



ECOLOOP

www.ecoloop-project.eu

D4.3 Decision support tool for farmers and foresters final version

31 March 2026



**Funded by
the European Union**

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. Horizon Europe Grant agreement N° 101118127.

Deliverable details

Title	WP	Version
D4.3 Decision support tool for farmers and forester's final version	4	V1.0

Contractual delivery date	Actual delivery date	Delivery type*	Dissemination**
March 2026 (M30)	March 2026 (M30)	R	PU

* Delivery type: R: Document, report; DEM: Demonstrator, pilot, prototype; DEC: Websites, patent fillings, videos, etc; DMP: Data management plan.

** Dissemination Level: PU - Public; SEN – Sensitive, limited under the conditions of the Grant Agreement.

Author(s)	Organization
Borja Albert Gramaje	ETRA I+D

Version	Date	Person	Action	Status***
V0.1	12/03/2026	Borja Albert Gramaje (ETRA)	Table of contents proposal	Draft
V0.1	16/03/2026	Borja Albert Gramaje (ETRA)	Chapter 1 & 2	Draft
V0.2	17/03/2026	Borja Albert Gramaje (ETRA)	Chapter 3	Draft

V0.3	20/03/2026	Borja Albert Gramaje (ETRA)	Chapter 4	Draft
V0.4	26/03/2026	Borja Albert Gramaje (ETRA)	Chapter 5	Draft
V0.5	27/03/2026	Borja Albert Gramaje (ETRA)	Chapter 6	Draft
V1.0	31/03/2026	Borja Albert Gramaje (ETRA), Sheila Sánchez Mínguez (ETRA)	Final review and submission of final version	Submitted

***Status: Draft, Final, Approved, Submitted (to European Commission).

Keywords

Internet of Things (IoT), sensors, analytics, big data, communication protocol, data, architecture, integration, hardware, decision support system, monitoring, renewable energy sources.



Index

1. INTRODUCTION	8
1.1. Purpose of the document	8
1.2. Scope of the document	8
1.3. Structure of the document	9
2. DECISION SUPPORT TOOL ARCHITECTURE	10
2.1. Microservices Architecure and Event-Driven communications	11
2.1.1. Microservices	12
2.2. Data Storage Architecture	14
2.2.1. Databases	14
2.2.1.1. MongoDB	14
2.2.1.2. InfluxDB	14
2.2.1.3. RustFS / S3 storage	15
2.2.2. Architecture	15
2.3. Application Framework and Presentation Layer	16
2.4. AI-Driven Predictive Analytics Architecture	17
2.4.1. Architectural Integration and Data Flow	18
2.4.2. Contextual Inputs and Operational Value	18
3. AI EVOLUTION AND HUMAN-IN-THE-LOOP DECISION	19
3.1. The ECOLOOP Approach: Bridging the Cognitive Gap	20
4. DECISION SUPPORT TOOL APPLICATION	21
4.1. Empirical Data Storage and Edge Processing	22
4.2. Views	24
4.2.1. Login View	24
4.2.2. Home View	26
4.2.2.1. Realm View	26
4.2.2.2. Site View	27

4.2.2.3. Zone View	29
4.2.2.4. Device View	30
4.2.3. Map View	33
4.2.4. Rules View	34
4.2.5. Digital Twin View	36
4.2.6. Notifications View	38
 5. PILOT SITE IMPLEMENTATION	 40
5.1. Picassent (l'Horta Sud, INDEREN)	40
5.2. Silla Biogas Plant (GENIA)	41
5.3. Finca Sinyent (Spain)	42
5.4. IC Jablje (KIS)	43
5.5. Bulgarian Pilot (Momchil / Albena)	45
5.6. Estonian Pilots	46
 6. NEXT STEPS	 47
 7. CONCLUSION	 48
 8. REFERENCES AND ACRONYMS	 50
8.1. References	50
8.2. Acronyms	52



List of figures

Executive Summary

This deliverable presents the development and implementation of a comprehensive big data analytics framework for the ECOLOOP IoT ecosystem, supporting the project's pillars in energy, agriculture/forestry, and energy community services. The framework is designed to integrate and harmonise data streams from a wide variety of IoT devices and monitoring systems deployed across multiple pilot sites, each representing distinct use cases and operational environments.

For each pilot site, a tailored set of IoT devices and sensors is deployed to capture key parameters relevant to local energy production, agricultural processes, and environmental conditions. These devices include, but are not limited to, sensors for solar irradiance, ambient temperature, relative humidity, air pressure, wind speed and direction, rainfall, soil moisture, and crop health. The collected data is transmitted through standardised protocols and aggregated within the ECOLOOP ecosystem, ensuring interoperability and scalability.

The integration process involves mapping the specific monitoring systems and data requirements of each use case to the broader ECOLOOP data architecture. This enables the seamless flow of information from field devices to centralised analytics platforms, where advanced big data techniques are applied. These analytics support real-time monitoring, predictive maintenance, resource optimisation, and informed decision-making for stakeholders across the energy and agricultural sectors.

By leveraging big data analytics across the entire value chain, the ECOLOOP ecosystem enhances operational efficiency, sustainability, and resilience. The deliverable outlines the methodologies for device integration, data harmonisation, and analytics deployment, providing a replicable model for digital transformation in rural and community-based energy and agricultural systems

Copyright statement

The work described in this document has been conducted within the ECOLOOP project. This document reflects only the ECOLOOP Consortium view, and the European Union is not responsible for any use that may be made of the information it contains.

This document and its content are the property of the ECOLOOP Consortium. All rights relevant to this document are determined by the applicable laws. Access to this document does not grant any right or license on the document or its contents. This document or its contents are not to be used or treated in any manner inconsistent with the rights or interests of the ECOLOOP Consortium or the Partners

detriment and are not to be disclosed externally without prior written consent from the ECOLOOP Partners.

Each ECOLOOP Partner may use this document in conformity with the ECOLOOP Consortium Grant Agreement provisions.

1. Introduction

1.1. Purpose of the document

Building upon the groundwork laid in Deliverable 4.1 (D4.1)—which established the foundation for harmonising and integrating data from diverse IoT devices across the ECOLOOP ecosystem using protocols like MQTT and REST APIs—this document (D4.3) addresses the next critical phase. While D4.1 successfully consolidated the big data analytics framework supporting the project’s energy, agricultural, forestry, and community service pillars, it was necessary to develop a platform to put this data into action. Consequently, this deliverable details the development of the ECOLOOP Decision Support Tool (DSS). This system is designed to integrate and visualise the previously harmonised data, providing professionals with the practical tools required to actively manage the ecosystem and make data-driven decisions.

1.2. Scope of the document

The scope of this document encompasses the design, development, and operation of the ECOLOOP Decision Support Tool (DSS). Building upon the harmonised big data foundation established in previous phases, this system provides the interfaces and analytical instruments necessary for professionals to leverage the data collected across multiple pilot sites. The document covers:

- The software architecture and user interfaces enabling the seamless visualisation, analysis, and management of the harmonised data previously integrated from distributed sensors and systems.
- The integration of advanced analytical modules, incorporating Artificial Intelligence (AI) and Machine Learning (ML) algorithms, to provide predictive analytics, intelligent alerts, and enhanced decision-making functionalities that empower professionals to interact with the unified ECOLOOP ecosystem.

- The integration of advanced analytical modules and decision-making functionalities that empower professionals to interact with the unified ECOLOOP ecosystem, ensuring practical interoperability and scalability of daily operations.
- The tool's support for ECOLOOP's holistic approach, which equips users to optimise renewable energy integration, sustainable agricultural protocols, and circular bioeconomy practices, while fostering regional development and environmental sustainability.

This document focuses on the technical, operational, and user-centric aspects of the Decision Support Tool and its integration within the wider ECOLOOP framework.

