DELIVERABLE

6.2 Report on implementation and status on trial sites

Version 1.2

Jan Weiler, SWW

Elena Malakhatka, Chalmers

David Steen, Chalmers

Manuel Pitz, RWTH

Rezeq Direya, RWTH

Adil Mukhtar, TU Wien

Lohiya Sathishkumar, FH JOANNEUM

Christof Sumereder, FH JOANNEUM

Radu Plamanescu, UPB

Matthias Primas, FH CAMPUS 02

Theresa Kohl, DiLT Analytics

INTERNAL REFERENCE

Deliverable No.:	D 6.2 (2025)
Deliverable Name:	**Report on implementation and status on trial sites**
Lead Participant:	SWW Wunsiedel
Work Package No.:	6
Task No. & Name:	T 6.2
Document (File):	
Issue (Save) Date:	19.12.2025

Document sensitivity

☒	**Not Sensitive**	Contains only factual or background information; contains no new or additional analysis, recommendations or policy-relevant statements
☐	**Moderately Sensit ive**	Contains some analysis or interpretation of results; contains no recommendations or policy-relevant statements
☐	**Sensitive**	Contains analysis or interpretation of results with policy-relevance and/or recommendations or policy-relevant statements
☐	**Highly Sensitive Confidential**	Contains significant analysis or interpretation of results with major policy-relevance or implications, contains extensive recommendations or policy-relevant statements, and/or contain policy-prescriptive statements. This sensitivity requires SB decision.

1. Introduction

1.1 Purpose and Scope of this deliverable

This deliverable documents the implementation progress and current status of all ECom4Future trial sites within Work Package 6 (Demonstration). It serves two main purposes: (i) to report what has been installed, commissioned and is operational at each site, and (ii) to compare the realised status against the plan defined in Deliverable 6.1 (baseline) across the WP6 activity chain—setup, monitoring, market integration and trial operations. The report consolidates partner inputs for each site (in particular the smart buildings and smart energy laboratories, WUNergy Energy Sharing Community, Chalmers Campus Energy Community, Jättesten, FDD/optimisation pilots, Energy Community Niedergrail, Energy Community Mariatrost, FH JOANNEUM & FH CAMPUS 02, and GreenMogo) and provides a portfolio view of schedule adherence, open issues/risks and mitigation actions.

1.2 Document structure

Section 2 captures the D6.1 baseline for all trial sites, summarising planned scope and objectives, planned activities (setup, monitoring, market integration, trial operation), and the planned technical assets and interfaces (DERs, metering, ICT platform connectors). Section 3 provides the implementation status per trial site, using a consistent template: actual implementation to date, measured variables with time resolution and coverage, and the site's open issues, risks and mitigation actions.

2. Baseline from D6.1

2.1 Planned scope and objectives per trial site

Deliverable D6.1 defines a coherent portfolio of complementary demonstrators that together cover residential, non-residential, laboratory and community contexts. The sites comprise: (i) smart buildings and smart energy laboratories at RWTH (heat-pump HiL benches and an AHU for real-time FDD), (ii) the WUNergy Energy Sharing Community (ESC) led by SWW, (iii) the Chalmers Campus Energy Community, (iv) the Jättesten community, (v) FDD and optimisation pilots by TU Wien/DiLT, (vi) the Energy Community Niedergrail, (vii) the Energy Community Mariatrost, (viii) FH JOANNEUM & FH CAMPUS 02 demonstrators, and (ix) GreenMogo as a prosumer site. Across these sites, D6.1 set four overarching objectives: establish compliant and operational community setups; instrument sites for high-quality, analysis-ready data; validate forecasting and physics-informed fault detection; and prepare grid-aware flexibility and local market use cases that can be demonstrated under real conditions. This D6.2 chapter restates those baselines to anchor the plan-versus-actual reporting that follows in Section 3.

2.2 Planned activities: setup, monitoring, market integration, trial operation

WP6 execution follows the four-step sequence set in D6.1 and reiterated in this report: T6.1 trial-site planning, setup and operational readiness; T6.2 monitoring of consumption/production with data-quality assurance; T6.3 market-schema integration and user-feedback analysis; and T6.4 trial operations and measurement campaigns. The plan foresees overlapping phases to allow iterative commissioning, data qualification and staged feature roll-out across 2024–2025. The Gantt view included in this document (page 26) reflects that phasing and shows the alignment to programme milestones. In D6.2, Section 3 reports each site's actual status against these planned phases to enable a clear plan-versus-actual comparison.

2.3 Planned technical assets and interfaces per site (DERs, metering, ICT platform connectors)

Assets and flexibility. The baseline technology stack across the portfolio comprises PV generation, battery energy storage (home- and lab-scale), heat pumps, EV charging, and HVAC/AHU assets. These provide the controllable and observable elements needed for the flexibility, market, and FDD use cases foreseen in D6.1 (e.g., RWTH's HP HiL for flexibility participation and voltage support; FHJ/CAMPUS 02 surplus-PV hot-water pilots; Niedergrail/Mariatrost household-level optimisation and forecasting).

Metering and data resolution. D6.1 specifies a two-tier metering approach. For community analytics and prospective energy-sharing settlement, data at 15-minute resolution are the baseline (e.g., WUNergy's smart-meter gateways for a core billing cohort plus research sub-metering across a larger cohort). For control-grade analytics, FDD and transient studies, selected sites operate high-resolution acquisition up to 1 s, notably via an "Unbundled Smart Meter" concept with Raspberry Pi gateways and direct inverter links; measurements are buffered locally (e.g., CSV) and forwarded to the platform. These provisions were designed to ensure both comparability to legacy smart-meter data and the fidelity required for physics-informed diagnostics and real-time market experiments.

ICT platform and interfaces. The common integration pattern is FIWARE-compatible and NGSI-based: a context-broker-centred architecture with standardised entities/attributes, long-term storage backends (e.g., MongoDB/InfluxDB/CrateDB, depending on site constraints), and visualisation/third-party integration layers (e.g., Chronograf, Grafana, or custom apps). RWTH operates a FIWARE instance in parallel with the building monitoring to host and compare FDD approaches; WUNergy exposes an authenticated API that separates configuration metadata (customers, assets, MPANs) from time-series measurements and delivers both live (quasi real-time) and historical data, enabling model back-testing and simulation-in-the-loop until full automation is in place; GreenMogo implements an updated NGSI data exchange and demonstrates high-rate analytics and third-party app interoperability. These interfaces were established by the end of 2025 and will produce data and findings through 2026.

3. Implementation Status per Trial Site

3.1 Smart buildings and smart energy laboratories (RWTH)

This pilot aims to extend the I-Greta setup to reach a higher TRL by incorporating flexibility markets and increasing the use of hardware-in-the-loop test benches. Furthermore, the fault detection algorithms developed in WP 3 will be demonstrated in a real building.

In contrast to the earlier report, the trial site setup has been refined to better align with the project's specific research goals. While the non-residential building remains a core component of the trials for Demo 4 (Fault Detection), its participation in the flexibility market (Demo 2) has been removed. As the project focus on Energy Communities evolved, the residential sector emerged as the primary driver for local flexibility markets. Consequently, we have concentrated all market-related demonstrations on the residential site (Demo 1). This allows us to avoid the significant overhead of deploying MPC-based market controls on the non-residential system. These resources are in turn spent deepening the analysis of residential flexibility and fault detection. Therefore, the planned trials are:

- Demo 1: Heat pump (HP) HiL thermal test bench (residential building): Participation in flexibility market
- Demo 2: HP HiL electric test bench(DC): voltage control in grids with a real HP and battery
- Demo 3: Air handling unit in a non-residential building: real-time fault detection

3.1.1 Actual implementation to date

Demo 1:

The HiL testbench is currently undergoing validation for the flexibility quantification methodology required for local market participation. Key activities include testing the cloud-based communication interface between the quantification algorithm and the HiL setup. We are also characterizing the custom-built heat pump's ramp-up, cycling behaviour, and start-up control sequences. Various control interfaces are being evaluated to optimize integration. The current measurement and control infrastructure is sufficient for the flexibility use case, though the modular design allows for hardware expansion if necessary. Modelling of the heat pump and HiL components is currently in the design and testing phase, with full testbench operations scheduled for the coming year.

Demo 2:

Waiting for additional equipment to be purchased.

Demo 3:

The preparation of the test hall is ready for first trials. This involved calibrating and replacing temperature sensors and identifying pre-existing system faults (e.g., stuck valves, fouled filters) to establish a baseline. A FIWARE platform has been successfully deployed to operate in parallel with the existing monitoring system, enabling fault detection algorithms to run across both environments. Currently, we are validating various fault detection approaches using system simulation models and historical data. Live operational tests on the air handling unit, utilizing real-time data streams, are scheduled for the upcoming year. For the final operational test, some equipment might be added.

3.1.2 Measured variables, time resolution and coverage

Demo 1:

Thermal demand is simulated using detailed Modelica models, allowing for data capture at flexible time resolutions. This simulated data, along with physical measurements from the HiL setup, is synchronized and stored via a cloud platform. Physical data points include environmental conditions from the climate chamber (ambient air, humidity, ...) and internal heat pump states (superheat, refrigerant cycle pressures, ...). The standard data acquisition resolution is set to ten seconds, though this can be adjusted to capture faster transient events if required.

