



HarvREST
Greener Farming with RES

D5.4

Methodology for farm based multisource energy management systems development

06 / 07 / 2026



Funded by
the European Union

www.harvrest.eu

PROJECT INFORMATION

ACRONYM	HarvRESt
PROJECT NAME	Harnessing the vast potential of RES for sustainable farming
PROGRAMME	Horizon Europe
TOPIC	HORIZON-CL6-2023-CLIMATE-01-7
TYPE OF ACTION	HORIZON Research and Innovation Actions
PROJECT NUMBER	101136904
START DAY	1 January 2024
DURATION	36 months

DOCUMENT INFORMATION

TITLE	D5.4 Methodology for farm based multisource energy management systems development
WORK PACKAGE	WP5
TASK	T5.4
LEAD PARTNER	SUITE5
CONTRIBUTORS	CIRCE, NORCE, VdV
DATE	06/07/2026
DISSEMINATION LEVEL	Public

DOCUMENT HISTORY

VERSION	DATE	CHANGES	RESPONSIBLE PARTNER
0.1	28/04/2026	Defined ToC	SUITE5
0.2	04/05/2026	Updated ToC	SUITE5, CIRCE
0.3	12/05/2026	Requirements for the HarvRESt Cloud Platform High-Level Architecture of the HarvRESt Cloud Platform	SUITE5
0.4	13/05/2026	Energy Performance — Self-Consumption Maximisation	SUITE5
0.5	29/05/2026	Validation on the HarvRESt Use Cases	SUITE5
0.6	19/06/2026	Energy Analytics Portfolio (PV, Wind, Biogas, Demand, Storage Flexibility) HarvRESt Dashboard	SUITE5, CIRCE, VdV
0.7	23/06/2026	Integration of multi-vector RES Executive Summary, Introduction, Conclusions	SUITE5, NORCE
0.8	24/06/2026	Candidate version submitted for internal peer review	SUITE5
0.9	30/06/2026	Implementing peer review comments	SUITE5, NORCE, CIRCE
1.0	06/07/2026	Final submission version	CIRCE

DISCLAIMER

The project is 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.

TABLE OF CONTENTS

1.	EXECUTIVE SUMMARY	7
2.	INTRODUCTION	8
2.1	Purpose and structure of the document	8
2.2	Connection with other tasks	8
3.	REQUIREMENTS FOR THE HARVREST CLOUD PLATFORM.....	9
3.1	Requirements Elicitation Approach.....	9
3.2	Functional Requirements.....	9
3.3	Non-Functional Requirements	11
4.	HIGH-LEVEL ARCHITECTURE OF THE HARVREST CLOUD PLATFORM.....	12
4.1	Upstream Data Sources	13
4.1.1	VdV centralised platform	13
4.1.2	ELEXIA platform	13
4.1.3	Farm DERs	14
4.2	Data Collection Layer	14
4.2.1	Data Collection Manager	14
4.2.2	Data Collection Pipeline Execution Service	15
4.3	Data Pre-processing and Quality Layer.....	15
4.3.1	Data Curation and Harmonisation.....	15
4.3.2	Data Observability Service	15
4.3.3	Data Access Control	16
4.4	AI Analytics Layer.....	16
4.4.1	Data Analytics Pipeline Execution Service	16
4.4.2	AI Models Catalogue	16
4.4.3	Self-Consumption Maximisation Engine	16
4.5	HarvRESt Dashboard	17
5.	ENERGY ANALYTICS PORTFOLIO	18

5.1	Overview of the HarvRESt Pre-Trained Analytics	18
5.2	PV Generation Forecasting	18
5.3	Wind Generation Forecasting	20
5.4	Biogas Generation Forecasting	21
5.5	Energy Demand Forecasting	21
5.6	Storage Flexibility.....	22
6.	HARVREST DASHBOARD.....	25
6.1	Energy Performance — Self-Consumption Maximisation	25
6.1.1	Problem framing	25
6.1.2	Inputs from the Energy Analytics Portfolio	25
6.1.3	Dispatch strategy: priority-based self-consumption rule	25
6.1.4	Output and Key Performance Indicators	26
6.1.5	Implementation footprint	27
6.1.6	Adaptation across the HarvRESt Use Cases	27
6.2	HarvRESt Dashboard	28
6.2.1	Energy Performance Overview	28
6.2.2	Next Action — Operational Guidance Messages	30
6.2.3	Strategy Panel — Energy Savings Summary.....	31
6.3	Validation on the HarvRESt Use Cases.....	31
7.	MANAGING INTEGRATION OF MULTI-VECTOR RES ON FARMS: COMPLEMENTARY MODELLING COMPONENTS.....	34
8.	CONCLUSIONS	36
9.	REFERENCES	37

ABBREVIATIONS

AI	Artificial Intelligence
API	Application Programming Interface
AVPP	Agro-based Virtual Power Plant
BESS	Battery Energy Storage System
CO ₂	Carbon Dioxide
DER	Distributed Energy Resource
DNI	Direct Normal Irradiance
DSO	Distribution System Operator
EF _{grid}	Grid Emission Factor (kg CO ₂ eq / kWh)
EMS	Energy Management System
GHI	Global Horizontal Irradiance
KPI	Key Performance Indicator
LGBM	Light Gradient Boosting Machine
MERRA-2	Modern-Era Retrospective Analysis for Research and Applications, Version 2
MQTT	Message Queuing Telemetry Transport
nMAE	Normalised Mean Absolute Error
PCC	Point of Common Coupling
PV	Photovoltaic
RBAC	Role-Based Access Control
REST	Representational State Transfer
SCR	Self-Consumption Ratio
SoC	State of Charge
SSR	Self-Sufficiency Ratio
TMY	Typical Meteorological Year
UC	Use Case
WP	Work Package
XGBoost	Extreme Gradient Boosting

1. EXECUTIVE SUMMARY

This report describes the development of the HarvRESt cloud platform and the HarvRESt Dashboard — the software deliverable of Task 5.4 in Work Package 5. The deliverable in hand covers platform design, AI model training and validation, and the first operational demonstration of the energy management interface across three HarvRESt demo sites: Viñas del Vero (VdV, Spain), Fattoria Solidale del Circeo (FSDC, Italy), and Grønn Gårdsenergi (GGE, Norway).

The HarvRESt cloud platform is organised in four layers. The Data Collection Layer ingests site data from three upstream sources: the VdV centralised platform in Spain, the data from various multi-vector sources at the Norwegian pilot that can be potentially integrated with the ELEXIA platform, and direct on-site monitoring exports for the remaining sites. The Data Pre-processing and Quality Layer standardises, aligns, and validates the incoming data. The AI Analytics Layer hosts five categories of pre-trained models: PV generation forecasting, wind generation forecasting, biogas estimation, energy demand forecasting, and battery storage flexibility simulation. The HarvRESt Dashboard presents the analytics outputs to farm operators through a web interface accessible to non-technical users.

The AI models are trained and evaluated on site-specific datasets. PV generation forecasting achieves explained variance scores between 0.82 (FSDC, Italy) and 0.92 (VdV, Spain). Demand forecasting ranges between 0.37 (VdV, Spain) and 0.89 (FSDC, Italy). Wind generation forecasting, operational for the Norwegian GGE site only, achieves an explained variance of 0.79. LightGBM is adopted as the primary production algorithm across all sites, with XGBoost evaluated as an alternative for the Norwegian site.

The Self-Consumption Maximisation Engine applies a priority-based dispatch rule to 48-hour generation and demand forecasts, producing battery charge/discharge schedules and five KPIs for each site: Self-Consumption Ratio, Self-Sufficiency Ratio, Energy Savings (kWh), Estimated Savings (EUR), and Avoided CO₂ (kg CO₂eq). The HarvRESt Dashboard presents these KPIs alongside a generation-demand timeline, a battery dispatch panel, and plain-language “Next Action” guidance messages. Pre-configured user accounts per active farm provide role-based access to site-specific data.

Chapter 7 documents the multi-vector energy systems advisor. The decision-support tool allows farmers to evaluate and simulate combinations of renewable energy technologies (solar PV, vertical-axis wind, hydropower, biogas CHP, wood-chip boilers, and energy storage) through a web-based interface, prior to making investment decisions. Together with the HarvRESt Dashboard, it constitutes a complete farmer-facing toolchain spanning pre-deployment RES planning through to day-to-day operational energy management.

2. INTRODUCTION

2.1 Purpose and structure of the document

D5.4 is the software deliverable of T5.4 in WP5. It documents the design, development, and initial validation of the HarvRESt cloud platform and the HarvRESt Dashboard, a web-based energy management interface for farm operators at the HarvRESt demo sites. The deliverable encompasses the technical architecture of the platform, the AI Analytics Portfolio trained on site-specific datasets, the Self-Consumption Maximisation Engine, and the full dashboard implementation including KPI methodology and operational guidance features.

The document is structured as follows:

Chapter 3 defines the functional and non-functional platform requirements;

Chapter 4 describes the four-layer cloud architecture (Data Collection, Data Pre-processing and Quality, AI Analytics, and HarvRESt Dashboard) and its upstream data sources;

Chapter 5 presents the Energy Analytics Portfolio, covering PV generation forecasting, wind generation forecasting, biogas estimation, energy demand forecasting, and battery storage flexibility simulation;

Chapter 6 describes the HarvRESt Dashboard in full: the Self-Consumption Maximisation algorithm, the output KPIs, the dashboard interface, and the validation of the platform on each active use case;

Chapter 7 presents the multi-vector energy systems advisor, a decision-support tool that addresses the pre-deployment dimension of T5.4 mandate to develop smart energy systems for agricultural settings, by enabling farmers to evaluate and size optimal combinations of renewable energy technologies before commissioning; and

Chapter 8 concludes the deliverable.