1.3. Structure of the document

This document is organised into six main sections:

- **Section 1: *Introduction*.** Outlines the purpose, scope, and structure of the deliverable, providing essential context regarding the transition from data harmonisation to data utilisation.
- **Section 2: *Decision Support Tool Architecture*.** Describes the overall software architecture of the DSS, building upon the integrated data foundation established in previous phases. It details the user interfaces, data visualisation mechanisms, and the practical tools provided to professionals for active ecosystem management.
- **Section 3: *Artificial Intelligence Integration*.** Dedicated to the AI and Machine Learning (ML) modules implemented within the DSS. This section explores how these technologies process the harmonised data to generate predictive analytics, intelligent alerts, and enhanced decision-making capabilities.
- **Section 4: *Decision Support Tool Application*.** Describes the ECOLOOP Decision Support Tool as the user-facing layer of the platform, presenting its application architecture and the main views through which users interact with the system, including hierarchical navigation, map-based exploration, rule management, digital twin representation, and device-level analysis and actuation capabilities.
- **Section 5: *Pilot Site Implementation*.** Describes how the ECOLOOP DSS has been progressively deployed and adapted across the different pilot environments. It explains the current integration status of each pilot, highlighting the technical developments incorporated at site. The chapter also reflects the heterogeneous nature of the pilots, showing how the platform has been implemented according to the operational maturity, constraints, and priorities of each location.

- **Section 6: *Next Steps*.** Outlines the forthcoming developments of the ECOLOOP Decision Support Tool, including the integration of additional empirical and model-based datasets through CSV and Excel import mechanisms, the deployment and validation of the platform across the remaining pilot sites, and the iterative refinement of the solution within WP6 based on real feedback from sector professionals and operational pilot environments.
- **Section 7: *Conclusion*.** Summarises the DSS development, its architectural and AI foundations, and outlines the upcoming deployment phases, future work, and expected enhancements.
- **Section 8: *References and Acronyms*.** Provides a list of referenced materials and a glossary of acronyms used throughout the document for clarity and ease of understanding.

This structure ensures a logical flow from the software and AI foundations to their strategic real-world application, highlighting the continuous improvement cycle that will optimise the ECOLOOP ecosystem.

2. Decision Support Tool Architecture

The ECOLOOP project is dedicated to optimising the integration of renewable energy sources—such as biogas, biomass, agri-photovoltaics (agri-PV), and geothermal energy—across rural areas. While reducing carbon footprints, enhancing sustainability, and fostering circular economy practices are the ultimate goals, achieving them requires more than just raw data collection; it requires actionable intelligence. To address this, this chapter details the architecture of the Decision Support Tool (DSS), the robust platform designed to bridge the gap between complex big data and practical, day-to-day decision-making.

Building upon the Big Data Analytics Framework established in earlier phases, the DSS architecture is specifically engineered to ingest, process, and visualise substantial volumes of diverse information. The system seamlessly integrates real-time inputs from multiple sources, ranging from renewable energy devices (such as solar inverters and biogas analysers) to agricultural sensors deployed across crop fields (measuring humidity, temperature, solar radiation, and more).

Rather than simply aggregating this telemetry, the DSS architecture transforms it into a unified, user-centric platform. By structuring this foundational layer with intuitive interfaces and advanced analytical modules, the tool directly empowers farmers, foresters, and energy managers. As the core focus of this deliverable (D4.3), this architecture ensures that every stakeholder has the necessary instruments to translate harmonised data into optimised operations and sustainable, real-world solutions. Below in Figure 1, this process described over this Chapter 2 is present.




Figure 1: ECOLOOP minified architecture diagram

2.1. Microservices Architecture and Event-Driven communications

As established in **Deliverable 4.1 (D4.1)**, the ECOLOOP DSS is built upon **CITRIC**, a highly scalable platform developed by ETRA. Originally designed to address communication requirements in previous European H2020 projects (such as X-FLEX [1], ASSISTANCE [2], and MATCHUP [3]), CITRIC provides the foundational architecture for the ECOLOOP ecosystem. This architecture operates on the RIVER principles—Reactive, Interoperable, Visible, Elastic, and Resilient—and adheres strictly to the IEC30182:2017 and UNE178104:2017 standards for functional layer orientation.

Unlike traditional monolithic systems, the CITRIC platform leverages a modern microservices architecture. This approach divides the platform's core functionalities into independent, standalone services. Each microservice is entirely decoupled from the others and is responsible for a specific

operational capability, such as ingesting agricultural sensor data, executing AI predictive models, managing user access, or triggering intelligent alerts. To ensure high availability, fault tolerance, and flexible deployments, all these components are containerised using Docker and orchestrated through Kubernetes [4], operating on Google Cloud Services. This deep modularity ensures that the DSS remains resilient, easy to maintain, and highly adaptable to the evolving demands of the diverse pilot sites. More details about this architecture based on microservices is explained at 2.1.1 Microservices.

To enable seamless interaction between these isolated components without creating rigid dependencies, the architecture relies on a highly efficient central communication bus. At the core of the CITRIC ecosystem is **NATS** [5], a high-performance message broker that functions as an Enterprise Service Bus (ESB) for the platform. Rather than relying on direct, point-to-point HTTP connections between services, the platform utilises an event-driven model where microservices [6] publish and subscribe to specific data streams asynchronously through secure, access-controlled topics on the NATS broker.

By positioning NATS as the central communication axis, the ECOLOOP platform achieves true technological decoupling. If a specific service requires an update, needs to scale due to increased data loads, or temporarily goes offline, it does not disrupt the rest of the platform's operations. Furthermore, this infrastructure allows for ultra-low latency data distribution. When a renewable energy inverter or a crop sensor transmits a new reading, the NATS broker instantly routes this telemetry to all relevant microservices—ensuring that the database logger, the AI analytics engine, and the frontend dashboard receive and process the information simultaneously in near-real time.

2.1.1. Microservices

Microservices are a software architecture approach in which a platform is divided into a set of small, independent services, each responsible for a specific function within the overall system. Instead of concentrating all functionalities in a single monolithic application, each microservice performs a well-defined task, such as data ingestion, transformation, storage, authentication, analytics, or visualisation support. These services communicate with each other through standardised interfaces and messaging mechanisms, allowing the platform to operate as a coordinated but modular ecosystem.

In the context of ECOLOOP, this approach has been adopted because the project integrates highly heterogeneous data sources, pilot environments, and functional requirements. The platform must simultaneously support agricultural sensors, renewable energy devices, empirical datasets, AI services, user management, and decision-support interfaces across different countries and operational scenarios. A monolithic design would make this ecosystem harder to maintain, scale, and adapt. By contrast, a microservices architecture allows each functional component to evolve independently without compromising the rest of the platform.

This provides several important benefits for ECOLOOP. First, it improves scalability, since the services that process more data or require greater computational capacity can be scaled independently according to the needs of each pilot or application. Second, it increases resilience, because the temporary failure or update of one component does not necessarily disrupt the operation of the entire system. Third, it facilitates interoperability, as different services can be connected to external tools, communication protocols, and pilot-specific infrastructures in a flexible way. Finally, it improves maintainability and extensibility, enabling the incorporation of new services, data flows, analytical modules, or visualisation capabilities over time without redesigning the complete platform.




Figure 2: Microservices vs Monolithic Architecture [7]

Based on Figure 2, a wide variety of microservices can be foreseen in the context of building the whole DTS. The key point is understanding that the DST is not a unique service or piece of code deployed but rather an architectural decision made, which several components have their own responsibility towards building a scalable, independent and efficient platform.

For ECOLOOP, the value of microservices is therefore not only technical, but also operational and strategic. They provide the flexibility required to support a multi-pilot and multi-domain project, while ensuring that the DST can progressively evolve as new requirements, tools, and feedback emerge. In this sense, the microservices approach is a key enabler for building a robust, adaptable, and future-oriented platform capable of transforming heterogeneous data into practical support for farmers,

foresters, and energy managers. Where NATS the service communication orchestrates all the workflow between each microservice.

2.2. Data Storage Architecture

The ECOLOOP DSS relies on a multi-layered data storage architecture designed to address different data access patterns, persistence requirements, and file management needs. Rather than depending on a single database technology, the platform adopts a specialised approach in which each storage component is responsible for a specific type of information. This design improves performance, scalability, and maintainability, while ensuring that both real-time and historical data can be efficiently managed across the system.