Demo 2:

Waiting for additional equipment to be purchased.

Demo 3:

The facility's energy system is monitored via a custom data logger that communicates directly to the monitoring platform. Coverage includes all relevant water and air system volume flows, temperatures, power consumption, humidity levels, and user-defined setpoints. As the automation system was developed in-house, the measurement framework is highly extensible. The current setup supports a time resolution of ten seconds for all variables.

3.1.3 Open issues, risks, and mitigation actions

Demo 1:

The primary technical risk involves the fidelity of the heat pump models; inaccuracies here could compromise the flexibility quantification results. To mitigate this, we are adopting an iterative validation approach using testbench data, scheduled for immediate implementation. A secondary risk is physical damage to the heat pump during flexibility activation, particularly if the system is pushed to extreme operating points. As no backup unit is available, extensive safety controls and limiters are being implemented and rigorously tested to prevent operation outside safe envelopes.

Demo 2:

Waiting for additional equipment to be purchased.

Demo 3:

The primary challenge involves managing user interaction within an occupied facility. Induced faults for testing purposes could lead to thermal discomfort. Mitigation strategies include continuous user communication regarding experiment schedules and validating all scenarios in simulation first to predict and minimize impact. There is also a risk of component failure when intentionally inducing faults. This is mitigated through pre-simulation impact analysis and the implementation of strict safety overrides in the control logic.

3.2 WUNergy Energy Sharing Community (SWW)

3.2.1 Actual implementation to date

In line with the D6.1 plan to establish a fully compliant Energy Sharing Community, the WUNergy ESC has been formally founded as a cooperative in mid-2025, with all legal requirements met under the applicable German cooperative and energy market regulations. The cooperative structure provides a suitable governance framework for member participation, decision-making and benefit sharing, and is compatible with the regulatory developments around citizen energy communities and energy sharing in Germany.

On the technical side, the build-up phase has focused on creating a core group of households with sufficiently detailed metering to support both operational monitoring and research activities:

- 14 households have been equipped with ImSys devices (Smart Meter Gateway plus Smart Meter). These installations enable legally compliant metering and secure data communication, forming the core of the energy sharing group in terms of billing-relevant data and potential future settlement.
- An additional 62 households have been equipped with “Shelly” devices (sub-metering / smart plugs) providing 15-minute resolution consumption data. These are used primarily for research and modelling purposes, allowing the project to analyse load profiles, intra-day

variability and flexibility potential, even where full gateway-based metering is not yet in place.

This two-layer approach – a smaller core of fully instrumented members and a larger set of research participants – matches the D6.1 intention to balance regulatory compliance, technical realism and practical deployment effort. It ensures that at least part of the community can, in principle, move towards real contractual energy sharing, while still giving the project broad visibility into residential demand patterns.

ESC matchmaking and optimisation model

A central element of the WUNergy trial site in both D6.1 and D6.2 is the development of an Energy Community matchmaking and optimisation model. The model is intended to support both:

- Operational decision support for energy communities (e.g. how to allocate local PV generation among members, when to charge/discharge batteries), and
- Strategic planning and scenario analysis (e.g. which households to include, how storage or tariff design affect self-consumption and costs).

The main problem the model addresses is the optimal matching of local generation, consumption and storage inside an energy community. More specifically, it aims to:

- Increase the self-sufficiency and self-consumption of locally generated energy,
- Reduce energy procurement costs for the community members compared to a standard supply model, and
- Support an operational strategy for storage systems (home batteries, community batteries) and flexible loads that is both economic and grid-friendly.

In the real-life application envisioned in D6.1, the model is to be provided as a web-based matching platform. Community operators, consultants and potentially even end users could use a browser-based interface to:

- Explore different community configurations and tariff structures,
- Compare the performance of alternative asset portfolios (e.g. with/without storage, different PV penetration levels), and
- Evaluate the economic impact of joining, leaving or restructuring communities ("interested parties can switch between or join communities").

At the time of D6.2, the model is operational in simulation mode and used for exploratory analysis and design support, with the following characteristics:

Inputs

- Metadata for all energy sources (e.g. PV systems, generation capacities, technical constraints, locations).
- Metadata for all energy consumers (e.g. household type, connection capacity, presence of flexible loads, potential for storage).
- Historical measurement data (load and generation time series from WUNergy and synthetic or external data where necessary).

- Live measurement data for consumption and feed-in, retrieved through the newly created API (see below), to enable near real-time analyses and back-testing.

Outputs

- Self-sufficiency level of the community for a given configuration or scenario.
- Resulting energy cost at community and member level, including network tariffs and (where applicable) local compensation schemes.
- Savings compared to a standard model (e.g. traditional individual supply without community sharing).
- Scenario analyses, such as changes in membership, load growth, or additional PV and storage investments.
- Optimal use of storage technologies, expressed as charging/discharging schedules or utilisation metrics.

Modelling approach and algorithms

The current prototype implementation uses a combination of:

- Simple statistical algorithms to analyse load and generation data (e.g. typical daily profiles, coincidence factors, diversity of consumption).
- A linear programming (LP) formulation that is under active development and validation. The LP approach allows the definition of an objective function (e.g. cost minimisation, maximisation of local self-consumption, or a weighted combination) subject to constraints such as energy balances, storage capacities, and technical limits on import/export.

In line with the D6.1 baseline, the model also includes a flexibility and grid impact perspective. It can incorporate flexible resources such as batteries and controllable loads, and is designed to be extended with a grid impact sub-model, which would estimate, for instance, the effect of community operation on local transformer loading or voltage profiles. This will be especially relevant when aligning the model with distribution network constraints and flexibility market mechanisms.

Market aspects and regulatory environments

The D6.1 baseline emphasised the need to make the WUNergy model applicable across different regulatory environments. To this end:

- The model already considers use cases in the Czech market (CZ) through advisory inputs from SWWs consulting activities, enabling the comparison of different national frameworks for energy communities.
- A dedicated representation of the German market will be developed once the forthcoming energy sharing regulations are finalised. This will allow the model to reflect German-specific tariff schemes, levies and incentives, and to evaluate the business case of energy communities under realistic national conditions.

Current status and roadmap

As of D6.2, the core script of the matchmaking and optimisation tool is planned to be ready by the end of 2025, with automated data flows from the WUNergy API targeted for early 2026. Until then, the tool is used in simulation mode, fed primarily with historical data exports and synthetic data where needed. This timeline is broadly aligned with the WP6 schedule that foresees continued trial operations and data collection throughout 2025, while leaving room for the tool to mature into a reusable asset beyond the project.

3.2.2 Measured variables, time resolution and coverage

A key requirement from D6.1 was the provision of a transparent, secure and well-documented data access interface that would allow WP6 and other work packages to consume both live and historical data from the WUNergy community. This has been addressed by the development of a dedicated API that interfaces with the "AURORA" aggregation software operated by SWW.

The API is designed to be accessible to all ECom4Future partners, while respecting data protection and security requirements. It exposes three main categories of information:

1. **Authentication and access control**

- Access is restricted to trusted partners and users through an authentication layer, ensuring that only authorised entities can query community data.
- This aligns with GDPR requirements and with the general ECom4Future approach to data governance as described in D6.1.

2. **Configuration data (static / slowly changing information)**

- Configuration data provides a structured description of ESC customers and assets, including identifiers for metering points, types of installed devices (e.g. PV, storage, controllable loads), and basic contractual attributes.
- At present, much of this configuration information resides in an Excel-based format, which is read by the AURORA system. From a software engineering perspective, it is considered advantageous to convert this into JSON or a similar structured format better suited for APIs. This would simplify data manipulation and integration with external tools and is therefore identified as a recommended enhancement.
- The configuration data is conceptually separate from time series data so that clients needing only measurements do not have to repeatedly download static metadata.

3. **Measurement data (live and historical)**

Live data API:

- Provides quasi real-time measurements for all customers in the ESC, obtained from smart meters and sub-metering devices.
- At the time of D6.2, meters are configured to send consumption data in 15-minute intervals, consistent with the D6.1 baseline for minimum time resolution. Other resolutions (30 minutes, 1 hour) are technically possible and may be used for aggregated or resampled views.

- The live data includes consumption and, where applicable, feed-in from local generation (e.g. PV), enabling near real-time analysis and back-testing of optimisation algorithms.

Historical data API:

- Allows the retrieval of time series data for individual customers over specified time ranges.
- To obtain data for multiple customers, the API is called multiple times, which keeps individual responses manageable and aligns with common REST design patterns.
- Historical data is used extensively for model calibration, scenario analysis and evaluation of flexibility and self-consumption potential.