2.2 Connection with other tasks

D5.4 builds on two predecessor deliverables in WP5. D5.1 established the machine learning methodology and the literature-based biogas substrate database that underpin the Energy Analytics Portfolio described in chapter 5. D5.2 delivered the first validated AI model prototypes; D5.4 extends those results with site-specific training on the full demo-site datasets and integrates the models into the production platform.

- The platform's Data Collection Layer is designed to interface with the HarvRESt data management infrastructure coordinated under WP3. Currently, this interface operates in batch mode, ingesting static dataset exports from partners. The architecture additionally supports real-time API and MQTT streaming connections.
- Live DER data for model training and platform validation are provided by WP7. Three upstream data pathways have been identified. The VdV centralised platform in Spain aggregates measurements from the winery's distributed monitoring equipment. The ELEXIA system in Norway can serve as a digital service platform for the GGE demo site after the installation of all planned components at the site. For the remaining sites (FSDC, VRT), data are delivered directly as on-site monitoring exports. At the time of writing, all datasets are static exports with mixed temporal granularities and non-overlapping time windows.

3. REQUIREMENTS FOR THE HARVREST CLOUD PLATFORM

This chapter presents the requirements for the HarvRESt cloud platform components developed under Task 5.4, as elicited from the project's Use Cases (UCs) and their technical counterparts. The requirements cover the data collection, data pre-processing and quality, AI analytics, energy optimisation and visualisation capabilities of the platform that together support the centralised, EU-replicable management of Distributed Energy Resources (DERs) at farm level.

3.1 Requirements Elicitation Approach

The requirements presented in this chapter were elicited through a structured process involving the HarvRESt demo-site partners and their technical counterparts within the consortium, guided by the operational context of the HarvRESt use cases and the characteristics of each site's distributed energy resources. The elicitation activities comprised desk research and iterative consultations with the technical and demonstration partners, progressively refined as the platform design advanced and as the data exchange arrangements with each demo site were consolidated.

The collected requirements were classified by area, namely data collection, data quality and pre-processing, data privacy and security, AI analytics and forecasting, and energy optimisation and strategy communication – and by type, distinguishing functional from non-functional requirements. Each requirement was prioritised to ensure that the HarvRESt cloud platform delivers, in its centralised cloud deployment, the core capabilities needed to maximise the use of RES in agricultural settings.

3.2 Functional Requirements

Table 1 lists the functional requirements of the HarvRESt cloud platform, grouped under the five areas of the platform: data collection, data quality and pre-processing, data privacy and security, AI analytics and forecasting, and energy optimisation and strategy communication. Each requirement is identified by a unique HarvRESt Cloud Functional (HCF) identifier, classified as functional, and assigned a priority.

Table 1. HarvRESt cloud platform functional requirements

Req ID	Description	Type	Priority
Data Collection			
HCF_001	The HarvRESt platform shall accept batch data files describing DER measurements (e.g., CSV, JSON, XML) submitted by authorised data providers.	Functional	High
HCF_002	The HarvRESt platform shall enable the scheduled collection of batch data files of significant volume from farm-side storage locations into the centralised HarvRESt cloud platform.	Functional	Medium
HCF_003	The HarvRESt platform shall enable scheduled retrieval of DER data assets via vendor or farm monitoring system APIs.	Functional	High
HCF_004	The HarvRESt platform shall support near-real-time and streaming data collection from DERs via publish/subscribe mechanisms (e.g., Kafka, MQTT).	Functional	High
HCF_005	The HarvRESt platform shall execute configured data collection pipelines automatically at the defined schedule and persist the collected DER data in the platform's cloud storage.	Functional	High
HCF_006	The HarvRESt platform shall enable mapping of raw DER data fields to a HarvRESt common data model during pipeline configuration.	Functional	High
Data Quality and Pre-processing			

Req ID	Description	Type	Priority
HCF_007	The HarvRESt platform shall allow the user to configure curation rules for handling incorrect, incomplete, inaccurate, irrelevant or missing parts of DER data.	Functional	High
HCF_008	The HarvRESt platform shall apply automated harmonisation rules to align raw DER data from heterogeneous sources (e.g., PV inverters, wind turbines, smart meters).	Functional	High
HCF_009	The HarvRESt platform shall maintain an audit log of pipeline executions, including success, failure and data-quality events, accessible to platform operators.	Functional	Medium
Data Privacy and Security			
HCF_010	The HarvRESt platform shall enforce role-based access control so that each data asset is accessible only to the farm stakeholders authorised by the corresponding data provider.	Functional	High
AI Analytics and Forecasting			
HCF_011	The HarvRESt platform shall provide an AI analytics layer in which DER data are read from the platform's storage and used to train and execute analytics pipelines whose outputs are made available to downstream components.	Functional	High
HCF_012	The HarvRESt platform shall support generation forecasting extraction at DER and farm levels for PV and wind sources.	Functional	High
HCF_013	The HarvRESt platform shall support energy demand forecasting extraction at farm level, covering combined farm load (e.g., irrigation, climate control, processing equipment).	Functional	High
HCF_014	The HarvRESt platform shall support the extraction of battery storage flexibility profiles and state-of-charge forecasts to be factored into self-consumption strategy computation.	Functional	High
Energy Optimisation and Strategy Communication			
HCF_015	The HarvRESt platform shall fuse the generation and demand forecasts produced by the pre-trained models to compute an optimal energy dispatch strategy for each farm across the scheduling horizon.	Functional	High
HCF_016	The HarvRESt platform shall optimise the dispatch strategy with the objective of maximising on-site self-consumption of locally generated renewable energy and minimising grid imports.	Functional	High
HCF_017	The HarvRESt platform shall consider battery storage state-of-charge limits, DER generation capacity and the forecasted farm demand profile as constraints during optimisation.	Functional	High
HCF_018	The HarvRESt platform shall produce time-series strategy signals indicating, for each scheduling step, when to store, when to consume and when to export excess generation.	Functional	High
HCF_019	The HarvRESt platform shall provide a dashboard with real-time and forecast views of generation, demand and storage state, the recommended strategy and key energy performance indicators (KPIs) at farm level.	Functional	High
HCF_020	The HarvRESt platform shall support a batch / replay execution mode in which historical DER data feed the analytics and dispatch pipelines on a configurable schedule, in the absence of streaming data sources.	Functional	High

3.3 Non-Functional Requirements

Table 2 summarises the non-functional requirements of the HarvRESt cloud platform, covering interoperability, security and access control, and scalability and performance. Each requirement is identified by a unique HarvRESt Cloud Non-functional (HCNF) identifier.

Table 2. HarvRESt cloud platform non-functional requirements

Req ID	Area	Description
Interoperability		
HCNF_001	Interoperability	The HarvRESt platform shall adopt standards-based data models for the representation of DER measurements, generation, demand and storage data, to support EU-wide replicability across diverse farm contexts.
HCNF_002	Interoperability	The HarvRESt platform shall expose well-documented APIs (REST, publish/subscribe) for upstream data collection from the upstream data sources at each site (the VdV centralised platform in Spain, the ELEXIA platform in Norway, and on-site farm DER monitoring systems).
HCNF_003	Interoperability	The HarvRESt platform shall support the configuration of new data collection pipelines for additional farms without changes to the platform code base, in order to ease replication beyond the project Use Cases.
Security and Access Control		
HCNF_004	Security and Access Control	The HarvRESt platform shall implement role-based access control distinguishing platform operators, farm data managers and DER owners, with privileges aligned to each role's responsibilities.
HCNF_005	Security and Access Control	The HarvRESt platform shall securely manage credentials (e.g., API keys, tokens, certificates) used to access upstream DER data sources, following industry good practice for secret management.
Scalability & Performance		
HCNF_006	Scalability & Performance	The HarvRESt platform shall be deployable on cloud infrastructure with horizontal auto-scaling for compute and storage to accommodate the data volumes produced by additional farms.
HCNF_007	Scalability & Performance	The HarvRESt platform shall support the concurrent execution of multiple data collection and analytics pipelines from different farms without degradation of service for any individual pipeline.
HCNF_008	Scalability & Performance	The HarvRESt platform shall execute pre-trained energy forecasting and self-consumption optimisation pipelines with latency compatible with the day-ahead and intra-day scheduling horizons used by farm operators.

The requirements collectively establish the functional and operational boundaries within which the HarvRESt cloud platform was designed and implemented. Taken together, they ensure that the platform is capable of ingesting and processing heterogeneous agricultural energy datasets, delivering accurate AI-driven forecasts, executing the self-consumption maximisation strategy, and presenting results through a secure, role-controlled, and scalable dashboard interface. The architectural choices described in chapter 4 and the analytics implementations described in chapter 5 are directly traceable to these requirements.

4. HIGH-LEVEL ARCHITECTURE OF THE HARVREST CLOUD PLATFORM

The HarvRESt cloud platform is designed as a centralised, layered architecture that turns raw DER measurements from farm sites into actionable energy management strategies for DER owners. The platform is organised into four layers – Data Collection, Data Pre-processing and Quality, AI Analytics, and the HarvRESt Dashboard – deployed centrally on the cloud. Each layer exposes well-defined inputs and outputs to the next, so that the addition of a new farm or DER type requires configuration only, with no changes to platform code.

Data flow follows a single direction. DER data (e.g., PV panels, wind turbines, battery storage systems) are collected from the upstream data sources described in section 4.1 – the VdV centralised platform in the Spanish UC, current on-site data collection in the Norwegian UC, which will be changed to the ELEXIA platform when all infrastructure is completed, and direct monitoring exports from on-site equipment at the Italian and remaining Spanish sites.

The Data Collection Layer brings these data into the platform; the Data Pre-processing and Quality Layer curates and harmonises them; the AI Analytics Layer applies the pre-trained forecasting models and executes the self-consumption maximisation engine; finally, the HarvRESt Dashboard presents the resulting strategies and KPIs to farm operators and DER owners through a user-friendly graphical interface.




