereignty	EU-DREAM Consortium	GA
GA-013	Define a common structure	The system shall define a common structure that allows data in various formats to be mapped and transformed into a single overarching linked-data space.	Data Integration	EU-DREAM Consortium	GA
GA-014	Integrate a semantic broker	The system shall integrate a semantic broker to combine data from various connectors, acting as an intermediary that processes and integrates data into a cohesive data lake.	Data Integration	EU-DREAM Consortium	GA

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

GA-015	Federated Query & SPARQL Access	Users shall be able to issue queries—including SPARQL queries—across multiple remote or local data sources.	Data Access	EU-DREAM Consortium	GA
GA-016	SAREF Ontology Adoption	The system shall integrate the SAREF ontology and the necessary extensions.	Data Integration	EU-DREAM Consortium	GA
GA-017	Comprehensive Data Catalogue	The system shall include a data catalogue that provides a comprehensive inventory of all datasets, including metadata.	Data Access	EU-DREAM Consortium	GA
GA-018	Offer visualization tools	The system shall offer visualization tools.	Data visualization	EU-DREAM Consortium	GA
GA-019	Handle streamed data.	The system shall be able to handle streamed data.	Data processing	EU-DREAM Consortium	GA
GA-020	Historical Data Access	The system shall allow queries to access historical data stored in the platform.	Data Access	EU-DREAM Consortium	GA
GA-021	Manage access to data	The system shall allow users to manage access to data generated by connected assets.	Data Security - Privacy and Sovereignty	EU-DREAM Consortium	GA
GA-022	Comply with GDPR regulation	The system shall comply with the GDPR.	Data Security - Privacy and Sovereignty	EU-DREAM Consortium	GA
GA-023	Implement CIM-compliant data models.	The system shall implement CIM-compliant data models.	Data Integration	EU-DREAM Consortium	GA
GA-024	Monitor performance metrics	The system shall monitor performance metrics to facilitate analysis of system scalability.	Data Platform Monitoring	EU-DREAM Consortium	GA
GA-025	Log data operation	The system shall capture and log every operation on data for auditability.	Data Platform Monitoring	EU-DREAM Consortium	GA
GA-026	Scalable, Replicable Architecture	The system shall provide a scalable, replicable technical architecture.	Data Infrastructure	EU-DREAM Consortium	GA
GA-027	ISO/IEC 27002 and IEC 62351 Compliance	The system shall comply with ISO/IEC 27002 and IEC 62351 security standards.	Data Security - Privacy and Sovereignty	EU-DREAM Consortium	GA

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

GA-028	Metadata and Naming Convention	The system shall use consistent naming conventions and metadata for datasets.	Data Integration	EU-DREAM Consortium	GA
DMP-001	Data Sensitivity Classification	Collect only essential data and label each item by sensitivity level.	Data Security – Privacy and Sovereignty	UPO	DMP
DMP-002	FAIR Data Governance	Apply the FAIR principles so every dataset is Findable, Accessible, Interoperable, and Reusable.	Data Integration	UPO	DMP
DMP-003	Non-Energy Data Management	Gather and serve all non-energy data needed alongside energy data for EU-DREAM users and components.	Data Ingestion	UPO	DMP
DMP-004	Role-Based Dataset Security	Limit dataset access to authorized users via strong authentication and authorization mechanisms.	Data Security – Privacy and Sovereignty	UPO	DMP
DMP-005	Energy Consumption Dashboard	Provide real-time and historical visualizations of energy use for individuals or groups.	Data Visualization	UPO	DMP
DMP-006	Be compliant with IEC 62351 Cybersecurity Standards	Secure energy-system communications by implementing IEC 62351 standards.	Data Security – Privacy and Sovereignty	UPO	DMP
DMP-007	Be compliant with ISO 27002 Security Standards	Protect data confidentiality, integrity, and availability by meeting ISO/IEC 27002 guidelines.	Data Security – Privacy and Sovereignty	UPO	DMP
DMP-008	Scalable Modular Platform	Design an energy-efficient, DLT-supported architecture that grows with data volume and participants.	Data Infrastructure	UPO	DMP
DMP-009	Consistent Dataset Metadata & Naming	Attach rich metadata and follow structured naming conventions for easy retrieval and context.	Data Integration	UPO	DMP
DMP-010	Personalized Preference Storage	Store user comfort limits and usage preferences to enable tailored services and optimization.	Data Storage	UPO	DMP

