AMAZING-6G
Amazing Large-Scale Trials and Pilots for Verticals in 6G

Deliverable 5.1

Energy domain intermediate developments

Co-funded by the European Union

AMAZING-6G project has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101192035.

Project Details

Call	HORIZON-JU-SNS-2024
Topic	HORIZON-JU-SNS-2024-STREAM-D-01-01
Project start date	01/01/2025
Duration	36 months
GA No	101192035

Deliverable Details

Deliverable WP:	WP5
Deliverable Identifier:	D5.1
Deliverable Title:	Energy domain intermediate developments
Editor(s):	Cristian Petrache (ORO)
Author(s):	Gabriele Scivoletto, Giada Landi, Gabriele Ferrigno, Marco Ruta (NXW) Bogdan Lebu, Dragos Anca (CSOFT) Mihai Neagu (SIM) Yohan Toh, Marco Turrini, Andre Schlindwein (TNO) Raul Barbosa (CAPG) Razvan Mihai, Alexandru Oprea, Cristian Petrache (ORO)
Reviewer(s):	Paulo Duarte (CAPG), Heikkinen Antti (VTT)
Submission Date:	31/05/2026
Dissemination Level:	PU

Table of Contents

List of Figures

List of Tables

List of Acronyms and Abbreviations

TERM	DESCRIPTION
5G	Fifth-generation
5GC	5G Core
6G	Sixth-generation
AMF	Access and Mobility Management Function
AM/SM	Access/Session Management
APC	Authentication and Policy Control
API	Application Programming Interface
AI	Artificial Intelligence
AI/ML	Artificial Intelligence / Machine Learning
APN	Access Point Name
AI-AAS	AI-as-a-Service
AUSF	Authentication Server Function
BVLOS	Beyond Visual Line of Sight
B5G	Beyond 5G
B5G/6G	Beyond 5G / Sixth-generation
CAAS	Compute-as-a-Service
CAPG	Capgemini
CL	Closed Loop
CPE	Customer Premises Equipment
CQI	Channel Quality Indicator
CSOFT	Customer Soft Solution SRL
DNN	Deep Neural Network
DSO	Distribution System Operators
DL	Downlink

TERM	DESCRIPTION
E3	Energy use case 3
E2E	End-to-End
EC	European Commission
EU	European Union
EMS	Energy Management System
MQTT	Message Queuing Telemetry Transport
FL	Federated Learning
FR	Failure Risks
GA	Grant Agreement
IBN	Intent-Based Networking
GNSS	Global Navigation Satellite System
gNB	next Generation NodeB
GNSS and RTK	GNSS + Real-Time Kinematic
HVAC	Heating, Ventilation and Air Conditioning
HITL	Human-In-The-Loop
HPA	Horizontal Pod Autoscaler
IoT	Internet of Things
IO&M	Inspection, Operation and Maintenance
KPI	Key Performance Indicator
KVI	Key Value Indicator
LLM	Large Language Model
MCDM	Multi-Criteria Decision Making
EuCNC	European Conference on Networks and Communications
FEM	Finite Element Method

GUI	Graphical User Interface
HTTP REST	Hypertext Transfer Protocol Representational State Transfer
I2C	Inter-Integrated Circuit
DT	Digital Twin
E1	Energy use case 1
E2	Energy use case 2
EMBB	Enhanced Mobile Broadband
NTRIP	Networked Transport of RTCM via Internet Protocol
NR	New Radio
NSA	Non-Standalone
S-NSSAI	Slice selection assistance information
ORO	Orange Romania
OPEN5GS	Open-source implementation of 5G Core
OPEX	Operational Expenditure
OSM	Open Source MANO
PCF	Policy Control Function
PDU	Protocol Data Unit
PM	Particulate Matter
PV	Photovoltaic
QOE	Quality of Experience
QOS	Quality of Service
REC	Renewable Energy Community
REC-EMS	Renewable Energy Community Energy Management System
REST	Representational State Transfer
RL	Reinforcement Learning
RPC	Remote Procedure Call

IMSI	International Mobile Subscriber Identity
NGC	Next Generation Core
NTN	Non-Terrestrial Network
MEC	Multi-access Edge Computing
ML	Machine Learning
SMF	Session Management Function
NLP	Natural Language Processing
SNS JU	Smart Networks and Services Joint Undertaking
SLA	Service Level Agreement
SNS Stream B	Stream B
TDD	Time Division Duplex
TM Forum	Telecom Management Forum
TLS	Transport Layer Security
TNO	Netherlands Organisation for Applied Scientific Research
UDM	Unified Data Management
UE	User Equipment
UDR	Unified Data Repository
URLLC	Ultra-Reliable Low-Latency Communications
UPB	Politehnica University of Bucharest
UPF	User Plane Function
UL	Uplink
VPN	Virtual Private Network
VPA	Vertical Pod Autoscaler
WP	Work Package
XAI	Explainable Artificial Intelligence
XGBoost	Extreme Gradient Boosting

RAN	Radio Access Network
FTDI	Future Technology Devices International
UART	Universal Asynchronous Receiver-Transmitter
UC	Use case
RUL	Remaining Useful Lifetime
SIM	Simtel
SA	Standalone
SB-EMS	Smart Building Energy Management System
S-NSSAI	Single Network Slice Selection Assistance Information
ZSM	Zero-touch network and Service Management

Executive Summary

This deliverable, D5.1 – Energy Domain Intermediate Developments, serves as the intermediate technical report of Work Package 5 (WP5) Energy Domain within the AMAZING-6G project. It consolidates the technical work conducted from M3 to M17, detailing how three energy-related use cases were defined, designed, preliminarily integrated, and evaluated through early prototype validation. D5.1 is intended to provide an implementation-oriented view of progress at the end of the intermediate engineering phase, where architectural choices and technology enablers are already sufficiently mature to be exercised in preliminary environments and linked to measurable performance criteria. It therefore covers not only the conceptual use case descriptions, but also the concrete integration steps performed, the interfaces connected across domain boundaries (field/edge/cloud/network), and the early evidence supporting readiness for the larger system demonstrations in WP7.

The scope and objectives of D5.1 are centered in three main directions. First, it aims to define and describe the three energy-related use cases, each addressing distinct operational challenges with different combinations of communication, edge processing, and data-driven models. Second, it documents the implementation and preliminary testing of relevant technology enablers, including edge/cloud computing orchestration, machine learning pipelines, federated processing approaches where applicable, and secure low-latency communication mechanisms based on Beyond Fifth-generation / Sixth-generation (B5G/6G). Finally, the deliverable contributes to preparing the foundation for full-scale integration and trial activities in WP7 by identifying which functional blocks have already reached integration readiness, and by providing initial Key Performance Indicators/Key Values Indicators (KPI/KVI) definitions that guide evaluation in upcoming demonstration phases. The expected outcomes of this intermediate phase include well-documented use case solutions, early test results from integration environments, and actionable feedback for WP2 (system requirements and architecture) and WP3 (core technology enablers), strengthening the technical base for large-scale pilots executed in the next project stages.

The deliverable also reports progress on the implementation and testing of relevant technology enablers, including edge/cloud computing orchestration, machine learning pipelines, federated data processing approaches where applicable, and secure low-latency communication mechanisms based on B5G/6G. These enablers are mapped to the specific operational needs of each use case, ensuring that evaluation can progress from component-level verification toward system-level demonstrations. Use cases introduction is presented below.

Renewable Energy Communities (Energy 1-E1) use case develops a cloud-native Smart Building Energy Management System (SB-EMS) and Renewable Energy Community (REC) extension that jointly optimizes user comfort and energy consumption/production. The system is designed to operate as a distributed application using Internet of Things (IoT) telemetry (environmental sensing and smart-meter data) together with predictive Artificial Intelligence/Machine Learning (AI/ML) functions for short/mid-term forecasting. Decision-making is implemented via configurable control loops (open recommendation or closed-loop actuation) and includes self-optimization mechanisms based on user feedback. At REC scope, buildings participate as federated learning clients, while a central aggregator coordinates federated model training and multi-building energy flow optimization.

In E1, a large-scale trial is planned to validate the system's scalability and efficiency by deploying the federated SB-EMS across a diverse network of interconnected residential and commercial buildings to demonstrate real-world energy savings within a live Renewable Energy Community.

The scope of the pre-trial and preliminary integration presented in this deliverable, focused on validating readiness of the B5G/6G-enabled infrastructure and aligning it with the SB-EMS application component model. On the network side, a private 5G Standalone (SA) environment was prepared and validated in the Orange 5G Lab premises, including isolation through a private Access Point Name (APN) and

preparation of network slicing resources for two traffic classes: energy monitoring (massive IoT-mIoT, like) and control/actuation (Ultra-Reliable Low-Latency Communications- URLLC like). On the compute side, Kubernetes-based edge/cloud continuum resources were instantiated (including building-level emulated edge nodes and centralized REC-level resources as future extension), and core SB-EMS platform components were deployed to support service orchestration, closed-loop governance, AI/ML model lifecycle, and monitoring. On the IoT side, the ThingsBoard Community Edition platform was installed and integrated with Message Queuing Telemetry Transport (MQTT)-based sensor telemetry available in the lab/campus environment.

Key enablers from the E1 platform were incorporated to ensure end-to-end operability under realistic conditions: compute-as-a-service orchestration in the device-edge-cloud continuum; AI-as-a-service (AI-aaS) and ML Operations (MLOps) for containerized ML onboarding and inference at the edge; IoT data ingestion/telemetry management; and communication enablers including network performance monitoring. Preliminary steps also integrated an intent processing chain with foundational integration to network telemetry and telemetry-driven corrective actions. Additionally, an AI-based traffic profile classification component was integrated as an early AI service to support predictive, service-aware orchestration during subsequent closed-loop trials.

Robotized Offshore Wind Turbine Blade Inspection and Maintenance (Energy 2-E2) use case addresses wind turbine Inspection, Operation and Maintenance (IO&M) by introducing robotized blade inspection supported by B5G connectivity, Digital Twin (DT) and edge computing. The goal is to reduce the time between inspection and decision-making, thereby lowering operational risk, cutting costs associated with turbine downtime, and improving safety for offshore technicians. The concept involves the use of an inspection drone equipped with non-destructive sensing (ultrasound or similar) and high-accuracy localization method to collect blade structural health measurements and stream them in near real-time to a DT system at the edge and to an onshore monitoring centre. The DT system transforms pre-processed measurement data into blade health indicators such as failure risk and Remaining Useful Lifetime (RUL). Based on these results, a decision can be made on what action to take such as do nothing, repair the blade, or even replace the blade. The pre-trial scope and preliminary integration activities focused on validating the readiness of the B5G communication and localization enabler which are prerequisite for the full end-to-end drone-to-DT workflow. Network components were initially set up in a lab environment for the pre-trial phase comprises of a B5G standalone radio access network using Amarisoft AMARI Callbox as Next-Generation Node B (gNB) and an open-source implementation of 5G Core (Open5GS 5G Core). In parallel, two Quality of Service (QoS)-related slice concepts were being prepared in the pre-trial: one dedicated slice for drone/inspection traffic and another for background traffic, demonstrating how network slicing can ensure consistent performance and service continuity under dynamic network conditions.

On the localization side, the integration effort targeted centimeter-level positioning accuracy needed for spatially-referenced structural assessment. A hybrid Global Navigation Satellite System (GNSS) and Real-Time Kinematic (RTK) approach was implemented. Instead of using a commercial RTK service, the use case opted to use its own GNSS base station to provide the RTK correction data to the rover mounted on the drone. Preliminary tests have verified that the RTK correction link was functional and a positioning accuracy in the centimeter-level has been achieved.

The final trial phase of the E2 use case is planned at the Zephyros Field Lab in Vlissingen, where several decommissioned wind turbine blades are available for research purposes. While the trial is planned to take place in a controlled indoor facility, its classification as a large-scale trial is fundamentally grounded in the use of a full-scale, decommissioned industrial blade sourced directly from an operational wind farm. Unlike lab-scale prototypes, this blade provides the actual physical dimension and real-world structural defects as found in an offshore wind farm.

Solar Energy Monitoring, Control and Predictions Using B5G/6G Communications and Edge-Cloud (Energy E3) develops and tests a smart solar energy management system that uses B5G/6G connectivity combined with edge–cloud computing to provide real-time monitoring, remote control, and short-term energy production forecasting. The trial targets a representative solar installation near Bucharest and validates an end-to-end architecture that connects field-level Photovoltaic (PV) inverters and telemetry sources to a cloud orchestration platform, while applying time-critical performance characteristics for operational control. The system is designed around a three-tier distributed IoT model: (i) industrial edge gateways that interface with inverters and sensors and buffer data during connectivity fluctuations, (ii) a B5G network with slicing-based separation of control and telemetry traffic, and (iii) a cloud/edge management layer hosting forecasting and optimization logic and exposing secure Application Programming Interfaces (APIs) to dashboards and partner systems.

The pre-trial and preliminary integration scope implemented in this deliverable focused on de-risking the core operational chain and establishing functional interoperability across field devices, connectivity, and the cloud ingestion/visualization/control plane. Integration activities covered: reliable Modbus-based telemetry acquisition and command execution between an edge gateway and a PV inverter, secure networking using 5G SA and RedCap connectivity for gateway reachability, and continuous telemetry ingestion into the platform with synchronized dashboard visualization and command acknowledgements. A key enabler validated early is AI-aaS readiness through the CAPG forecasting pipeline (implemented in the Capgemini-CAPG testbed for pre-trial), including model training/serving workflows and MLOps-style model management concepts which are planned to be integrated into the Romanian testbed in the next phase.

In the pre-trial, the system exercised the monitoring and control loop using a laboratory PV setup (one inverter with a 10-panel string) connected through an edge gateway using Modbus Remote Terminal Unit/ Transmission Control Protocol (TCP/RTU) and RedCap over 5G. The primary validation objective was the end-to-end chain: telemetry acquisition at the field, transfer to the cloud, correct mapping into a unified schema, and execution of remote-control actions with near real-time reflection in telemetry. This baseline supports the subsequent full E3 trials in H1 2027, where URLLC-like slice characteristics, automated control/forecasting workflows, and integrated AI-driven prediction and optimization are to be validated in the operational solar park environment. This use case is considered a large-scale trial, as the targeted solar park spans hundreds of square meters, integrates and manages more than 100 PV panels, and involve continuous data acquisition and analysis over several months of operation.

The results and integration progress reported in D5.1 are intended to inform the planning and execution of system level system-level demonstrations in WP7. The deliverable provides a traceable mapping between each use case’s operational workflow and the enabling technology components that were validated early (e.g., edge/cloud orchestration, slicing preparation, IoT onboarding/control interfaces, and localization/DT readiness where applicable), enabling WP7 to focus on closing end-to-end gaps rather than re-establishing foundational connectivity and data/control interoperability. In addition, the initial KPI/KVI definitions and the preliminary test outcomes support the refinement of the performance measurement methodology, clarifying what has already been validated under controlled conditions and what must be re-validated under the more representative trial configurations and workloads.

1 Introduction

The Energy Domain WP5 focuses on defining, implementing, and testing use cases that demonstrate how B5G-enabled technologies can address current and emerging challenges in the energy sector. WP5 aligns with the project's broader goals of integrating technological enablers with operational energy systems, targeting measurable benefits for both performance and sustainability. This document, D5.1, is an intermediate milestone that summarizes the developments achieved up to month 17 and provides a technical baseline for subsequent work in WP7.

This document presents a structured overview of what has been defined and built so far in WP5. The first milestone of WP5 covered the definition and design of three representative energy use cases, selected for their relevance to energy efficiency and transition goals, their ability to demonstrate technical innovations, and their capacity to be validated in realistic integration environments. After the use case definitions, work was initiated on supporting technologies and testing infrastructure. The results of these activities are organized to feed into the trial and demonstration activities planned in WP7, including the validation of KPIs and the refinement of system-level integration plans.

D5.1 is structured to support engineering traceability and the reuse of intermediate outcomes, and it follows the main sections outlined below.

Section 2 describes the use cases, including their scope, deployment approach, and technology enablers used to support the intended workflows.

Section 3 provides use case planning that captures the next stage preparation, including how testing activities evolve toward demonstrations.

Section 4 presents preliminary integration work and early results for each use case, including the current integration status, initial test outcomes, and progress relative to target KPIs. Section 4 also includes the concrete evidence base that guides integration decisions during WP7.

Section 5 concludes the report by summarizing the overall progress made and the next steps required to complete final field integration and trial execution.

A references section and annex materials are also included to support traceability and provide additional context where needed for implementation continuity and cross-work-package alignment.

2 Use Case Definition and Design

This section describes the definition and design of the three Energy use cases developed in WP5. It covers not only the goals and scope of each use case but also their implementation schemes, key components, and how they interact with the 5G/B5G infrastructure and other technical enablers.

Each use case addresses important challenges in the energy sector by applying new B5G technologies. The use cases focus on different areas: collaborative energy management in renewable energy communities, robotized inspection and maintenance of offshore wind turbines, and solar energy monitoring and control using smart edge devices.

Defining and designing the use cases in detail is the first step to enable the start of their implementation. By clearly outlining the technical solutions and integration with communication networks and enablers, this section sets the foundation needed to move forward with development and testing activities.

2.1 Renewable Energy Communities (Energy 1-E1)

The E1 use case proposes a cloud-native Smart Building Energy Management System (SB-EMS) solution enabled by B5G/6G private networks for joint comfort and energy optimization in smart buildings and Renewable Energy Communities (RECs). The SB-EMS performs real-time monitoring, prediction and user-centric optimization of ambient conditions, energy consumption and generation from Renewable Energy Sources (RES), leveraging IoT platforms integrated with edge/cloud continuum, AI-as-a-Service (AI-aaS) and MLOps with support for Federated Learning (FL) mechanisms, combined with B5G/6G private network connectivity.

Each smart building operates as a self-contained intelligent energy node with local IoT sensing and actuation, integrating edge computing resources organized in Kubernetes clusters for local AI/ML inference and optional training. At the REC level, buildings participate as federated learning clients, while a central aggregator coordinates global model training and multi-building energy flow optimisation. The E1 testing area, hosted within the Orange 5G Lab premises (Figure 1), provides a controlled environment for validating the cloud-native SB-EMS enabled by B5G/6G private networks, integrating IoT sensing, edge/cloud AI processing, and federated learning for joint energy and comfort optimization in smart buildings and REC.

Figure 1. E1 testing areas in Orange 5G/6G laboratory premises.