Figure 1. HarvRESt cloud platform – layered architecture

The design follows three guiding principles. First, interoperability: all components rely on standards-based data models and open APIs to support replicability beyond the project UCs. Second, modularity: each component is responsible for a well-scoped function and can evolve independently provided its interface contracts are preserved. Third, replicability: the platform configuration is parameterised so that the same architecture supports diverse farm contexts across the EU.

Although the HarvRESt Dashboard is the only user-facing interface of the platform, the upstream Data Collection, Data Pre-processing and Quality, and AI Analytics layers are equally part of the HarvRESt cloud platform. They run unattended on the cloud, continuously transforming raw DER measurements into the forecasts, strategies and KPIs that the dashboard presents to its users; together with the dashboard they constitute a complete solution covering data collection, data management, analytics and energy performance optimisation.

4.1 Upstream Data Sources

The HarvRESt cloud platform ingests DER data from three categories of upstream sources: the VdV centralised platform at the Spanish UC, the ELEXIA platform associated with the Norwegian UC, and direct on-site monitoring exports from the Italian and remaining Spanish sites. In the current implementation all datasets are delivered as static exports — comma-separated files at sub-hourly to hourly resolution — and are ingested via the batch-upload mode of the Data Collection Manager (section 4.2.1).

4.1.1 *VdV centralised platform*

The VdV centralised platform is an energy management and monitoring installation at the Viñas del Vero (VdV) winery estate in Barbastro, Aragon, Spain. The Spanish VdV UC operates a 40.95 kWp single-axis tracking PV array paired with a battery energy storage system, managed by a site-level platform that records and exports the farm energy balance at 15-minute resolution. Data are delivered as daily CSV files, each carrying six channels expressed in kilowatts: self-supply, grid consumption, total farm consumption, grid injection, self-consumption, and PV generation. The HarvRESt Data Repository holds 242 such daily exports spanning 13 January 2024 to 10 September 2024 (approximately eight months of continuous operation). The format and daily naming convention of these exports are consistent with the systematic export function of an EMS-class monitoring platform; the specific software product has not been formally identified.

4.1.2 *ELEXIA platform*

The ELEXIA platform is an energy monitoring service which will be deployed at the Grønn Gårdsenergi (GGE) farm at Buevegen, Norway. Currently available data files cover PV production power at 5-to-10-minute resolution over the period February 2023 to January 2024, together with aggregated daily PV energy production in kWh. Electricity consumption data for the farm's total electrical load are delivered separately at hourly resolution and originate from grid-operator smart-metering infrastructure. The association between these files and the ELEXIA platform will be inferred from project documentation and naming conventions. In addition to the measured time-series, the Norwegian portion of the HarvRESt Data Repository includes modelled generation profiles: a 50 kW and a 200 kW wind-turbine scenario and a 2,100 kW PV expansion scenario, all calibrated to 2019 as the model year and using MERRA-2 reanalysis weather data. These modelled profiles are used exclusively for the wind generation forecasting pipeline (section 5.3) and are explicitly labelled as model-derived in the HarvRESt analytics outputs.

4.1.3 *Farm DERs*

For the Italian and remaining Spanish UCs, DER data are delivered directly from on-site monitoring equipment rather than through a named centralised platform. At the Fattoria Solidale del Circeo (FSDC) site in Lazio, Italy, three PV installations are present: a 5.5 kWp glass-glass bifacial single-axis-tracked greenhouse array (15-minute interval power exports in watts, February 2024 to February 2025, with a gap from September to December 2024 due to an inverter failure), an 80 kWp rooftop installation (hourly generation data including net flows to battery storage, January 2025 to April 2025), and a 70 kWp ground-mounted array for which no generation data have been delivered at the time of writing. Electricity consumption at FSDC is supplied from Distribution System Operator (DSO) smart-meter readings at hourly resolution covering January 2024 to April 2025. These data originate directly from PV inverter monitoring software and national grid-operator metering and do not pass through a named EMS. At the Viñedos del Río Tajo (VRT) estate in Extremadura, Spain, a 40.95 kWp single-axis tracking PV installation is in operation; the Data Repository contains hourly electricity consumption readings from the CUPS meter (warehouses, offices, and ancillary loads) but does not include a measured PV generation time series for this site at the time of writing. At the Sorigué (ACSA) biogas plant in Catalonia, Spain, daily operational records covering digester feed flows, biogas yield, and biomethane output are available for 2024; these data support the biogas generation forecasting module (section 5.4) and are not incorporated into the HarvRESt dashboard.

4.2 Data Collection Layer

The Data Collection Layer brings DER data from farm sites into the centralised HarvRESt cloud platform in a reliable, scheduled and configurable manner. It comprises two backend components: the Data Collection Manager, which handles pipeline configuration and onboarding, and the Data Collection Pipeline Execution Service, which executes the configured pipelines at runtime. The layer is designed to support three collection modes — batch file upload, scheduled API pull, and near-real-time streaming via MQTT or Kafka — enabling the platform to accommodate both historical batch exports and live telemetry feeds as the data pipelines of each UC mature. In the current HarvRESt implementation, the upstream sources described in section 4.1 have delivered data exclusively as static batch exports; the API and streaming modes are fully architected and can be activated through pipeline configuration without changes to platform code.

4.2.1 *Data Collection Manager*

The Data Collection Manager is the backend configuration component of the HarvRESt cloud platform. It captures, for each DER asset or group of assets, the operational arrangements agreed with the UCs during onboarding and translates them into the end-to-end pipeline specifications executed at runtime. The configuration is performed by HarvRESt platform developers in cooperation with the data providers and is not exposed as a separate interface to the end users of the platform.

The Data Collection Manager supports three collection modes. Batch upload accepts CSV or JSON files exported from farm monitoring systems or EMS platforms and uploaded directly by the farm operator or platform engineer; this is the primary mode used in the current HarvRESt implementation, given that all UC datasets are delivered as static exports (see section 4.1). API pull schedules periodic retrieval of DER data from vendor or farm monitoring system APIs. Streaming supports near-real-time collection via publish/subscribe mechanisms such as MQTT and Kafka, for assets that produce high-frequency telemetry.

For each pipeline, the configuration steps are: (a) selecting the collection method (batch, API or streaming); (b) mapping raw DER data fields to a HarvRESt common data model using the field inventory observed for each upstream source (section 4.1); (c) configuring curation and pre-processing rules to address missing values,

outliers and unit conversions; (d) defining the execution schedule. The output of the configuration step is a versioned pipeline specification stored centrally on the platform, ready to be executed by the Data Collection Pipeline Execution Service.

Overall, the Data Collection Manager standardises the way farm-side DER data are brought into the HarvRESt cloud platform across heterogeneous sources, facilitating both the project UCs and the replication of the platform to additional farms.

4.2.2 *Data Collection Pipeline Execution Service*

The Data Collection Pipeline Execution Service is the runtime component that executes the data collection pipelines defined through the Data Collection Manager. It is triggered automatically on the schedule set at design time for each pipeline, and runs centrally on the HarvRESt cloud. For each execution, it applies the configured collection method, mapping and curation rules in sequence.

The result of each execution is a collected, staged and stored DER data asset, persisted in the HarvRESt cloud platform's data storage and made available to the downstream pre-processing and analytics layers. Execution outcomes (success, failure, data-quality alerts) are logged through the Data Observability Service and surfaced to the platform operators for follow-up.

4.3 Data Pre-processing and Quality Layer

The Data Pre-processing and Quality Layer ensures the quality, consistency and security of the collected DER data before they are used by the AI Analytics Layer. It comprises Data Curation and Harmonisation rules, the Data Observability Service.

4.3.1 *Data Curation and Harmonisation*

Data curation and harmonisation are applied during the execution of each data collection pipeline. Curation rules address common DER data-quality issues: detection and treatment of missing values, identification of outliers (e.g., implausible PV output values during night-time, abnormal wind generation peaks) and filling of short gaps via interpolation where appropriate.

Harmonisation translates raw measurements from heterogeneous sources, be it PV inverter monitoring systems, EMS platform exports, wind turbine controllers, battery storage management systems, or grid-operator smart meters, into a HarvRESt common data model, aligning units, time stamps, sampling rates and naming conventions. The harmonised data assets become the canonical input for the AI Analytics Layer and for any KPI computed by the dashboard.

4.3.2 *Data Observability Service*

The Data Observability Service captures the operational status and data quality of each data collection pipeline at runtime. It is scheduled to run alongside the execution of every pipeline and continuously assesses a set of pipeline-level and data-level metrics: pipeline execution status (success or failure, with classified error codes), data completeness (missing records, gaps, missing values), data freshness and the actual versus scheduled update rate of each DER data asset.

When anomalies or threshold breaches are detected, the Data Observability Service emits log entries with a pertinent criticality level and notifies platform operators. This supports troubleshooting of data collection pipelines, minimises data downtime, and safeguards the quality and reliability of the data feeding the AI Analytics Layer.

4.3.3 *Data Access Control*

The HarvRESt platform enforces data access control to ensure that farm-level DER data are accessible only to authorised stakeholders. Access to each data asset is governed by the role assigned to the requesting user: platform operators may manage pipeline configurations and review observability logs; farm data managers may access and export the data assets of their own farm; DER owners may view the data and KPIs associated with their own installations. These controls are applied uniformly across data collection, pre-processing and analytics layers.

4.4 *AI Analytics Layer*

The AI Analytics Layer executes the energy forecasting models and computes the self-consumption maximisation strategy for each farm. It is structured as a set of configurable, executable pipelines that run on the centralised HarvRESt cloud and combine pre-trained AI models from the platform catalogue with the optimisation engine. The pre-trained AI models that run on this layer are described in detail in chapter 5.

4.4.1 *Data Analytics Pipeline Execution Service*

The Data Analytics Pipeline Execution Service is responsible for executing the data analytics pipelines configured for the HarvRESt cloud platform. It invokes the appropriate AI framework or library for each step of the pipeline (data preparation, model application, post-processing and output storage) based on the configuration parameters supplied at design time.