2.2.1. Databases

The storage layer of the ECOLOOP DSS is therefore composed of several complementary technologies, each one integrated into the platform according to its particular strengths and intended function. This separation allows the system to manage heterogeneous information in a more efficient and structured way, from operational records and time-dependent measurements to file-based empirical resources. The following subsections describe the main storage components adopted in the platform and their specific contribution to the overall architecture.

2.2.1.1. MongoDB

At the operational level, *MongoDB* [8] is used to store the latest values received from the different devices and services integrated into the platform. MongoDB is a NoSQL document-oriented database designed to store information in flexible, JSON-like documents rather than in rigid relational tables. It is especially suitable for applications that manage heterogeneous or evolving data structures.

This database acts as the primary source for current system state, enabling fast access to the most recent measurements, statuses, and processed outputs required by dashboards, APIs, and decision-support interfaces. By keeping only the latest relevant values in this layer, the system ensures efficient queries and rapid response times for day-to-day monitoring and interaction.

2.2.1.2. InfluxDB

For long-term persistence and time-series analysis, the platform uses *InfluxDB* [9] to store historical data. InfluxDB is a time-series database specifically designed to store, organise, and query data

indexed by time. It is particularly well suited for environments where measurements are continuously generated, such as sensor readings, telemetry, or operational indicators collected at regular intervals.

This layer is specifically designed to manage continuous telemetry streams generated by sensors, renewable energy systems, and other monitoring devices over time. Through this approach, the DSS can support trend analysis, temporal visualisation, performance evaluation, and the analytical processing required for forecasting and AI-driven functionalities.

2.2.1.3. RustFS / S3 storage

In addition to structured and time-series data, the ECOLOOP architecture also requires a dedicated solution for storing empirical datasets and manually uploaded telemetry-related files. For this purpose, the platform uses RustFS, an open-source S3-compatible object storage interface. S3-style object storage organises data as objects within buckets, making it particularly suitable for managing unstructured or semi-structured resources in a scalable and efficient manner [10] [11]. Unlike databases focused on structured or time-series information, RustFS is intended for storing files such as documents, images, datasets, and exported resources, thereby complementing MongoDB and InfluxDB within the broader ECOLOOP microservices architecture.

This storage layer is intended for files such as Comma-separated values (CSV) datasets, images, manually entered telemetry records, and other supporting documents that do not fit naturally into conventional database schemas but remain essential for validation, traceability, and extended analysis.

2.2.2. Architecture

Following the architecture where the ECOLOOP platform is structured around a microservices approach, the database layer must also be understood from the same perspective. In this context, **MongoDB**, **InfluxDB**, and **RustFS** are not treated as isolated technical components, but as specialised services within the broader ecosystem, each one focused on its own well-defined responsibility. This separation of responsibilities is aligned with the overall CITRIC philosophy, where each component is designed to fulfil a specific role in a modular and scalable way.

On top of these storage technologies, additional microservices have been developed to manage the **ingestion**, **adaptation**, **retrieval**, and **exposure** of data from the different data sources. Rather than allowing applications to interact directly with each database in an ad hoc manner, the ECOLOOP architecture introduces dedicated intermediary services that are responsible for validating incoming data, transforming it into the ECOLOOP data model, routing it to the appropriate storage backend, and later exposing it to higher-level services such as analytics modules or the Decision Support Tool.

This abstraction layer improves maintainability, reinforces security and access control, and allows the platform to evolve without tightly coupling business logic to any specific storage engine.




Figure 3: ECOLOOP Storage architecture

This design provides flexibility to support heterogeneous data types, facilitates the future incorporation of new storage backends or retrieval services, and ensures that the DST can access harmonised information through controlled and reusable interfaces. Together, these three storage technologies followed by a microservice architectural design form a coherent and complementary persistence strategy within the DSS. **MongoDB** supports real-time access to the latest operational values, **InfluxDB** ensures robust management of historical time-series information, and **RustFS** provides flexible storage for unstructured and manually supplied files.

2.3. Application Framework and Presentation Layer

To provide a seamless, real-time user experience, the ECOLOOP DSS application layer is developed using Meteor.js [12]. Meteor is an open-source, full-stack JavaScript framework designed for the rapid development of responsive and highly interactive web applications. Its isomorphic nature allows developers to use the same codebase on both the client and the server, reducing development overhead and ensuring consistent logic across the platform. Furthermore, Meteor natively supports

real-time data synchronization via WebSockets (utilising the Distributed Data Protocol, DDP [13]), making it the ideal technological choice for a platform that relies on live telemetry and continuous ecosystem monitoring.

Within the ECOLOOP architecture, the Meteor server-side acts as a secure intermediary between the user interface and the core backend infrastructure. Rather than directly accessing the databases, the application's server communicates with the previously described CITRIC microservices. It retrieves harmonised data, AI predictions, and historical time-series records, while also dispatching user-generated inputs, configurations, or manual data entries. This communication is primarily handled through RESTful APIs and event-driven updates, ensuring that the presentation layer remains fully decoupled from the raw data storage (MongoDB, InfluxDB and RustFS) and complex background processing. This design guarantees that data retrieval and storage operations are managed securely, efficiently, and in strict alignment with the platform's microservices principles.

Leveraging this robust and real-time backend integration, the primary purpose of the Meteor application is to deliver comprehensive data analytics and interactive dashboards to the end-users. The DSS processes the continuous data streams received from the microservices and translates them into actionable visualisations. Professionals—such as farmers, foresters, and energy managers—can access customised dashboards that display Key Performance Indicators (KPIs), real-time sensor readings, and predictive AI alerts in an intuitive, user-friendly format. By providing these advanced analytical tools and graphical interfaces, the application empowers stakeholders to actively monitor the ECOLOOP ecosystem, evaluate historical trends, and make data-driven decisions to optimise their daily operations across all pilot sites.

2.4. AI-Driven Predictive Analytics Architecture

Within the ECOLOOP DSS, predictive capabilities are provided by Delorean, a specialised AI forecasting microservice integrated into the CITRIC architecture. This service is based on machine learning; a branch of artificial intelligence focused on the development of algorithms capable of learning patterns and relationships from historical data in order to make predictions on new observations. Unlike traditional rule-based logic, where the reasoning is explicitly predefined, machine learning models infer predictive behaviour from examples observed during a training process.

In the case of ECOLOOP, this approach is particularly suitable because many relevant variables, such as energy production, energy consumption, or environmental conditions, evolve over time and can therefore be forecasted from past measurements and contextual inputs. Rather than operating as a monolithic analytical engine, Delorean functions as an independent, fully decoupled module dedicated exclusively to time-series forecasting. It is designed to process historical data and estimate future

trends across any time-varying magnitude critical to the ECOLOOP pilot sites, such as energy consumption, photovoltaic production, or agricultural temperature variations.

2.4.1. Architectural Integration and Data Flow

Architecturally, the predictive engine relies on the platform's multi-layered storage and event-driven communication infrastructure. To generate accurate forecasts, the microservice retrieves historical telemetry seamlessly from the established storage layers—primarily querying InfluxDB for time-series metrics and MongoDB for operational state—as well as pulling complementary data from external APIs when necessary.

The service interacts with the broader ECOLOOP ecosystem entirely through the NATS message broker. By utilising this asynchronous messaging system, it receives forecasting requests and dispatches its predictions without blocking or degrading the performance of the user interface or other platform components. Furthermore, it operates invisibly in the background, continuously broadcasting its health metrics to the Kubernetes orchestration layer to ensure system reliability.

2.4.2. Contextual Inputs and Operational Value

To maximize the accuracy of its forecasts, which can be generated for intervals as granular as 5 minutes, the AI engine does not rely solely on past numerical values. It is configured to automatically ingest and evaluate multidimensional contextual inputs. This includes calendar events, such as national holidays and days of the week, and external environmental variables, such as solar radiation or cloud cover, that significantly influence energy and agricultural behaviour.