2.1.1 Description of the use case

The E1 use case implements an AI-driven and user-centric Energy Management System (EMS) for smart buildings (SB-EMS) and Renewable Energy Communities (REC-EMS) for collaborative and scalable optimization of energy production and consumption in B5G/6G-enabled smart cities and smart districts. The SB-EMS application performs a joint analysis of multiple variables describing environmental conditions, building layouts, and user preferences, applying trustworthy AI techniques for decision making to identify the optimal compromise between energy savings and user comfort, with the capability to scale from single buildings to REC scope (see Figure 2). The proposed solution applies user-centric paradigms, adopting configurable open or closed control loops depending on the user preferences. The system can either issue well-explained suggestions on best practices and recommendations on smart building configurations (e.g., Heating, Ventilation and Air Conditioning -HVAC or lighting settings) asking the user to manually accept or reject them, or automatically apply the most convenient configurations limiting the interactions with the user. Moreover, Reinforcement Learning (RL) [1] techniques for self-optimization and self-tuning based on user feedback allow to tailor the behaviour of the system to different categories of buildings (smart homes, smart offices, smart campuses, hotels, etc.) and user profiles, without requiring complex settings or customizations.

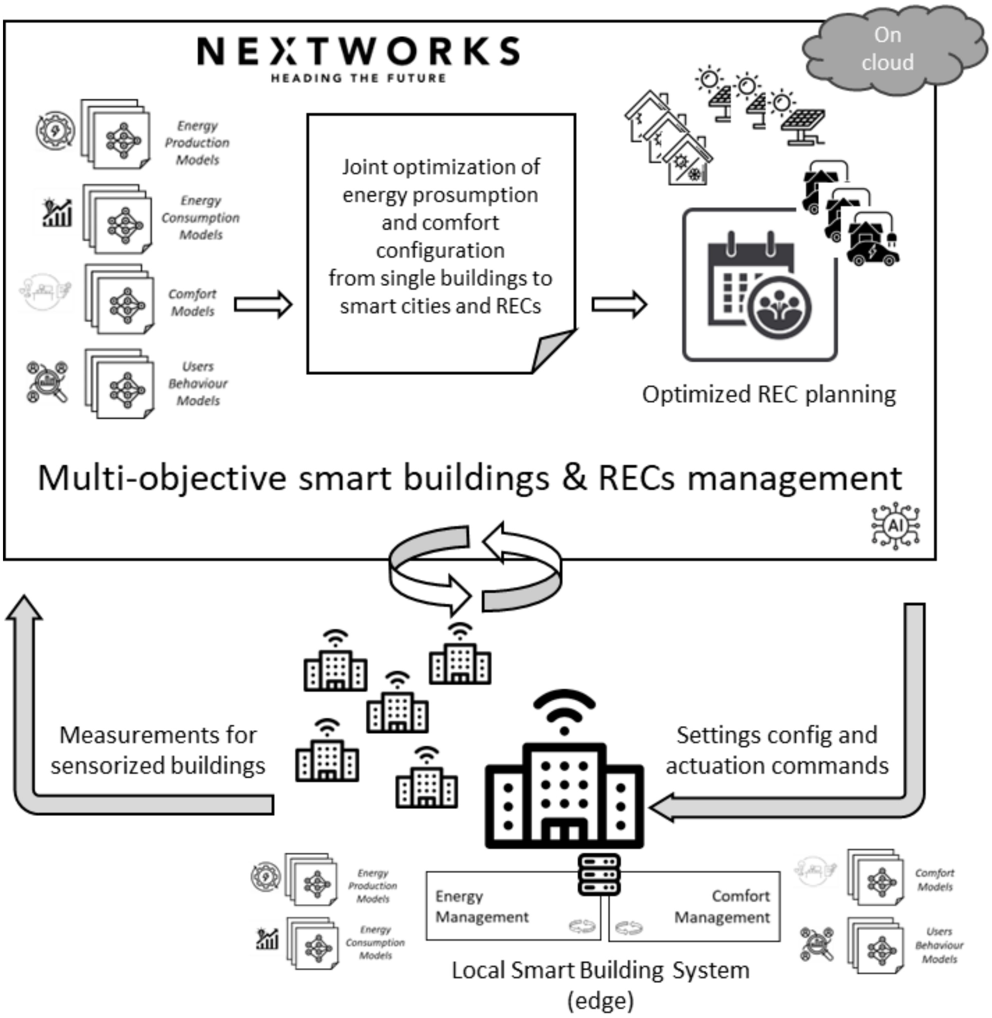

Figure 2. E1 Use case concept.

From the technical perspective, the E1 use case introduces several innovations, not only at the application level, but also in terms of more features enabled by the integration with B5G/6G technical enablers developed in WP3. The main innovations in E1 can be summarized as follows:

- Global energy consumption optimization, scaling from single buildings to RECs, considering renewable energy sources and energy storage capabilities, making use of cloud-native distributed applications distributed in the edge/cloud continuum;
- Self-tuning energy management system capable to adapt itself to different types of building (smart homes, smart offices, smart hotels) by applying RL techniques and requiring minimal manual configuration or customization;
- Joint optimization of energy savings and user comfort, to achieve optimal energy consumption while improving the living conditions. This is enabled through the combined processing of heterogeneous datasets, collected from pervasive IoT devices (energy meters, environmental sensors), also used to derive users' preferences (e.g., rooms occupancy, preferred HVAC settings), as well as external services (e.g., weather forecasts);
- AI/ML models for joint energy consumption and user comfort optimization including Explainable Artificial Intelligence (XAI) [2] and FL techniques for efficiency and privacy in scalable and collaborative scenarios;
- Intent-based service deployment and configuration, with integrated monitoring and configuration of network slices in B5G/6G infrastructures, for optimized traffic flows related to the collection of IoT sensing data and the actuation of commands on smart building devices;
- IoT, monitoring and MLOps platforms integrated with device/edge/cloud continuum and B5G/6G networks, with end-to-end and zero-touch service and resource orchestration.

The SB-EMS and REC-EMS solutions have been designed following an integrated approach that combines together the vertical application components related to IoT data processing, AI/ML models for user and energy consumption/production profiles, and user-friendly interactions with the technical enablers offered by E1 platform and foreseen as key features of future B5G/6G private networks. This methodology leads to the design of a "6G-native" SB-EMS/REC-EMS solution, resulting on one hand in a vertical application ready to fully exploit the capabilities of future 6G infrastructures and, on the other hand, in the possibility to propose a bundled offer for advanced smart buildings and RECs that integrates in a single product an easy-to-extend platform with private B5G/6G and edge capabilities with smart IoT devices and EMS applications. This integrated system jointly addresses several technical challenges at both infrastructure and vertical application level:

- Coordinated management of IoT platforms controlling multi-vendor IoT devices and distributed computing resources in the device-edge continuum through a single platform for 6G-enabled smart buildings;
- Unified monitoring and control of heterogeneous IoT sensors and actuators guaranteeing interoperability and application portability across multiple buildings with minimal customizations;
- Optimization of placement strategies for AI/ML functions in single buildings or among multiple buildings (at the REC level), with joint optimization of traffic flows and slice configuration, for scalability and adaptability to scenarios where buildings have different computing capabilities;
- Self-learning and self-tuning trade-off strategies for joint optimization of energy consumption and user comfort, with accurate interpretation of users' intents and preferences, for a solution that is capable to adapt itself to the preferences of different profiles of users and buildings;
- Scalability, data privacy, and trustworthy AI when scaling from single buildings to RECs, fundamental to comply with European Union (EU) regulations on data privacy and AI.

The E1 use case is designed around two scenarios: energy and comfort optimization in single buildings for the first scenario, addressed in the first part of the project, which is then extended to the full REC scope for the second scenario, developed in the second part of the project. The baseline technical enablers adopted in the two scenario are the same (orchestration of virtual computing resources, integration between IoT platforms and extreme edge resource orchestration, network slicing for IoT traffic, MLOps and AI-aaS for automated lifecycle management of ML models), but with customizations specialized to the requirements of each of them. In particular, the REC scenario extends the first one introducing further logic to handle FL [3] for collaborative training among smart buildings belonging to the same REC and management of resources in the continuum edge/cloud.

The first scenario is focused on the concept of a SB-EMS that collects real-time data from heterogeneous and multi-vendor IoT sensors connected via a private B5G/6G network installed in the building, also equipped with an IoT platform for the unified management of different IoT devices, both sensors and actuators. The buildings are represented by the lab in ORO premises in Bucharest), already equipped with a variety of sensors exposing data via Message Queuing Telemetry Transport (MQTT). The data collected include indoor and outdoor environmental data (temperature, humidity, air quality), information on the presence of people in the 5G Lab room, the use of appliances and devices, together with the related power consumption retrieved via smart meters.

These datasets are used to feed different ML models designed for profiling the behaviour and the preferences related to the usage of the smart building and predict energy consumption and production. The first type of model aims to build per-building profiles related to the use of specific devices or rooms on daily, weekly, and monthly basis, as well as related preferences on rooms' settings, e.g., temperature, humidity, use of natural vs artificial light, etc. The models can be either initially trained in the cloud using public datasets available for similar categories of buildings or trained locally on the real data using edge nodes on premises for data confidentiality (personal data are not collected in any case). The models for the prediction of energy consumption and generation are mostly using data on weather conditions, geographical area, building position and layout. They were initially based on public datasets or even pre-trained ML models, with provision for runtime validation and retraining in the event of model drift.

The storyline for the single building scenario can be summarized as follows. Data are collected at the smart building level through the IoT platform, and here they are aggregated, unified, and anonymized as needed to properly feed the AI/ML pipelines. The initial training phase creates the ML models for profiling building usage and preferences. Then, using the MLOps platform, the models are deployed at the edge for the inference tasks, to predict the short/mid-term evolution of factors like rooms occupancy, comfort preferences, energy consumption and RES generation. This inference constitutes the analysis phase of a closed loop that follows with decisions on local per-building optimization (see Figure 3). The decision takes as input the predictions elaborated at the analysis phase and finds the best compromise between energy and comfort, according to objectives defined by policies through configurable weights (e.g., maximization of comfort level, minimization of energy consumption or cost, maximization of use of locally generated - or stored – renewable energy). The closed loop proceeds with the execution phase where, depending on the system configuration, the SB-EMS can either actuate autonomous configurations on the smart building (e.g., adjusting HVAC, lighting, and blinding settings) or give suggestions to the users via mobile applications, enriched with explanations on the motivation of the recommendation or the potential benefits that could be gained.

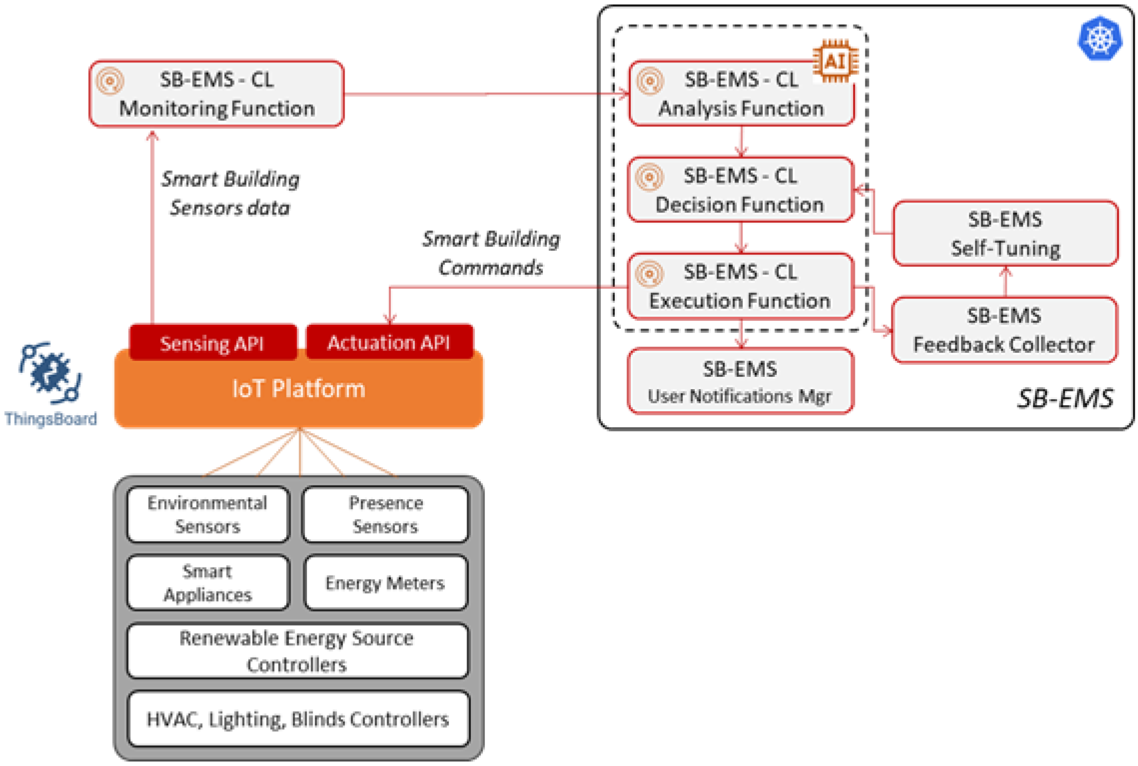

Figure 3. High-level architecture of E1 SB-EMS.

In parallel to the continuous execution of the closed loop for the comfort/energy optimization, the SB-EMS gathers feedback from the users, either explicit (acceptance/rejection of suggestions) or implicit (manual override of building settings automatically configured by the SB-EMS). This approach allows the system to automatically improve the learning of the users' preferences, for self-adaptation and self-tuning, e.g., adjusting the weights assigned to the different objectives of the decision function. Moreover, comparisons between predicted and actual data are continuously performed to detect possible drifts at the ML models' level and trigger re-training procedures if needed. The mechanisms for self-tuning and re-training are planned to be implemented and integrated at the beginning of the second half of the project (M19-M24).

The REC scenario extends the SB-EMS concepts to the entire REC level, thus introducing additional challenges in terms of data confidentiality and scalability of the solution. These challenges are addressed introducing FL to build more accurate models through the collaboration of different buildings participating in the REC and the adoption of resource allocation algorithms specialized for edge/cloud continuum, capable to migrate tasks between edge and cloud and to select target buildings for AI/ML functions based on their computing and sensing capabilities. In particular, the use of computing resources for AI/ML tasks (both training and inference) is handled via the AI-aaS mechanisms offered by MLOps Platform, considering edge nodes capabilities and related energy sources of the buildings to be selected as FL clients in the FL groups, together with their energy consumption constraints. Furthermore, the distribution of training tasks among edge/cloud nodes is designed to account for aspects related to data confidentiality and required level of trustworthiness.

The storyline for the REC scenario is quite similar to the single building scenario, assuming a replication of the same workflow in each building and the adoption of FL at the training phase with an FL aggregator operating at the REC level. Moreover, additional tasks are performed in parallel. At the monitoring phase, data on shared RES production at collected at the REC level together with aggregated data on per-building energy consumption and RES generation. These data feed an additional closed loop operating at the REC level for global multi-building optimization. This results in a multitude of closed loops working in parallel, some of them with per-building scope and other with per-REC scope, all collaborating together under hierarchical coordination mechanisms for conflict management. This approach is also designed to enrich the per-building optimization address in the first scenario with new decisions at the REC-level, e.g., for scheduling of energy storage, energy exchange, etc.

The intent-based service deployment and network slice optimization framework represents a key cross-scenario extension to the B5G/6G infrastructure. It provides a high-level abstraction layer where application and user intents, expressed in terms of reliability, latency constraints, or guaranteed data rates, are automatically translated into network-level configurations and energy-resource service deployment actions. Through this mechanism, users, control systems, and decision-making engines can dynamically request and adapt QoS profiles, subscriber configurations, and resource allocations. For example, intents associated with real-time sensor streams may trigger the provisioning of low-latency, high-reliability slices, while periodic energy-metering traffic can be mapped onto more cost-efficient connectivity profiles. As illustrated in Figure 4, the Intent-Based Networking (IBN) system is integrated with the testbed, interfacing with the ORO core APIs, network telemetry mechanisms and with service orchestrator to take and trigger decisions on scaling based on AI workloads.

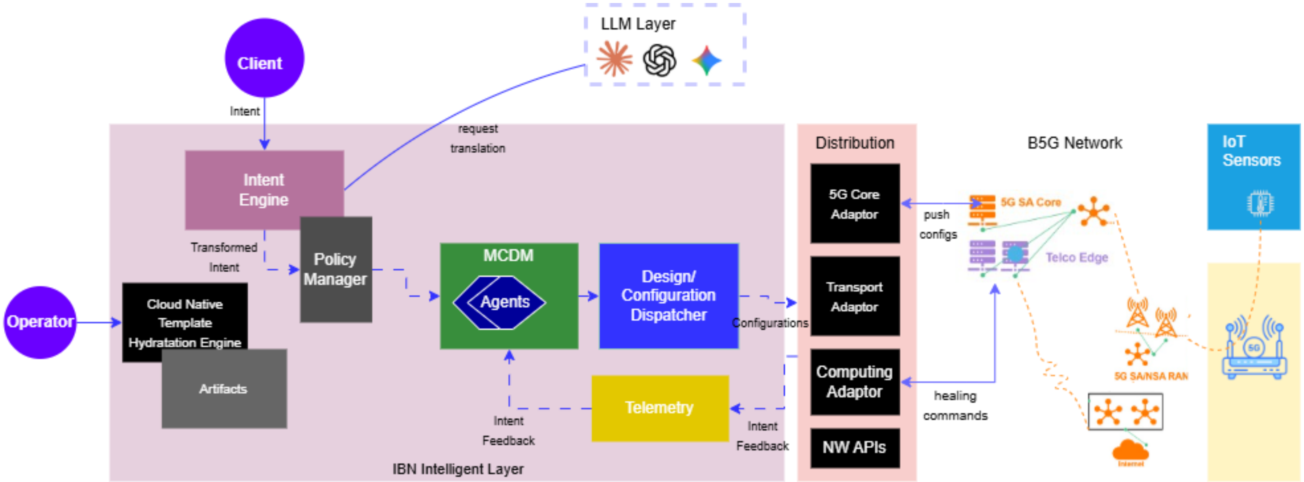

Figure 4. Intent-Based Networking Framework E2E integration within the testbed.

The framework incorporates Multi-Criteria Decision Making (MCDM) algorithms combined with GenAI-based multi-agent orchestration for slice recommendation and QoS optimization. It is composed of three main modules: (a) Intent Processing module, where the operator or user express intents in natural language. A GenAI-driven reasoning layer interprets these intents, extracts contextual information, and identifies domain-specific requirements. The intents are validated, normalized and translated into actionable network operations, such as subscriber configuration, slice instantiation, modification of QoS parameters or resource provisioning requests or resource provisioning requests. (b) the decision matrix module estimates the expected performance of available network slices or configurations using historical KPI measurements. It aggregates metrics such as latency, packet loss, jitter availability, computational load and energy footprint to build a dynamic performance matrix that reflects current network conditions. (c) MCDM algorithms module evaluate, and rank candidate slices or configurations based on the requirements extracted from the intent and the performance matrix. The outcome is a recommended slice or configuration profile that best satisfies the request while optimizing network efficiency to support the IoT traffic. Additionally, the system integrates a intelligent monitoring and anomaly detection modules that continuously analyse the network metrics.