All pipelines run on the centralised HarvRESt cloud infrastructure. The service reads input DER data from the platform's data storage – the output of the Data Collection and Data Pre-processing layers – and applies the configured analytics steps to produce generation and demand forecasts for downstream use by the Self-Consumption Maximisation Engine and the HarvRESt Dashboard.

4.4.2 *AI Models Catalogue*

The AI Models Catalogue is the HarvRESt-specific, energy-domain repository of pre-trained models that are ready to be executed as analytics building blocks on the platform: PV and wind generation forecasting (per asset), farm-level energy demand forecasting, and battery storage flexibility analytics. The catalogue is the single point of reference for the analytics that feed the Self-Consumption Maximisation Engine.

Each model in the catalogue carries with it the metadata necessary for its execution (i.e., the expected input schema, the training data lineage, the validation metrics, the version of the underlying framework and any retraining cadence) so that the Data Analytics Pipeline Execution Service can invoke it deterministically and reproducibly. New models can be added or existing models can be re-trained without changes to the platform code base: registering a new entry in the catalogue and pointing the corresponding analytics pipeline to it is enough. The pre-trained models held in the catalogue are described in detail in chapter 5.

4.4.3 *Self-Consumption Maximisation Engine*

The Self-Consumption Maximisation Engine is the algorithmic component that fuses the generation and demand forecasts produced by the pre-trained models and computes the optimal energy dispatch strategy for each farm across the scheduling horizon. It is a HarvRESt-specific adaptation of the energy management optimisation approach to the farm scale: instead of operating at community or grid level, it is configured around the assets, loads and constraints of an individual farm. For each forecast input —PV generation, wind generation and farm demand— the platform exposes a forecaster interface that can be implemented either by a pre-trained AI model or by a deterministic baseline (e.g., persistence with day-of-week and hour-of-day

seasonality); the Engine is agnostic to the specific implementation, provided that the forecaster returns a time series at the scheduling resolution.

The optimisation objective is the maximisation of on-site self-consumption: the engine schedules the use of locally generated renewable energy first and minimises imports from the grid, while taking the forecasted farm demand profile into account. The constraints considered include battery storage state-of-charge limits, charge/discharge rates, the rated capacity of each DER and the forecasted demand of farm processes such as irrigation, climate control and on-site processing equipment.

The output of the Self-Consumption Maximisation Engine is a time-series strategy that, for each scheduling step, indicates when to consume locally generated energy, when to charge or discharge the battery storage and when excess generation is available for export. The strategy and the associated KPIs (e.g., self-consumption rate, estimated cost savings, excess generation) are made available to the HarvRESt Dashboard for communication to farm operators and DER owners.

The Self-Consumption Maximisation Engine is exposed to the rest of the platform as a pure function: it consumes the per-farm forecasts and asset configuration, it returns a dispatch schedule and the associated KPIs in a structured payload, and it has no side effects. The HarvRESt Dashboard reads this payload directly and renders it; no other component needs to know how the schedule was computed. This decoupling is what allows the same engine to drive the dashboard for any HarvRESt farm with no code change.

4.5 HarvRESt Dashboard

The HarvRESt Dashboard is the primary interface between the HarvRESt cloud platform and the human users of the system, be them farm operators, energy managers and DER owners. It is equipped with farm-specific views and KPI panels that turn the outputs of the AI Analytics Layer into actionable information.

The dashboard provides the user with real-time and forecast views of generation, demand and storage state, a recommended energy strategy panel sourced from the Self-Consumption Maximisation Engine, and a set of KPIs covering the self-consumption rate, estimated cost savings and excess generation that may be available for trading. The dashboard reads the Self-Consumption Maximisation Engine's output payload directly, which makes the same dashboard layout usable for any HarvRESt farm regardless of its asset mix. Screenshots and a detailed walkthrough of the dashboard are provided in chapter 6.

5. ENERGY ANALYTICS PORTFOLIO

5.1 Overview of the HarvRESt Pre-Trained Analytics

The HarvRESt Energy Analytics Portfolio comprises five categories of pre-trained analytics deployed on the centralised cloud platform: PV generation forecasting, wind generation forecasting, biogas generation forecasting, energy demand forecasting, and battery storage flexibility simulation. Of these, the first four are AI-driven predictive models; the fifth is a deterministic state-of-charge simulator that translates the generation and demand forecasts into battery scheduling boundaries. Together they constitute the input layer of the Self-Consumption Maximisation Engine described in section 6.1.

All forecasting models in the portfolio share a common design philosophy: they are trained on the static datasets delivered by the HarvRESt demo-site partners, they expose a uniform forecast (history, horizon) interface so that the underlying implementation is interchangeable without changes to the dispatch engine or the dashboard, and they are configured to produce 48-step-ahead (48-hour) forecast horizons at hourly resolution. Weather covariates (irradiance components, wind speed, temperature, humidity, dew point) are sourced from the Open-Meteo API using the geographic coordinates of each site. The primary modelling algorithm is LightGBM (LGBM), selected for its efficiency on tabular time-series data; XGBoost (XGB) is evaluated alongside LGBM for the Norwegian UC, where data availability allowed a model comparison. The three active UCs in the HarvRESt Dashboard are the Italian UC (FSDC), the Norwegian UC (GGE), and the Spanish UC (VdV).

5.2 PV Generation Forecasting

PV generation is forecasted across all three active UCs: the Italian UC (a 5.5 kWp single-axis-tracked greenhouse array), the Norwegian UC (a multi-array rooftop system at Buevegen farm), and the Spanish UC (a single-axis tracked array). In all cases the primary driver features are solar irradiance components — Direct Normal Irradiance (DNI) and Global Horizontal Irradiance (GHI) — sourced as 48-hour lead forecasts from Open-Meteo and enriched with cyclical calendar encodings for hour of day and, where seasonality spans are long enough, month of year. For the Norwegian and Spanish UCs, autoregressive lags of the PV generation signal itself are also included to capture site-specific cloud-shadow and tracking patterns.

The models are trained on the non-overlapping portion of each site's available history, with a common held-out test set of approximately 14 days used to evaluate 48-step-ahead performance. Table 3 summarises training windows, test windows, and performance metrics on the test set for each site and model variant.

Table 3. PV generation forecasting — training configuration and test-set performance

Use Case	Model	Feature Sets	Train Window	Test Window	Explained Variance	nMAE
Italy (FSDC)	LightGBM	DNI leads 0–48h, GHI leads 0–48h, Hour (cyclical)	2024-01-02 to 2025-03-16	2025-03-17 to 2025-03-31	0.824	0.324
Norway (GGE)	LightGBM	PV Gen lags 0–24h, DNI leads 0–48h, GHI leads 0–48h, Month / Hour (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.806	0.347
Norway (GGE)	XGBoost	DNI leads 0–48h, GHI leads 0–48h, Month / Hour (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.806	0.348
Spain (VdV)	LightGBM	PV Gen lags 1–168h (selected), Shortwave / DNI / GHI leads 1–48h, Month / Hour (cyclical)	2024-01-13 to 2024-07-31	2024-08-01 to 2024-09-08	0.920	0.160

The Spanish VdV site achieves the highest accuracy (explained variance 0.92, nMAE 0.16), benefiting from a denser dataset spanning eight months and a more regular daily irradiance cycle driven by its single-axis tracker. The Italian and Norwegian sites achieve solid but lower explained variance (around 0.82), consistent with shorter training windows and greater cloud variability at those latitudes. For Norway, LightGBM and XGBoost achieve virtually identical performance; LightGBM is adopted as the production model for consistency with the other UCs.

Dataset scope and preprocessing: Across the three active UCs, PV generation training data vary substantially in duration and temporal density. The Italian FSDC dataset covers 74 days (21 January – 4 April 2025; 1,755 hourly rows), with three partially recorded days identified on 21 January (12 hours), 1 March (23 hours), and 4 April (16 hours). The Spanish VdV dataset is the densest, spanning approximately eight months (13 January – 10 September 2024; 5,808 rows, including 58 missing generation values). The Norwegian GGE dataset covers 212 days (9 February – 9 September 2023; consolidated hourly entries). All datasets are converted to UTC, resampled to a common 1-hour granularity, and augmented with weather covariates from the Open-Meteo API at the respective site coordinates.

Feature engineering: PV generation models are trained on irradiance-forward feature matrices. The Italian and Norwegian models use DNI and GHI leads from 0 to 48 hours ahead, combined with cyclical hour-of-day encodings. The Spanish model extends the feature set with generation lags at 1–24, 32, 40, 48, 72, 96, 120, 144, and 168 hours, shortwave radiation leads (1–48 h), and cyclical month and hour encodings. The train/test split uses a chronological cut-off: approximately the final two weeks of each site’s data window form the held-out evaluation set, with all earlier records used for training.

5.3 Wind Generation Forecasting

Wind generation forecasting is modelled for the Norwegian UC only. No other HarvRESt demo site has an operational wind turbine within the evaluation scope. An important data provenance note applies to this UC: the wind generation time series used for training and evaluation is a modelled output generated by a simulation platform using MERRA-2 reanalysis wind resource data for the Buevegen site, rather than a measured production series from an installed turbine. This is because the GGE wind turbine is planned but not yet operational. The dashboard flags this accordingly: wind generation figures for the Norwegian UC are labelled as modelled estimates.

The feature set is designed around the physical relationship between wind speed and power output. Beyond the wind speed forecast itself (Open-Meteo, 48-hour lead), the cube of wind speed is included as an explicit feature, reflecting the cubic relationship between wind speed and turbine power in the region below rated capacity. Autoregressive lags of the generation signal at both dense (hourly) and sparse (4- to 8-hour) intervals capture persistence and diurnal wind patterns. Both LightGBM and XGBoost are evaluated; their performance on the test set is reported in Table 4.