Once the predictions are calculated, they are routed back through NATS to be stored in the InfluxDB to be later on visualised on the Meteor-based frontend dashboards or consumed by other automated services. Within the ECOLOOP framework, this continuous flow of predictive intelligence enables critical operational functionalities, allowing site managers to anticipate energy peaks, detect systemic anomalies, and optimise resource planning before physical events occur.

3. AI Evolution and Human-in-the-Loop Decision

To fully grasp the artificial intelligence strategy implemented within the ECOLOOP DSS, it is essential to contextualise how Artificial Intelligence (AI) paradigms have evolved and how they mirror human cognitive processes. Human decision-making is famously categorised by psychologist Daniel Kahneman into two distinct modes: *System 1* and *System 2* thinking.

System 1 is fast, automatic, intuitive, and subconscious. It excels at immediate pattern recognition—for example, instantly recognising a colleague's face or instinctively hitting the brakes when a car stops suddenly ahead. Conversely, System 2 is slow, deliberate, analytical, and logical. It requires conscious effort and situational context, such as solving a complex calculus equation or formulating a long-term agricultural strategy based on shifting economic and environmental factors.

From the "AI Winter" to the Current Boom

The history of AI closely reflects this cognitive divide. In its early decades, researchers focused heavily on replicating System 2 thinking through Symbolic AI (also known as rule-based or expert systems). The goal was to explicitly programme human logic and domain rules into computers. However, the sheer complexity of real-world environments made it impossible to hardcode every possible scenario. Coupled with a lack of computational power and insufficient digital data, this approach eventually stalled, leading to periods of severe funding cuts and diminished interest known as the "AI Winters" (primarily in the 1970s and late 1980s) [14] .

Today, Artificial Intelligence is experiencing an unprecedented boom. This resurgence is not driven by Symbolic AI, but rather by Statistical AI (Machine Learning and Deep Learning). Thanks to the modern explosion of Big Data (generated by IoT ecosystems like ECOLOOP) and massive increases in computational power (GPUs), Statistical AI can process millions of data points to find hidden patterns. Essentially, the current AI boom represents the triumph of **System 1 AI: it is incredibly fast and highly accurate at pattern matching and forecasting, but it often operates as a "black box," lacking the true reasoning, logic, and situational awareness of System 2** [15].

3.1. The ECOLOOP Approach: Bridging the Cognitive Gap

In the ECOLOOP project, we recognise that relying exclusively on Statistical AI (System 1) is insufficient—and potentially risky—for managing complex rural microgrids, biogas plants, or agricultural fields. While the DSS's predictive microservices (such as the CITRIC forecasting microservice engine) excel at rapidly analysing vast amounts of IoT telemetry to forecast energy peaks or detect anomalies, they do not inherently understand why these events occur or how external socio-economic factors might impact a local pilot site.

To provide these predictive capabilities, ECOLOOP integrates **Delorean**, a dedicated forecasting microservice built upon the **XGBoost** machine learning algorithm [16], widely recognised for its accuracy and efficiency with structured, time-series data. Rather than simply looking at raw past values, Delorean enriches its analysis by automatically incorporating several layers of contextual information. It considers recent historical trends through a sliding window of past readings, accounts for recurring weekly behavioural patterns (distinguishing, for instance, a Monday from a Saturday), and factors in national public holidays—days which typically produce significantly different energy and agricultural consumption profiles. Furthermore, where available, the service can integrate external environmental variables such as solar irradiance or ambient temperature, but only when a statistically meaningful correlation with the target variable is confirmed. This ensures that the model's predictions remain grounded in relevant, high-quality information rather than noise.

To bridge this cognitive gap, the ECOLOOP Decision Support Tool is built around a Human-in-the-Loop philosophy. Rather than attempting to fully automate operations with opaque algorithms, the platform leverages the strengths of both systems:

- The AI provides the System 1 capabilities: instantly processing historical data to deliver rapid forecasts, trend analyses, and anomaly alerts through interactive dashboards.
- The sector professional provides the System 2 capabilities: farmers, foresters, and energy managers inject their deep domain knowledge, intellect, and situational context to interpret the AI's suggestions and make the final operational decisions.

Crucially, to fully integrate human intellect into the automated processes, **the DSS incorporates a transparent Rule Engine**. Acting as a modern implementation of Symbolic AI, this engine empowers experts to define, view, and modify explicit logical rules (e.g., "IF soil moisture drops below X AND solar radiation is Y, THEN trigger an irrigation alert"). By exposing these rules rather than hiding them

behind complex mathematics, professionals can clearly understand exactly how and why specific outputs are generated.

This synergistic approach ensures that ECOLOOP users maintain full control, trust, and contextual accuracy over their daily operations, combining the computational speed of modern AI with the irreplaceable reasoning of human experts.

4. Decision Support Tool Application

The ECOLOOP DSS application represents the user-facing layer of the project, transforming the harmonised data and analytical services developed in previous phases into a practical and accessible working environment for professionals. Through this application, users are able to explore real-time and historical information, monitor operational indicators, and interact with the analytical outputs generated within the ECOLOOP ecosystem. In this sense, the DSS acts as the bridge between the technical data infrastructure and the decision-making processes required in the project's energy, agricultural contexts.

From a technical perspective, the application has been developed using **MeteorJS** as the full-stack framework and **React** for the implementation of the user interface. This architecture enables the development of dynamic, responsive, and near real-time web interfaces adapted to the operational needs of the different pilot sites. The application communicates with the underlying platform through a combination of **REST HTTP calls** and **WebSocket-based communication**, allowing it to retrieve, update, and synchronise information efficiently across the different services integrated into the ECOLOOP ecosystem.

A **REST HTTP call** is a standard request-response communication mechanism widely used in web applications. In practice, it allows the application to request specific information from a server, such as the latest measurements of a device, a list of registered assets, or a historical dataset for a selected time period. The server processes the request and returns the corresponding response, typically in a structured format such as JavaScript Object Notation (JSON). This approach is especially suitable for controlled data retrieval, configuration management, and operations where an immediate but non-continuous exchange of information is required [17].

A **WebSocket**, by contrast, enables persistent, bidirectional communication between the client application and the server. Instead of opening a new connection for every request, the connection remains active, allowing data to flow continuously in both directions. This makes WebSockets particularly suitable for near real-time dashboards, live status updates, notifications, and dynamic monitoring environments where the information displayed to the user must be refreshed as soon as

new data becomes available [18]. Within the ECOLOOP DSS, this communication model supports a more interactive and reactive user experience.

The combination of MeteorJS, React, REST, and WebSockets provides a flexible and modern application architecture, capable of supporting both standard data queries and live operational visualisation. On top of this technological foundation, the DSS offers a set of dedicated views designed to address different user needs, ranging from asset monitoring and historical analysis to predictive insights and operational support. The following sections describe these views in detail and explain their role within the overall application.

As the project is now entering its final phase, the DSS is progressively being validated through its deployment in the pilot sites. This stage is especially relevant, as the exposure of the application to real operational environments may reveal the need for further refinements, usability improvements, or site-specific adaptations. Such adjustments are considered part of the natural evolution of the platform and will be addressed within the scope of **WP6**, which focuses on pilot deployment, validation, and iterative improvement of the ECOLOOP solutions.

4.1. Empirical Data Storage and Edge Processing

In addition to real-time telemetry and structured operational data, the ECOLOOP DSS must also support a large volume of **empirical datasets** generated across the different pilots and project activities. These include, for example, measurements collected in **living labs**, field observations registered by professionals, and historical empirical datasets such as those provided by the Estonian partners for forestry and land-use analysis. Unlike telemetry streams integrated through standard device protocols, these resources are often heterogeneous in structure, format, and granularity, as each partner may already work with its own established templates, spreadsheets, file naming conventions, and domain-specific data collection methodologies.

To address this challenge, the platform adopts an S3-compatible object storage approach as the main solution for managing empirical data assets. This model is especially suitable because it allows the platform to store not only a very large number of files, but also a wide variety of file types and structures without imposing a rigid database schema. In practice, this means that ECOLOOP can preserve partner-provided datasets in a scalable and flexible manner while still integrating them into the broader decision-support environment. This approach is particularly valuable in a project such as ECOLOOP, where empirical knowledge is highly relevant but not always produced through homogeneous pipelines.