2.1.2 Deployment and implementation details

2.1.2.1 Design of the application

The system architecture design for the single building scenario of E1 use case is represented in Figure . The figure shows the SB-EMS application components, represented in the light grey boxes and running in Kubernetes containers deployed in the edge nodes at the buildings, and the infrastructure and platform components implementing 6G technological enablers developed in E1.

The infrastructure level, represented with orange boxes, includes computing resources in edge nodes, extreme edge nodes or devices (e.g., CPEs, IoT gateways) and a private B5G/6G network that interconnects the smart building sensors and actuators. All the computing nodes can host containers to deploy the application software and they are organized in a Kubernetes cluster. The IoT devices are handled through an IoT platform, exposing sensing APIs (to read data from the sensors) and actuation APIs (to issue configuration commands to the smart building).

The E1 platform components, represented with blue boxes, provide mechanisms for orchestrating the end-to-end SB-EMS service and the computing resources in the device/extreme edge/edge continuum, for monitoring the system and all the related resources, and to manage the execution of the SB-EMS application in terms of closed loop functionalities and lifecycle of their ML models.

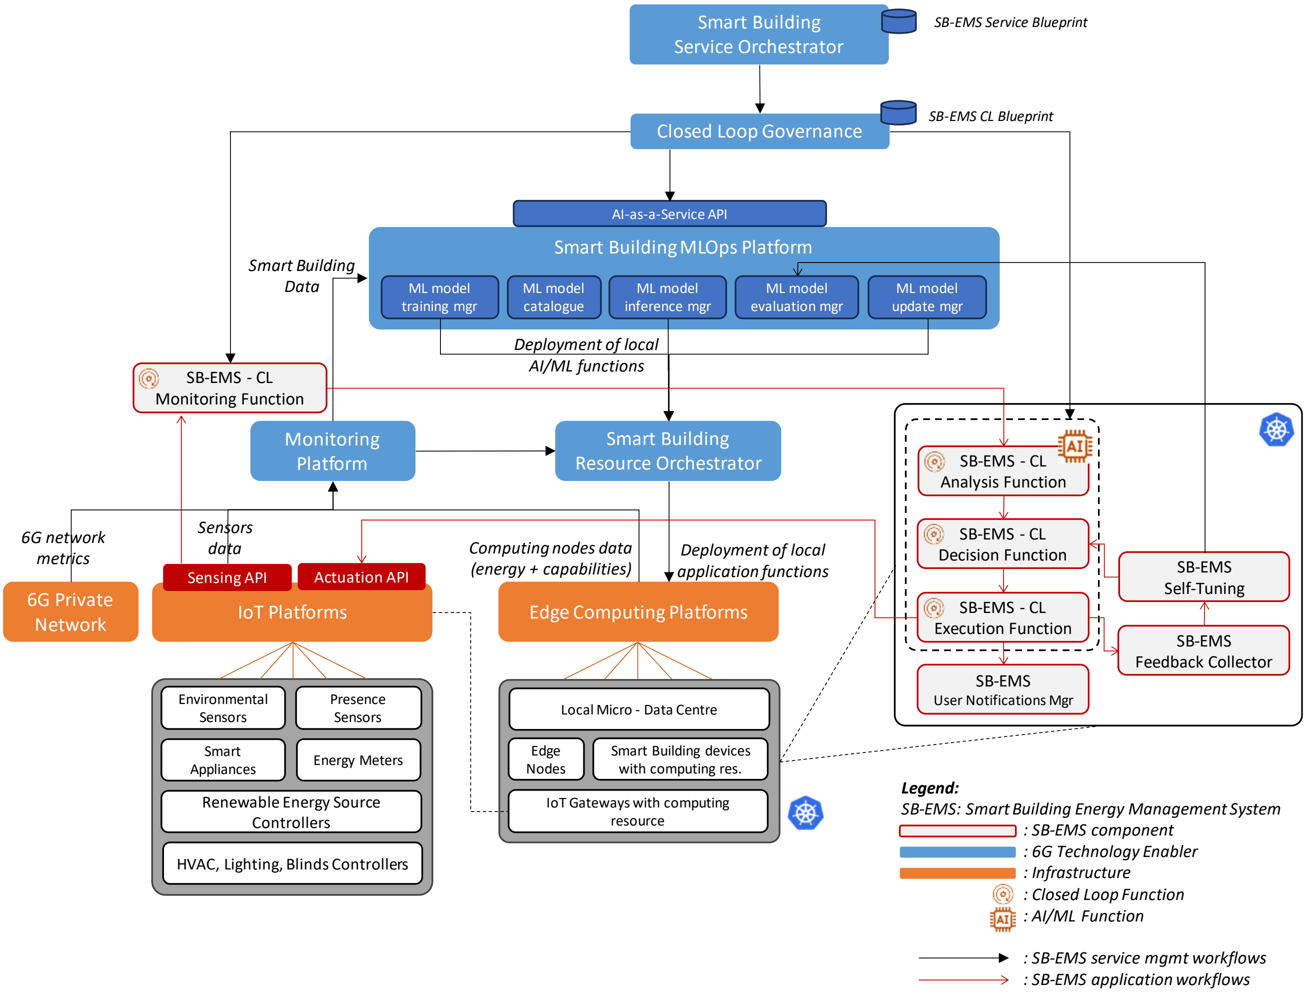

Figure 5. SB-EMS system architecture.

The platform components have been defined in D3.1 and they can be summarized as follows. The Resource Orchestrator is specialized to handle the computing resources in the local edge and extreme edge nodes, considering their capabilities, the interconnection to the private network, and the connectivity to the IoT devices attached to the IoT platform. The Service Orchestrator manages the workflows for the provisioning, termination, and runtime operation of the SB-EMS service, interfacing with the rest of the platform components. In particular, the instantiation and management of the Closed Loop (CL) functions of the SB-EMS application are delegated to the Closed Loop Governance component. The CL functions relying on AI tasks are managed through the MLOps Platform that offers an AI-aaS interface to select ML models from its catalogue and deploy the functions for the inference tasks. These models are used in the analysis CL function for prediction of energy consumption and production, rooms occupancy or related comfort preferences. The MLOps Platform is also used offline for training tasks and, at runtime, for model tuning or re-training to achieve better accuracy. The decisions on service function placement, including ML workloads, are handled through a resource allocation engine that is synchronized with the Resource Orchestrator for the real-time availability of computing resources in the continuum and the characteristics of the nodes, including their mobile connectivity and their association with IoT devices. Real-time information on usage of computing resources, network performance, and status of IoT devices is collected through a Monitoring Platform, with drivers to interact with Kubernetes, the private network telemetry system in Orange lab, and the IoT platform. All these platform components are implemented by Nextworks in the context of WP3, as customization and extension of software assets developed in previous research projects (mainly Smart Networks and Services - SNS Stream B Adroit-6G project).

The IoT platform, instead, is based on the open-source ThingsBoard[1] - Community Edition (Apache 2.0 license). The interface with the IoT sensors deployed in the buildings in Bucharest (Orange lab) is based on the MQTT protocol. The integration between ThingsBoard and the Monitoring Platform, where aggregated IoT data are stored and pre-processed, is based on the REST-based query APIs exposed by ThingsBoard, while commands are issued through the ThingsBoard Remote Procedure Call (RPC) API.

The system architecture design for the REC-EMS scenario is the same for the deployment at each smart building. However, the MLOps platform introduces the support of Federated Learning management and additional application components are deployed on the cloud for the REC management, as shown in Figure .

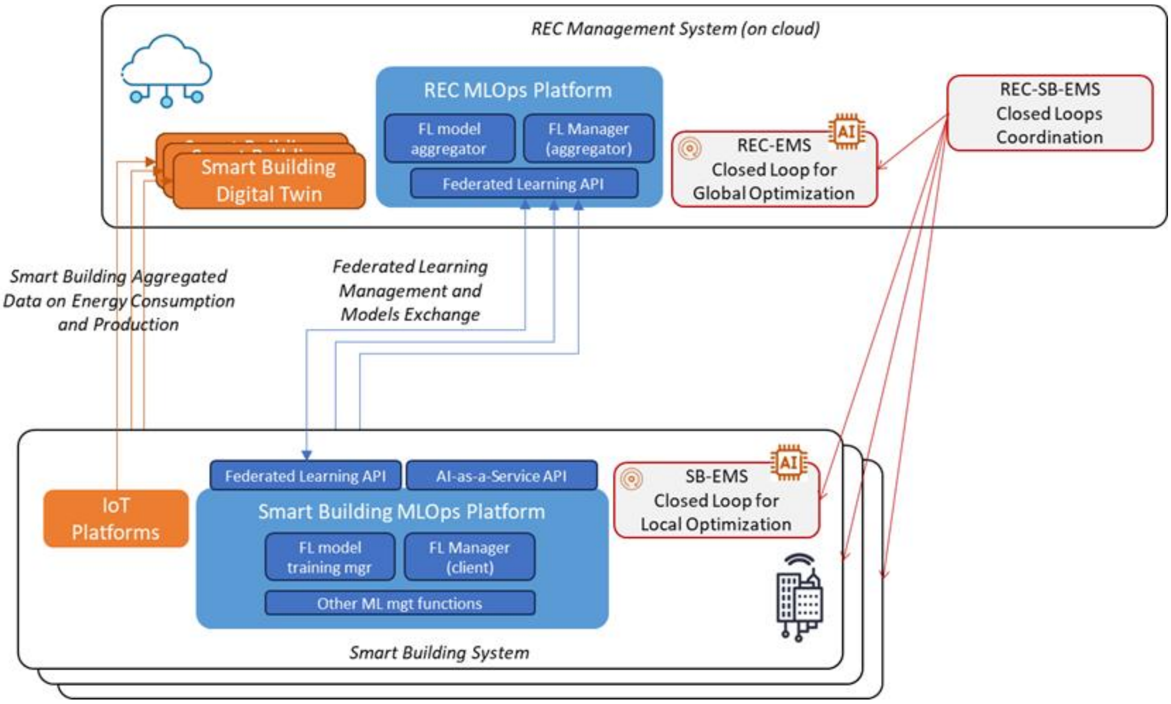

Figure 6. E1 – REC-EMS system architecture.

The MLOps Platform at the buildings handles the management of FL clients, while a centralized MLOps Platform instance in the cloud coordinates the FL server management and the procedures for the creation of FL groups, the selection of FL clients, and the aggregation of FL models, working with REC

[1] https://thingsboard.io/

scope. The interaction between MLOps platforms uses FL APIs, enabling advertisement of FL client's capabilities, federation management, exchange of FL models trained by clients and redistribution of aggregated models.

Figure 7 depicts the end-to-end interaction between the IBN framework and the underlying 5G infrastructure, illustrating how high-level user intents are automatically translated into concrete network actions such as slice configuration and computing plane such as AI workloads management. At the top the cognitive layer handles intent parsing, policy enforcement and Ai-driven decision making, and multi-criteria slice selection. The resulting configuration actions are passed through core adaptor and computing with triggers the corresponding configurations of slice at the core and or at the computing plane the scaling up or down of the AI workloads. At infrastructure level the framework interacts directly with 5G Core UDR API, with active URLLC and mIoT slices. A feedback loop from the infrastructure back to the intent engine closes the control cycle enabling continuous and autonomous adaptation. This architecture enables the automation of the network configuration, where the IoT sensor traffic and energy monitoring applications require dynamic and policy-driven slice management.

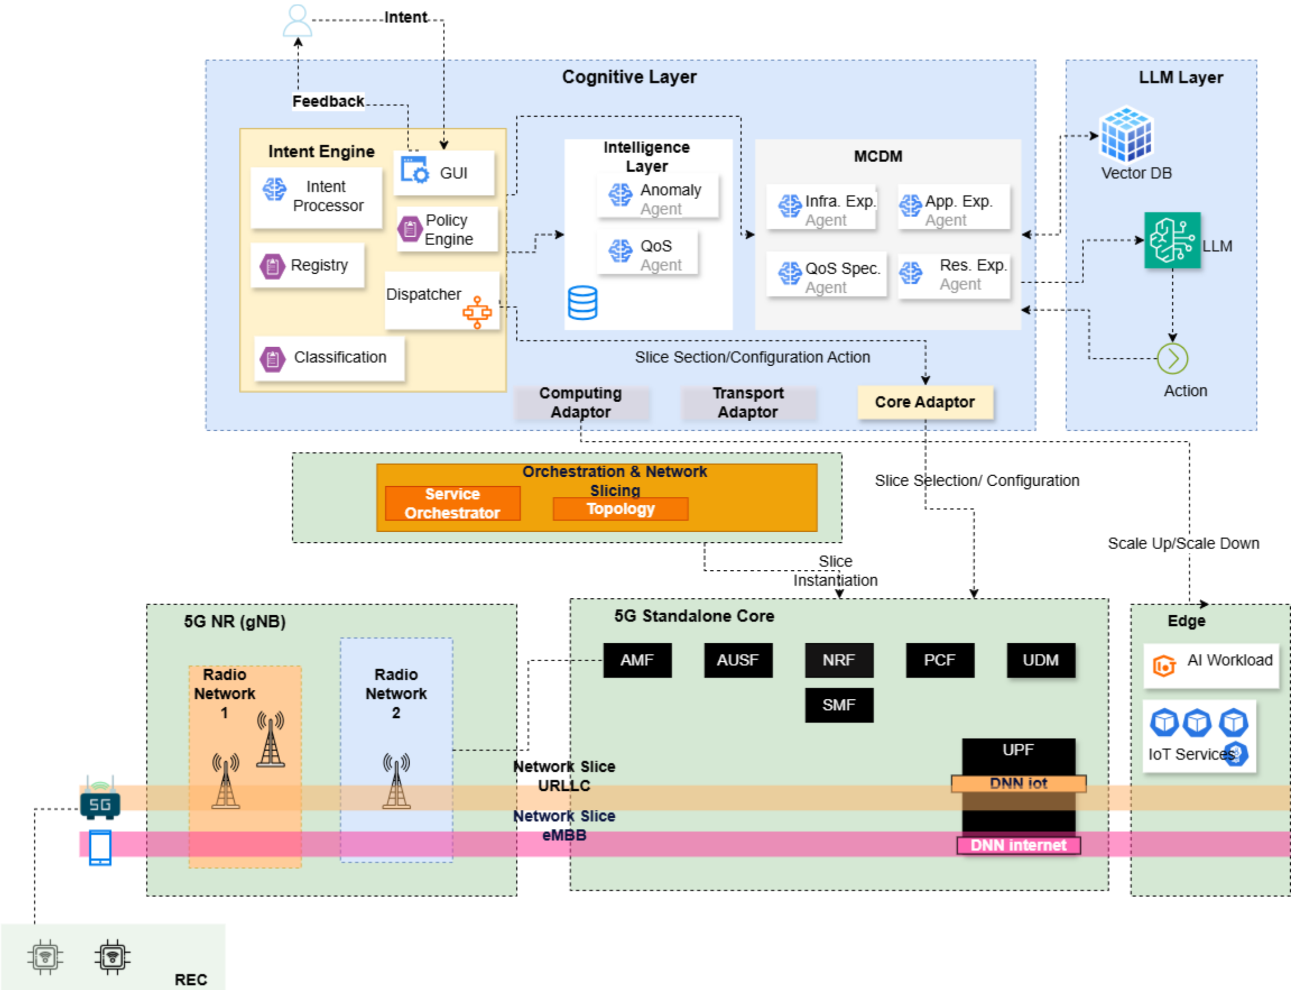

Figure 7. E2E interaction between the IBN framework and underlying infrastructure.

The framework processes each intent through a sequential pipeline, illustrated end-to-end in Figure .

The Natural Language Processing (NLP) processor acts as the entry point, performing semantic parsing to extract key service parameters, latency targets, reliability levels, throughput requirements, time windows, and service classification, and producing a structured intermediate representation regardless of how the intent was originally phrased. In the trial context, inputs range from explicit operator commands to automated triggers generated by the traffic pattern classifier.

The structured intent is forwarded to the Large Language Model (LLM) agent, which performs deeper semantic reasoning to produce a normalized intent object containing QoS parameters, service type classification, and contextual constraints. The agent decomposes the intent into the expected network configuration goals and constraints that guide downstream decision-making.

Before any provisioning action is taken, the normalized intent is validated by the Policy Engine against operator-defined policies, regulatory constraints, and current infrastructure capabilities. Rejected intents are returned with a structured explanation and written to the audit log — a critical guardrail in autonomous provisioning scenarios.

Once validated, the intent is passed to the MCDM engine, which selects the optimal execution strategy by scoring candidate slice configurations and resource allocations against a dynamic performance matrix of real-time and historical KPIs. The deliberation is structured around four specialised LLM-based agents, a network infrastructure expert, a QoS specialist, an application expert, and a resource efficiency analyst, whose collective output is an optimised orchestration plan specifying slice selection, compute placement, resource allocation, and policy enforcement actions.

These agents are supplied by the AI-aaS platform, which acts as the horizontal intelligence layer providing reusable models for anomaly detection, predictive analysis, and optimisation. A multi-agent orchestration layer operates continuously in the background, feeding observations back into the MCDM cycle and triggering re-evaluation whenever deviations from the desired network state are detected.

The IBN framework integrates with the ORO testbed, which exposes URLLC and Enhanced Mobile Broadband (eMBB) slices to support the different traffic types generated by IoT devices, energy management services, and AI workloads. Based on the MCDM output, the IBN selects the most appropriate slice via the ORO API and coordinates the selection of Kubernetes scaling strategies for AI workloads in parallel.