Table 4. Wind generation forecasting — training configuration and test-set performance
(Norwegian UC only; generation data are modelled)

Use Case	Model	Feature Sets	Train Window	Test Window	Explained Variance	nMAE
Norway (GGE)	LightGBM	Wind Gen lags 0–24h, Wind speed leads 0–48h, Wind speed lags (sparse), Wind speed ³ leads 0–48h, Month / Hour (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.792	0.230
Norway (GGE)	XGBoost	Wind Gen lags (sparse 0–48h), Wind speed leads 0–48h, Wind speed ³ leads 0–48h, Month / Hour (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.780	0.236

LightGBM achieves a slightly higher explained variance (0.792 vs 0.780) and lower nMAE (0.230 vs 0.236) and is adopted as the production model. The level of explained variance is consistent with the nature of the target: wind output from a modelled series at a single site is inherently smoother than measured data from a turbine, which tends to reduce the upper bound on model accuracy metrics. In production, when real turbine data become available, the model can be retrained without changes to the forecasting interface.

Dataset scope and preprocessing: The Norwegian wind dataset covers the same 212-day window as the PV and demand datasets (9 February – 9 September 2023), with PV generation, wind power, and electricity demand consolidated into a single unified hourly matrix. Missing values and timeline inconsistencies were addressed during initial offline processing before upload to the platform. Cyclical calendar features are extracted for month, day-of-week, and hour-of-day.

Feature engineering: Two feature configurations are evaluated. For XGBoost, wind generation lags are provided at 0, 4, 8, ..., 48-hour intervals (step 4 h); for LightGBM, lags cover the full 0–24-hour window. Both configurations include Open-Meteo wind speed leads (0–48 h) and wind speed cubed leads (0–48 h) — the cubic relationship captures the non-linear dependency of aerodynamic power on wind speed encoded in the

Betz turbine power curve. The train window runs from 12 February to 31 July 2023; the evaluation window from 1 August to 7 September 2023.

5.4 Biogas Generation Forecasting

The HarvRESt biogas generation forecasting functionality targets long-term annual biogas production estimation for anaerobic digestion systems, with a specific focus on agro-industrial feedstock streams. Unlike solar and wind generation, biogas production is not primarily driven by meteorological conditions, but by the availability, composition and biodegradability of the organic substrates fed into the digester. The module therefore estimates the expected production potential of raw or mixed feedstocks, supporting stakeholders in assessing alternative feeding strategies for existing biogas plants. The proposed methodology follows a feedstock-driven modelling approach based on the characterization of organic substrates and their conversion potential. This approach is based on the characterisation of organic substrates and their conversion potential, using parameters such as TS, VS, BMP and biodegradability as the core descriptors of production.

The work reported in D5.1 established the basis of the biogas prediction tool by compiling a literature-based database covering the most relevant substrate categories for the project context and by deriving representative parameter values for each of them. These values were synthesised and constitute the reference basis for long-term biogas production calculations. D5.1 identifies this work as the completion of Phase 1 of the tool, i.e. the methodological stage that enables long-term production estimation and prepares the subsequent software integration steps.

On top of this methodological basis, the AVPP developed in T6.1 incorporates the biogas generation calculation by implementing the logic derived from D5.1 in a user-oriented software component. The user provides the required inputs related to the feedstock or feedstock mixture, and the tool computes the corresponding biogas production estimate. In this way, the biogas functionality has been integrated into the AVPP as a practical calculation module that allows users to explore the production potential associated with different substrate configurations.

5.5 Energy Demand Forecasting

Energy demand forecasting is deployed for all three active UCs. The target variable is the combined farm electrical load, i.e., irrigation pumps, climate control, post-harvest processing, and auxiliary equipment, measured at the farm's point of common coupling. For the Spanish VdV UC, the 15-minute energy balance exports are resampled to hourly resolution to match the granularity of the demand model; for Norway and Italy, hourly readings from the DSO meter or grid-operator smart meter are used directly. Demand is inherently harder to forecast than solar generation because it reflects human and agronomic decisions that do not follow a simple physical law. The feature engineering therefore relies primarily on autoregressive lags of the demand signal itself, complemented by weather covariates (temperature, humidity, dew point) that correlate with irrigation and climate control loads, and cyclical calendar encodings.

Table 5 summarises the model configuration and test-set performance for each site.

Table 5. Energy demand forecasting

Use Case	Model	Feature Sets	Train Window	Test Window	Explained Variance	nMAE
Italy (FSDC)	LightGBM	Demand lags 0–48h, Hour (cyclical)	2024-01-02 to 2025-03-16	2025-03-17 to 2025-03-31	0.894	0.189

Use Case	Model	Feature Sets	Train Window	Test Window	Explained Variance	nMAE
Norway (GGE)	LightGBM	Demand lags 0–24h, Dew point leads 0–48h, Relative humidity leads 0–48h, Month / Hour / Day of week (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.384	0.155
Norway (GGE)	XGBoost	Demand lags 0–48h, Dew point leads 0–48h, Relative humidity leads 0–48h, Month / Hour / Day of week (cyclical)	2023-02-12 to 2023-07-31	2023-08-01 to 2023-09-07	0.368	0.156
Spain (VdV)	LightGBM	Demand lags 1–168h (selected), Temperature / Humidity lags (selected), Month / Hour (cyclical)	2024-01-13 to 2024-07-31	2024-08-01 to 2024-09-08	0.370	0.330

The Italian UC achieves the strongest demand forecast performance (explained variance 0.894, nMAE 0.189), reflecting a relatively stable and repetitive facility load at the FSDC greenhouse and packhouse. The Norwegian and Spanish UCs show substantially lower explained variance (0.37–0.38), though the normalised error for Norway is modest (nMAE 0.155), indicating that while the model fails to capture high-amplitude peaks, it tracks the general level of demand reasonably well. For the Spanish VdV site, both metrics are weaker (EV 0.37, nMAE 0.33), consistent with the highly irregular, irrigation-dominated load profile of a vineyard operation where peak demand events are driven by crop management decisions rather than a predictable schedule. In the absence of tariff and operational schedule data from partners, these baseline LightGBM models represent the best achievable accuracy with the available data; the forecasting interface remains unchanged should richer contextual features be provided in future iterations.

Dataset scope and preprocessing: The Italian FSDC demand dataset is the longest available, covering 458 days (1 January 2024 – 2 April 2025; 11,064 hourly rows). A data quality issue on 27 October 2024 (DST changeover) produced 96 sub-hourly entries; these were resolved by averaging to standard hourly values. The Spanish VdV demand dataset mirrors the generation window (13 January – 10 September 2024), with 96 missing values imputed by linear interpolation. The Norwegian GGE demand dataset is co-registered with the generation data (212 days, 9 February – 9 September 2023) and cleaned for missing data points and timeline shifts.

Feature engineering: Demand forecasting relies primarily on autoregressive features. The Italian and Spanish models use demand lags at 1–24, 32, 40, 48, 72, 96, 120, 144, and 168 hours. The Spanish model additionally includes temperature and humidity lags at selected horizons (4, 8, 16, 24, 43, 40, 48, 72, 96, 120, 144, 168 h). The Norwegian model is evaluated with two configurations: XGBoost uses demand lags at 0–48 h plus dew-point temperature and relative humidity leads at 0–48 h; LightGBM uses lags at 0–24 h with the same weather leads. All models include cyclical calendar features (month, hour-of-day, day-of-week where relevant). The evaluation window is consistently the final ~15 days of each site’s data window.

5.6 Storage Flexibility

The Storage Flexibility module provides the battery state-of-charge (SoC) simulation that the Self-Consumption Maximisation Engine uses to determine when and how much energy to store or release. Unlike the forecasting

models, this module is deterministic: given the net residual power at each time step and the current SoC, it computes the physically feasible dispatch range and updates the SoC accordingly.

Site-specific battery hardware specifications have not been received from partners for any of the three active UCs. The simulation therefore uses representative market-average parameters as a baseline configuration, listed in Table 6. These values are defined as configurable constants in the asset configuration layer.

Table 6. Battery storage baseline configuration parameters (hypothetical market-average values)

Component	Parameter	Value
Battery storage	Nominal capacity (E_{bat})	100 kWh
Battery storage	Max charge / discharge rate	50 kW
Battery storage	Operational SoC range	10 % – 90 %
Battery storage	Initial SoC (SoC_0)	50 %
Inverter	Round-trip efficiency (η)	95 %

The simulation logic proceeds in six steps at each discrete time step t of the 48-hour forward horizon.

Step 1 / Net residual power. The net balance between generation and load is computed as:

$$P_{net}(t) = P_{gen}(t) - P_{load}(t)$$

where $P_{gen}(t) = P_{pv}(t) + P_w(t)$ is the total on-site generation. A positive P_{net} indicates a surplus that can charge the battery or be exported; a negative P_{net} indicates a deficit that the battery or grid must cover.

Step 2 / Dynamic SoC power boundaries. Before dispatching, the module calculates the power limits imposed by the current SoC to prevent over-charge or over-discharge:

$$P_{absorb_SoC_limit} = (SoC_{max} - SoC(t)) * E_{bat} / (dt * \eta)$$

$$P_{supply_SoC_limit} = (SoC(t) - SoC_{min}) * E_{bat} * \eta / dt$$

Step 3 / Absolute physical constraints. The hardware inverter and charge controller impose additional upper bounds on charge and discharge rates:

$$P_{max_absorb} = \min(\text{charge_rate}, P_{absorb_SoC_limit})$$

$$P_{max_supply} = \max(-\text{discharge_rate}, -P_{supply_SoC_limit})$$

Step 4 / Actual battery dispatch. The actual power routed through the battery, positive for charging and negative for discharging, is bounded by both the SoC limits and the hardware constraints:

$$P_{actual}(t) = \max(P_{max_supply}, \min(P_{max_absorb}, P_{net}(t)))$$

Step 5 / SoC state update. The SoC transitions to the next step with efficiency applied asymmetrically to charging and discharging:

$$SoC(t+dt) = SoC(t) + (P_{actual}(t) * dt * \eta) / E_{bat} \quad [\text{charging}]$$

$$SoC(t+dt) = SoC(t) + (P_{actual}(t) * dt) / (E_{bat} * \eta) \quad [\text{discharging}]$$

Step 6 / Grid flexibility metrics. The module derives two flexibility indicators that express how much additional charge or discharge headroom remains after the current dispatch:

Upward flexibility = $|P_{\text{max_supply}} - P_{\text{actual}}|$ [remaining discharge capability]

Downward flexibility = $|P_{\text{max_absorb}} - P_{\text{actual}}|$ [remaining charge capability]

The module outputs, for each of the 48 forecast steps: PV generation forecast, wind generation forecast, demand forecast (and their actuals where available for historical replay), SoC trajectory, battery status (charging / discharging / idle), P_{actual} , $P_{\text{max_absorb}}$, $P_{\text{max_supply}}$, grid import, grid export, upward flexibility, and downward flexibility. This structured output is consumed directly by the Self-Consumption Maximisation Engine as the storage input block.