In order to make this storage strategy scalable, the platform also incorporates edge-side data transformation and preprocessing before publication. Rather than sending all raw files directly to the

server for processing, ECOLOOP makes use of DuckDB executed in the browser, allowing substantial parts of the transformation pipeline to be performed closer to the user. DuckDB is an analytical in-process database engine designed for fast execution of data-intensive queries, especially over tabular formats such as CSV and Parquet [19]. Its lightweight architecture and strong support for analytical workloads make it especially suitable for running directly in client-side environments, where it can filter, reshape, aggregate, and convert data before upload.

This browser-based use of DuckDB brings an important architectural advantage to ECOLOOP. By moving part of the processing workload to the edge, the platform avoids overloading the central backend with repetitive transformation tasks, reduces unnecessary data transfer, and makes ingestion workflows more scalable as the number of users and uploaded files grows. In this sense, DuckDB is not only a technical convenience, but a strategic component that helps keep the core of the platform lighter and more efficient while still supporting data-intensive empirical workflows from multiple partners.

For tabular datasets, one of the key optimisations applied before storage is the conversion from raw formats such as CSV into **Parquet**. Parquet is a columnar storage format widely adopted in big data and analytics environments because it is highly efficient for both storage and query performance [20]. By organising data by columns rather than by rows, Parquet allows systems to read only the variables required for a given analysis, significantly reducing I/O operations. It also supports efficient compression and encoding mechanisms, which typically result in smaller file sizes than equivalent CSV files while preserving the original information. For these reasons, Parquet has become a standard format in large-scale analytical ecosystems and is particularly well suited for ECOLOOP, where empirical datasets may be large, multidimensional, and repeatedly queried for comparison or historical analysis.

A similar optimisation strategy is applied to image-based empirical resources. When uploaded images are too large, the platform reduces their storage weight through controlled optimisation processes that preserve visual quality while lowering bandwidth and storage costs. In addition, thumbnails are generated automatically to facilitate fast preview and navigation within the user interface. This ensures that the platform can efficiently manage photographic evidence, field documentation, and other visual records without compromising usability or overloading the storage infrastructure.

Overall, this combination of S3-compatible storage, edge-side processing with DuckDB, Parquet conversion for tabular data, and optimised image handling provides ECOLOOP with a practical and scalable strategy for empirical data management. It enables the platform to accommodate the diversity of formats used by different partners, preserve the value of already existing datasets, and integrate empirical knowledge into the DSS without placing unnecessary pressure on the central

services. In this way, the platform becomes better prepared to support large-scale, heterogeneous, and professionally generated data assets as the project evolves.

4.2. Views

The DSS is structured around a set of complementary views that allow users to interact with the platform at different levels of detail and from different operational perspectives.

The **Home View** provides the main hierarchical navigation of the application, enabling users to move from the **realm level**, which represents the full project ecosystem, down to the **site, zone**, and **device** view levels, thus supporting both aggregated and asset-specific analysis. In addition to this general navigation view, the application also includes a **Map View** at site level, offering a spatial perspective of the pilot infrastructure; a **Rules View** at site level, where users can review and manage the logical rules associated with decision-support processes; and a **Digital Twin View**, also at site level, which provides a more advanced and integrated representation of the pilot environment. Together, these views form the core user interface of the DSS and support monitoring, analysis, and informed decision-making across the ECOLOOP pilots.

The navigation model based on a tree-navigation model, has been designed to support progressive exploration of the available information. Users may begin with a high-level project perspective and then move step by step towards site-specific, zone-specific, and device-specific details, depending on the operational question they need to address. This approach improves usability and ensures that the same application can serve both strategic monitoring needs and more detailed technical analysis tasks.

4.2.1. Login View

The **Login View** provides the authentication entry point to the ECOLOOP DSS, ensuring that access to the platform is restricted to authorised users and aligned with the security requirements of the overall ecosystem. This view is designed to support both usability and secure access management, acting as the initial control layer before users can interact with the different monitoring, analytics, and decision-support functionalities offered by the application.

From an authentication perspective, the DSS supports both **traditional username and password login** and more advanced **standardised secure authentication protocols**. In particular, the platform is compatible with Single-Sign-On, SSO **OpenID-based authentication** through the **Keycloak** identity and access management service deployed under the ECOLOOP project.

This approach enables centralised and secure user authentication, facilitates the integration of single sign-on mechanisms, and supports more robust access control policies across the platform. By

combining conventional login mechanisms with modern identity standards, the Login View ensures both accessibility for end users and compliance with the security and interoperability principles defined for the ECOLOOP ecosystem. Below at Figure 4, the login view is displayed.




Figure 4: ECOLOOP DSS login view

4.2.2. Home View

The Home View constitutes the main navigation environment of the ECOLOOP DSS. It allows users to explore the platform through a hierarchical structure that progressively moves from a global overview of the project to the detailed inspection of individual monitored assets. Through this approach, the application supports both high-level monitoring and fine-grained operational analysis.

Within the Home View, the information is organised into four hierarchical views:


- **Realm View**
- **Site View**
- **Zone View**
- **Device View**

Each of these levels provides a different perspective on the ECOLOOP ecosystem, enabling users to navigate the project structure in a consistent and intuitive way.

4.2.2.1. Realm View

The Realm View provides the highest-level perspective within the Home View. It offers a global overview of the ECOLOOP ecosystem, allowing users to visualise and access the different pilot sites participating in the project. This view is intended to support aggregated monitoring across the full project scope, enabling users to understand the overall structure of the platform and identify the different operational environments available for exploration.

From this level, users can initiate navigation towards the individual pilot sites, moving from a general project perspective to more site-specific information. As such, the Realm View (Figure 5) acts as the main gateway to the hierarchical organisation of the DSS.






Figure 5: ECOLOOP DSS realm view

4.2.2.2. Site View

The Site View focuses on a specific pilot site and provides access to its local structure, assets, and operational context. At this level, users can inspect the configuration of the selected site and understand how its infrastructure is organised within the platform. This view is particularly important in ECOLOOP, as each pilot has its own technical characteristics, environmental conditions, and operational requirements.

This view can be seen under Figure 6, which showcases one of the ECOLOOP pilot sites, in this case the Bulgarian pilot site, located at Balchik.



Figure 6: ECOLOOP DSS site view

In addition to visualisation and navigation, the Site View also supports management functionalities for authorised users. These include the possibility to edit the site configuration, modify the hierarchical structure of the site, and manually add empirical data associated with the pilot when this information is not acquired automatically through connected devices or integrated services.

As such, Figure 7 displays how the contextual layout of a given site can be changed. This functionality allows end-users to setup their own hierarchical instance of a given site, to logically represent a physical site in a digital way. Allowing from creating new “virtual logical” zones representing a physical location inside a site, as well as assigning devices to those zones.




Figure 7: Edit layout in DST

4.2.2.3. Zone View

The Zone View (Figure 8) introduces an intermediate hierarchical level within each pilot site. It allows users to organise and interpret the site infrastructure through logical or functional subdivisions, which may correspond to operational areas, subsystems, agricultural sectors, energy units, or any other meaningful grouping defined for the pilot.

This intermediate layer improves the readability and scalability of the application, especially in pilots with complex infrastructures. By structuring the site into zones, users can navigate large environments more efficiently and better contextualise the assets and measurements associated with each part of the pilot.




Figure 8: ECOLOOP DSS zone view

4.2.2.4. Device View

The Device View (Figure 9) provides the most detailed level of inspection within the Home View. It is centred on an individual monitored asset and gives users direct access to its available operational information. This includes the current measurements associated with the device as well as more detailed historical information relevant for analysis and validation.




Figure 9: ECOLOOP DSS device view with chart

At this level, the application also provides advanced data access functionalities. Users can perform custom queries on historical data, selecting the required temporal range and variables of interest customizing dataset display options such as chart or table mode (Figure 10). In addition, the resulting datasets can be exported in CSV format, enabling external analysis, reporting, and further processing outside the platform.