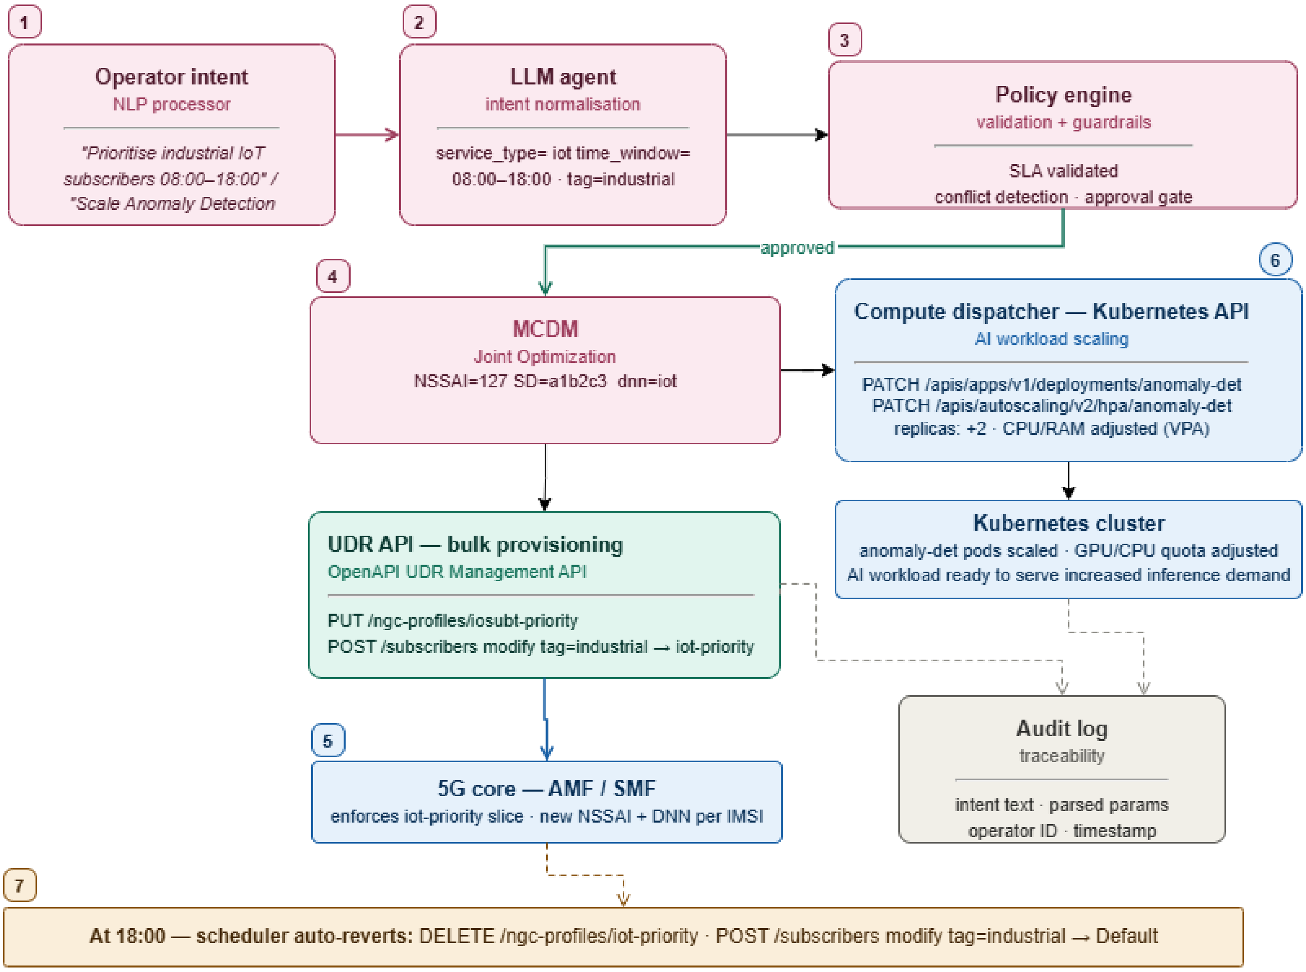

Figure 8. Scenario Intent Operator NLP.

Slice decisions are enforced through the Nokia Unified Data Management (UDR) Management API. Next Generation Core (NGC) profiles — defining Single Network Slice Selection Assistance Information (S-NSSAI), Deep Neural Network (DNN), and Access/Session Management (AM/SM) policy parameters — are created or updated via PUT */udr/ngc-profiles/{profId},* and subscriber bindings are adjusted via PUT */udr/subscribers/{subscriberId}*. Bulk reprovisioning across multiple subscribers is handled via POST */udr/subscribers* in a single API call.

A key capability, of the integration is the use of Radio Access Network (RAN) metrics and Customer Premises Equipment (CPE)-side data, illustrated in Figure 9, to predict traffic patterns and trigger pre-emptive slice provisioning. An AI model continuously analyses 60-second snapshots of throughput measurements collected from CPE devices via TR-069, complemented by RAN-side KPIs, and classifies the current traffic profile as either normal or active. The moment an active profile is detected, the system automatically selects the appropriate NGC profile, updates it via the UDR API, and binds the subscriber International Mobile Subscriber Identity (IMSI) before user demand fully materializes, ensuring the correct slice is in place ahead of demand and avoiding the session degradation of a reactive, post-peak response.

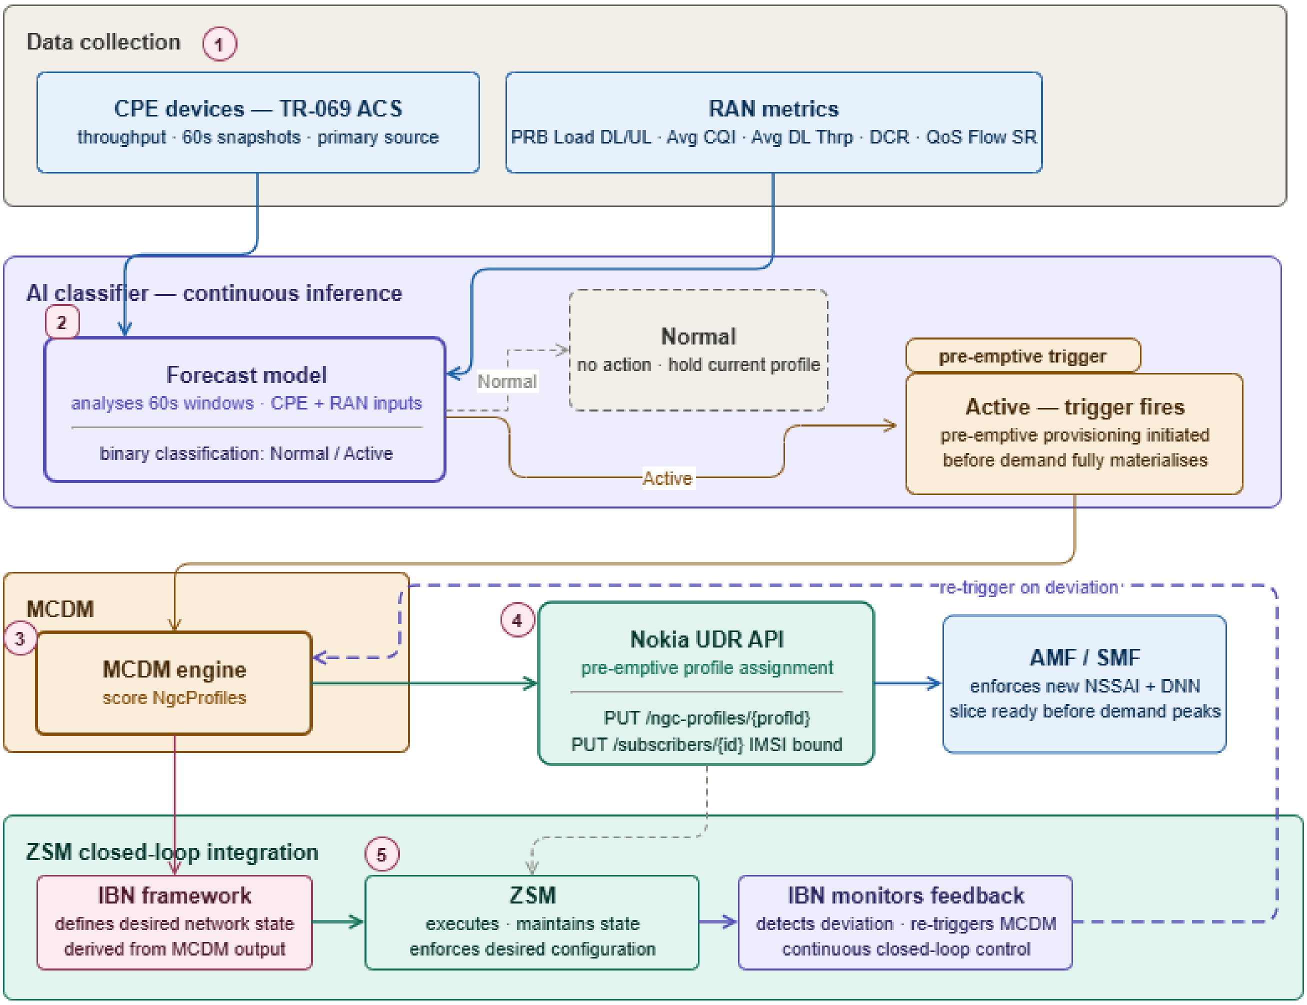

Figure 9. Scenario Forecast Model Based in CPE side and RAN data.

The framework integrates with Zero-touch network and Service Management (ZSM) through a clean separation of responsibilities: the IBN defines the desired network state derived from the MCDM output, while ZSM executes and maintains it. The IBN monitors the feedback continuously, re-triggering the MCDM decision cycle whenever deviations between the intended and actual network state are detected.

In the initial phase, development focuses on the core intent processing pipeline — NLP parsing, LLM normalization, policy validation, and MCDM-driven profile selection — and its integration with the Nokia UDR API. The AI-aaS platform is deployed with the first set of trained models, including the traffic pattern classifier trained on collected CPE datasets. By Q2, the first functional prototype was operational, capable of interpreting natural language intents, selecting the optimal NGC profile via MCDM scoring.

In the second phase, full closed-loop control is planned to be introduced through continuous RAN metric monitoring, enabling automatic corrective re-provisioning without operator intervention. The ML model catalogue is expanded through the AI-aaS platform to cover predictive slice selection, anomaly detection, and resource optimisation. The Kubernetes Horizontal Pod Autoscaler/ Vertical Pod Autoscaler (HPA/VPA) scaling extension is integrated into the dispatcher layer, enabling joint optimisation of network slice assignments and compute resource allocations within a single MCDM decision cycle.

2.1.2.2 Devices, sensors and supporting equipment -ORO &NXT (virtual devices)

Table 1 presents the list of devices used for the deployment of the E1 use case, considering both IoT devices and 5G network devices.

Table 1. Devices used in E1 deployment.

Category	Component	Technical Details	Units
5G Connectivity Module	Quectel RM500Q-AE [5] 	Industrial-grade 5G module enabling wireless connectivity for the robots. Supports 5G NR, LTE-A Pro, LTE-A, and WCDMA. Download speed up to 2.5 Gbps and upload up to 660 Mbps. Integrated GNSS (GPS, GLONASS, BeiDou, Galileo) for precise positioning. Operates in harsh environments (-40°C to +85°C, shock and vibration resistant). Designed for smart transportation, IIoT, video surveillance, and robotic applications.	2
IoT Sensor	Industrial power meters [6] 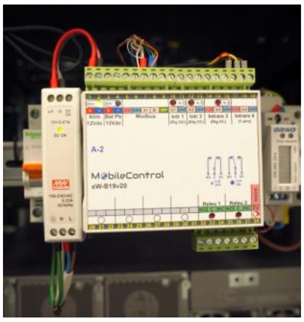	Used for measurement of power consumption for 5G Core Network, 5G gNodeB, IPFABRIC L2/L3 networking, firewall, edge computing.	2
IoT Sensor	Industrial temperature sensors [7] 	Used for measurement of air temperature around the 5G Core Network, 5G gNodeB, IPFABRIC L2/L3 networking, firewall, edge computing.	4
IoT Sensor	Air Quality Sensors [8] 	Sensors measuring the following metrics: Particulate Matter (PM2.5, PM1, PM10), Ozone, Formaldehyde, Carbon Dioxide, Volatile Organic Compounds (VOC), temperature, barometric pressure, air humidity and noise	4

2.1.3 Network infrastructure and functional components

The network infrastructure designed for the E1 and E3 use cases represents a fully integrated B5G implementation that extends ORO's experimental and commercial 5G SA environments toward advanced energy management applications.

Figure illustrates an E2E Network configuration tailored for a Smart Campus scenario. This setup leverages a 5G SA architecture with a dedicated edge computing infrastructure to achieve ultra-low latency and dedicated computing resources [9]. The Smart Campus is covered with both indoor and outdoor radio units, integrating sensors and devices while benefiting from the secure, isolated, and performance-guaranteed network slicing.

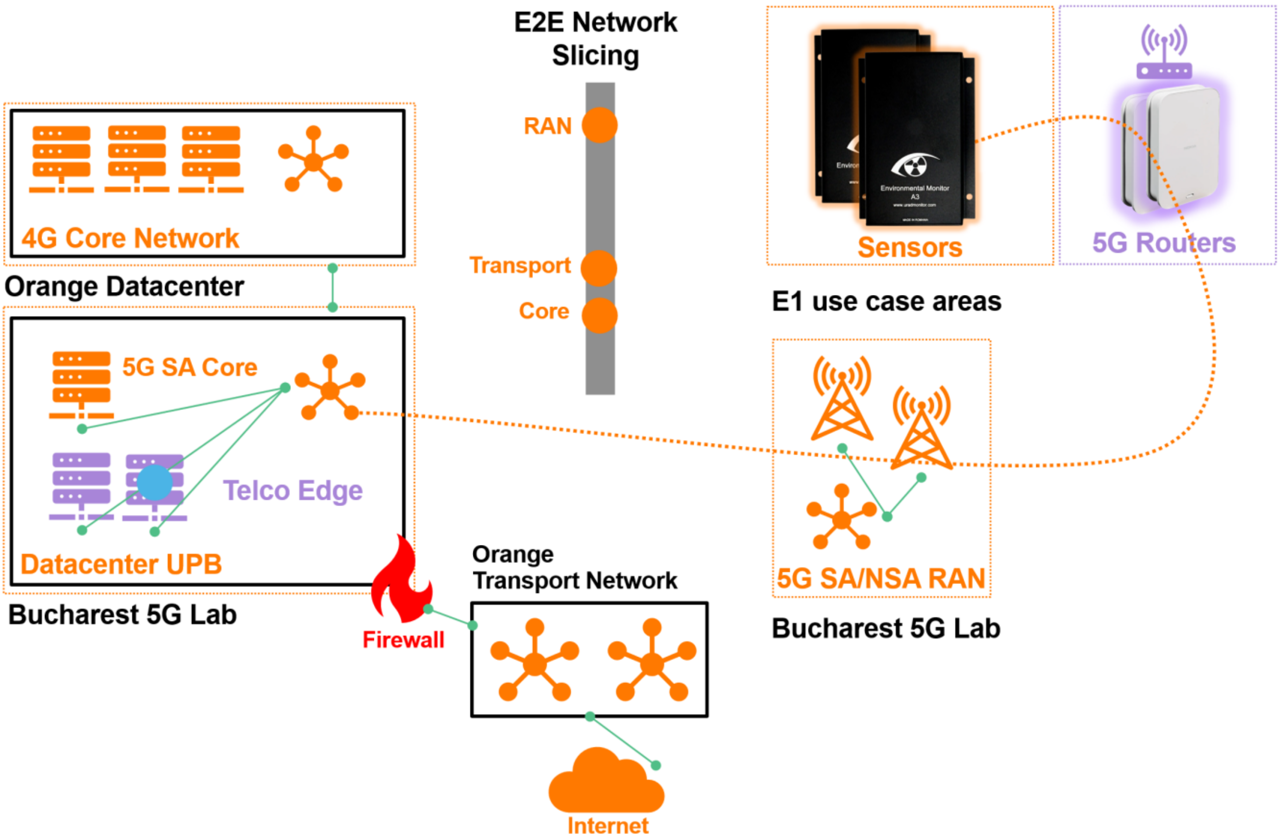

Figure 10. B5G Architecture for Smart Campus implementation in E1.

The architecture relies on a distributed setup that interconnects the Bucharest 5G Lab with multiple locations in Romania (see Figure 11, showcasing the system expandability. Together, these form a resilient end-to-end platform for energy monitoring, renewable energy optimization, and real-time analytics across geographically separated domains. These sites serve as key validation nodes within the broader European 6G ecosystem, adhering to the architectural principles for vertical industries established by the Smart Networks and Services Joint Undertaking (SNS JU) [10].

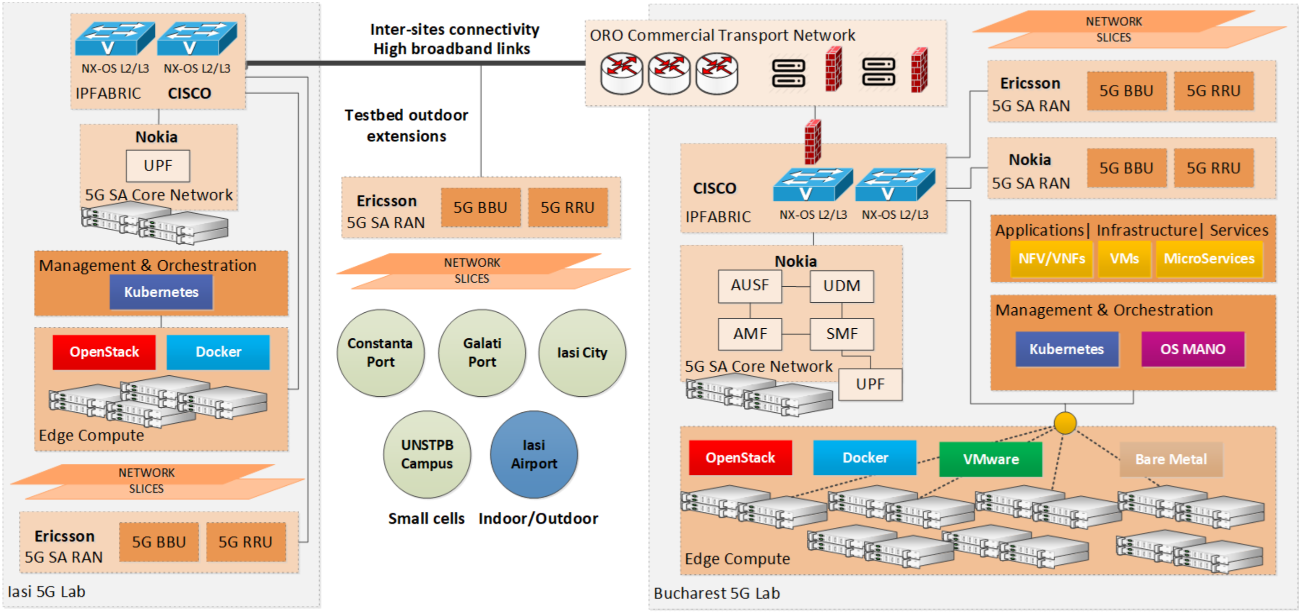

Figure 11. Orange 5G Labs Architecture.

At the core of the infrastructure is the 5G Core Network (5GC) deployed on a virtualized cloud infrastructure hosted in the Bucharest 5G Lab datacenter. The 5GC follows (see Figure 12) a service-based architecture that integrates the complete 3GPP suite of network functions, including the Access and Mobility Management Function (AMF), the Session Management Function (SMF), the User Plane Function (UPF), and the Authentication Server Function (AUSF). Together with the Unified Data Management (UDM) and Policy Control Function (PCF), these elements enable dynamic session management, policy-based traffic prioritization, and low-latency routing across multiple service domains. The Bucharest Lab also integrates the Authentication and Policy Control (APC) cluster and the Cloud Mobility Manager (CMM), which provide centralized control over session lifecycles.