Pipeline execution outputs: The storage block executes iteratively per time step across the 48-hour forward horizon. It produces a structured output matrix with the following column groups for each time step: energy input benchmarks (PV generation forecast, actual PV generation, wind power forecast, actual wind power, demand forecast, actual demand); battery metrics (SoC, battery status, current dispatch); operational boundaries ($P_{\text{max_absorb}}$, $P_{\text{max_supply}}$); grid interchange (grid import, grid export); and flexibility analytics (upward flexibility, downward flexibility). Depending on the execution mode — historical evaluation or live deployment — the ‘actual’ columns are populated either from the site’s historical records or left as NaN pending measurement ingestion.

Flexibility analytics: Two flexibility indicators complement the core dispatch outputs. Upward flexibility (supply capability) represents the remaining headroom to safely discharge additional power into the local load ecosystem at each time step: Upward Flexibility = $|P_{\text{max_supply}} - P_{\text{actual}}|$. Downward flexibility (absorption capability) represents the remaining capacity to absorb additional surplus from generation: Downward Flexibility = $|P_{\text{max_absorb}} - P_{\text{actual}}|$. These indicators are included in the structured pipeline output and are available for downstream consumption by the Self-Consumption Maximisation Engine.

6. HARVREST DASHBOARD

6.1 Energy Performance — Self-Consumption Maximisation

This section documents the methodology of the Self-Consumption Maximisation Engine, the algorithmic component of the HarvRESt cloud platform that consumes the demand and generation forecasts produced by the pre-trained AI models described in chapter 5 and translates them into an actionable energy dispatch strategy for each farm. The strategy is then surfaced to DER owners through the HarvRESt Dashboard (section 6.2).

6.1.1 *Problem framing*

Each HarvRESt farm is modelled as a prosumer microgrid behind a single point of common coupling (PCC) with the distribution network. The microgrid integrates DERs of heterogeneous nature (e.g., PV panels, wind turbines and battery energy storage systems -BESS), alongside the farm's electrical load, which combines climate control, irrigation pumps, post-harvest processing equipment and ancillary auxiliaries. While the wider HarvRESt vision encompasses agro-community-level optimisation, T5.4 deliberately scopes the Self-Consumption Maximisation Engine to the farm boundary: the optimisation is performed per farm, on the centralised HarvRESt cloud, using the data collected for that farm. This choice mirrors the operational reality of the UCs, where DER ownership and metering are individual, and keeps the algorithm tractable and replicable across the diverse topologies of the HarvRESt UCs.

The Self-Consumption Maximisation Engine targets the maximisation of on-site self-consumption of locally generated renewable energy, the most informative single metric introduced by the Renewable Energy Directive (RED II) [1] and consolidated in the prosumer literature by Luthander et al. [2].

6.1.2 *Inputs from the Energy Analytics Portfolio*

The Self-Consumption Maximisation Engine consumes the three forecast analytics and the storage analytics described in chapter 5 plus the static asset configuration registered for the farm. The three forecast time series, each at the same temporal resolution (typically 15 min or 1 h), span the scheduling horizon (typically the next 24 h):

- PV generation forecast $\hat{P}_{pv}(t)$ (see section 5.2).
- Wind generation forecast $\hat{P}_w(t)$ (see section 5.3).
- Energy demand forecast $\hat{P}_{load}(t)$ (see section 5.5).

The static asset configuration adds the storage flexibility parameters delivered by section 5.6 (BESS nominal capacity E_{bat} , charge/discharge rate limits P_{ch_max} and P_{dis_max} , allowed state-of-charge band $[SoC_min, SoC_max]$, round-trip efficiency η), the grid import/export caps imposed by the farm's connection contract, and the time-of-use tariffs $\pi_{imp}(t)$ and $\pi_{exp}(t)$.

6.1.3 *Dispatch strategy: priority-based self-consumption rule*

The Self-Consumption Maximisation Engine implements the so-called Self-Consumption Maximisation rule, the de-facto baseline strategy in the prosumer literature [2][3]. The rule applies a fixed priority order to every time step of the horizon and keeps the model strictly linear and free of any binary decision. Compared with optimisation-based controllers, it has three operational advantages relevant to the HarvRESt platform: (i) no mathematical-programming solver is required, so the Self-Consumption Maximisation Engine can be packaged as a module without any commercial or heavyweight dependency; (ii) the strategy is transparent and

reproducible, which matches the auditability expectations of an EU-funded project; (iii) the per-step computation is trivial, so the Self-Consumption Maximisation Engine can comfortably be re-executed whenever a new forecast becomes available, with no impact on platform latency.

The priority order is fixed and reads:

- First, cover the farm load with the local renewable generation (PV + wind).
- If a surplus remains, charge the battery up to its maximum charge power and to the upper SoC bound.
- If a surplus still remains, export the remainder to the grid (within the contracted export cap).
- If a deficit appears, discharge the battery down to its maximum discharge power and to the lower SoC bound.
- If a deficit still remains, import the missing energy from the grid (within the contracted import cap).

Formally, with the net residual power $r(t) = \hat{P}_{pv}(t) + \hat{P}_w(t) - \hat{P}_{load}(t)$, the schedule is computed at every time step as follows:

```
# Inputs at step t: r(t), SoC(t), asset parameters
if r(t) >= 0:                                # surplus generation
    P_ch = min(r(t), P_ch_max,
               (SoC_max - SoC(t)) * E_bat / (eta_ch * dt))
    P_dis = 0
    P_exp = min(r(t) - P_ch, P_exp_max)
    P_imp = 0
else:                                         # deficit
    deficit = -r(t)
    P_dis = min(deficit, P_dis_max,
                 (SoC(t) - SoC_min) * E_bat * eta_dis / dt)
    P_ch = 0
    P_imp = min(deficit - P_dis, P_imp_max)
    P_exp = 0
SoC(t+1) = SoC(t) + (eta_ch * P_ch - P_dis / eta_dis) * dt / E_bat
```

Two refinements are applied on top of the base rule when the use case justifies them. First, when the time-of-use tariff exposes a clear cheap/expensive band, a simple price-aware override is enabled: the battery is allowed to charge from the grid during the cheap band only if its state of charge is below a configurable threshold, and to discharge preferentially during the expensive band. Both refinements remain rule-based, preserve the linearity of the model and add only a handful of lines of code.

6.1.4 *Output and Key Performance Indicators*

For every recomputation, the Self-Consumption Maximisation Engine returns a structured object that the HarvRESt Dashboard consumes directly. The payload comprises two blocks. The first is the dispatch schedule itself: a time-indexed array of P_{ch} , P_{dis} , P_{exp} , P_{imp} and SoC for each of the 48 forecast steps. The second is the KPI block: Self-Consumption Ratio (SCR), Self-Sufficiency Ratio (SSR), Savings (kWh and EUR) and Avoided CO₂ values computed over the scheduling horizon. Because the HarvRESt platform processes static datasets, the 'current date' used in all KPI computations is a pre-configured reference date rather than the wall-clock

time; KPI values shown under the 'Today' timeframe reflect the full-day predictions generated at 00:00 of that reference date.

$$SCR = \sum_t (P_{gen}(t) - P_{exp}(t)) \cdot \Delta t / \sum_t P_{gen}(t) \cdot \Delta t$$

$$SSR = \sum_t (P_{load}(t) - P_{imp}(t)) \cdot \Delta t / \sum_t P_{load}(t) \cdot \Delta t$$

$$Savings_kWh = \sum_t (P_{imp_base}(t) - P_{imp}(t)) \cdot dt \quad | \quad Savings_EUR = Savings_kWh \cdot p_static$$

with $P_{gen}(t) = P_{pv}(t) + P_w(t)$. SCR measures the share of local generation consumed on site; SSR measures the share of farm demand covered by local generation and battery storage. Savings_kWh counts every kWh of grid import avoided relative to a no-BESS baseline; Savings_EUR multiplies this by a static representative retail electricity price p_static per country (Italy EUR 0.32/kWh, Spain EUR 0.21/kWh, Norway EUR 0.16/kWh [4]). A fifth indicator, Avoided CO₂, is defined as:

$$Avoided\ CO_2 = \sum_t (P_{imp_base}(t) - P_{imp}(t)) \cdot dt \cdot EF_{grid} \quad [kg\ CO_2eq]$$

where EF_{grid} is the national grid emission factor (kg CO₂eq/kWh) derived from the country's annual electricity generation mix. This KPI quantifies the greenhouse gas emission reduction attributable to on-site renewable generation and battery storage operation relative to a grid-only baseline.