Figure 10: ECOLOOP DSS device view with table format

When empirical data is available, the **Device View** also provides access to a historical repository of all manually uploaded files associated with the selected asset within the ECOLOOP platform. This functionality, reflected in Figure 11 allows users to review the full record of empirical datasets over time, facilitating the comparison of different uploaded files and the analysis of their evolution across successive observations or campaigns. In addition, these empirical records can be contrasted with real telemetry data acquired automatically from the field, enabling users to identify correlations, validate observed trends, and draw more robust conclusions from the combined interpretation of manual and sensor-based information




Figure 11: ECOLOOP DSS empirical analysis view

In addition to monitoring and data exploration, the **Device View** also supports operational interaction with those assets that allow actuation. For compatible devices, the platform is capable of issuing control or command actions through the communication mechanisms supported by the underlying infrastructure, including **SCADA-based integrations**, **MQTT**, **REST APIs**, and **AMQP**. This enables the DSS to go beyond passive observation and become an active operational interface, allowing authorised users to interact directly with field equipment or control systems when such functionality is available and permitted by the pilot configuration. In this way, the platform reinforces its role as a practical decision-support environment that not only visualises and analyses data, but can also contribute to the execution of informed actions in real operational contexts.

4.2.3. Map View

The Map View provides a spatial representation of the selected pilot site, enabling users to visualise its assets and infrastructure in their geographical context. This view showcased in Figure 12 and Figure 13, is particularly valuable for understanding the physical distribution of devices, zones, and relevant operational elements within the pilot environment. By complementing the hierarchical navigation of the Home View with a georeferenced perspective, the Map View improves situational awareness and facilitates a more intuitive interpretation of the monitored ecosystem.



Figure 12: ECOLOOP DSS map view

Through this view, users can identify the location of relevant components and better understand the relationship between infrastructure layout and operational behaviour. This is especially useful in distributed rural environments, where the physical positioning of sensors, energy systems, and agricultural assets may be directly related to the interpretation of the collected data and the decisions derived from it.



Figure 13: ECOLOOP DSS map view with device details

4.2.4. Rules View

The **Rules View** provides access to a set of preconfigured logical rules associated with the decision-support processes implemented within the ECOLOOP platform. At site level, this view allows users not only to inspect these predefined rule sets, but also to adapt them to the specific operational context of each pilot. In practice, these rules define how certain outputs, alerts, or automated responses are generated from the available data, establishing an explicit and understandable logical layer on top of the broader analytical capabilities of the platform. In this way, the Rules View promotes a more transparent and explainable interaction with the intelligent functionalities of the DSS and translates the conceptual AI approach described in Chapter 3 into a practical operational interface.

This view is particularly relevant because it represents the concrete implementation of the hybrid reasoning model pursued in ECOLOOP, combining **System 1** and **System 2** approaches. On the one hand, the platform incorporates fast, data-driven intelligence through predictive models, automated analytics, and near real-time processing of operational data. On the other hand, the Rules View enables domain experts to apply their own deliberate reasoning, professional judgement, and contextual understanding by reviewing and modifying the prebuilt logical rules proposed by the system. Rather than relying exclusively on opaque outputs generated by statistical models, users can

intervene directly in the symbolic layer of the platform, adjusting thresholds, conditions, and decision logics so that the system behaviour better reflects the local reality of the site.




Figure 14: ECOLOOP DSS rules view

As a result, the Rules View, shown in Figure 14 and Figure 15, becomes a key element within the ECOLOOP Human-in-the-Loop strategy. It ensures that automated intelligence is not applied in isolation, but is instead complemented and refined by expert knowledge from professionals working in agriculture, energy, and rural management. This combination strengthens trust in the platform, improves contextual adaptation, and gives end users a more active role in shaping how recommendations and alerts are produced. In this sense, the Rules View is not only a configuration

interface, but also a practical mechanism for aligning advanced analytics with the real decision-making needs of sector professionals.




Figure 15: ECOLOOP DSS rules view with detailed rule

4.2.5. Digital Twin View

The Digital Twin View provides an integrated and advanced representation of the selected pilot site, bringing together its relevant assets, data flows, and operational context within a unified digital environment. This view is intended to offer a more holistic understanding of the pilot by connecting the physical infrastructure with its digital counterpart, thereby supporting a more comprehensive interpretation of the monitored ecosystem.

By consolidating multiple sources of information into a coherent representation, the Digital Twin View, shown in Figure 16 and Figure 17, enables users to better analyse the current state of the pilot and the interactions between its different components. This type of view is particularly useful for complex environments in which isolated data points are not sufficient to fully capture the operational reality. In the ECOLOOP context, the Digital Twin View strengthens the capacity of the DSS to support monitoring, analysis, and future optimisation of the pilot sites through a more connected and system-oriented perspective.




Figure 16: ECOLOOP DSS digital twin view




Figure 17: ECOLOOP DSS digital twin top-view

4.2.6. Notifications View

The Notifications View of the ECOLOOP DSS has been designed as a centralised interface for monitoring relevant operational activity across the different pilot sites. Its main purpose is to provide users with a clear and structured overview of platform-generated notifications, helping them identify notable events, detect incidents, and follow the recent activity associated with each monitored environment. In this way, the view contributes to improving situational awareness and supports a more proactive interpretation of the system status.

From a user-interface perspective, as Figure 18 shows, the view incorporates a navigation panel inspired by repository-style layouts, similar to GitHub, which allows users to browse the available pilots in a simple and intuitive manner. This side panel provides a hierarchical structure where the different sites can be accessed quickly, facilitating movement across the platform and improving usability when multiple pilot environments are being monitored simultaneously. Such an approach is especially useful in ECOLOOP, where the DSS must support several heterogeneous pilots while preserving a coherent and unified user experience.




Figure 18 ECOLOOP DST Notifications View

Within the main content area, the Notifications View presents the activity associated with the selected pilot through a visual and chronological organization. One of its key components is the activity summary, which gives users an immediate perception of the temporal distribution of notifications

over time. This makes it possible to identify periods with higher operational activity and to better contextualize the incidents and events registered in the system.

A particularly relevant feature of this view is the distinction between events and incidents, both of which can be consulted from the same interface. Events represent relevant operational records or noteworthy system actions, while incidents correspond to situations that may require attention, such as abnormal sensor readings, threshold exceedances, or adverse environmental conditions. By grouping both categories in a unified dashboard, the platform allows users to understand not only what has happened in the system, but also which situations may have operational relevance or require further assessment.




Figure 19 ECOLOOP DST Notifications Detail View

In addition, the interface presents notification entries together with metadata such as severity level, description, associated device, service, status, and date (Figure 19). This structured presentation helps users quickly assess the importance of each notification and understand its origin and current handling

state. As a result, the Notifications View becomes an important functional component of the ECOLOOP DSS, supporting traceability, usability, and operational awareness within the monitoring environment.

5. Pilot Site Implementation

This chapter presents the implementation status of the ECOLOOP DSS across the main pilot sites, describing the current level of integration achieved in each case and the specific developments carried out according to the technical and operational realities of each pilot. As already outlined in D4.1, the ECOLOOP pilots represent heterogeneous environments in Spain, Slovenia, Bulgaria, and Estonia, each one with different renewable energy assets, monitoring needs, and integration constraints.

Rather than following a fully uniform deployment path, the implementation has progressed in a pilot-specific manner, adapting the platform to the maturity, infrastructure availability, and domain priorities of each site. In this sense, the following subsections describe the practical integration work performed so far, with particular attention to sensor deployment, empirical data incorporation, and the progressive evolution of each pilot within the ECOLOOP DSS.

5.1. Picassent (l'Horta Sud, INDEREN)

At the Picassent pilot site, the current integration efforts have focused primarily on **field sensing infrastructure**, which has already been connected to the ECOLOOP DSS and incorporated into the platform's monitoring and visualisation workflows. This decision responds to the current implementation status of the pilot, where the integration of the planned photovoltaic installation remains subject to external administrative requirements and licensing procedures. As a result, the most relevant progress to date has been achieved through the incorporation of environmental and field-level data streams, which already provide a valuable operational basis for the pilot.