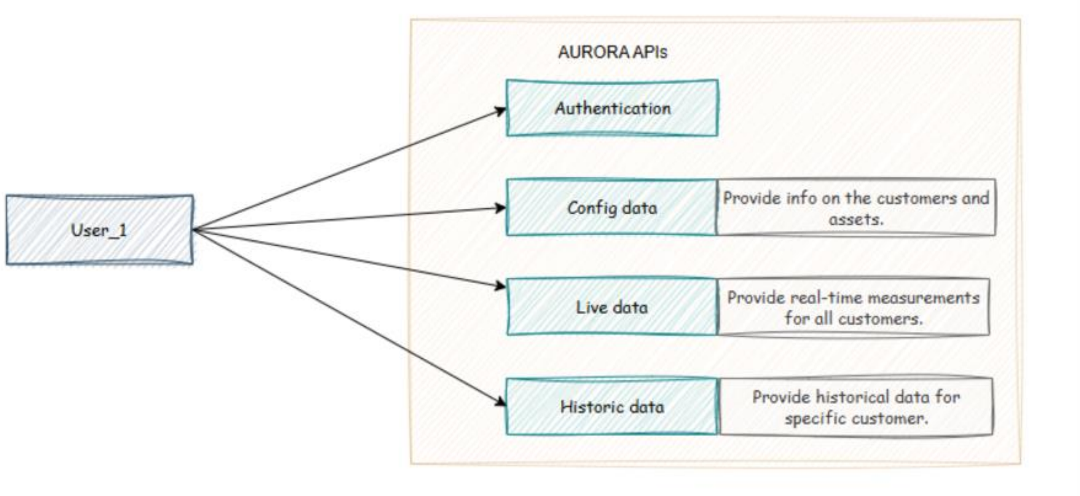

Figure 1 AURORA API at SWW

Through this API, the models and software solutions developed within ECom4Future can be fed directly with both live and historical data from the WUNergy ESC. This closes the loop between real-world field data and algorithmic development, in line with the objectives of WP6 and the baseline described in D6.1. It also supports cross-pilot re-use, as other partners can treat WUNergy as a "data-providing" site for testing their own methods on realistic residential community data.

3.2.3 Open issues, risks, and mitigation actions

1) Timeline for full automation of data flows and model integration

- As noted above, the matchmaking/optimisation script is expected to be ready by the end of 2025, with automated integration into the WUNergy data pipeline foreseen for early 2026.
- Given that WP6 trial operations run through 2025, there is a risk that fully automated, near real-time model-in-the-loop operation will be available only in the later stages of the project or mainly post-project.

Mitigation:

- Use manual or semi-automated data exports for intermediate testing and validation.
- Focus during the project runtime on robust simulation-based evaluation with real data feeds, ensuring that the model is sufficiently mature for post-project operationalisation and exploitation

2) Data management and interoperability

- While the API is operational, some configuration data is still maintained in Excel files, which is not ideal for long-term scaling and integration with diverse tools.
- This creates a potential interoperability and maintainability risk, especially if multiple tools and partners need to access up-to-date asset and customer metadata.

Mitigation:

- Gradually migrate configuration information to JSON-based or similar structured formats, as recommended.

3) Regulatory uncertainty and market design

- The planned inclusion of a German energy sharing market representation in the model depends on the finalisation of national regulations and detailed tariff structures.
- This introduces a regulatory dependency risk, as some aspects of the economic evaluation and business case analysis can only be finalised once the legal framework is clear.

Mitigation:

- Adopt a modular modelling approach that separates generic energy community mechanics from country-specific regulations.
- Use the Czech and other existing markets as reference cases to test the modelling framework.
- Update the German-specific modules as regulations become available, with clear documentation of assumptions used in interim analyses.

3.3 Chalmers Campus Energy Community (Chalmers)

The demonstration at Chalmers campus includes the demonstration of an EMS utilizing different flexibility units such as BES and V2G devices. Different market and tariff structures will be tested as well as different LEC alternatives.

3.3.1 Actual implementation to date

The real estate owner Akademiska Hus collects real time data from all their buildings in the local building SCADA system Webport developed by Kiona. To allow access for researcher a new Webport instance was set up at a local Chalmers server to mirror the most important measurement and control points. The Webport include an API for real time access of data and possibility to post commands such as charge/discharge commands to the batteries. Figure XX presents the interface of the Webport system for one of the buildings at the campus.

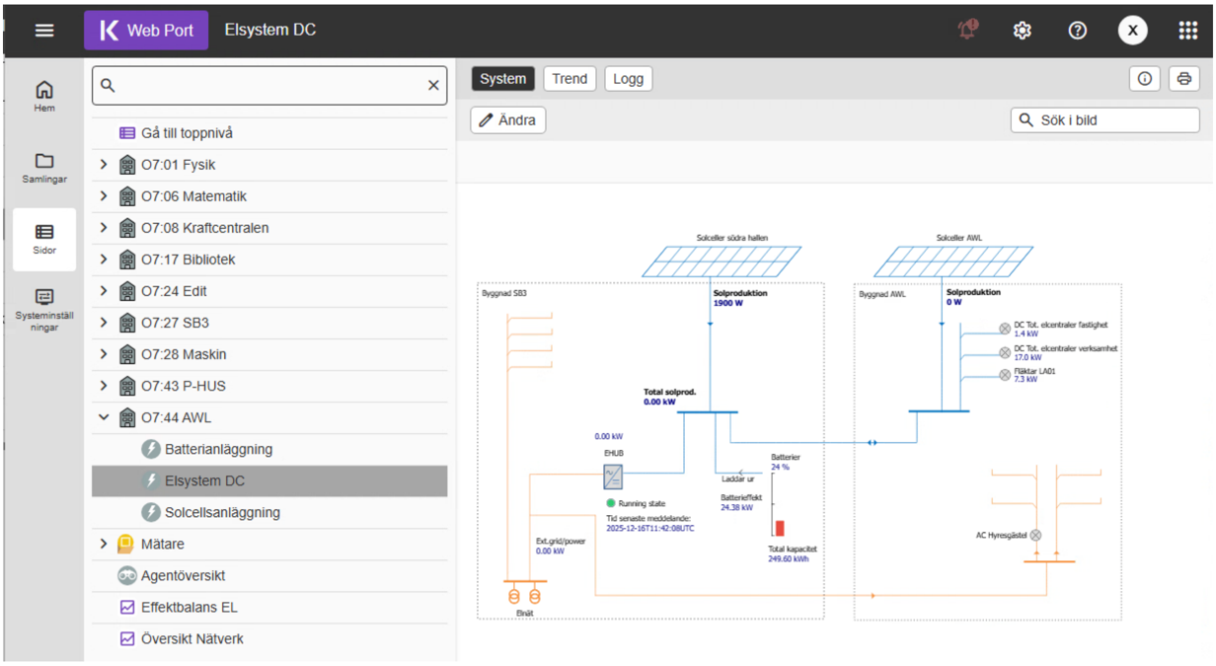

Figure 2 Screenshot of the frontend of Webport interface for one of the campus buildings.

In addition, a database with building data has been developed to collect other building essential parameters such as footprint, height etc. The database also includes an API for easy access to the data. Figure XX presents a screenshot of the frontend of the database.

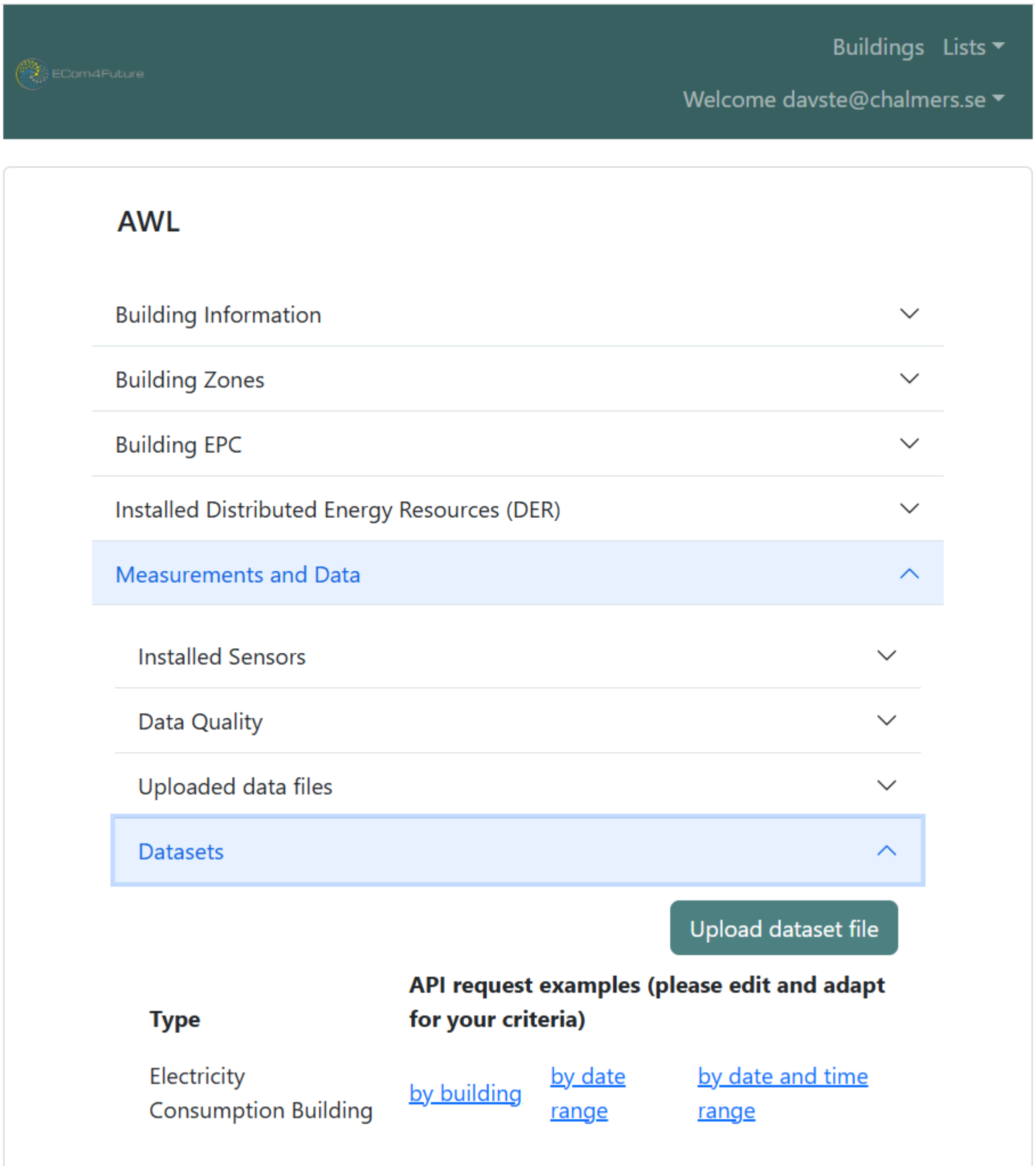

Figure 3 Screenshot of the frontend of Chalmers campus building database (illustrating building named AWL)

For the specific demo cases is presented below:

Demo 1: A LEC optimization model have been developed to dispatch the flexibility resources within a LEC, with the possibility to simulate different combinations of buildings and different market/tariff structures. The model is utilizing a rolling time horizon approach and the aim is to use the model to dispatch the flexibility resources in the campus in real time.