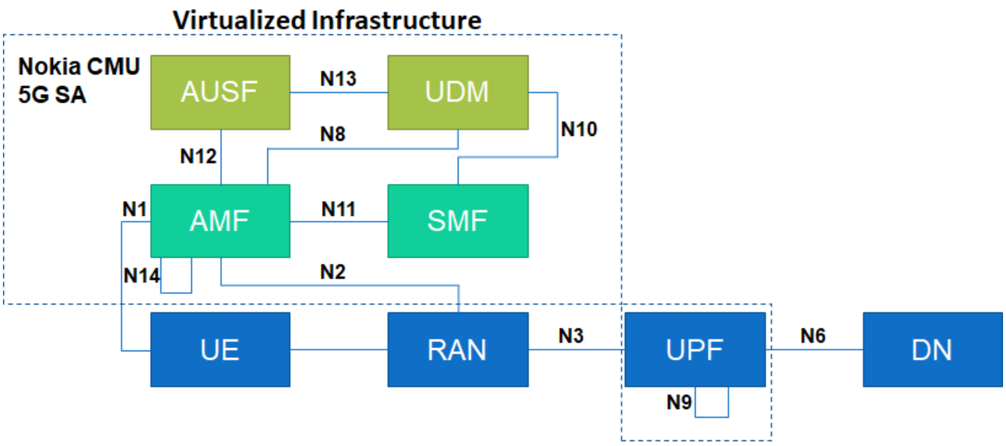

Figure 12. 5G SA Implementation in the Romanian Testbed.

The radio access layer combines both indoor and outdoor 5G SA and NSA deployments, using the n78 band (3.5 GHz) for enhanced capacity and coverage. The Bucharest 5G Lab operates a Ericsson-based 5G SA RAN which extends to both indoor and outdoor environments.

The Smart Campus test area also connects, through a Private APN, to the ORO public network and hosts a local radio unit that supports both indoor energy monitoring sensors and outdoor energy management gateways.

Network slicing (see Figure 13) is a central element of the Romanian testbed. Dedicated slices are foreseen for energy monitoring, energy control, and analytics-intensive applications, each to be characterized by its own QoS parameters and routing rules.

The eMBB slice provides high-bandwidth connectivity for data collection, visualization, and cloud synchronization of large datasets from solar farms and building management systems. The URLLC slice, by contrast, is optimized for ultra-reliable, low-latency communications, suitable for real-time control or automated switching.

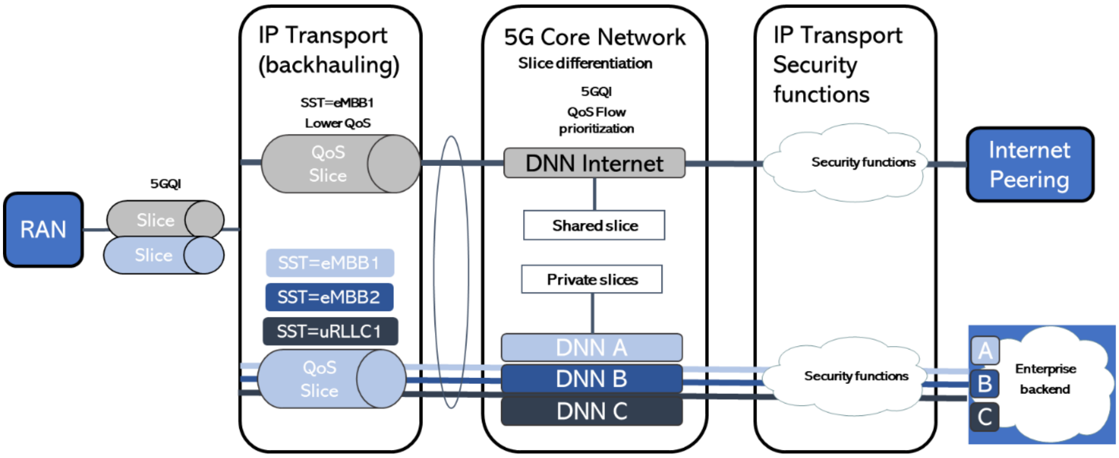

Figure 13. Romanian Testbed Slicing Architecture.

The Orange Datacenter provides the backbone connectivity and long-term data retention for both E1 and E3 use cases. It hosts the application onboarding platform initially derived from the VITAL-5G framework [11] and extended with new modules for energy management. This platform enables service instantiation, resource allocation, and monitoring of active deployments through a user-friendly graphical interface.

Applications such as demand-response coordination, photovoltaic performance analytics, and distributed energy optimization can be automatically deployed on either central cloud resources or edge computing nodes, depending on the latency and reliability requirements.

The integration of the Smart Campus within this ecosystem allows for testing of intelligent energy management systems under realistic conditions. E1 focuses on the optimization of energy consumption across multiple buildings by combining IoT sensor data, predictive control algorithms, and real-time actuation through 5G URLLC communication.

E3 (see Section 2.3) extends this concept toward distributed renewable energy monitoring, where IoT and sensors continuously measure solar production, inverter efficiency, and grid injection levels. The collected data is processed locally at the edge for rapid response and simultaneously transmitted to the datacenter for historical trend analysis and forecasting. This dual-level processing ensures both operational continuity and predictive energy balancing.

2.1.4 Use case enablers

As introduced in deliverable D3.1, E1 enablers are presented in Table 2.

Table 2 List of enablers and sub-enablers for E1.

Category	Technology enablers	Sub-enablers
Communication enablers	Communication resource management	Zero-touch network and service management

Communication enablers	Communication resource management	Network performance monitoring and control
Communication enablers	Network slicing	
Compute-as-a-Service enablers	Compute resource management	Intelligent service and resource orchestration in extreme edge/edge/cloud compute continuum
Compute-as-a-Service enablers	Compute resource management	Intent-based service and resource management
Compute-as-a-Service enablers	Compute continuum	Distributed user-edge-cloud compute continuum
Application enablers and AI	AI-as-a-Service framework	AI-as-a-Service
Application enablers and AI	AI-as-a-Service framework	ML models catalogue
Application enablers and AI	AI-as-a-Service framework	MLOps
IoT and localization enablers	Data Ingestion and Telemetry	Data Collection from sensors and devices
IoT and localization enablers	Data Ingestion and Telemetry	Monitoring Platform
IoT and localization enablers	IoT Service Platforms and Management	IoT Sensors Orchestration
IoT and localization enablers	IoT Service Platforms and Management	Service and Resource Lifecycle Manager
IoT and localization enablers	IoT Service Platforms and Management	Message queue fabric
IoT and localization enablers	IoT Service Platforms and Management	IoT Exposure Function

In E1 use case, ZSM [4] is adopted as a supervisory and automation layer operating on top of the existing infrastructure and management stack provided by ORO. ZSM complements the underlying orchestration and control mechanisms, by introducing intent-based automation and control-loop service assurance, while respecting the configuration constraints imposed by the pre-existing infrastructure. High level intents expressed by users, or applications are interpreted and mapped into feasible network configurations. The ZSM is designed to handle the requirements and enforce policies, such as enforcing real-time, ultra-reliable connectivity, during peak periods, or consumption periods. Also, when a demand response event is activated, an intent is placed to prioritize the service's subscribers and the ZSM temporarily configures the network slices for the specific subscribers supporting demand response signaling.

Close-loop control mechanisms are responsible for continuously monitoring the realization of user intents by collecting relevant metrics from the underlying infrastructure. These measurements are evaluated against the expected service level objectives. When deviations, degradations, or anomalies are detected, an alarm is triggered, and the control layer automatically triggers corrective actions.

Network slicing in the use case operates based on pre-defined slices already instantiated by the underlying infrastructure, focusing on intelligent slice selection and optimization. User or operator intents are expressed in natural language requirements and translate them into network operations such as slice selection, subscriber management and association, and QoS parameter adjustment. MCDM algorithms rank candidate slices and configuration profiles to identify an option that satisfies the intent while optimizing network efficiency.

In context of networking ML catalogue provides a structured repository for AI/ML models supporting ZSM and service optimization functions. ML catalogue includes models for KPI prediction, anomaly detection and intent validation support.

At the infrastructure level, Compute-as-a-Service enablers facilitate intelligent service and resource orchestration across a distributed extreme edge/edge/cloud compute continuum. This is supported by an intent-based management framework that optimizes service deployment and network slice configurations to handle specific traffic profiles. Application and AI enablers provide an AI-aaS framework, which includes an automated MLOps platform and a centralized ML models catalogue. These tools manage the full lifecycle of the system's predictive models, including occupancy inference, thermal dynamics, and power generation, enabling zero-touch optimization and the implementation of federated learning for scalable, multi-building renewable energy communities

2.1.5 KPIs definition

Table 3 presents the set of KPIs defined for measurement and validated during the final trial phase of the E1 use case, planned for Q1–Q2 2027. The KPIs are defined to assess the performance of the B5G/6G private network supporting the SB-EMS under an indoor coverage scenario, focusing on latency, throughput, availability, reliability, coverage, and connection density. KPI measurements are envisaged to be obtained using a combination of network-native monitoring functions, traffic generation and probing tools at the user and application layers, and application-level logging integrated within the SB-EMS platform.

Table 3. E1 defined KPIs.

KPI name	Description/KPI definition	Scenario	KPI target
Reliability of actuation commands (KPI_VA_1)	% of actuation commands successfully executed		>99.999%
End-to-end latency for actuation commands (KPI_VA_3)	Time between command sent and its execution		<20 ms
End-to-end latency for sensors data (UL) (KPI_VA_4)	Time between sensor data sent and its storage in the monitoring platform		<50 ms
Uplink throughput (KPI_VA_11)	Data rate required to support IoT traffic		>10 Mbps
Energy consumption of edge computing infrastructure (KPI_VA_16)	Power consumed by edge nodes in the building during runtime operations		<20% energy savings at building
Service provisioning latency (KPI_VA_5)	Time for service provisioning from request to service up and running		<30 s
Inference speed (KPI_AI_1)	Latency incurred in the computation of an AI model in inference mode	CAPG Testbed	<2 s

ML accuracy (KPI_AI_2)	Accuracy of ML classification model	CAPG Testbed	>0.9

2.1.6 KVIs definition

Table 4 presents the set of KVIs defined for assessment during the final trial phase of the E1 use case, planned for Q2 2027, with the objective of capturing the perceived value and impact of the SB-EMS solution beyond pure technical performance. The KVIs focus on stakeholder perceptions related to energy consumption reduction, user comfort, trust in AI-driven decision making, scalability to multi-building and REC scenarios, and expected operational cost savings. During the trial execution, relevant stakeholders, including building users, operators, and technical partners, are envisaged to attend live trial demonstrations and actively participate in the evaluation process by completing structured questionnaires. The questionnaires are designed using standardized Likert-scale metrics and tailored to different stakeholder profiles, ensuring consistency and traceability of the collected feedback. The qualitative and quantitative responses are to be aggregated and analysed to derive the KVIs reported in the table, enabling a direct correlation between the technical behavior of the SB-EMS, its integration with the B5G/6G infrastructure, and the perceived benefits at building and REC level. In addition, the KVIs collectively support the assessment of environmental sustainability, by capturing stakeholder-perceived reductions in energy consumption, improved efficiency, and optimized resource usage, contributing to lower carbon footprint at both building and REC levels.

Table 4. E1 defined KVIs.

KVI name	Description/KVI definition	KVI target
Reduction in Energy Consumption	Measures stakeholder perception of reduced overall energy consumption achieved through AI-driven optimization at building and REC level.	At least 70% of respondents rate perceived energy consumption reduction as “good” or “very good” (≥4 on a 5-point Likert scale.
User Comfort and Quality of Life Satisfaction	Evaluates user perception of indoor comfort (thermal comfort, air quality, lighting) and its impact on well-being while energy optimization is active.	An average score of ≥4.0/5 on comfort and well-being questions, with ≥75% of users reporting maintained or improved comfort levels.
Trust in Automated and AI-driven Energy Decisions	Measures user and operator trust in AI-based recommendations and automated control actions (including explainability and predictability of system behavior).	An average score of ≥4.0/5 on trust, transparency, and explainability questions, with ≥70% of respondents willing to rely on automated or semi-automated control modes.
Confidence in Scalability to Multi-building and REC Scenarios	Measures stakeholder confidence that the SB-EMS solution can scale from single buildings to multiple buildings and REC-level operation with limited additional effort.	At least 75% of respondents rate scalability and replicability as “good” or “very good” (≥4/5)
Perceived Reduction in Energy-related Opex	Captures operator perception of operational cost savings resulting from reduced energy expenditure and more efficient building energy management.	At least 60% of operators report perceived Opex reduction corresponding to ≥10% savings (≥4/5)

2.1.7 Description of the pre-trial and trial setup

Table 5 summarizes the information about the trial setup for E1 use case.

Table 5. E1 pre-trial and trial setup.

Trial ID & Name		Trial 12/E1 - Renewable Energy Communities
Trial Infrastructure / Venue	ORO 5G lab in Bucharest	
Trial Description	Joint optimization of energy consumption and user comfort in (i) single smart buildings and (ii) REC scenarios.	
Technical Setup	**System Components**	• Platform components: Service Orchestrator with Closed Loop Governance; Resource Orchestrator working on Kubernetes clusters; MLOps Platform; AI-aaS Plaform; Monitoring Platform; IoT Platform. • SB-EMS application functions (deployed as containers on Kubernetes). • Intent-based system for Network Slicing • MCDM for network slicing and resource management (AI Workloads) • IoT sensors. • 5G private network.
	System Configuration	• Platform components deployed as containers on VMs in ORO datacenter in Bucharest. • IoT platform controlling IoT devices via 5G private network. • 5G network with UDR reachable and subscribers pre-configured with default network slices • SB-EMS deployed on Kubernetes worker nodes running in VMs in ORO datacenter in Bucharest (emulating edge nodes for the buildings). • IBN for 5G network slice selection based on NLP intents integrated with ORO APIs • MCDM system for slice decision making integrated with the intent-based system
Pre-Trial execution (initial phase, included in this deliverable, section 4.1)	**Pre-conditions**	• IoT devices are configured to expose HTTP REST endpoints for telemetry data retrieval. • The IoT Platform is fully operational and hosted on a dedicated Kubernetes cluster. • The Monitoring Platform is deployed and active to ensure system-wide visibility. • Intent pipeline deployed and smoke-tested (with LLM locally deployed) • ORO 5G Core UDR reachable with base profiles
	Trial Steps	*5G Connectivity & Network Slicing* • Connect devices to 5G CPEs; validate via ping + iperf • Verify subscriber registration in 5G core • Validate RAN metrics access via BigQuery *IoT Telemetry & Monitoring Pipeline* • Validate device connectivity and payload schemas via Postman • Configure IoT dashboard for device mapping and real-time metrics • Set up monitoring plugins to persist measurements into InfluxDB *Intent-Driven Intelligence Layer* (in CAPG testbed)

		• Deploy LLM in lab; validate intent parsing on sample commands; Run IBN pipeline • Train pattern classifier on collected datasets • Document model performance and scoring
	Methodology	Network KPIs: • UL data rate for IoT sensors traffic (from telemetry data of 5G private network in ORO lab) • End-to-end latency for IoT sensors data and actuation commands (from timestamps in logs) Service KPIs: • Service provisioning time (from orchestrator logs) Resource usage (average, peak) at edge nodes (from Kubernetes APIs)
	Technical Enablers	Communication Enablers: 5G network with pre-configured network slices. IBN framework for high-level SLA definition and MCDM engine for slice (at Lab) IoT and Localization enablers: • Data Ingestion and Telemetry: • Monitoring Platform Compute enablers: AI-aaS
	Expected Outcomes	Validation of the end-to-end data pipeline, from IoT sensor ingestion via Hypertext Transfer Protocol Representational State Transfer (HTTP REST) to real-time visualization on the IoT Platform dashboards.
Trial Execution (final phase, to be included in final deliverable D5.2)	**Pre-conditions**	• PV system operational • Sensors calibrated and synchronized • 5G/B5G network coverage active at the site • Edge–cloud integration fully validated • Monitoring dashboard linked to the edge pipeline • Intent-Based Networking pipeline fully tested and integrated with 5G APIs • Intelligence layer with prediction agents and AI workload management integrated • Subscribers bound to default profile in UDR
	Trial Steps	• Deployment Baseline: provision NW slices and deploy SB-EMS via orchestration on K8s worker nodes • Collect KPIs (latency, THP, resource usage) and IoT data via 5G private network. • SB-EMS operation and validation: Validate SB-EMS energy consumption visualization, and personalized comfort recommendations; measure power consumption reduction. • AI-Driven optimization: run IBN pipeline for traffic-driven slice optimization with real-time classification and MCDM; SLA monitoring and joint optimization of network slicing with AI workload scaling via k8s and AI-aaS; • Monitoring and Evaluation: Monitor dashboards (SB-EMS actions, slice decision); evaluate resource efficiency and user comfort impact. • Document findings
	Methodology	Network KPIs: • UL data rate for IoT sensors traffic (from telemetry data of 5G private network in ORO lab) • End-to-end latency for IoT sensors data and actuation commands (from timestamps in logs)

		Application KPIs: • Acceptance rate for reconfiguration suggestions (from SB-EMS application logs) Comfort satisfaction index indicating > 85% of satisfied users (from questionnaires to trial users)
	Complementary Measurements	Service KPIs: • Service provisioning time (from orchestrator logs) • Resource usage (average, peak) at edge nodes (from Kubernetes APIs)
	Calculation Process	Collect logs, traces.
Technical Enablers	**Communication Enablers**	• *Communication enablers*: Full B5G network slicing (eMBB and URLLC), zero-touch network and service management, continuous network performance monitoring, QoS enforcement, and prioritized traffic handling for telemetry and control. IBN framework for high-level intent for network slicing. Policy engine. MCDM for slicing ponderation. AI/ML models for network performance prediction.
	Compute-as-a-Service (CaaS) Enablers	• IBN framework for high-level intents, and scaling mechanisms for AI workloads.
	Application enablers and AI	• AI-aaS Platform • MLOps • ML models catalogue
	IoT and Localization Enablers	• Monitoring Platform • IoT Sensors Orchestration • Service and Resource Lifecycle Manager • Message queue fabric • IoT Exposure Function • IoT Sensors
Expected Outcomes	• Reduction of power consumption > 10% compared to the power consumption without SB-EMS, with > 85% users satisfied with comfort conditions.	