Figure 20 represents the **main dashboard of the Picassent pilot**, where several key environmental and soil-related measurements can be visualised in an integrated way. Among the monitored variables, the dashboard includes **soil temperature**, **soil humidity**, **vapour pressure**, **atmospheric pressure**, as well as **ambient temperature** and **ambient humidity**. Through this interface, the ECOLOOP DSS

provides a clear overview of the field conditions at the pilot site, supporting continuous monitoring and facilitating the interpretation of local environmental behaviour.



Figure 20 Picassent Pilot Site

This phased approach allows the platform to deliver useful monitoring capabilities even before the full renewable energy infrastructure becomes available. Once the photovoltaic installation is completed and the corresponding licences are obtained, the pilot will be further extended so that the DSS can also incorporate the associated energy data and complete the intended integration between field conditions and local renewable generation.

5.2. Silla Biogas Plant (GENIA)

In the case of the **Silla biogas plant**, the current work has focused on the preparatory phase required for a correct future integration into the ECOLOOP DSS. More specifically, the viability of the sensorisation strategy has already been analysed, including the study of the devices to be installed and the monitoring points most relevant to the pilot. This means that the technical groundwork for integration has already been defined, even if the full deployment has not yet taken place.

The actual installation of the planned sensing infrastructure will proceed once the Silla plant is fully operational. Therefore, this pilot should currently be understood as a site in an advanced planning and readiness stage, where the integration logic has been prepared but the effective deployment is dependent on the operational availability of the plant itself. This staged implementation is consistent with the broader ECOLOOP strategy of adapting deployments to the real status of each pilot environment.

5.3. Finca Sinyent (Spain)

At **Finca Sinyent**, the ECOLOOP DSS has already incorporated external and empirical data sources that are especially relevant for agricultural monitoring. In particular, the platform integrates weather-related information through the **governmental SIAR service**, which provides the meteorological basis required to support site monitoring and analysis. This integration is particularly valuable because it enables the platform to consume trusted external environmental data even when that information is not produced directly by local field devices.




Figure 21: Weather station from finca sinyent using SiAR

Figure 21 presents the **weather station data integrated for the Finca Sinyent pilot**, showing the environmental information used by the ECOLOOP DSS to monitor local weather conditions, while the accompanying Figure 22 provides a comparison between **empirical data** and **real telemetry data captured by field sensors**. Together, these visualisations illustrate how the platform is able to combine external meteorological information, manually collected or empirically generated datasets, and real-time sensor measurements within the same analytical environment, thereby supporting a more complete interpretation of field conditions and strengthening the monitoring capabilities of the pilot.




Figure 22: Empirical data compared to telemetry data

In addition to this external data source, the DSS also includes the **CSV datasets generated in the Living Lab context**, thereby extending the pilot beyond real-time or automated telemetry and incorporating empirical records collected through project activities. This combination of institutional weather data and manually produced Living Lab datasets makes Finca Sinyent a relevant example of how ECOLOOP can merge heterogeneous sources into a single decision-support environment.

5.4. IC Jablje (KIS)

At the **IC Jablje pilot site**, integration efforts have continued to evolve beyond the initial monitoring setup. A particularly relevant development has been the incorporation of **new sensing infrastructure related to biomethane**, extending the original scope of the pilot and enriching the set of variables available in the ECOLOOP DSS. This is especially significant because it reflects not only a deployment effort, but also an expansion of the monitored domain according to the operational evolution of the pilot itself.



Figure 23 SWG1000 analyzer that analyzes both biogas and biomehtane

As shown in Figure 23, the pilot is able to monitor both **biogas** and **biomethane** streams through the integration of the **SWG100**, a dedicated gas analysis device used to measure and characterise the composition of the processed gases. This equipment provides the necessary analytical information to distinguish the quality and properties of the gas output, making it possible to track the transformation from raw biogas to upgraded biomethane within the ECOLOOP DSS.

Figure 24 illustrates the main monitoring dashboard of the **KIS Jablje pilot**, where the ECOLOOP DSS integrates key variables related to the **biogas and biomethane process**, together with additional measurements captured by the installed sensor network. In particular, the dashboard allows the visualisation of gas-related indicators associated with biomethane upgrading, as well as **temperature values obtained from different sensors** distributed across the system.



Figure 24: KIS dashboard where both biogas and biomethane data can be displayed

In parallel, the pilot remains under active extension, and **new devices continue to be incorporated** into the platform as they become available. This makes Jablje one of the clearest examples of the modular and incremental nature of the ECOLOOP architecture, where pilot integration is not treated as a one-time action but as a progressive process of enrichment and adaptation.

5.5. Bulgarian Pilot (Momchil / Albena)

For the **Bulgarian pilot**, ECOLOOP is currently progressing with the integration of **field sensing infrastructure**, extending the data acquisition capabilities of the DSS at site level. This includes the deployment and incorporation of sensors intended to capture relevant operational and environmental variables, allowing the platform to improve monitoring and support future analytical services across the pilot.

In parallel, the platform is also receiving **green energy data** associated with the pilot, particularly in relation to **photovoltaic generation** and **biogas-based energy production**. The integration of these

renewable energy streams is especially important because it allows the ECOLOOP DSS to combine field-level sensing with energy-related monitoring, thereby providing a more complete representation of the pilot's operational context and sustainability performance.




Figure 25: Bulgarian pilot site dashboard with data integrated from sensors

Figure 25 shows the main monitoring dashboard of the Bulgarian pilot, where the ECOLOOP DSS displays a set of environmental and soil-related variables captured by the installed field sensors. These measurements include soil temperature, soil moisture, air temperature, air humidity, and light intensity (lux), providing a broad overview of the local conditions affecting the pilot site. By combining soil and atmospheric measurements with luminosity data, the platform supports a more complete interpretation of field dynamics and enables continuous monitoring of the environmental context.

As a result, the Bulgarian pilot now reflects the implementation of the sensing layer that had originally been foreseen during the planning stage. This progress is important because it closes a previous deployment gap and ensures that the pilot can now contribute more consistently to the ECOLOOP data ecosystem, both in terms of monitoring and future analytical exploitation.

5.6. Estonian Pilots

The **Estonian pilot activities** are especially relevant in relation to the management of empirical and professional measurement datasets, which differ in nature from conventional IoT telemetry. In this

context, ECOLOOP is preparing and supporting the ingestion of empirical datasets generated by professionals, including forestry- and land-related measurements that are central to the project's carbon sequestration and land assessment objectives. These datasets are particularly valuable because they bring long-term expert knowledge into the platform and complement the more automated streams available in other pilots.

From the DSS perspective, this reinforces the need for flexible ingestion and storage mechanisms capable of handling partner-specific formats and large empirical repositories. The Estonian contribution therefore represents a key use case for the empirical-data workflows of the platform, especially in relation to CSV/Excel ingestion, historical repository management, and future analytical exploitation of professionally curated measurements.

6. Next steps

The next phase of the ECOLOOP DSS will focus on extending the platform to support additional domain-specific tools and methodologies developed within the project, particularly those related to **carbon sequestration assessment (T4.2)** and **renewable-based agricultural protocols for sustainability and circularity (T4.3)**, both detailed in Deliverable 4.2. In contrast to the real-time telemetry and automated sensor data already integrated into the platform, these new contributions will rely primarily on **empirical datasets**, calculated outputs, and externally generated results derived from field measurements, laboratory analyses, and scientific modelling processes.

In this context, the DSS will support the incorporation of such information through its existing **CSV and Excel import capabilities**, enabling project partners to upload and manage structured datasets generated by these tools. This approach is especially relevant for the carbon sequestration calculator and the agricultural protocol methodologies, where many of the results will originate from empirical observations, model-based estimations, and manually processed data rather than directly from connected IoT infrastructure. To ensure flexibility and practical interoperability, the platform supports the development of **custom importers** adapted to the specific structure and requirements of each tool, pilot, or data provider. This will allow the ECOLOOP ecosystem to accommodate heterogeneous data sources while maintaining a unified visualisation and analysis environment.