6.1.5 *Implementation footprint*

The Self-Consumption Maximisation Engine is implemented as a module that is registered with the Data Analytics Pipeline Execution Service described in section 4.4.1 and invoked once per scheduling cycle (e.g., once a day after the AI forecasts of chapter 5 are refreshed). Its runtime on the centralised HarvRESt cloud is dominated by the interaction with the platform's data storage rather than by the computation itself. The Self-Consumption Maximisation Engine operates in batch / replay mode (HCF_020): each invocation selects a historical window of input data from the platform's data storage, executes the priority-based dispatch over that window, and persists the resulting schedule and KPIs for downstream consumption by the HarvRESt Dashboard. Scheduled or streaming activation, once streaming feeds become available, is a configuration change and does not require code changes.

6.1.6 *Adaptation across the HarvRESt Use Cases*

The Self-Consumption Maximisation Engine is validated on the HarvRESt demo-site datasets. These cover three electrical-dispatch use cases — VdV (Spanish UC, PV + BESS), FSDC (Italian UC, multi-PV-installation + BESS) and GGE (Norwegian UC, PV + BESS with planned wind turbines, currently represented by a modelled wind series). One biogas-plant operational dataset (ACSA) is documented as a separate research output of WP5 (see section 5.4) and is not surfaced in the dashboard, since the biomethane produced at ACSA is injected into the gas grid rather than into a farm electrical bus.

The electrical UCs exercise materially different asset mixes against the same engine and the same output schema: VdV stresses the irrigation-driven seasonality of the load and the densest data feed (15-min, eight months on a single time grid); FSDC exercises the multi-PV-installation breakdown in the dashboard chart; GGE exercises the wind branch (currently modelled), the cold-climate battery efficiency curve and the heterogeneous-granularity case. Switching between use cases requires only re-parameterising the asset bounds, the tariffs and the optional price-aware override; the rule itself, the KPI definitions and the output schema remain identical, which is what allows the platform to claim EU-wide replicability. Detailed per-UC validation results are reported in section 6.3.

6.2 HarvRESt Dashboard

The HarvRESt Dashboard is a web application that translates the platform's analytics outputs into an at-a-glance energy management interface for farm operators, energy managers and DER owners. User accounts, per active UC (Italy, Norway, Spain VdV), render the farm-specific view upon login via basic email-and-password authentication. Since the platform processes static datasets, the 'current date' shown in the dashboard is a pre-configured reference date rather than the wall-clock time, ensuring full reproducibility of analytics results across all three UCs.




Figure 2. HarvRESt Dashboard

6.2.1 Energy Performance Overview

The farm view centres on an Energy Performance Overview composed of five KPI summary cards, a generation-demand-storage timeline chart, a battery storage dispatch panel, and a recommended Next Action banner. A timeframe selector (Today / Next 24h / Next 48h / Last 7 days) at the top of the page updates the KPI cards and the timeline chart.




Figure 3. HarvRESt Dashboard, VdV UC: Energy Performance Overview

The five KPI summary cards at the top of the page present the key energy performance indicators. All values for the 'Today' timeframe are derived from the full-day forecast generated at 00:00 of the reference date, ensuring consistency between the historical and forecasted portions of the day:

Self-Consumption Ratio (SCR): The share of total on-site renewable generation consumed locally rather than exported to the grid. A higher SCR indicates that the farm is making effective use of its generation assets.

Self-Sufficiency Ratio (SSR): The share of total farm electricity demand covered by local renewable generation and battery storage. A higher SSR indicates greater energy autonomy and reduced reliance on the utility grid.

Energy Savings (kWh): The cumulative energy saved over the last 7 days by substituting grid imports with local generation and battery dispatch, relative to a no-BESS baseline in which all unmet demand is covered by grid import.

Estimated Savings (EUR): The monetary equivalent of the 7-day energy savings, computed as total kWh savings x the static representative retail electricity price for the farm's country (Italy EUR 0.32/kWh, Spain EUR 0.21/kWh, Norway EUR 0.16/kWh [4]).

Avoided CO₂ (kg CO₂eq): The cumulative greenhouse gas emissions avoided over the last 7 days by substituting grid imports with on-site renewable generation and battery storage, relative to the no-BESS baseline. Emissions are calculated by applying the national grid emission factor (EF_{grid}, kg CO₂eq/kWh) to the energy savings quantity: $\text{Avoided CO}_2 = \text{Savings_kWh} \times \text{EF_grid}$. This KPI contextualises the farm's energy performance in terms of its environmental impact, enabling operators to communicate the carbon-reduction value of their investment in renewable assets to stakeholders and policy audiences. EF_{grid} values are drawn from national grid carbon intensity databases and are treated as static parameters for the current deployment.

The central chart panel displays an hourly timeline of PV generation forecast, energy demand forecast, battery charge and discharge schedule, and the resulting net grid import/export balance. The timeframe selector at the top of the page switches the displayed window between Today, Next 24h, Next 48h and Last 7 days. For the 'Today' view, all values reflect the full-day forecast generated at 00:00 of the reference date.

The battery storage panel complements the main chart by showing the current day's dispatch schedule at hourly resolution. Each cell is colour-coded by dispatch priority (discharge > charge > import > export). Hours

that have elapsed show historical actuals; upcoming hours show the forecasted schedule. The panel always covers the full 24-hour window of the reference date, independently of the selected timeframe.

6.2.2 Next Action — Operational Guidance Messages




Figure 4. HarvRESt Dashboard, GGE UC: Next Action guidance banner and battery dispatch panel

A distinctive feature of the HarvRESt Dashboard is the Next Action banner, which surfaces operational guidance in plain language to support farm operators who do not have a background in energy systems optimisation. The banner allows operators to act on the platform's recommendations immediately, e.g., by shifting heavy loads to coincide with generation surpluses or peak export windows.

The message is determined by the dominant dispatch mode in the nearest upcoming forecast hour, following the priority order: battery discharging > battery charging > grid import > grid export. Representative messages for each category are listed below; the dashboard selects the most contextually appropriate message based on the current dispatch state and SoC trajectory:

Grid Import: 'Farm is importing grid power — consider pausing non-essential heavy loads (e.g., milling, drying) to minimise utility costs.' | 'The system is drawing grid power during a peak pricing window; check equipment schedules.'

Grid Export: 'Storage is full and farm load is fully met — excess green energy is being returned to the grid.' | 'System is exporting surplus solar; this is the ideal time to run heavy irrigation or pump water to storage tanks.'

Battery Charging: 'Surplus solar generation is redirecting to charge the farm battery bank.' | 'Ideal conditions — the energy storage system is charging at maximum speed.' | 'Storage is nearing full capacity; the farm is well-positioned for evening ventilation and automated feeding cycles.'

Battery Discharging: 'The sun has set — the farm is now running entirely on stored energy.' | 'Battery discharging has activated to protect the farm from high on-peak grid tariffs.' | 'High operational loads are depleting the battery quickly; estimated full-discharge runtime is dropping faster than forecast.'

6.2.3 *Strategy Panel — Energy Savings Summary*

The Savings bar chart shows the daily energy savings in kWh for each of the seven days preceding the reference date. A summary figure at the top of the panel converts the cumulative 7-day total into an estimated monetary value, expressed in euros, by multiplying by the static country price (Italy EUR 0.32/kWh, Spain EUR 0.21/kWh, Norway EUR 0.16/kWh [4]). The chart is independent of the timeframe selector and always reflects the 7-day window ending on the reference date, giving operators a stable weekly savings narrative regardless of which forecast horizon they are viewing.

6.3 Validation on the HarvRESt Use Cases

This section reports on how the Self-Consumption Maximisation Engine and the surrounding analytics were exercised on each of the HarvRESt UCs. The following paragraphs describe the dataset characteristics, the asset mix validated, and the key energy behaviours observable in the HarvRESt dashboard for each of the three active electrical-dispatch use cases — VdV, FSDC, and GGE. Specific KPI figures are not cited as these are computed dynamically from the site's data and depend on the reference date and the forecast horizon filter (24-hour or 48-hour view) selected by the user on the dashboard.

VdV (Spain) — anchor demo. The VdV dataset is the densest and most internally consistent of the five: 242 daily files spanning 13 January 2024 to 10 September 2024 (approximately eight months) at 15-minute resolution, covering PV generation and demand at the winery. After resampling to hourly granularity, the dataset feeds a LightGBM PV forecast model (EV 0.92, nMAE 0.16) and a LightGBM demand forecast model (EV 0.37, nMAE 0.33). The BESS is simulated with the platform's representative market-average parameters (100 kWh, 50 kW charge/discharge rate, $\eta = 95\%$, SoC window 10–90%). The dashboard for VdV (**Figure 3**) shows a clear daily charge cycle driven by the winery's high midday PV surplus, with the irrigation-linked demand spikes visible in the late-afternoon discharge window. The seven-day savings panel reflects the seasonal consistency of the Mediterranean irradiance profile over the dataset window.

FSDC (Italy) — multi-PV secondary demo. Two PV installations contribute real generation data: a 5.5 kWp bifacial greenhouse array (15-minute samples, February 2024 – February 2025, with a Sep–Dec 2024 gap) and an 80 kWp rooftop installation (hourly samples, January – April 2025). After alignment and resampling to a 74-day joint window (January – April 2025), the merged generation signal feeds the LightGBM PV forecast (EV 0.82, nMAE 0.32). Demand is modelled with LightGBM (EV 0.894, nMAE 0.189) on the 458-day consumption history of the Circeo cooperative greenhouse facility, making the FSDC demand forecast the most accurate in the portfolio. The battery is simulated with the same representative parameters as VdV. On the dashboard, the FSDC view (**Figure 5**) demonstrates the Self-Consumption Maximisation Engine operating under a shorter and more variable generation signal; the demand forecast captures the repetitive facility load pattern well, and the Next Action messages reflect the cooperative's predominantly import-reduction use case.