2.2 Robotized offshore wind turbine blade inspection and maintenance (Energy 2- E2)

The E2 use case aims to address the challenges in wind turbine Inspection, Operation and Maintenance (IO&M). The offshore wind sector is developing very rapidly, while the current IO&M of wind turbine blades is still suboptimal and requires significant efforts of technicians under difficult and dangerous offshore conditions. Furthermore, a future projected shortage of technicians may pose a challenge to the ambitions for offshore wind energy. Consequently, the sector is moving towards robotized IO&M, where drones equipped with various sensors are used to assess the structural health of the turbine blades. This development has the potential to lower risks (less people offshore), increase sustainability (less boat transfers), and lower costs (less downtime).

The E2 use case is planned be tested and demonstrated at the Zephyros Field Lab in Vlissingen (as shown in Figure 14), where several decommissioned wind turbine blades are located and regularly used for research projects.

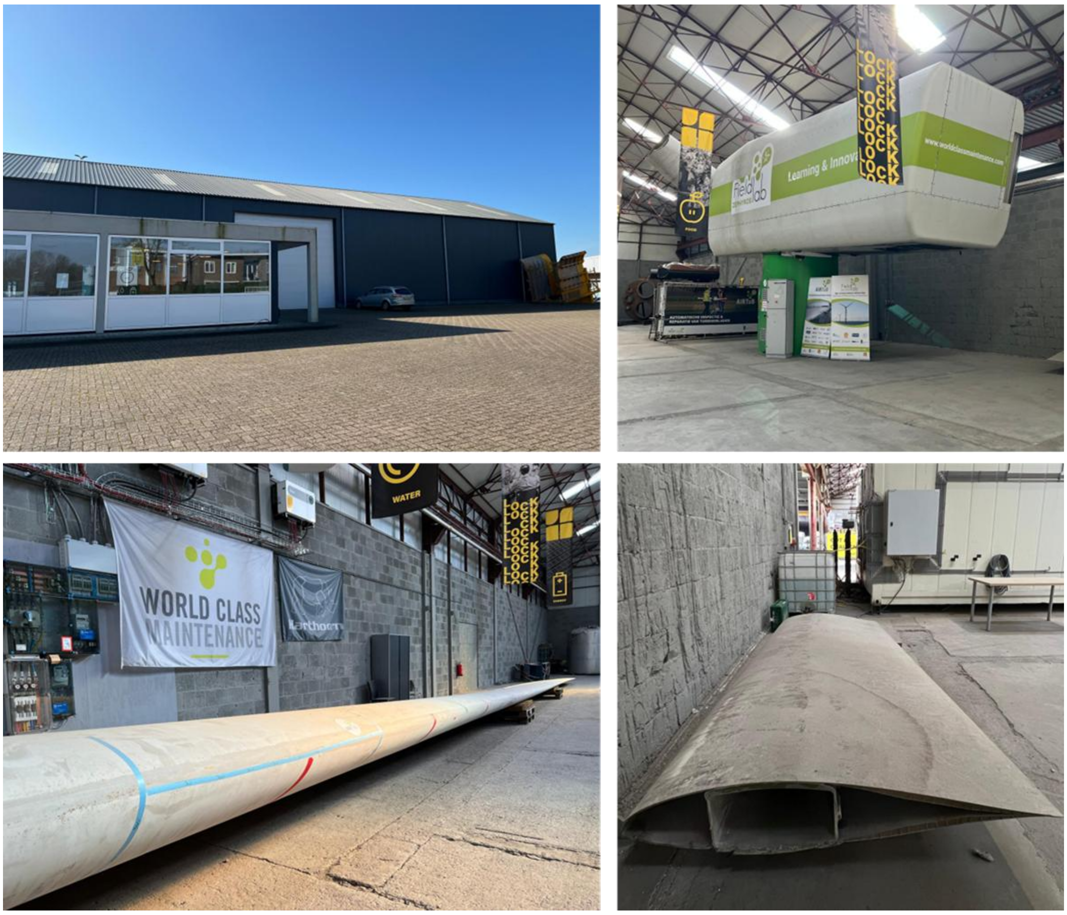

Figure 14. E2 testbed located at the Zephyros Field Lab in Vlissingen, The Netherlands.

2.2.1 Description of the use case

The E2 use case demonstrates how B5G connectivity and edge computing can significantly improve the wind turbine IO&M. Wind turbines operate continuously in harsh environments such as salt spray, rain, hail, dust, ultraviolet exposure, lightning, and high cyclic loads. Over time this leads to a range of degradations and faults that reduce energy yield, raise loads on components, and increase safety risks if left unaddressed. Typical blade defects, as shown in Figure 15 include leading-edge erosion, which increases surface roughness and aerodynamic drag, as well as gelcoat and paint cracking, skin-to-core debonding, and laminate delamination, particularly around adhesive joints and trailing edges. Additional issues involve lightning strike damage and bond line defects that may propagate under fatigue, along with structural problems such as root and spar cap degradation, skin cracks, and core failures.

Using the B5G connectivity, the information collected by the drone during an inspection session from ultrasound or similar sensor measurements can be transferred in real-time to the onshore monitoring centre. The aim is to reduce the time between inspection and decision-making and provide inspection data to onshore control groups faster than currently possible. The data gathered can be analysed with the help of DT and numerical models, which can result in a decision on what action to take: do nothing, repair the blade, or even replace the blade.

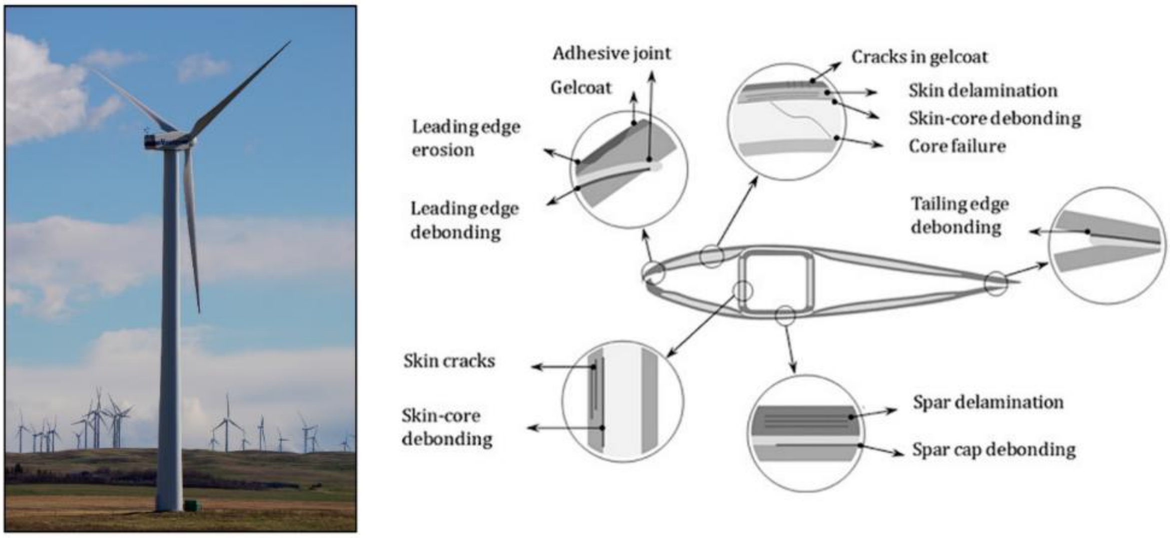

Figure 15. Types of structural faults in a wind turbine blade.

The drone envisaged for use in the E2 use case is a drone with an onboard sensor capable of performing non-destructive inspections and deliver measurements to the network. The original plan was to use the drone belonging to the Automated Inspection and Repair of wind Turbine Blades – Resident Offshore Monitoring and Inspection (AIRTuB-ROMI) project [12]. However, delays in the AIRTub-ROMI project made it no longer feasible to rely on this asset within AMAZING-6G timeline. Currently TNO is investigating the possibility to use another drone, namely DJI Matrice M600. The envisioned drone (as shown in Figure) is equipped with, among other things, a sensor (ultrasound or similar), a communication module for B5G connectivity and localization hardware to allow for remote sensing with high positioning accuracy. The B5G module is foreseen to provide new possibilities of data transfer towards the so called "DT system", as well as to the onshore control support. The collected data are intended to be used by the DT system to estimate the wind turbine blade health and potentially predict its upcoming risk of failure. The DT system is planned to take the form of a numerical support tool aimed at processing information obtained by the measurements for improved decision making. This information is designed to be made available within the time required, at maximum, for one drone flight inspection, enabling near real-time decision making from onshore.

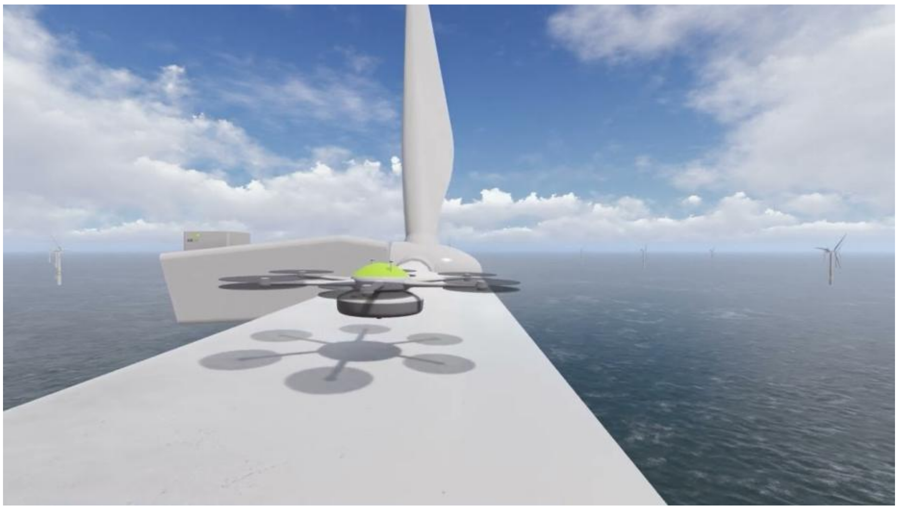

Figure 16. Illustration of the envisioned drone to be used for the turbine blade measurement.

2.2.2 Deployment and implementation details

2.2.2.1 Design of the application

The high-level application architecture is shown in Figure 17. The drone lands on a wind turbine blade and performs ultrasound measurement. In addition, a positioning module installed on the drone allows the drone to determine its position during the measurement. The drone is connected to the B5G network via a B5G modem.

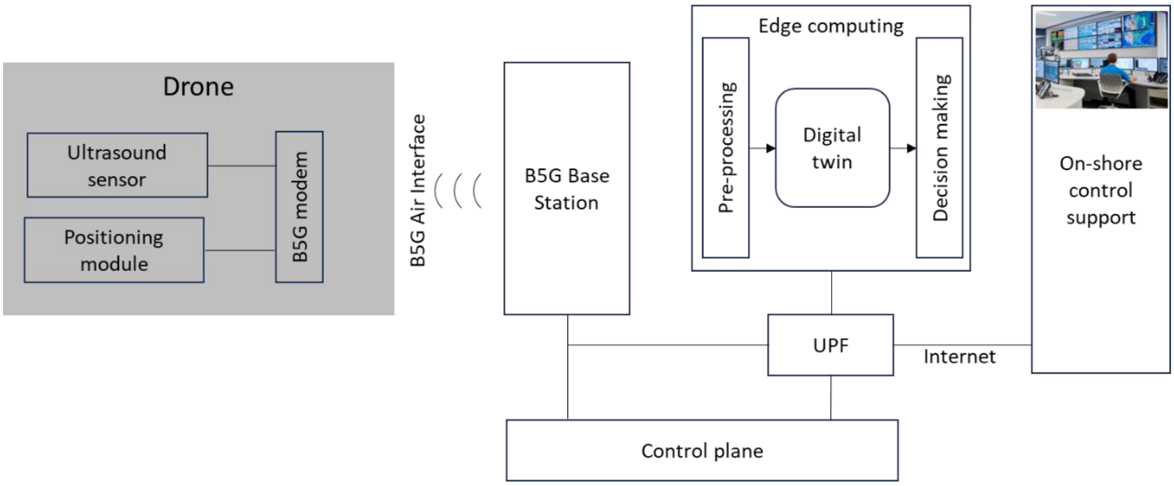

Figure 17. High level application architecture.

The B5G network is used to transfer the measurement data to two destinations. The first one is the DT system located at the edge and the second is a remote desktop emulating the control room where human operators can receive and process the data. The DT system is a model capable of delivering information about the blade status and structure as a result of the measurements. One candidate method is a surrogate Finite Element Method (FEM) blade structure simulation tool obtained through ABAQUS software suite [13]. It is a simplified FEM model capable of estimating the loads on the blades in different operating conditions as depicted in Figure 18. In the step of the DT, the loads distribution from a FEM model of a blade is used to estimate internal efforts, Remaining Useful Lifetime (RUL) and Failure Risks (FR). Assessing the internal efforts such as stress distributions, bending moments, and shear forces are critical for determining structural integrity of a blade. Another option considered, depending on the data and measurements that will be available, is a logic-based tool aimed at translating the information gained from UT measurements (crack depth, size, orientation) on a list of "failure hotspots" to be inspected next.

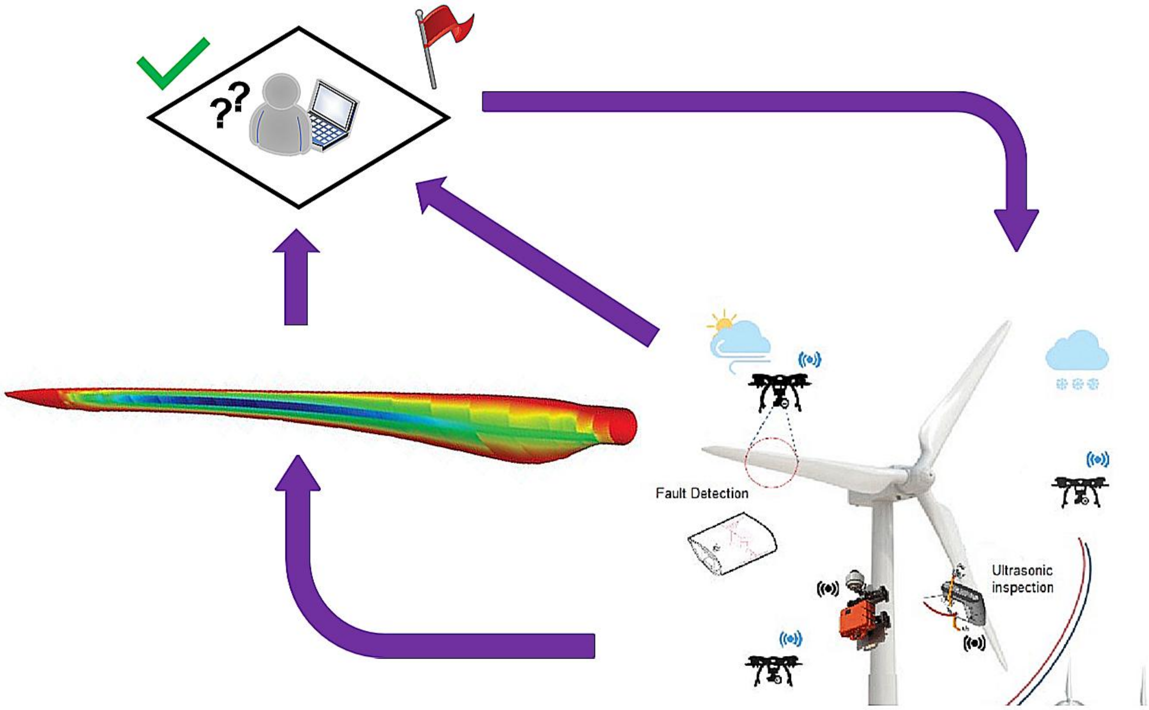

Figure 18. Data workflow from inspection to DT to decision making.

The input for the DT numerical processing is the pre-processed measurement data. The measurement and their spatial coordinates are used to estimate the existence of cracks or other faults in the blade (delamination, erosion). These faults are converted in structural changes from the factory and the observed structural properties. The output of the DT system is then transferred to the remote operators for inspection and processing alongside the raw measurements which can be used if necessary for cross-validation. The onshore control centre may use the collected data for decision making during the inspection.

During the preliminary phase, the B5G connectivity, the measurement sensor, the positioning module, and the B5G network system were installed and tested individually. In the trial phase, the objective is to operate the drone as an integrated system, hovering and placing itself on the blade in order to measure the blade's structural properties. The drone would be capable of flying onto the wind turbine blade when the blade is at a standstill and placed horizontally. The original planned drone was equipped with a Sonatest rolling ultrasound sensor [14], which is designed for inspection of thick composite materials up to 110 mm thickness. This ultrasound sensor has a scan speed of 40-60 mm/s with a scan width of 85 mm.

2.2.2.2 Devices, sensors and supporting equipment

Table presents the list of devices used for the deployment of the E2 use case.

Table 6. Devices used in E2 deployment.

Category	Component	Technical Details	Units
5G connectivity module	Quectel RM500Q-GL [15]	RM500Q-GL is a 5G module optimized for IoT and eMBB applications. It supports both 5G NSA and SA modes. The module is equipped with Qualcomm Snapdragon X55 modem.	2

			
UAV	DJI Matrice M600 [16] 	The drone to be equipped with ultrasound or similar sensor, 5G communication modem and positioning module. This drone is planned to be used for the final trial.	1
UAV	DJI Matrice 4T [17] 	Additional drone used for evaluating the positioning module during the pre-trial phase.	1
Positioning module	u-blox EVK-X20P [18] 	High precision GNSS module with all-band GNSS positioning and built-in RTK technology providing centimeter level accuracy.	2

2.2.3 Network infrastructure and functional components

The network infrastructure and functional components are depicted in Figure 19. For the B5G Radio Access Network (RAN), the AMARI Callbox Classic from Amarisoft [19] is deployed in the Zephyros lab in Vlissingen. This setup provides B5G connectivity for the Quectel modems. For the B5G Core, the Open-source 5G System (Open5GS [20]) implementation is used. The Open5GS Core, installed at the TNO lab in The Hague, provides the key network functions required for 5G connectivity, including:

- Access and Mobility Management Function (AMF): Handles registration, connection, and mobility management for 5G devices.
- Session Management Function (SMF): Manages Protocol Data Unit (PDU) sessions and allocates Internet Protocol (IP) addresses for user equipment.
- Authentication Server Function (AUSF): Performs secure authentication of subscribers.
- Network Repository Function (NRF): Maintains service discovery and registration for network functions.
- User Plane Function (UPF): Routes and forwards user data traffic, enabling high-throughput and low-latency communication.

In addition to the UPF located at TNO lab, a local UPF is placed at the Zephyros lab to function as the edge server.