Once imported, these datasets will be made available within the DSS so that they can be visualised, explored, and compared alongside the rest of the project information. This will allow users to work not only with live telemetry and historical operational data, but also with empirical and model-derived

indicators related to carbon uptake potential, biodiversity, avoided emissions, soil behaviour, erosion dynamics, and other sustainability-related variables addressed within ECOLOOP.

Significant efforts have already been made towards the deployment and validation of the tool in several pilot environments. At the time of writing, the DSS has already been deployed and tested in pilots such as **Picassent (INDEREN)**, **Jablje (KIS)**, and the **Albena biogas plant**, among others. These deployments have provided valuable insight into the operational behaviour of the platform and its applicability in real-world contexts. However, not all pilot integrations have progressed at the same pace. In some cases, delays in pilot readiness or infrastructure availability have postponed the full deployment of the solution, as is the case for the **Silla biogas plant** led by GENIA.

For this reason, the final stage of the project will place particular emphasis on completing the deployment of the DSS in the remaining pilot sites where integration is still pending or partial. These efforts will be carried out within the scope of **WP6**, which will provide the operational framework for pilot deployment, validation, and refinement. Beyond the purely technical rollout, this phase will be essential for gathering **real feedback from sector professionals and end users**, whose experience will guide the next iteration of the platform.

This iterative validation process is expected to play a central role in improving the ECOLOOP DSS. Feedback from professionals working in energy, agriculture, forestry, and rural management contexts will help identify practical needs, usability challenges, pilot-specific requirements, and opportunities for functional enhancement. As a result, the next steps of the project are not limited to deploying the existing solution, but also to **reiterating and improving the tool** based on direct interaction with real users and operational scenarios. This will ensure that the final ECOLOOP DSS is not only technically robust, but also truly aligned with the needs of the professionals it is intended to support.

7. Conclusion

This deliverable has presented the design and development of the ECOLOOP Decision Support Tool (DSS), building upon the harmonised data foundation established in Deliverable D4.1. While D4.1 defined the integration of heterogeneous IoT data sources across the ECOLOOP ecosystem, this document has focused on the implementation of the system that transforms such data into practical, operational, and user-oriented decision-support capabilities.

The DSS has been conceived as a scalable and modular platform, based on a microservices architecture and event-driven communication principles, enabling the integration of real-time telemetry, historical records, empirical datasets, and advanced analytical services within a unified environment. Its

technological foundation combines interoperable communication mechanisms, specialised storage layers, and a flexible application framework capable of supporting dynamic dashboards, historical analysis, spatial visualisation, rule-based logic, and digital twin representations.

Particular emphasis has been placed on the integration of Artificial Intelligence as an enabling component of the platform rather than as an isolated black-box solution. In this context, the ECOLOOP approach combines predictive analytics with human expertise through a Human-in-the-Loop philosophy, ensuring that automated insights are complemented by contextual reasoning and sector-specific knowledge. This is further reinforced through transparent rule-based mechanisms that allow professionals to understand, validate, and refine the logic behind certain outputs and recommendations.

From the user perspective, the deliverable has described how the application materialises these architectural and analytical capabilities into an accessible interface structured through multiple views, including hierarchical navigation, map-based exploration, rule management, and digital twin visualisation. In this way, the DSS is not only a technical integration layer, but also an operational environment in which harmonised data can be explored, interpreted, and transformed into informed decisions across the energy, agricultural, and forestry domains addressed by ECOLOOP.

As the project enters its final phase, the deployment of the DSS in the pilot sites will provide the necessary real-world validation of the solution. This stage is expected to generate valuable feedback from professionals and end users, allowing further refinements, adaptations, and usability improvements to be incorporated within the scope of WP6. Therefore, this deliverable should be understood both as a description of the current state of the ECOLOOP DSS and as the foundation for its final validation and consolidation in the operational environments of the project.

8. References and acronyms

8.1. References

- [1] X-FLEX, [Online]. Available: <https://xflexproject.eu/>.
- [2] assistance, [Online]. Available: <https://assistance-project.eu/>.
- [3] MATCHUP, [Online]. Available: <https://www.matchup-project.eu/>.
- [4] kubernetes, “kubernetes,” [Online]. Available: <https://kubernetes.io/docs/concepts/overview/>.
- [5] NATS, “NATS,” [Online]. Available: <https://nats.io/>.
- [6] C. Richardson, “Microservices.io,” [Online]. Available: <https://microservices.io/>.
- [7] V. Nadar, “Monolithic Architecture Vs. Microservices,” [Online]. Available: <https://www.clariontech.com/blog/monolithic-vs.-microservices>.
- [8] MongoDB, [Online]. Available: <https://www.mongodb.com/>.

- [9] InfluxData, “InfluxData,” [Online]. Available: <https://www.influxdata.com/>.
- [10] “RustFS,” [Online]. Available: <https://rustfs.com/>.
- [11] “What is Amazon S3?,” [Online]. Available: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/Welcome.html>.
- [12] Meteor, “Github,” [Online]. Available: <https://github.com/meteor/meteor/blob/devel/packages/ddp/DDP.md>.
- [13] “DDP Specification,” [Online]. Available: <https://github.com/meteor/meteor/blob/devel/packages/ddp/DDP.md>.
- [14] “AI Winter,” [Online]. Available: https://en.wikipedia.org/wiki/AI_winter.
- [15] “Daniel Kahneman Explains The Machinery of Thought,” [Online]. Available: <https://fs.blog/daniel-kahneman-the-two-systems/>.
- [16] XGBoost, “XGBoost,” [Online]. Available: <https://xgboost.readthedocs.io/en/stable/>.
- [17] “What is a REST API?,” [Online]. Available: <https://www.redhat.com/es/topics/api/what-is-a-rest-api>.
- [18] R. Hat, “Red Hat,” 8 May 2020. [Online]. Available: <https://www.redhat.com/en/topics/api/what-is-a-rest-api>.
- [19] “DuckDB,” [Online]. Available: <https://duckdb.org/>.
- [20] “Parquet,” [Online]. Available: <https://parquet.apache.org/docs/>.
- [21] Opentunity Grant Agreement. EC, 2022.
- [22] “WebGL,” [Online]. Available: <https://get.webgl.org/>.

- [23] “MQTT,” [Online]. Available: <https://mqtt.org/>.
- [24] Mozilla, “MDN,” [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/HTTP/Guides/Overview>.
- [25] “SailPoint,” 9 February 2025. [Online]. Available: <https://www.sailpoint.com/identity-library/access-token>.
- [26] FIWARE, “FIWARE,” [Online]. Available: <https://www.fiware.org/smart-data-models/>.
- [27] JSON, “JSON,” [Online]. Available: <https://www.json.org/json-en.html>.
- [28] KEYCLOAK, “KEYCLOAK,” [Online]. Available: <https://www.keycloak.org/>.
- [29] G. A. B. J. D. G.-C. K. R. Klemen Peter Kosovinc, “Deliverable 2.1 Pilot site analysis and use cases, requirements and KPIs definition,” [Online]. Available: https://ecoloop-project.eu/wp-content/uploads/2025/03/ECOLOOP_D2.1.pdf.
- [30] D. Specification. [Online]. Available: <https://github.com/meteor/meteor/blob/devel/packages/ddp/DDP.md>.

8.2. Acronyms

Table 1: Table of acronyms.

Acronym	Description
---------	-------------

Agri-PV	Agri-photovoltaics
DSS	Decision Support Tool
RIVER	Reactive, Interoperable, Visible, Elastic and Resilient
ESB	Enterprise Service Bus
HTTP	Hypertext Transfer Protocol
API	Application Programming Interface
CSV	Comma-separated values
DDP	Distributed Data Protocol
KPI	Key Performance Indicator
AI	Artificial Intelligence
IoT	Internet of Things
JSON	JavaScript Object Notation
SSO	Single Sign On