Figure 5. HarvRESt Dashboard, FSDC UC

GGE (Norway) — multi-asset and mixed-data demo. The dataset is the most heterogeneous of the five: real PV production at Buevegen (5–10 minute resolution, February 2023 – January 2024), a daily battery log, and modelled wind generation data over the same 212-day analysis window (February – September 2023) after resampling to hourly. Three AI models run in parallel: a LightGBM PV forecast (EV 0.806), a LightGBM wind generation forecast (EV 0.792, nMAE 0.230), and a LightGBM demand forecast (EV 0.384, nMAE 0.155). The Self-Consumption Maximisation dispatch formula $r(t) = \hat{P}_{pv}(t) + \hat{P}_w(t) - \hat{P}_{load}(t)$ incorporates both generation sources, making GGE the only use case where the wind power forecast directly influences the battery dispatch schedule. On the dashboard, the GGE view (**Figure 6**) exercises the full multi-asset interface and demonstrates the Next Action messages in a mixed wind-and-solar operational context.




Figure 6. HarvREST Dashboard, GGE UC

Across the three UCs, the HarvREST platform successfully ingested, pre-processed, and modelled site-specific energy data, and the dashboard rendered complete energy performance views for each active site. The Self-Consumption Maximisation Engine is demonstrated across materially different asset configurations — PV-only with BESS, dual-PV with BESS, and multi-vector PV-wind with BESS — confirming the engine's adaptability. Full operational KPI values for each site are accessible directly on the dashboard and are updated whenever the underlying pipeline is re-executed against new data.

7. MANAGING INTEGRATION OF MULTI-VECTOR RES ON FARMS: COMPLEMENTARY MODELLING COMPONENTS

The transition towards sustainable energy in agriculture can be significantly enhanced through the widespread adoption of multi-vector energy systems. Modelling and simulation are essential tools for the design, management, and optimisation of renewable energy systems, as they enable the technical assessment of RES technologies under different farm conditions and operational assumptions.

Within HarvRESt, this work has been developed as an exploratory technical activity focused on complementary multi-vector modelling components. Its purpose is to extend the range of renewable energy technologies that can be represented and simulated at farm level, enriching the exploratory work already performed in T5.2 and T5.3 and providing a potential basis for future integration with the AVPP developed under WP6. This component should not be interpreted as a standalone AVPP or DSS. While the AVPP focuses on the generation, sizing and multi-criteria assessment of alternative RES scenarios, this activity addresses a different level of the decision process: the operational assessment of predefined multi-vector energy configurations. In this sense, it does not generate or rank alternative RES sizing scenarios, but evaluates how a predefined configuration can be operated under specific demand, generation and constraint assumptions.

The technologies considered include solar PV systems, vertical-axis wind turbines, hydropower, biogas combined heat and power (CHP) systems, wood-chip boilers, and energy storage systems. To make these models easier to test and visualise, they have been integrated into a web-based demonstration environment.

The user interface of the prototype is shown in Figure 7.




Figure 7. The UI, where a combination of RES technologies can be selected

Each technology is represented by a models developed in Python that estimates energy production based on system characteristics and environmental conditions. These models allow users to explore the potential behaviour of different RES technologies and combinations under predefined assumptions. The main contribution of this activity is therefore methodological and technical: it expands the portfolio of RES models available within HarvRESt and provides a basis for future pre-assessment and operational simulation functionalities.

To explore the coordinated behaviour of multiple technologies, a prototype economic optimisation routine inspired by Model Predictive Control (MPC) principles has been implemented. In general terms, MPC uses a dynamic system model over a predefined prediction horizon to anticipate future system behaviour and identify suitable control actions while satisfying operational constraints. In the current version, this routine is applied in simulation mode to coordinate simulated technology outputs under predefined input data and constraints. It is not connected to farm assets and does not replace the Self-Consumption Maximisation Engine implemented in the HarvRESt Dashboard.

The predicted energy outputs from the technology-specific models can be compared against the farm's energy demand to provide a first consistency check of the modelling approach and to assess the potential suitability of different multi-vector configurations for future real-world applications. The results of this modelling process are shown in Figure 8.




Figure 8. The UI with the results of modelling

The web-based environment includes the ability to insert Excel files in xlsx or csv format containing data such as farm energy demand, meteorological station data, data obtained from the farm's wind and solar atlas, and other data relevant to the system's operation. It provides an intuitive environment for configuring renewable energy technologies, assessing their indicative technical and economic performance, and visualising energy flows within the farm energy system.

Overall, this prototype provides complementary multi-vector modelling capabilities that can support future AVPP developments by offering additional operational indicators for predefined candidate configurations. By combining technology-specific Python models with an MPC-inspired optimisation routine, it enables the simulation of coordinated multi-vector operation under defined demand, generation and constraint assumptions, with the aim of maximising system performance and economic benefits. Some additional integration capabilities are still under development.

8. CONCLUSIONS

The deliverable at hand documents the development and initial validation of the HarvRESt cloud platform and its associated AI Analytics Portfolio, culminating in the operational HarvRESt Dashboard — a web-based energy management interface for farm operators across three HarvRESt demo sites.

The platform's four-layer cloud architecture (Data Collection, Data Pre-processing and Quality, AI Analytics, and Dashboard) has been designed for extensibility. New demo sites and data sources can be onboarded by registering additional pipeline configurations without changes to the core analytics engine. Although the current deployment operates on static datasets provided by partners, the architecture is ready for real-time API and MQTT streaming.

The Energy Analytics Portfolio has been trained and evaluated across three national contexts (Italy, Spain, and Norway) covering PV generation forecasting, wind generation forecasting, energy demand forecasting, and battery storage flexibility simulation. PV forecasting consistently achieves high accuracy across all three sites (explained variance 0.82–0.92, with the Spanish VdV site performing best). Wind generation forecasting for the Norwegian GGE site achieves an explained variance of 0.79, consistent with the known intermittency characteristics of small-scale wind at the Buevegen coastal location. Demand forecasting is strongest at the Italian FSDC site (EV 0.89), where the greenhouse facility load profile is relatively stable, and weakest for the Spanish VdV site (EV 0.37), where irrigation scheduling introduces pronounced demand variability.

The Self-Consumption Maximisation Engine applies a priority-based dispatch rule to the 48-hour generation and demand forecasts, producing charge/discharge schedules for the on-site BESS and five KPIs: Self-Consumption Ratio, Self-Sufficiency Ratio, Energy Savings (kWh), Estimated Savings (EUR), and Avoided CO₂ (kg CO₂eq). The HarvRESt Dashboard makes these outputs accessible to non-technical farm operators through a generation-demand timeline, a battery dispatch panel, plain-language Next Action guidance messages, and a seven-day savings summary. Role-based access control with three pre-seeded user accounts (one per active farm) ensures that operators see only their own site data.

Chapter 7 presents a multi-vector energy systems advisor, which enables farmers to simulate and evaluate combinations of renewable energy technologies prior to deployment. This tool uses an Economic Model Predictive Control framework to optimise multi-technology portfolios. Together with the HarvRESt Dashboard, it completes a farmer-facing toolchain that spans from pre-investment RES planning through to day-to-day operational energy management.

9. REFERENCES

- [1] European Parliament and Council. Directive (EU) 2018/2001 of 11 December 2018 on the promotion of the use of energy from renewable sources (recast – RED II). Official Journal of the European Union, L 328/82.
- [2] Luthander, R., Widén, J., Nilsson, D., & Palm, J. (2015). Photovoltaic (PV) self-consumption in buildings: A review. *Applied Energy*, 142, 80–94. <https://doi.org/10.1016/j.apenergy.2014.12.028>
- [3] Angenendt, G., Zurmühlen, S., Axelsen, H., & Sauer, D. U. (2018). Comparison of different operation strategies for PV battery home storage systems including forecast-based operation. *Applied Energy*, 229, 884–899. <https://doi.org/10.1016/j.apenergy.2018.08.058>
- [4] Eurostat. (2025). Electricity prices for household consumers — bi-annual data (from 2007 onwards) [Dataset nrg_pc_204]. European Commission Statistics. Available at: https://ec.europa.eu/eurostat/databrowser/view/nrg_pc_204. Representative retail prices used in this deliverable for the first half of 2025: Italy EUR 0.32/kWh, Spain EUR 0.21/kWh, Norway EUR 0.16/kWh (Eurostat band DC, medium-sized household consumers, all taxes included).




HarvRESt
Greener Farming with RES

The project

The HarvRESt project aims to enhance the sustainable production of renewable energy at farm-level. This approach not only makes farms climate-neutral but also optimizes production, reduces their impact on natural resources and biodiversity, and provides energy services to communities, thereby diversifying economic income. However, deciding how best to integrate renewable energy sources (RES) on a farm is not without its challenges. The decision is a complex one, with many factors to consider. Due to this, HarvRESt seeks to identify, understand, and overcome the existing barriers hindering the widespread adoption of this innovative approach. Current initiatives often overlook the complex interactions and factors within the farming and RES context, resulting in ineffective support for decision-making based on accurate projections, estimations, and forecasts. HarvRESt will therefore consolidate and enhance existing knowledge, creating an Agricultural Virtual Power Plant capable of running diverse scenarios and farm configurations. This tool will determine the best operational procedures for a given RES solution, providing valuable data to a decision support system. This system will weigh trade-offs and key indicators, offering tailor-made recommendations to farmers and policymakers.

PARTNER		SHORT NAME
	CIRCE Research Centre	CIRCE
	BETA Technological Centre	UVic-UCC
	NORCE	NORCE
	Tecnoliment	TCA
	WHITE	WR
	Suite5 Data Intelligence Solutions Ltd.	Suite5
	EnGreen	EnG
	ConTerra	CT
	Confagricultura	CONFAGRI

	Fattoria Solidale del Circeo	FSDC
	Viñas del Vero	VdV
	Viñedos del Rio Tajo	VRT
	Sorigué	ACSA
	Grønn Gårdsenergi AS	GGE
	Food & Bio Cluster Denmark	FBCD
	Climate KIC	CKIC
	HEC Paris	HEC

Contact us

www.harvrest.eu

 <https://linkedin.com/harvRESt>

 https://twitter.com/HarvRESt_eu