Demo 2: The LEC model developed will be linked to an ADMM model where the flexibility requirement from the DSO will be exchanged to find the most suitable solution for all partners. The model has been validated in simulation studies but will be demonstrated using the real data from Akademiska Hus.

Demo 3: A Bidirectional DC charger has been installed at the campus and will be used as a flexibility resource in the campus. The charger is however not connected to the Webport system but is controlled through a local MQTT broker.

3.3.2 Measured variables, time resolution and coverage

The measurement data in Webport is collected with a time resolution of 10 minutes. It includes the majority of the buildings of the campus owned by Akademiska Hus. For other buildings owned by other real estate companies' historical data have been collected with an hourly resolution.

3.3.3 Open issues, risks, and mitigation actions

Currently the BES located at the campus is connected to the Webport system. However, the control has not yet been implemented in Webport but is done locally through an MQTT broker. The aim is to connect the MQTT broker to the Webport platform and control it through that interface. If it for some reason will not be possible, we will control the BES directly through the MQTT broker.

3.4 Jättesten Energy Community (Chalmers)

The Jättesten use case has been executed throughout the year. The main outcome from the use case was presented in a paper presented at the Nordics EIA conference. The use case investigated the benefits of combining different building types into a local energy community.

3.4.1 Actual implementation to date

The use case relies on historical data from the Jättesten area in Gothenburg in combination with the data from Chalmers Campus. As there is limited possibility to control flexibility resources in Jättesten this use case is only based on simulations. The benefit from installing BESS and PV in one of the evaluated LEC combination is shown in Figure XX, where the potential cost reduction is higher for larger PV plants.

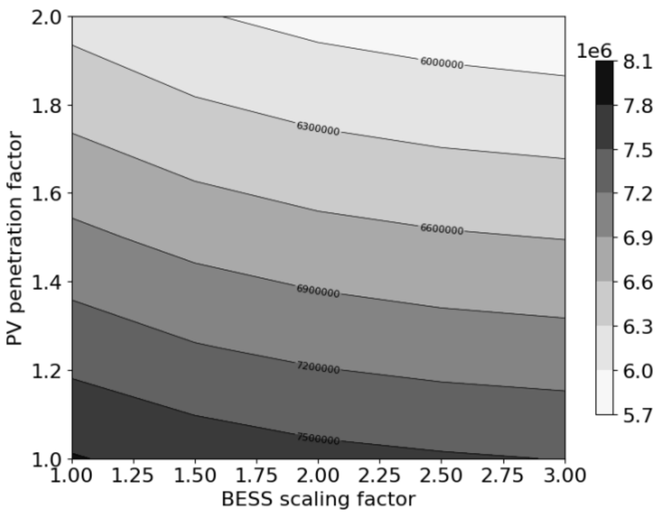

Figure 4 Results on total operational costs for varying PV penetration and BESS

3.4.2 Measured variables, time resolution and coverage

The measurement data used in this study include historical hourly electricity consumption data collected for the year 2024.

3.4.3 Open issues, risks, and mitigation actions

Due to the limited possibility to control the flex resources within Jättesten this use case was conducted as a simulation study with real data. No further evaluations are planned.

3.5 FDD and Optimization of Energy Community (TUW / DiLT)

3.5.1 Actual implementation to date

Energy communities consist diverse operational devices such as heatpumps, batteries, electrical appliances, and HVAC or chiller systems (particularly in mixed-use districts with facilities exhibiting high energy demand). There devices are prone to faults, and timely intervention is essential to prevent performance degradation. Consequently, reliable fault detection is critical to maintaining efficient and uninterrupted operations. In this task, we developed a framework for fault detection and diagnosis (FDD) in energy-efficient buildings, focusing on air handling units (AHUs) and validated with synthetic data. We propose One-class Classification Cluster ENsembles (OCCEN), an innovative approach for multivariate time series data. OCCEN employs an unsupervised, physics-informed artificial intelligence (AI) methodology that identifies anomalies based on deviations from physical system boundaries, without requiring prior fault knowledge. The framework integrates one-class classification for anomaly detection, clustering, and binary classification to identify fault features. It further enhances temporal accuracy using explainable AI and dynamic time warping (DTW).

For the task of anomaly detection, the objective is to learn a function $f : T \rightarrow \{0, 1\}$ where $f(t_i) = 0$ represents a non-anomalous instance and $f(t_i) = 1$ indicates anomalous data point. Furthermore, it is assumed that characteristics of T represent normal operations, and deviations from normal patterns are indicative of faults. Labeling high-resolution time series data is challenging due to the time required and the sparsity of class labels, particularly anomalies. This imbalance often biases models toward majority class labels. One-class classification (OCC) methods address this by focusing on learning the distribution of a single class, typically normal observations, to establish boundaries distinguishing inliers from outliers, making them effective for anomaly detection. For the anomaly detection task, various One-Class Classification models are evaluated, and the best one is considered to generate a ranked list of diagnosis through XAI and Dynamic time warping technique.

We evaluated the efficacy of different OCC methods in detecting anomalies and considered the best-performing method for the fault diagnosis analysis. To assess fault diagnosis performance, we established a simple baseline by employing explainable AI models immediately following the OCC methods, unlike OCCEN. The explainable model generates a list of ranked features, i.e., diagnoses, for a given time step. The top features ranked by the XAI model are then retained for over m time steps and further re-ranked based on distance measures calculated by FastDTW. The Elliptic Envelope model consistently showed high precision and recall across all fault types, resulting in strong F1 scores, particularly in detecting Cooling Coil Valve Stuck (Cvlv_stk) anomalies. Its precision above 93% and recall above 90% in most cases reflect its robustness and reliability in anomaly detection. Diagnosis metrics such as Overlap@P, HitRate@K and RootCause@k showed that features ranked within top places are highly relevant to the faults identified during the detection phase.

3.5.2 Measured variables, time resolution and coverage

We employ a synthetic dataset of a single-duct air handling unit (AHU). The dataset comprises one year of operational data, incorporating both faulty and nominal system behaviours, with a total of 525,600 time samples representing each operational condition.

3.5.3 Open issues, risks, and mitigation actions

When time series data is modelled with methods that do not consider time, several problems can occur. Important patterns such as trends, seasonality, and delayed effects may be lost, which can lead to poor predictions and missed changes in system behaviour. Models may also become sensitive to noise or perform poorly when real conditions change. These issues can be serious in systems where quick and correct decisions are needed. To reduce these risks, useful time-based features such as lag values and rolling averages should be added. It also helps to test models under different operating conditions and use hybrid methods that include some form of temporal information to maintain context and improve reliability. Additionally, the method is evaluated using publicly available synthetic datasets, which may not fully represent real-world conditions. A key open issue is the need for real data from operational energy communities to properly test performance and confirm whether similar results can be achieved in practice. However, access to such data is difficult due to privacy concerns and strict data-sharing rules that limit the availability of system indicators and time series patterns. Although this remains a challenge, the current approach is designed in principle to detect faults in similar types of input data.

3.6 Energy Community Niedergrail (TU WIEN / DiLT / dwh)

3.6.1 Actual implementation to date

Before building any models, the data was analysed using stochastic and exploratory analysis techniques to gain an understanding of the dataset and to identify potential optimization opportunities. The first analysis revealed multiple inefficiencies in battery operation and charging behaviour. For several users, batteries were charged using grid energy even though the state of charge (SoC) was not low, indicating avoidable costs and suboptimal control. Similar inefficiencies were observed in Wallbox (Wattpilot) charging, where grid charging occurred despite full batteries or incomplete battery discharging.
These findings highlight substantial potential for optimization through improved load shifting and charging strategies.

Forecasting models for consumption and PV production were developed. External weather data (solar irradiance, temperature) and time-related features were integrated into the dataset, with cyclic time features encoded using sine–cosine transformations. Various statistical and machine-learning approaches were evaluated, with XGBoost identified as the most suitable model. Both stepwise and direct 24-hour forecasts were tested, and daily 24-hour forecasts are now available for use in the optimization processes in WP4.