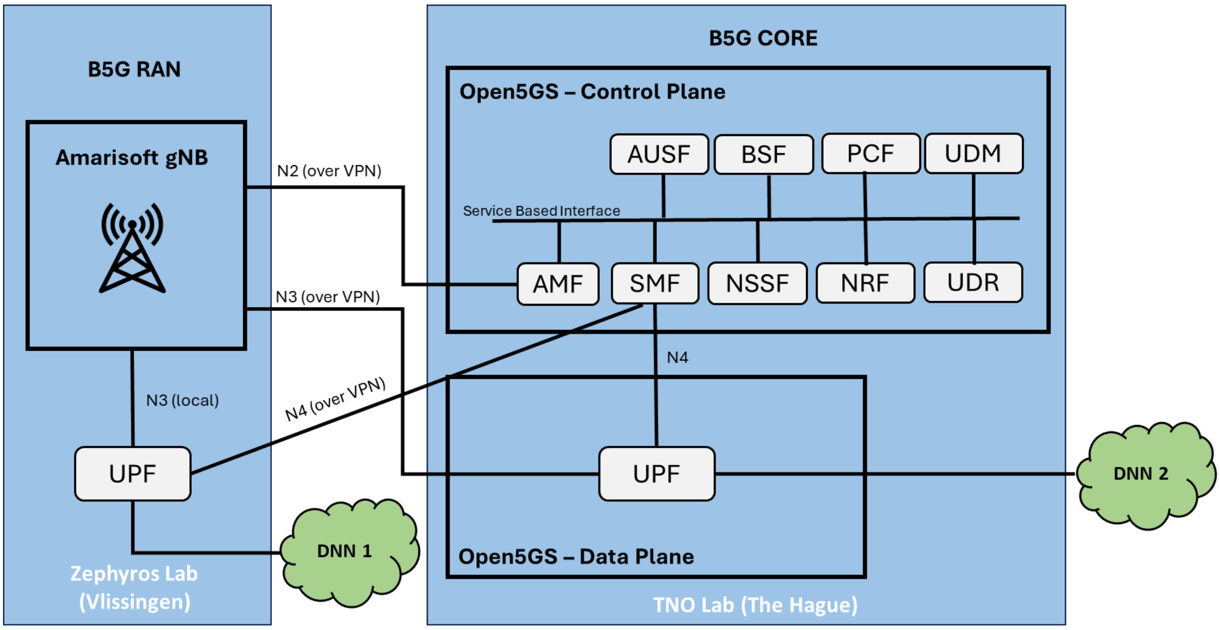

Figure 19. Network infrastructure and functional components.

2.2.4 Use case enablers

Table 7. presents a summary of the major E2 technology enablers, as introduced in the deliverable D3.1.

Table 7. List of enablers and sub-enablers for E2.

Category	Technology enablers	Sub-enablers
Communication enablers	Network slicing	
Communication enablers	Public-private network integration	
Communication enablers	Multiple-RAT connectivity	Non-Terrestrial Network (NTN) integration
Compute-as-a-Service enablers	Compute continuum	Distributed user-edge-cloud compute continuum
Compute-as-a-Service enablers	Compute continuum	Multi-party cloud
Application enablers and AI	OpenAPIs for Network Exposure	CAMARA APIs for Quality on Demand
Application enablers and AI	OpenAPIs for Network Exposure	CAMARA APIs for Edge/Cloud
IoT and localization enablers	Advanced Localization and Positioning Systems	Hybrid Localization (5G + GNSS)
IoT and localization enablers	Advanced Localization and Positioning Systems	GNSS and RTK Positioning

While all these enablers listed in Table 7. are relevant to the E2 use case, not all of them is foreseen to be implemented and integrated in the E2 pilot due to the availability of hardware and software. For example, NTN integration is a potential solution for extending coverage in offshore locations where there is limited or no coverage by the terrestrial network. However, in order to evaluate this enabler, access to NTN is required such that e.g. a satellite can be configured to act as a B5G base station. This is however unrealistic, considering the budget limit and currently limited availability of NTN. The NTN enabler is therefore not foreseen for evaluation within the E2 use case, though it remains identified as an important capability for future development. In the following text, only the enablers foreseen for implementation and integration in the E2 pilot are described in futher details.

Network slicing: In the E2 use case, drones are deployed to perform offshore wind blade inspections. This drone operation requires reliable connectivity, low latency, and guaranteed throughput to support real-time data transmission and remote control. Using network slicing, a dedicated slice can be deployed to meet the specific QoS for the traffic related to the drone operations, which goes over a network that handles other types of traffic too. Network slicing will ensure consistent performance and service continuity even under dynamic or congested network conditions.

The design of network slicing for the E2 use case will consist of two slices, i.e. one slice for the drone-related traffic and the other slice for the background traffic. As shown in Figure 20, the network components will be located at the Zephyros lab in Vlissingen and at the Netherlands Organisation for Applied Scientific Research (TNO) lab in The Hague. Specifically, the User Equipment (UE) and the gNB will be located at the Zephyros lab in Vlissingen and the core network will be located at the TNO lab in The Hague. The connection between the two locations will be provided via a VPN. This design is chosen to more accurately replicate the scenario where the drone and access network are located offshore and the core network is located onshore. However, to allow for lower latency, one UPF (shown as UPF1 in Figure 20) will be deployed at the Zephyros lab, specifically to serve the slice handling the drone-related traffic. All other background traffic from the second slice will be routed to the Internet via another UPF (shown as UPF2 in Figure 20), which will be deployed at the TNO location in The Hague. This design aims to showcase the capability of network slicing in achieving lower latency and replicates the scenario where an edge server (hosting the UPF1) is deployed offshore.

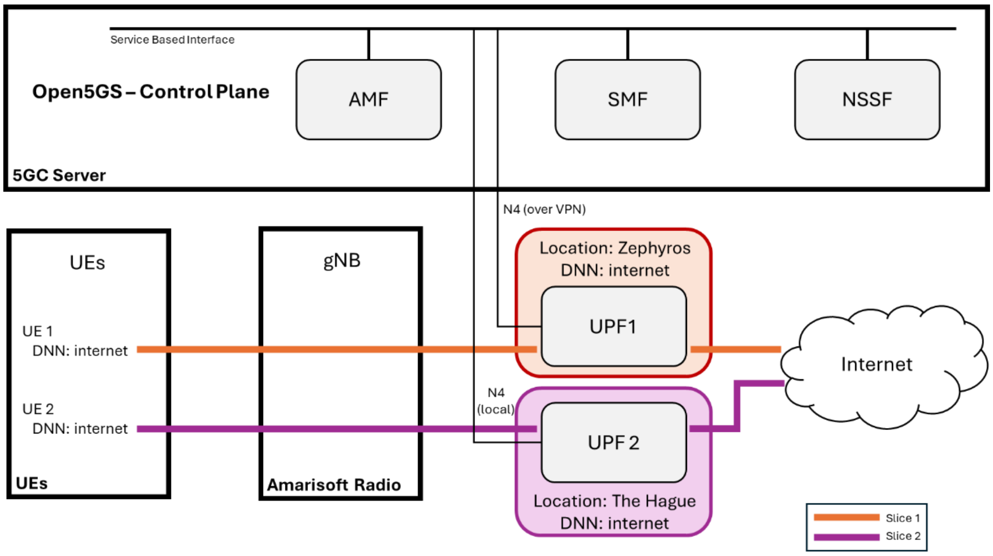

Figure 20. Network slicing design for E2 use case.

Public-private network integration: In the E2 use case, connectivity needs to be maintained and guaranteed between the drone and the base station located at the offshore platform. Due to the nature of the use case, it could be the case that there are limited deployed networks at the offshore location. Often the connectivity at the offshore location is provided by either standalone private network or by one network operator. In the latter case, the different enterprises operating offshore can establish a customer relationship with the network operator and have their own private network based on the public network, operated and maintained by the network operator.

The architecture as shown in Figure 21 uses a centralized Open5GS control plane to manage two separate UPFs, which effectively isolates private and public traffic at the data plane level. When a UE requests a PDU session, the SMF performs selection based on the DNN and steers traffic via the N4 interface to the designated UPF.

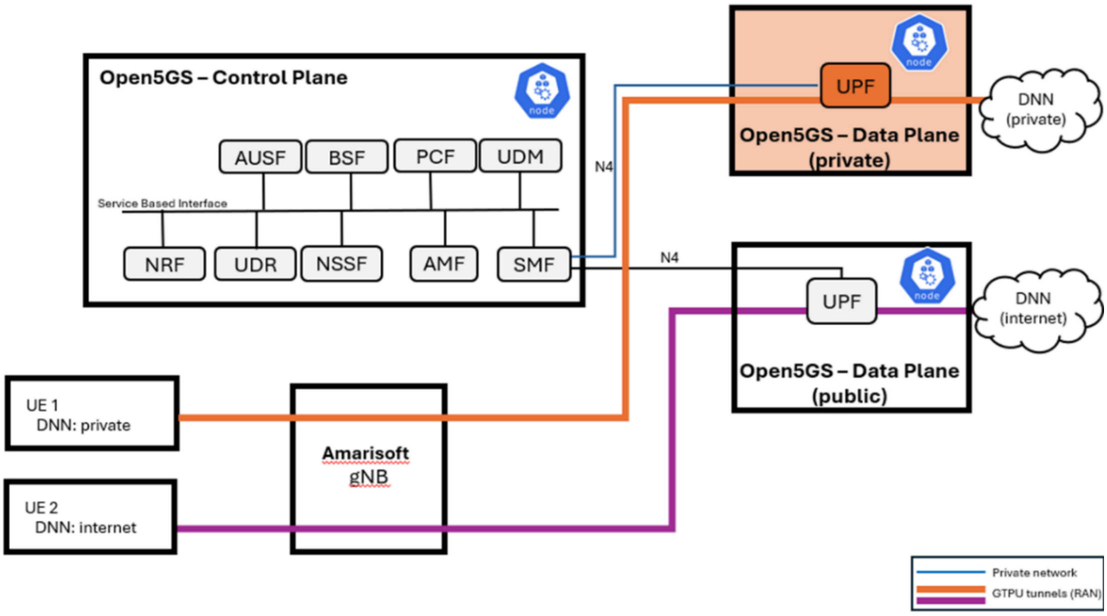

Figure 21. Public-private network integration design for E2 use case.

The Amarisoft gNB manages the radio side by handling the GPRS Tunneling Protocol for User plane (GTP-U) tunnels that connect the UEs to their respective data planes.

Advanced localization and positioning system: In the E2 use case, precise localization of the drone on the blade surface during inspection is achieved by using a combination of GNSS and RTK techniques. The typical accuracy of a GNSS positioning is between 1 – 5 m. By combining the GNSS positioning with RTK correction services, the positioning accuracy can be enhanced to the centimeter level. The RTK correction data can be obtained in two ways. The first approach is by subscribing to commercial RTK services (such as PointPerfect [21]) which rely on a broad network of static base stations. A base station is a stationary GNSS receiver placed at a precisely known location. However, because these base stations may be distant from the drone, this approach generally limits the accuracy to 3 – 6 cm. The second approach involves deploying a dedicated local base station on-site to provide the RTK correction data to the drone. In this case, the drone acts as a rover, which is a mobile GNSS receiver. Due to the close proximity of this local base station, the positioning accuracy improves to approximately 1 – 2 cm. Consequently, this second approach has been selected for the E2 use case. The test setup is depicted in Figure 22.

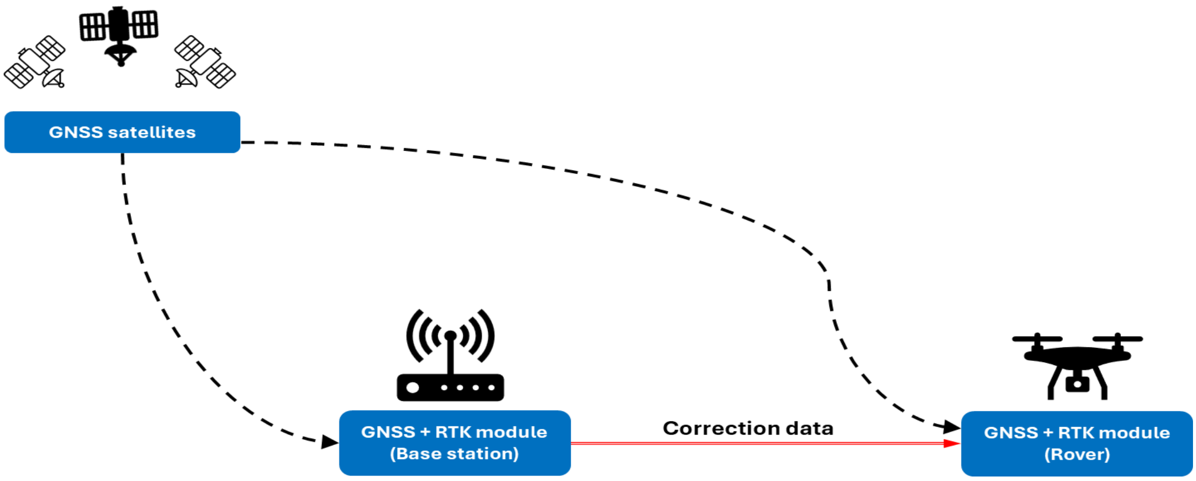

Figure 22. GNSS and RTK positioning test setup for E2 use case.

2.2.5 KPIs definition

Table 8 presents the set of KPIs that will be measured and validated during the final trial phase of the E2 use case, planned for Q1–Q2 2027. The KPIs are defined to assess the performance of the B5G network supporting blade inspection in a wind farm environment. The B5G network should be able to provide continuous and uninterrupted connectivity to the B5G modem mounted on the drone. It should also support reliable data transfer from the drone to the DT system and onshore control support. Furthermore, accurate localization of the drone during inspection is critical.

Table 8. E2 defined KPIs.

KPI code/name	Description/KPI definition	Scenario	KPI target
User experienced data rate (KPI_VA_11)	Uplink data rate	Wind farm environment	> 50 Mbps
End-to-end latency (KPI_VA_3)	Time between when the source node sends a data packet and when the destination node receives it	Wind farm environment	< 100 ms
Service availability (KPI_VA_1)	The system needs to be almost always available to provide the network connection	Wind farm environment	> 99%
Positioning accuracy (KPI_VA_24)	Difference between the calculated position and the actual position of the device	Wind farm environment	< 10 cm

2.2.6 KVIs definition

Table 9 summarizes the set of KVIs that will be evaluated during the final trial stage of the E2 use case, planned for Q1–Q2 2027, with the purpose of assessing how B5G connectivity and edge computing can significantly improves the wind turbine IO&M. The KVIs focus on stakeholder perceptions related to inspection time reduction, environmental sustainability and improved safety for inspection personnel. During the trial execution, relevant stakeholders, including energy-sector representatives, wind farm operators, will be invited to attend the live trial demonstrations and to actively participate in the evaluation process by completing structured questionnaires using standardized metrics to ensure consistent and reliable feedbacks. The collected feedbacks will be analyzed to derive the KVIs reported in the table, providing a direct assessment on the KVIs target. Overall, the KVIs support the evaluation of sustainability impact by capturing reductions in emissions, minimized operational interventions, and safer, more efficient maintenance practices enabled by digitalization and remote inspection capabilities.

Table 9. E2 defined KVIs.

KVI name	Description/KVI definition	KVI target
Reduced inspection time	Decrease in time required for blade inspection	At least 70% of respondents report a shorter inspection time and necessary decisions can be made within one drone flight window.
Reduced environmental footprint	Reduction of CO_2 emissions due to less boat transfers for personnel travelling to windfarm	At least 60% of respondents report fewer boat transfers are

		required per inspection campaign.
Improved safety for inspection personnel	Reduce risky tasks need for inspection (e.g. climbing on the blade)	At least 80% of respondents report the needs for technicians to physically access the turbine blades during inspection has decreased significantly.

2.2.7 Description of the pre-trial and trial setup

Table 10 summarizes the information about the trial setup for the E2 use case.

Table 10. E2 pre-trial and trial setup.

Trial ID & Name		Trial 13/E2 – Robotized offshore wind turbine blade inspection and maintenance
Trial Infrastructure / Venue	Zephyros Fieldlab in Vlissingen and TNO lab in The Hague (NL)	
Trial Description	Wind turbine blade inspection with drone. The data collected by the drone from an ultrasound or similar sensor will be transferred in real-time via B5G network to the onshore monitoring centre. The data is then analysed with the help of digital twins and numerical models, which can result in a decision on what action to take. The accurate localization of the drone during inspection is critical, which will be provided with GNSS and RTK modules.	
Technical Setup	**System Components**	• 5G SA RAN located in Vlissingen. One cell operating in the n77 band with 50 MHz bandwidth. • 5G Core network located in The Hague. • 5G UE (modem) • Edge node in Vlissingen. • Drone equipped with ultrasound or similar sensor and GNSS/RTK localization module. • DT system with numerical models
	System Configuration	• Network slicing with dedicated slide for drone traffic. • Edge node configured for low latency KPI. • Public-private network integration mimicking the typical configuration at offshore location.
Pre-Trial execution (initial phase, included in this deliverable, section 4.2)	**Pre-conditions**	• 5G network is up and running. For the pre-trial, the 5G network is set up at TNO lab in The Hague. • 5G UE is equipped with a SIM card provisioned with the correct credentials to access the 5G network.
	Trial Steps	Validation of 5G connectivity • Connect the 5G User Equipment (UE) to the 5G network • Perform KPI measurements (i.e. ping for latency and iperf for throughput measurements) • Document the results Initial assessment of the localization accuracy

		• Setting up the localization modules. One module is used for rover, and another module acts as base station to provide correction data. • Make sure the GNSS antennas have an unblocked view to the sky. • Wait until the base station module shows the "Time-Fix" message and the rover shows the "3D-Fix" message which indicates the rover is receiving a stable RTK correction stream. • Read out the 2D and 3D localization accuracy.
	Methodology	• Validation of 5G connectivity: Collect the network KPIs: • Initial assessment of the localization accuracy: Derive the localization accuracy from the GNSS/RTK modules
	Technical Enablers	*IoT and localization enabler*: GNSS/RTK positioning. In pre-trial phase, only one enabler is tested. Other enablers will be tested in the next phase.
	Expected Outcomes	• Validation that the 5G connectivity works as expected. • Validation that hybrid localization method using GNSS/RTK results in a significant improvement in accuracy compared to GNSS only.
Trial Execution (final phase, to be included in final deliverable D5.2)	**Pre-conditions**	• 5G network is up and running. For the final trial, the 5G network is set up at the field lab in Vlissingen. • 5G modem and GNSS/RTK module are mounted on the drone • The drone is equipped with an ultrasound or similar sensor capable of streaming the measurement data to the DT system.
	Trial Steps	• Connect the drone to the 5G network in the Vlissingen field lab. • Fly the drone towards the wind turbine blade. Note that the turbine blade is placed on the ground and not in the actual wind farm. • The drone lands on the blade and performs measurements on the blade's surface. • The measurement data is transferred via the 5G network to the DT. • Based on the output of DT, the onshore control center will make the inspection decision.
	Methodology	Validation of 5G-connected drone performing wind turbine blade inspection. The measurement data is transferred in real-time to the DT where analysis and decision is performed.
	Complementary Measurements	• Measurement of relevant network KPIs, such as uplink data rate and latency. • Measurement of DT processing latency.
	Calculation Process	The calculation of the blade health is performed by the DT system.
	Technical Enablers	*Communication enabler*: Network slicing Communication enabler: Public-private network integration.