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

DMP-011	End-to-End Data Encryption	Encrypt sensitive data—especially AI-related datasets—both at rest and in transit.	Data Security – Privacy and Sovereignty	UPO	DMP
DMP-012	Privacy-Preserving Anonymization	Anonymize or pseudonymize personal data whenever required to safeguard sensitive information.	Data Security – Privacy and Sovereignty	UPO	DMP
D1.2-001	Modular and Scalable Architecture	The architecture shall be built in self-contained modules that can be extended or modified over time. It should allow new services, tools, or pilot sites to be added without impacting existing features and be able to scale across various regions, user groups, and energy infrastructures.	Data Infrastructure	LINKS	D1.2
D1.2-002	Data Interoperability	The platform shall ingest and exchange data seamlessly from diverse sources—such as legacy IT systems, IoT sensors, building management platforms, weather APIs, and smart meters—by leveraging open standards and APIs to integrate internal and third-party components.	Data Integration	LINKS	D1.2
D1.2-003	Simulation and Forecasting Tools	The solution shall enable advanced analytics to predict future energy consumption, generation, and price movements, empowering stakeholders to plan strategically, make informed operational choices, and engage proactively with the grid.	Data Processing	LINKS	D1.2
D1.2-004	Role-Based Access Control	The platform shall enforce a robust access-control framework that grants users specific permissions and privileges according to their assigned roles.	Data Security - Privacy and Sovereignty	LINKS	D1.2
D1.2-005	Real-Time Monitoring	The platform shall continuously ingest and process live data streams, offering operators and end-users up-to-the-second	Data Ingestion	LINKS	D1.2

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

		visibility into energy generation, usage, and flexibility metrics.			
D1.2-006	Energy Flexibility Services	The platform shall include features for managing demand response, shifting loads, and orchestrating distributed energy assets to optimize flexibility services automatically.	Data Processing	LINKS	D1.2
D1.2-007	Integration with Market Platforms	The platform shall connect seamlessly to external trading venues—such as wholesale markets, local flexibility exchanges, or aggregator platforms—to enable market participation.	Data Integration	LINKS	D1.2
D1.2-008	Compliance with Regulatory Frameworks	Every element of the platform shall comply with applicable EU and national laws and standards—covering data protection (e.g., GDPR), energy market regulations, cybersecurity, and interoperability.	Data Security - Privacy and Sovereignty	LINKS	D1.2

Annex B

Req. ID	Requirement Title	Description	Category	Author	Source
LINKS-026	Broker endpoint for the Digital Twin	Offer a documented message-broker where the Digital Twin registers, publishes telemetry and subscribes to control commands.	Data Access	LINKS	T2.2
LINKS-027	Read queries pushed on Kafka channels	Configure Kafka topics for incoming query requests; the platform consumes these messages, routes them to the correct data service and prepares a reply.	Data Access	LINKS	T2.2
LINKS-028	Databases queryable through GraphQL	Provide a GraphQL layer so clients can issue flexible, typed queries against operational databases and receive JSON responses with only the requested fields.	Data Access	LINKS	T2.2