A realistic Reinforcement Learning (RL) environment was implemented using the Deep Deterministic Policy Gradient (DDPG) method. Initially, the optimization focuses on a single participant, aiming first to minimize grid load and later to minimize energy costs by incorporating grid price information. The learned policies were evaluated against rulebased baselines and actual operational data. Hyperparameter optimization is performed using Optuna to improve stability and reduce runtime. The next step is the extension to a multi-agent setting covering the entire energy community.

3.6.2 Measured variables, time resolution and coverage

Household-level variables

- Grid import and export per household
- PV production per household (if available for the participant)
- Battery system data (if available for the participant):
    - State of Charge (SoC)
    - Charging and discharging power
- Wallbox (Wattpilot) charging data (if available for the participant)

External / auxiliary variables

- Weather data (solar irradiance, temperature)
- Time features (hour, weekday, seasonal cycle via sine/cosine transformation)
- Grid prices (for import and export)

Time resolution and coverage

Data was collected from 01.01.2023 to 01.10.2024, providing more than 1.5 years of usable data. This dataset was split into a training and a test set, with the last 12 weeks reserved for testing and the remaining period used for training.

Depending on the source, different granularities are available:

- PV dataset from smart meters: 5-minute resolution
- EDA-report dataset: 15-minute resolution
- Weather data (Geosphere): available in 10-minute or hourly intervals

3.6.3 Open issues, risks, and mitigation actions

1. Inefficient battery and charging operations in RL

The RL agent currently struggles to identify the long-term benefit of charging the battery at specific times, because the reward associated with maintaining a high SoC is only observed several time steps later. As a result, battery usage is not yet optimal. This delayed reward structure remains one of the central challenges in the current RL setup.

2. Heterogeneous participant behaviour and objectives

Transitioning from a single-agent to a multi-agent optimization strategy requires capturing diverse user profiles, technical configurations, and household-specific objectives. One potential approach is to treat the entire community as a single aggregated entity rather than applying different strategies for each participant. Whether this is feasible depends on the legal and organizational structure of the community and may vary between ECs.

3. High computational demand of RL methods

RL requires extensive training data, significant computational resources, and careful hyperparameter tuning. Long training times and potential instability make it difficult to identify an optimal policy, especially when less than two full years of data are available. Additionally, it is challenging for the agent to learn seasonal patterns when each seasonal period is only observed once in the dataset.

3.7 Energy Community Mariatrost (TU Wien / DiLT)

3.7.1 Actual implementation to date

The Renewable Energy Community (EEG) in Mariatrost is centered around utilizing photovoltaic (PV) systems to generate solar power. Initially, the district's PV installations produced more electricity than local households consumed, sending the surplus back to the public grid. Thanks to Austria's Renewable Expansion Act (EAG) of 2021 and amendments to the Electricity Industry and Organization Act (ElWOG), this surplus electricity can now be shared across property boundaries with reduced network fees, making locally produced PV electricity cheaper than traditional public providers. Established as a formal association in July 2022 and fully organized by March 2023, the EEG Mariatrost includes diverse members such as private individuals, small businesses, and public institutions. It aims to expand by adding more PV installations, an electric vehicle charging station, and potentially an e-car sharing service, while fostering partnerships in the community through a grassroots approach.

In this analysis, we combine the community's overall electricity consumption with individual participants' usage profiles to identify key insights and explore the potential for forecasting electricity demand for each participant.

Energy Consumption Analysis

The total energy consumption of a community over time analysed from January to August 2024, revealed a clear seasonal trend. In the initial months (January to February), energy consumption is high, peaking between 6 and 8 kWh, likely due to increased heating needs during winter. From February to March, there is a noticeable decline in consumption, suggesting reduced energy requirements as temperatures warm. By April, consumption stabilizes at a lower level, fluctuating between 1 and 3 kWh, reflecting more consistent energy use through the spring and summer months. Occasional spikes in May, June, and late August suggest temporary increases in energy demand, possibly due to specific events or external factors. Overall, the consumption profile of the community reveals a significant reduction in energy consumption from winter to summer, indicating effective seasonal adaptation or conservation efforts within the community.

Energy Demand Prediction

A key objective of this project is to predict energy demand, in this case, electricity consumption, for the targeted energy community. For this purpose, we utilized data from the Mariatrost energy community, which includes electricity consumption records of individual participants. Given the data's structure, we approached this task with a multi-model learning strategy, creating separate models for each participant to capture the unique consumption patterns of each. The rationale behind this approach is that each participant exhibits distinct characteristics in their consumption profile, independent of others within the community. Since we only have individual consumption data and lack other contextual information (like demographic or appliance data), a personalized model for each participant can be trained using only the time-series data. This approach can avoid the need for cross-participant feature engineering, simplifying the modelling process. Additionally, each participant's model can be tuned specifically to their behaviour, leading to more precise predictions. This learning framework helps with optimizing load predictions and energy management strategies for each user, which is valuable for both participants and the community as a whole.

We implemented a simple baseline data-driven (machine learning) model for predicting energy demand. Baseline models are typically designed with a minimal parameter space required for fine-tuning. For each participant, the model follows a one-step forecasting approach, predicting consumption at time t+1 based on the current consumption at time t. In other words, it involves developing a function f(x) that maps the input variable x (energy consumption at time t) to minimize the average prediction error over training instances for forecasting consumption at t+1. Once trained for each participant, the model is applied iteratively on the test set, meaning each prediction is fed back as input to forecast subsequent steps. In the initial experiments we trained n models (estimators) where n=number of participants. The underlying model is classical machine learning method called linear regression. The data is resampled to hourly sum of energy consumption for each participant; hence the prediction is done for hourly energy demands.

The evaluation of the baseline model is performed through 10-fold cross validation for each participant. Mean Squared Error (MSE) distribution is analysed for training and testing sets across different participants in an electricity demand prediction model. The majority of participants show low and consistent MSE values for both training and testing, indicating that the model generally performs well in capturing demand patterns. However, a couple of participants significantly higher MSE, particularly in the testing phase, with a noticeable spread and outliers, suggesting that the model struggles to generalize effectively for these individuals. This discrepancy aligns with the findings and analysis that indicates unique or irregular demand patterns for these participants, which the model fails to predict accurately.

3.7.2 Measured variables, time *resolution* and coverage

Electricity consumption data for the energy community participants, covering approximately eight months, has been supplied to the consortium by one of our partners, DiLT Analytics. This data has a 15-minute time resolution. There are 19 participants in total, but we only have information on five (participants IDs: 4, 10, 13, 14 and 19) who have PV systems installed, although we currently lack access to their electricity generation profiles. The data includes participants who lack available measurements, so they were excluded from the analysis because imputation is not feasible without additional predictor variables, such as non-operational factors (like demographics or routines) or operational data. The only operational data we have is electricity consumption for this community. For participants with missing measurements in certain months, we used Multivariate Imputation by Chained Equations (MICE) to estimate and fill in those missing values.

3.7.3 Open issues, risks, and mitigation actions

To address this issue related to high error rates, additional information will be integrated into the learning feature space to enhance the model's effectiveness. One potential approach involves incorporating non-operational variables, such as weather data, obtained from open-source APIs. This data will be collected for the periods corresponding to the Mariatrost Energy Community’s measurement timeline and made available for analysis. Nonetheless, initial findings suggest that electricity consumption alone serves as an effective indicator, and the model demonstrates sufficient capacity for predicting energy demand in the baseline setup. Additionally, other baseline models, including support vector machines, XGBoost, random forests, and advanced deep learning methods, will be adapted and trained. Following this, data-driven learning methods and strategies will be developed to assess and compare performance improvements over the baseline model.

3.8 FH Johanneum & FH Campus 02

3.8.1 Actual implementation to date

Within the Energy Analytics and Solution Lab (EAS-Lab cooperation between FHJ and C02), three planned demonstrators could be carried out as pilots. The documentation took place as part of several scientific theses. In the thesis "Optimizing Energy Demand and Supply with Smart Heating and Cooling Solutions in Multi-Family Buildings through a Simulation Platform", an energy demand forecast of a single-family home was created by means of remote access to weather data, which can be used for energy management independently of the platform. Measured variables from the current PV generation, the electricity consumption and the heating demand serve as input variables. The calculation tool reads these values and determines the parameters for the energy and storage requirements in the next few hours in an optimization loop. Completion 100%

The other two demonstrators were carried out at Campus 02: the focus of the work so far has been on using surplus PV (Photovoltaic) energy efficiently for domestic hot water heating. This builds on a master's thesis, whose technical analyses and preliminary experiments were directly used in the project. A flexible system for automatic load control was developed, set up, and tested. A practical test setup with a small balcony PV system with a maximum output of 800 W was built, and measurement data from the FH CAMPUS 02 PV system were also used. This setup allowed to continuously record real PV generation and household consumption data and to evaluate different control strategies. Using real data ensures that the findings are practical and applicable to actual household situations rather than only theoretical models.

First, the classic on/off control method, also known as two-point control, was examined. Using an oscilloscope, it was clearly visible that switching the water heater fully on or off creates sudden jumps in power consumption. These sudden load changes are safe for the heater, but they reduce the efficiency of PV self-consumption because surplus energy is not used smoothly. Measurements confirmed that this method is not ideal for smaller PV systems because small surplus energy amounts cannot be efficiently captured. For larger PV systems, this simple control works better, but it still lacks flexibility in adjusting power to match instantaneous PV generation.

Next, phase-angle control was tested. Oscilloscope readings clearly showed the characteristic "cut" sine wave, and it became visible how power can be adjusted continuously. In practical experiments with water heating, this method was fast and responsive, making it possible to match power supply more precisely to PV generation. However, the measurements also revealed harmonics, or electrical distortions, in the system. These harmonics can cause small errors in energy meters and may, in some cases, affect the power quality in the grid. Understanding these effects allows the team to document the limits of phase-angle control and plan improvements in regulation algorithms.

3.8.2 Measured variables, time resolution and coverage

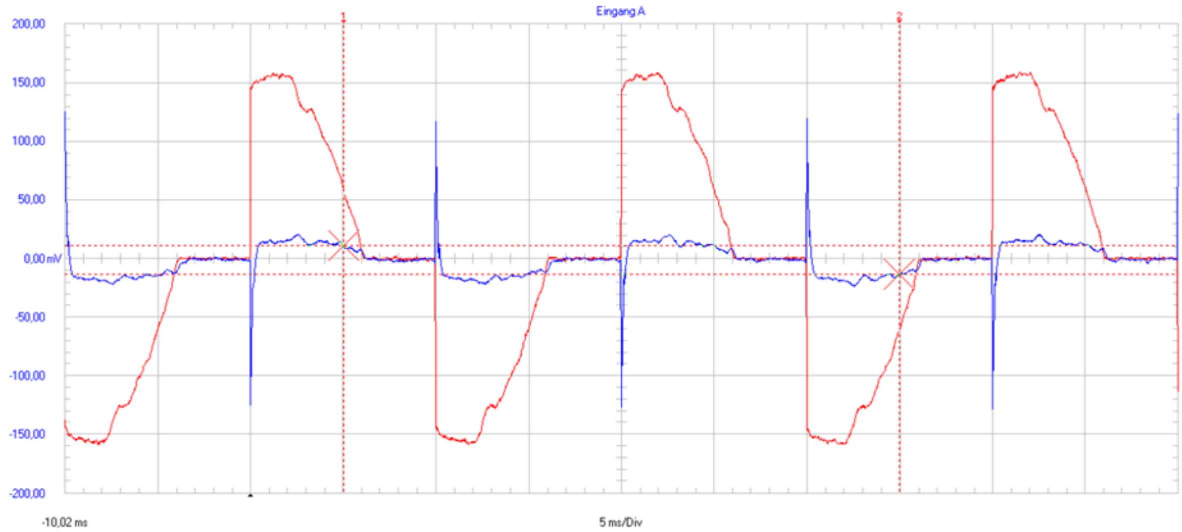

Figure 5 Cuts in the sine wave due to phase-angle control measured with the oscilloscope

Pulse-width control, also called burst firing or package control, was examined only theoretically. According to literature, it can maintain a nearly perfect sine wave and reduce harmonics, which would improve measurement accuracy and reduce potential grid issues. However, no practical measurements have been conducted yet, so it is unclear how well this method would perform in real water heating scenarios or how fine the control of power would be. This theoretical study is important for planning future steps, comparing different control approaches, and guiding improvements for water heating applications in households.

The mobile load control system implemented in this project uses phase-angle control to direct surplus PV energy to a water heater. Tests so far show that this approach can significantly increase self-consumption of solar electricity. How sudden changes in household loads affect system performance, as they naturally create short-term fluctuations that the controller must balance, have also been investigated. Measurements using two different energy meters helped identify potential deviations caused by phase-angle control. Different hot water systems, ranging from small electric boilers to larger storage tanks, were studied in theory, and all could benefit from intelligent load control. The modular design of the system allows flexible adaptation to various types of heaters and measurement setups.

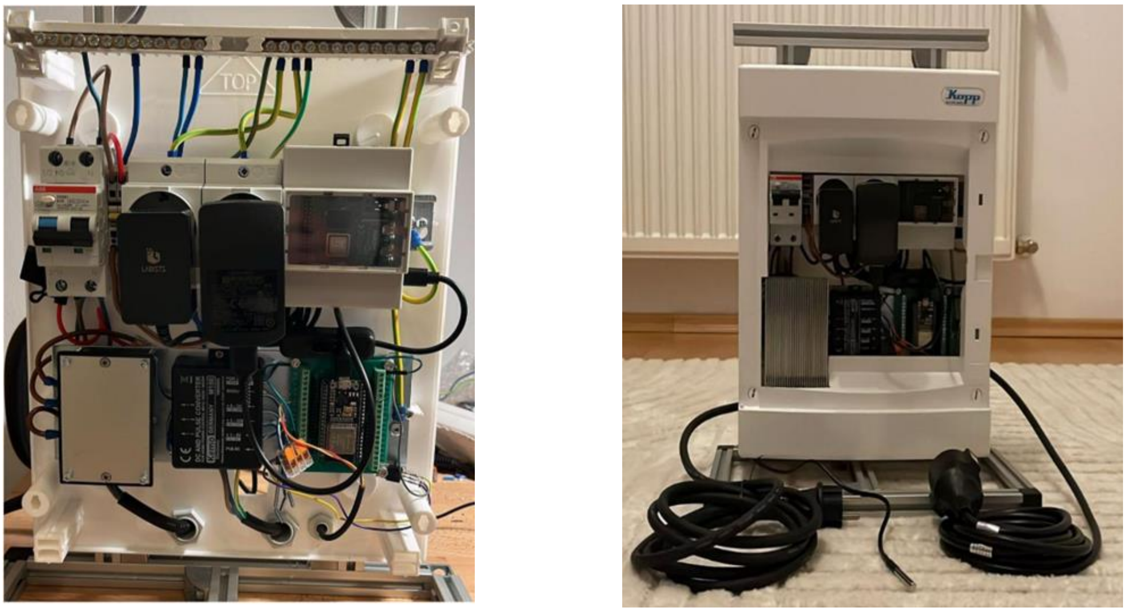

Figure 6 Interior of the mobile load control system for flexibel load control

In parallel, a simulation model for heat pumps and hot water preparation was developed using Node-RED. This model allows to reproduce the dynamic behaviour of heating system realistically without having to perform hardware tests for every system change. Node-RED's graphical, modular interface makes it possible to model energy flows, control strategies, and household consumption profiles in real time. The model includes the thermal house behaviour, PV generation, storage behaviour, heater operation, and different control strategies. Household heating profiles, water temperatures, and building characteristics can all be adjusted, enabling safe testing of different scenarios, extreme cases, and potential optimizations. In future project steps, this simulation model will be integrated into the Energy Analytics and Solution Labs process controlling system XAM Control. This allows control algorithms to be tested and improved in a controlled virtual environment before implementation on real hardware. Using this approach reduces risks, saves time, and supports practical, scalable solutions for efficient and flexible hot water heating using surplus solar energy.

Figure 7 Node-RED Dashoard for parameterizing and controlling the heatingsystem

3.8.3 Open issues, risks, and mitigation actions

Further investigations will focus on load control within energy communities. Both Node-RED and the evon Smart Home system will be used for this purpose. The goal is to develop and evaluate a decentralized and flexible load control concept that distributes surplus energy across the community in an optimal way, increasing self-consumption and improving the independence of all members. First tests on smart meter data analysis and on programming the required system interfaces have already been carried out.

3.9 GreenMogo Energy Community

3.9.1 Actual implementation to date

An overview of the Energy Community enabled Prosumer

GreenMogo is a zero-energy building operating in different scenarios to teach various target groups about more sustainable ways of living and building. Within this paper, the premise acts as a regulatory restricted prosumer (following the operation of a UniRCon) providing insights on energy needs for modern local energy communities. The general grid topology of this UniRCon is presented in the figure below and consists of a family residential building (with regular home appliances) with PV generation and BESS interfaced with a Victron inverter and a heat pump for thermal energy transfer. The measurement layer consists of high-reporting rate information (1 frame/second), which is provided by energy meters operating using the Unbundled Smart Meter (USM) concept for energy use, and also provided by the Colour Control GX TCP/IP integrated with the Victron Inverter for energy generation (PV and BESS). The grid configuration of the prosumer system (demo site at GreenMogo) is depicted in the figure below. It encompasses a residential structure equipped with conventional household devices and a heat pump for thermal energy exchange also featuring photovoltaic (PV) generation and a battery energy storage system (BESS) – all elements being integrated with a Victron inverter.