LINKS-029	MQTT broker for the Digital Twin	Deploy a dedicated MQTT broker to carry high-frequency telemetry and control traffic between the Digital Twin and platform services.	Data Access	LINKS	T2.2
LINKS-030	Publish query results on MQTT	After executing a query, publish the results to a designated MQTT topic so Digital Twin clients receive near-real-time updates.	Data Access	LINKS	T2.2
LINKS-031	Return Living-Lab (LL) data with DLT metadata	Whenever LL data are returned, attach provenance metadata (origin, timestamp, hash) issued by the Distributed Ledger Technology layer.	Data Integration	LINKS	T2.2
LINKS-032	MQTT broker for DLT traffic	Run a separate MQTT channel specifically for ledger-related messages, keeping DLT events isolated from regular operational data.	Data Access	LINKS	T2.2
LINKS-033	Write LL data to DLT topic (add UUID)	On ingestion, publish each LL record to the DLT MQTT topic and append a universally unique identifier to guarantee traceability.	Data Security - Privacy and Sovereignty	LINKS	T2.2
LINKS-034	Read DLT messages from MQTT	Subscribe to the DLT topic, parse incoming ledger events and forward verified updates to downstream components.	Data Integration	LINKS	T2.2
LINKS-035	Update stored data with DLT data	When ledger updates arrive, synchronize the platform's persistent storage to maintain consistency with the authoritative DLT record.	Data Storage	LINKS	T2.2
LINKS-036	MQTT broker for LL (LL)	Provide an MQTT endpoint that LL devices/services use to stream sensor data into the platform and receive control commands.	Data Access	LINKS	T2.2
LINKS-037	REST APIs for LL push	Expose RESTful endpoints so LL can push batch or ad-hoc records when MQTT is unsuitable.	Data Access	LINKS	T2.2

Horizon-CL5-2023-D3-03
Grant Agreement 101160614

LINKS-038	Pull LL data via their APIs	Implement client routines that periodically call Living-Lab REST APIs to retrieve data when those labs publish information passively rather than pushing it.	Data Access	LINKS	T2.2
D2.1-001	CSV & Image File Handling	The platform shall support reading and writing CSV files and storing additional image files generated by agents.	Data Ingestion	AKKO	D2.1
D2.1-002	Simulation Results Import/Export	The platform shall allow ingestion and export of simulation results in time-series table formats as well as related visual outputs.	Data Ingestion	AKKO	D2.1
D2.1-003	Unified Publish/Subscribe Message Bus	The platform shall provide a publish/subscribe message bus to enable communication between appliance models, DT agents, EMS, and NLP interfaces.	Data Access	AKKO	D2.1
T2.4-001	Secure API Gateway	The data platform shall expose a single, well-guarded entry point that routes every request, applies central security rules, and keeps traffic organised.	Data Access	AKKO	T2.4
T2.4-002	REST APIs	The data platform shall deliver REST endpoints that answer different components calls	Data Access	AKKO	T2.4
T2.4-003	Real-Time Message Broker	The data platform shall queue and forward messages fast enough to keep chats smooth during traffic, while letting back-end services scale independently.	Data Processing	AKKO	T2.4
T2.4-004	Right storage for each data type	The data platform shall store heterogenous data in optimised engines to ensure fast searches	Data Storage	AKKO	T2.4
T2.4-005	Plug-and-play chat flows	The data platform shall support modular conversational workflows so new language-model	Data Processing	AKKO	T2.4

	(LangGraph)	features can be deployed or rearranged without full redeployment.			
T2.4-006	MCP server for many parallel streams	The data platform shall process multiple chat sessions at the same time, spreading workloads across channels and keeping spares ready to avoid downtime.	Data Infrastructure	AKKO	T2.4
T3.2-001	Open API & SPARQL Service	Offer queries that return results in JSON or CSV via standards-based endpoints.	Data Access	AKKO	T3.2
T3.2-001	Cloud-Native Horizontal Scaling	Scale out with open-source, cloud-native components for elastic capacity.	Data Infrastructure	AKKO	T3.2
T3.2-001	Containerised Connector Onboarding	Let new connectors be onboarded plug-and-play via containerised micro-services.	Data Infrastructure	AKKO	T3.2

Horizon-CL5-2023-D3-03
Grant Agreement 101160614