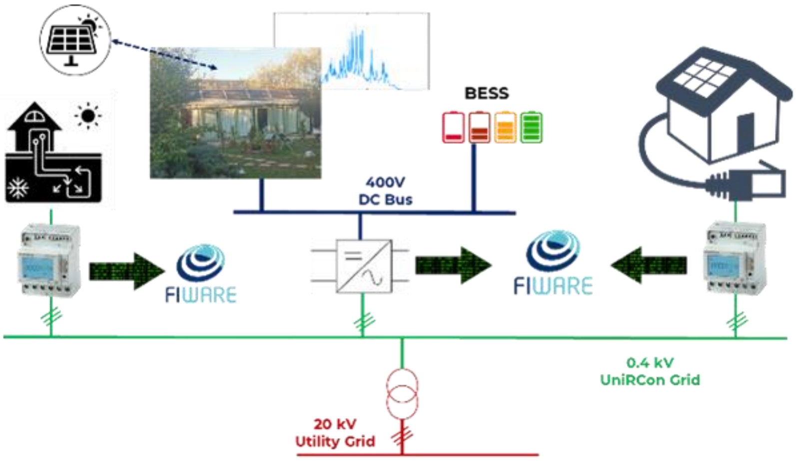

Figure 8 Architecture GreenMogo

Architecture and Implementation of the ICT Platform for prosumer environment

The below picture provides an overview of the open-source framework and its integration layers. At this stage, the Data Harvesting Layer is key, enabling unified data collection from various sources. The main goal is to standardize data exchange and storage. Above it, the Temporary Logic Layer includes Smart Converters (APIs) that translate heterogeneous source data into NGSI-like format, sending it to the central broker or database. The framework uses an updated version of the FIWARE NGSI format — a structured, JSON-based format—to ensure seamless data exchange and interoperability. It represents data as entities with attributes (values, types, and optional metadata) and relationships, enabling real-time updates, integration with various components, and support for scalable, smart system design. The Persistent Logic Layer is based on a context broker (configured per local specifications) that receives NGSI-like formatted data for processing and integration. The Data Storage Layer ensures long-term storage using technologies like MongoDB, InfluxDB, or CrateDB, depending on the trial site. Finally, the Presentation Layer supports integration with third-party services and visualization tools, using platforms such as Chronograf, Grafana, or custom applications in HTML, Java, etc. This platform integrating multiple sources of information (load side, generation side and ambient sensors) using an updated NGSI format and test the interoperability with third-party applications. The demo platform’s real-world cybersecurity performance depends solely on the local area network (LAN provided by the internet carrier) and on a dedicated virtual private network (VPN) for establishing connection between the entities at prosumer level and the hardware server running the applications.

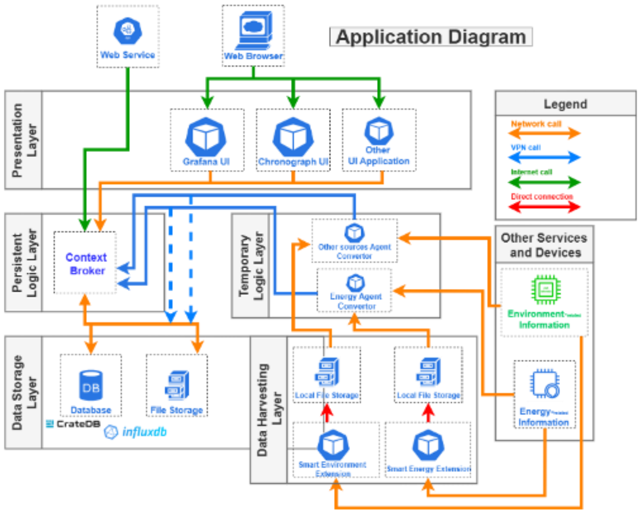

Figure 9 Application Diagram GreenMogo

3.9.2 Measured variables, time resolution and coverage

The measurement layer providing information to be integrated from the prosumer side comes from multiple sources of information including load and generation assets, different sensors and devices, derived from local measurements set on 1 frame/s reporting rate.

High-reporting rate information for loading

The information for the load side is derived from local measurements using the concept of the Unbundled Smart Meter (USM) set on 1 frame/second reporting rate. This is achieved using Raspberry Pi boards (RPIs) connected to different types of smart meters.

Figure 10 Unbundled Smart Meter

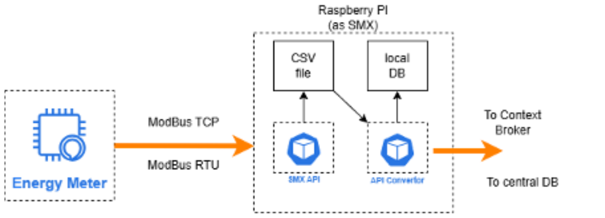

Figure 11 Setup Smart Meter with Raspberry PI

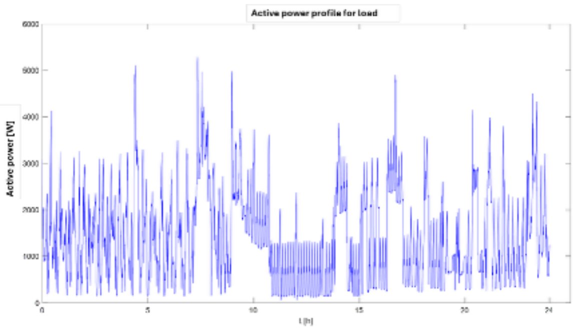

Figure 12 High-reporting rate information for generation

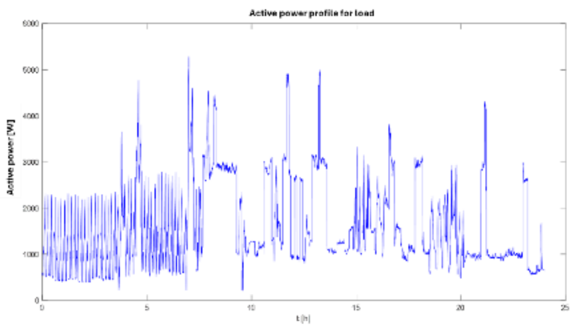

Figure 13 Active power profiles for load

High-time resolution measurement information (1s resolution) for PV generation and Battery Energy Storage System (BESS) is achieved by connecting a RPI via TCP/IP protocol to the inverter and performing dedicated software configuration based on Python. On the RPI, the script is running to collect the data locally in .csv files and forward the readings to the platform.

Figure 14 BESS GreenMogo

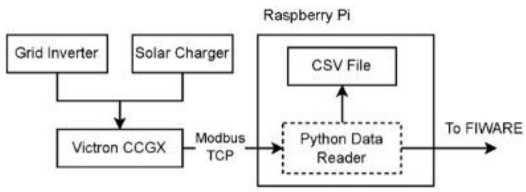

Figure 15 RPI Setup, GreenMogo

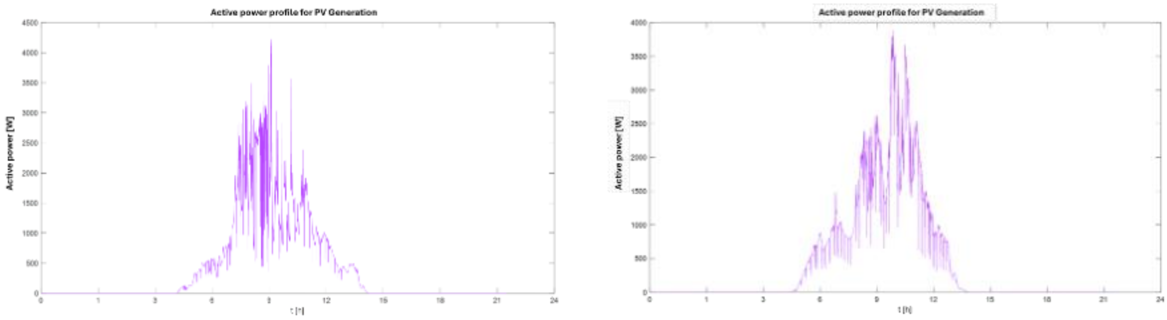

Figure 16 Active power profile for PV Generation

Third-party applications integration

The platform demonstrates a high degree of flexibility in analysing prosumer behaviour by supporting multiple reporting rates, diverse statistical models, and high-resolution measurement information. Through tools such as CV(RMSD), Goodness of Fit, and Matrix Profile–based anomaly detection, it can adapt its analytics to different temporal windows—from second-level resolution to legacy smart-meter intervals—while maintaining consistent performance. This makes it suitable for prosumers with highly dynamic profiles, such as those combining PV generation, storage, and no-power-injection constraints. External forecasting engines, variability classifiers, anomaly-detection systems, or behavioural-pattern analysis tools can be plugged directly into the data pipeline. This allows service providers to deploy customised algorithms at scale and prosumers to benefit from tailored insights, operational optimisation, and enhanced flexibility management. By combining high-reporting-rate data acquisition with an extensible analytical ecosystem, the platform creates a robust environment for real-time monitoring, intelligent control, and data-driven decision-making across heterogeneous prosumer setups.

Some examples of results for histograms, CVRMSD and GoF computation and matrix profile assessment are given below.

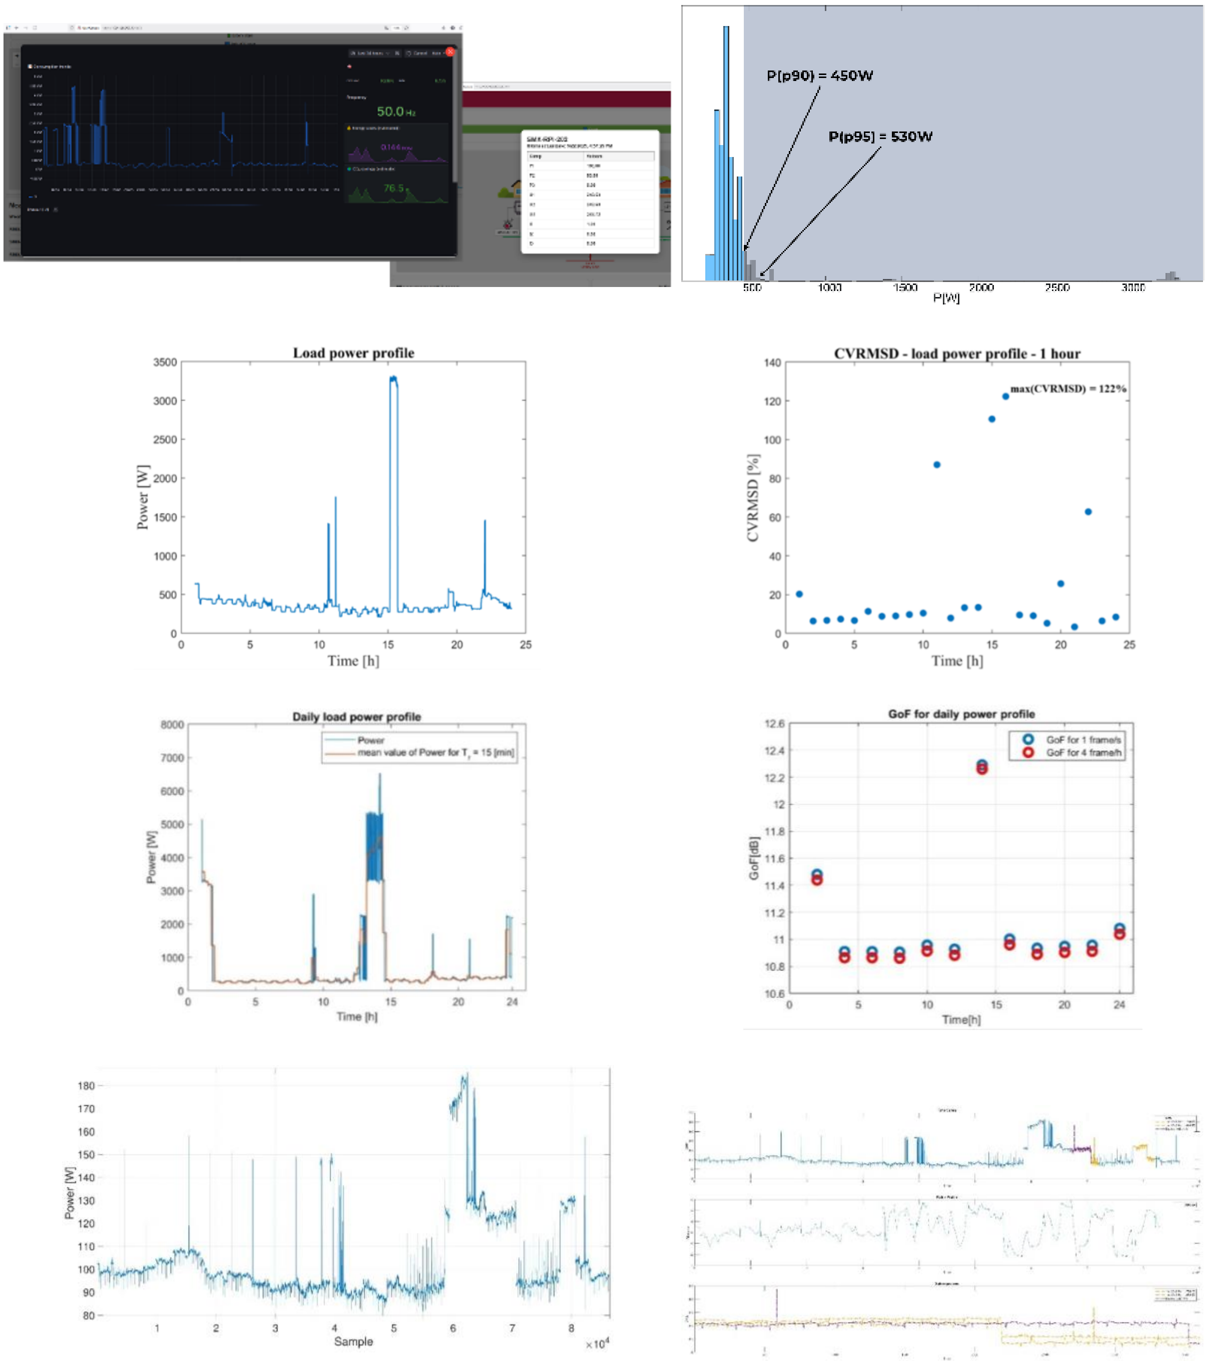

Figure 17 Third party integration

3.9.3 Open issues, risks, and mitigation actions

Several open issues remain regarding the deployment of the platform in a prosumer environment. Interoperability challenges persist, as third-party applications rely on different data models and interfaces, making seamless integration difficult. At the same time, the high–reporting-rate measurement information raise concerns about scalability, storage capacity, and processing delays, especially when large numbers of prosumers are connected. Variability metrics and forecasting models still require fine-tuning to cope with diverse prosumer behaviours, and the lack of standard thresholds complicates consistent interpretation. Finally, the coexistence of legacy and advanced metering infrastructures introduces inconsistent data quality, while cybersecurity and data-privacy concerns remain critical, particularly when external apps are granted access to granular energy data.

These issues introduce several risks, including potential system bottlenecks from data overload, inaccurate analytics caused by poor model-to-data fit, and reduced flexibility due to vendor-specific integrations. Security vulnerabilities or regulatory non-compliance could undermine trust, leading to low prosumer adoption.

To mitigate these challenges, adopting open standards and well-defined APIs is essential to ensure smooth integration of third-party tools. A scalable architecture—combined with edge processing— will contribute to big-data processing. Establishing clear performance indicators for variability metrics will support consistent interpretation across deployments. Harmonising legacy and high-resolution measurement data, along with strong cybersecurity, encryption, and privacy-preserving mechanisms, will help ensure reliable operation. Finally, transparent communication with prosumers about data usage can increase trust and encourage participation in the platform ecosystem.

4. Timeline

The timeline for WP6 encompasses the critical phases of trial site planning, setup, and operation, beginning in early 2024 and continuing through 2025. Initial activities focus on establishing infrastructure and aligning technical frameworks (T6.1) to support subsequent tasks. Monitoring of energy consumption and production (T6.2) will follow, running in parallel with the integration of market schema and user feedback analysis (T6.3) to ensure functionality. Full-scale trial operations (T6.4) are scheduled to commence in late 2024, allowing sufficient time for iterative adjustments and data-driven refinements throughout 2025. With this deliverable T6.2 is documented and on track to support T6.3/T6.4 in 2025.

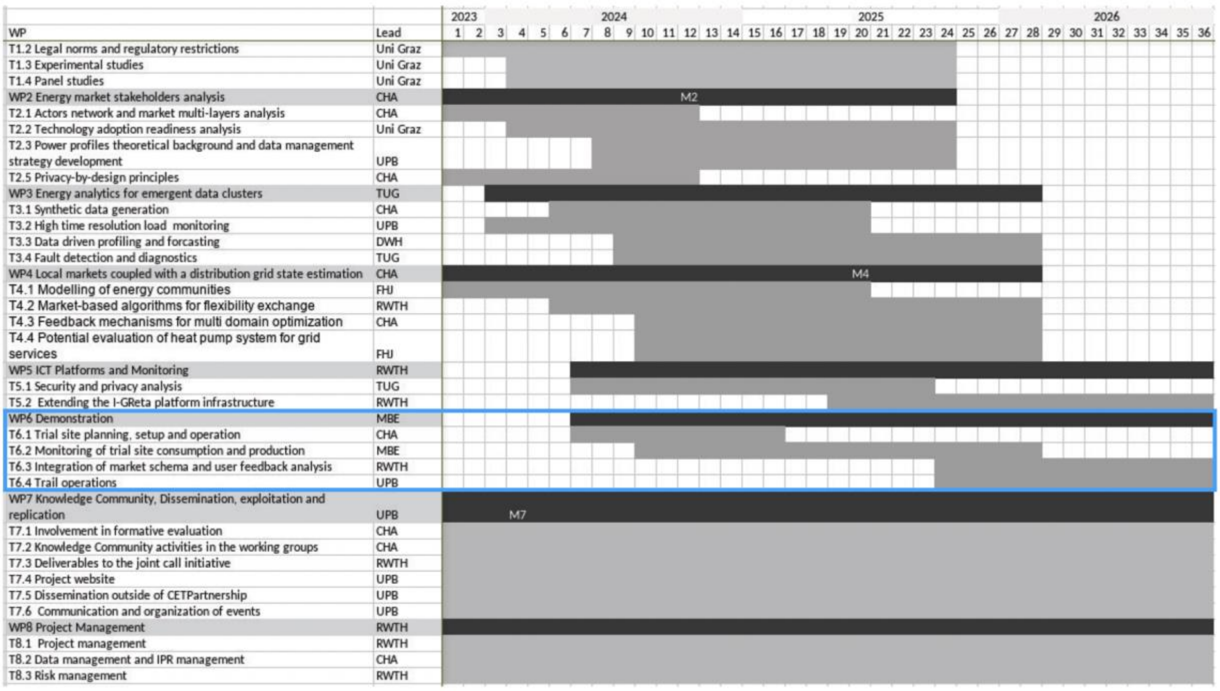

Figure 18 GANTT WP 6

I. List of Figures