		IoT and localization enabler: GNSS/RTK positioning *Application enablers and AI*: CAMARA APIs for Quality on Demand
Technical Enablers	**Communication enablers**	• Network Slicing • Public-private network integration.
	Application Enablers and AI	• CAMARA APIs for Quality on Demand
	IoT and localization enabler	• GNSS/RTK positioning
Expected Outcomes	• B5G connectivity enables the transfer of inspection data from the drone to the Digital Twin system onshore faster than currently possible. • The integration of DT reduces the time between blade inspection and decision-making. • Hybrid localization with GNSS and RTK provides better accuracy than just relying on GNSS positioning.	

2.3 Solar energy monitoring, control and predictions using B5G/6G communications and edge-cloud (Energy 3-E3)

This use case focuses on building and testing a smart solar energy management system using B5G/6G and edge-cloud computing. The goal is to make solar energy systems more efficient, easier to control remotely, and able to predict energy production more accurately. The trial will be implemented in a solar park located near Bucharest (Figure 23), a medium-scale photovoltaic installation equipped with modern string inverters and on-site monitoring infrastructure, providing a representative environment for validating the E3 smart solar energy management system built on B5G/6G connectivity and edge-cloud integration.

Figure 23. Final trial location for E3.

2.3.1 Description of the use case

The E3 use case aims to demonstrate how distributed solar energy infrastructure can be enhanced through real-time monitoring, control, and prediction capabilities enabled by B5G/6G communications and edge-cloud architectures. The objective is to validate an E2E industrial solution that integrates field-level photovoltaic data acquisition with cloud-based management and forecasting, while ensuring that control operations meet reliability and latency requirements typically associated with national energy coordination mechanisms. The system will operate using a three-tier distributed IoT architecture that combines industrial edge computing, advanced mobile network slicing, and cloud-hosted machine learning services. This architecture presented in Figure 24 is designed to support intelligent, resilient, and scalable energy operations across solar power plants.

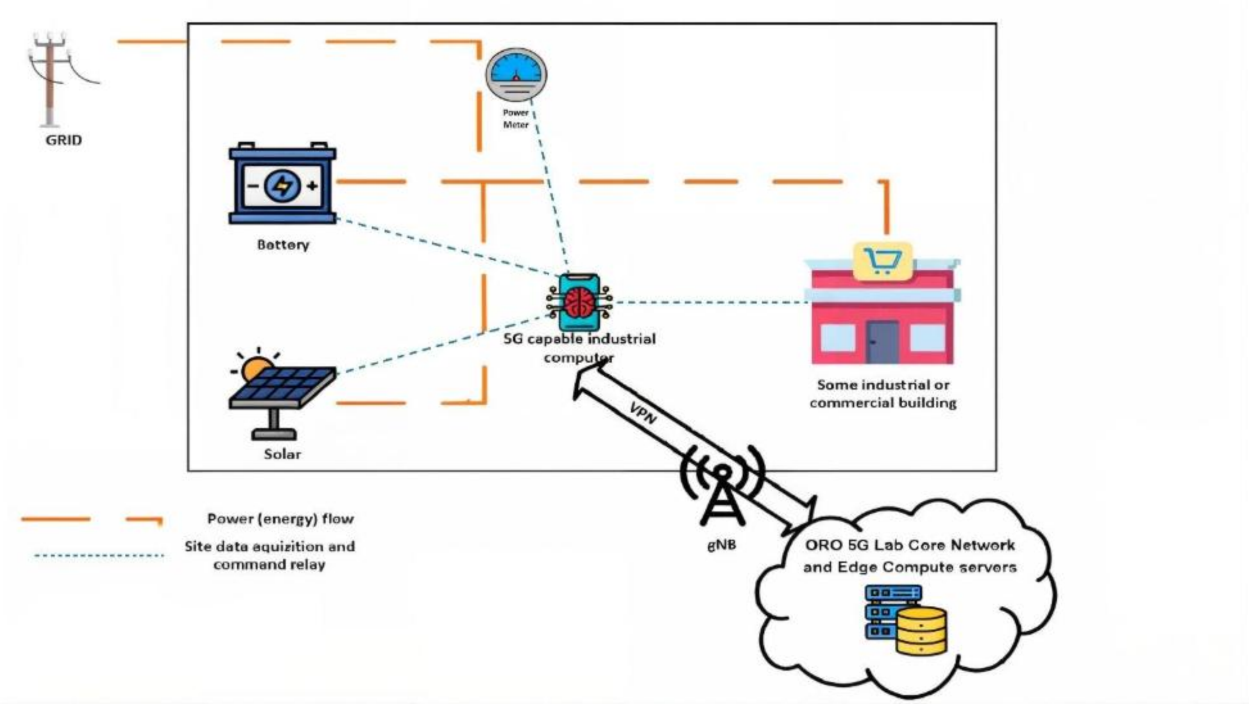

Figure 24. E3 use case architecture.

The field layer consists of industrial-grade edge gateways installed at solar plants in Bucharest. These devices, developed by CSOFT, are equipped with 5G NR RedCap modems [22] compliant with 3GPP Release 17 [23]. They integrate with photovoltaic inverters and energy storage units through Modbus TCP and Modbus RTU interfaces [24], enabling real-time telemetry collection of metrics such as power, voltage, and current. The gateways perform local preprocessing, execute inverter control commands, and incorporate buffering capabilities for resilience during fluctuating connectivity conditions. Communication between the field equipment and the cloud is secured through encrypted protocols, ensuring authenticated and tamper-resistant data flows.

The network layer relies on ORO's B5G infrastructure, which provides macro gNBs and network slicing tailored to the operational demands of the use case. A dedicated URLLC slice supports the transfer of time-critical control commands, enabling sub-20ms network latency, and 99.999% service reliability. A separate enhanced broadband slice is used for bulk telemetry uploads and non-time-sensitive data transfers. This configuration ensures that the system maintains deterministic performance even as network load fluctuates, and that regulatory-driven commands reach the field devices without delay.

The cloud and management layer are built around CSOFT's energy management platform, deployed on an edge-cloud environment with containerized services for telemetry ingestion, device orchestration, and data processing. Here, ML models will generate short-term predictions of solar energy production using merged data streams combining real-time field telemetry with external inputs such as weather forecasts. The cloud layer disseminates recommendations and predicted values to operator dashboards and partner systems through secure APIs. It also manages remote diagnostics, software updates, and the full device lifecycle. The architecture is designed to accommodate scalability and future enhancements, including predictive maintenance, automated grid-response mechanisms, and extended analytics components.

Two operational scenarios will be validated during the field trials. The first scenario focuses on regulatory compliance and real-time command execution. In this test, a grid adjustment request originating from the National Energy Authority will be simulated by the project partners and processed through the energy management platform. Upon receiving the simulated request, the platform will generate and dispatch a control instruction, such as a temporary power reduction, through the URLLC network slice. The edge gateway, presented as central component in Figure 24, is expected to receive the instruction with ultra-low latency, execute the required inverter adjustment, and report back the new operating conditions to the cloud platform. This scenario validates the platform's ability to provide reliable, ultra-low-latency control paths, ensuring that solar assets can respond within the five-minute regulatory window required for national grid balancing.

The second scenario highlights the forecasting intelligence of the platform. The machine learning engine running in the edge-cloud environment will continuously integrate real-time telemetry with external weather forecast data to generate near-term production estimates. Based on these predictions, the system will compute recommended inverter settings to optimize energy yield, manage storage usage, and improve operational planning. Operators may choose to apply these recommendations manually or allow the platform to execute them automatically. This testing scenario evaluates the accuracy of predictions, the responsiveness of cloud-to-edge command propagation, and the capacity of the architecture to support continuous adaptation based on environmental and operational changes.

The E3 use case addresses several technical challenges inherent to IoT deployments in the energy sector, including the need for resilient field equipment, stable connectivity under variable radio conditions, and secure and efficient data exchange. The use of RedCap modems in field devices is particularly relevant for validating 5G performance in industrial deployments. Achieving predictable URLLC behavior over a B5G test network is fundamental for ensuring grid-compliant control operations. The integration of ML adds complexity requiring careful orchestration between compute resources, data pipelines, and model management workflows.

The upcoming implementation and testing phases will validate the interoperability of all components across the architecture. Focus areas include verifying the latency and reliability of the network slices, confirming edge-cloud synchronization accuracy, and evaluating the operational impact of predictive forecasting.

2.3.1.1 Design of the application

The E3 application architecture showcased in Figure 25 is a distributed energy monitoring and control system designed to optimize the operation of solar power plants by combining edge computing, B5G communication, and cloud-based orchestration.

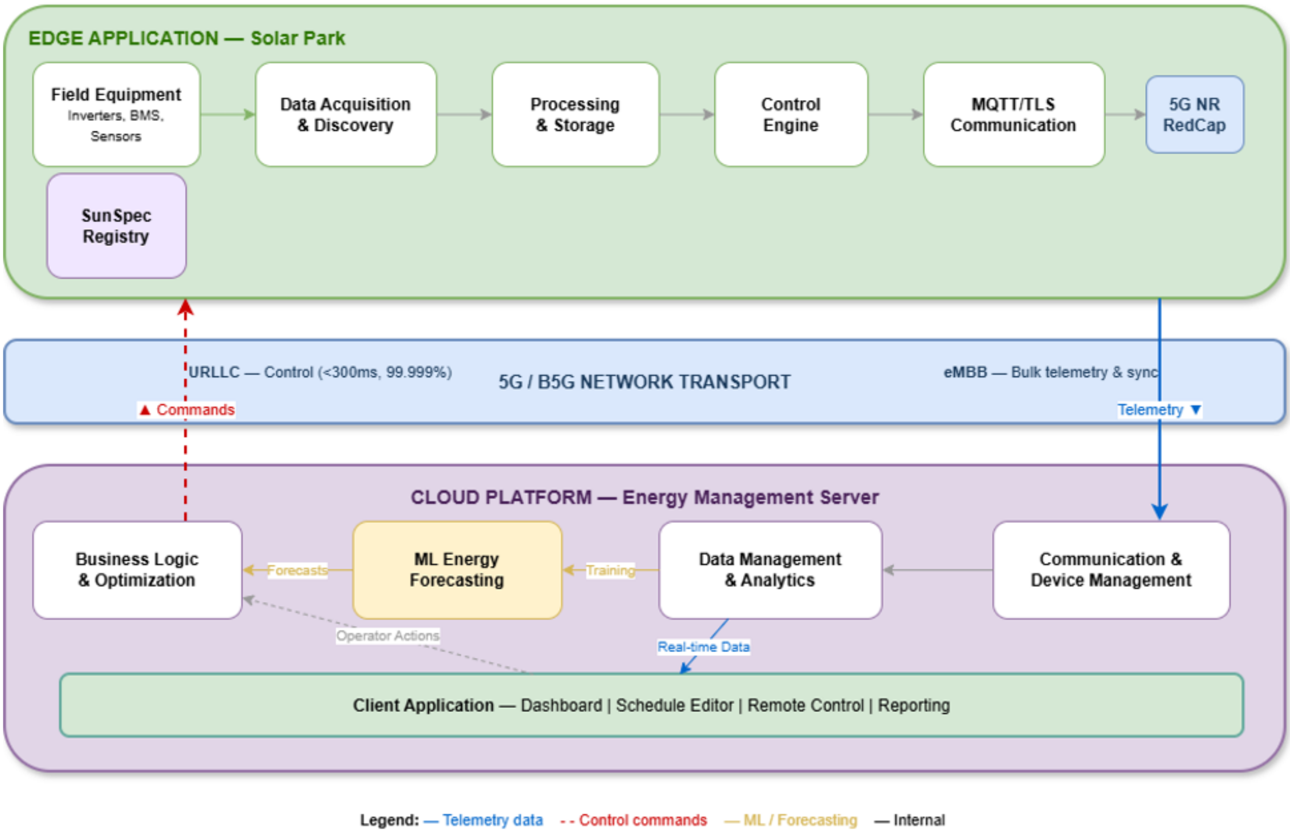

Figure 25. E3 application architecture.

The architecture shows how field data from the solar plant is collected, transmitted, processed, and used to generate both operational commands and production forecasts. At the field level, the solar inverters and monitoring equipment send measurements to the embedded controller. The controller integrates a RedCap-based communication module, which establishes the wireless connection to the cloud platform through the mobile network.

Once transmitted, the data enters the cloud application, which is organized into three functional layers. The communication layer manages device connections, handles incoming telemetry, and ensures secure and consistent data transfer. The data layer stores measurements, maintains device states, and exposes the required interfaces for internal processing. The command layer generates control actions based on operator input or automated logic, routing these commands back to the field equipment along the same communication path.

In parallel, forecast handling uses external weather data and machine-learning models to compute short-term production predictions. These outputs support both operator decision-making and automated adjustments within the command layer. The client application allows users to visualize real-time measurements, inspect historical data, and trigger manual control actions.

The full chain enables bidirectional communication, real-time visibility, and predictive functions, connecting the solar plant infrastructure with cloud intelligence and user-facing tools in a coordinated architecture.

At the field level, each edge device collects telemetry data from PV inverters, batteries, and auxiliary sensors installed at the site. The devices are multi-protocol capable and support both standard SunSpec registry maps and vendor-specific maps, which can be configured during deployment. Telemetry and control communication between edge devices and solar equipment are performed using industry-standard protocols such as Modbus TCP and RTU over RS485. Edge devices incorporate local buffering and embedded databases to store time-serialized telemetry during short- or long-term network outages, ensuring continuity of service and historical data availability for trend analysis. The devices are equipped with 5G NR RedCap (Reduced Capability) modems to provide reliable connectivity where full 5G capability is not required, while local processing reduces network load by handling non-critical tasks at the edge.

The edge gateways also implement SunSpec-based model discovery and runtime telemetry polling [25]. Automatic device discovery allows the gateway to register new inverters or batteries with minimal configuration, creating detailed device profiles with capabilities, registry data, and firmware versions. Command routing, event broadcasting, and firmware-aware synchronization are supported, enabling centralized management to interact seamlessly with all registered devices. Data interrogation is executed using a decoupled polling strategy, where high-priority energy prediction parameters (AC power, cumulative energy, DC input, operating state) are collected first, while less critical data is aggregated in intelligent batches to optimize performance. Communication with the cloud uses MQTT over Transport Layer Security-TLS for bidirectional telemetry, remote configuration, health monitoring, and command execution.

The application's main functions are implemented across the edge, cloud, and client components. Telemetry is collected at the edge (via MQTT) and forwarded to the Cloud Server, where a central repository stores time-series data and device states, and where business logic and machine learning models perform forecasting and optimization.

The Cloud Server exposes APIs for both data provisioning to the dashboard and managed control commands back to the network gateway for coordinated actuation. The Client Application then consumes the processed live and forecasted data for visualization and enables manual control actions through the cloud control interfaces, closing the feedback loop between analytics and field execution.

Figure 26 illustrates the control and forecasting flow. The cloud-based server acts as the central orchestration engine, aggregating telemetry from all gateways, maintaining a virtual device registry, and providing REST APIs for external and internal applications. Through integration with AI-aaS frameworks, ML models analyze historical inverter performance and real-time telemetry combined with external weather data to generate predictive energy production forecasts. These models produce recommendations for energy storage, consumption, or grid injection, which are then pushed to the edge gateways for execution.

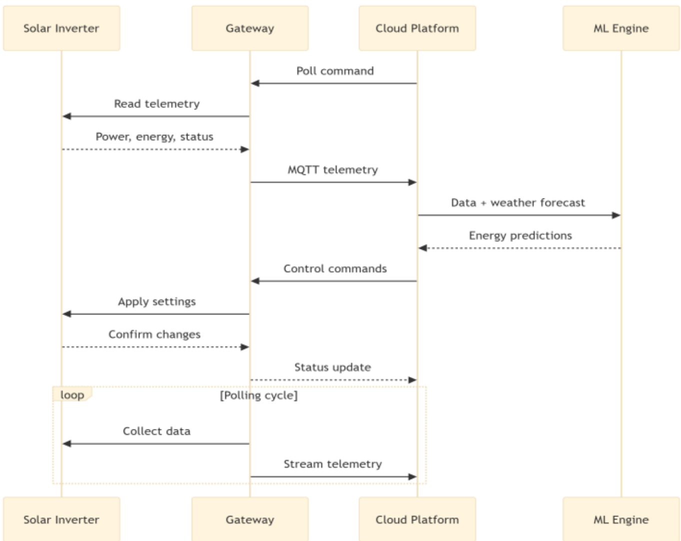

Figure 26. Solar energy forecasting & control flow.

Edge devices incorporate local intelligence to autonomously handle rule-based decisions, temporary network outages, and data buffering. Data collected at the edge is stored chronologically, indexed for trend detection, and continuously fed back to the cloud server for analysis and visualization. The cloud platform provides dashboards for operators, displaying real-time site status, device health, and forecasted energy production. Feedback loops are implemented to update AI/ML models based on observed deviations between predicted and actual energy production, ensuring continuous refinement of forecasting and control strategies.

Security is enforced through encrypted communication channels, authenticated access, and compliance with privacy requirements, ensuring that sensitive operational and sensor data remains protected. The architecture is scalable to support hundreds of distributed solar sites, and the system is designed for standard off-the-shelf edge hardware, allowing easy deployment and maintenance.

2.3.1.2 Devices, sensors and supporting equipment

Table 11 summarizes the hardware components used in the E3 environment and their role within the solar park monitoring and control prototype. The listed equipment supports validation of inverter telemetry polling, command execution, communication reliability, and end-to-end forecasting workflows, ensuring functional alignment with the hardware planned for field deployment.

Table 11. E3 devices.

Category	Component	Technical Details	Units
5G Connectivity Module	Quectel RM500Q-AE [26] 	Industrial-grade 5G module enabling wireless connectivity for the robots. Supports 5G NR, LTE-A Pro, LTE-A, and WCDMA. Download speed up to 2.5 Gbps and upload up to 660 Mbps. Designed for smart transportation, IIoT, video surveillance, and robotic applications.	2
Solar park inverter	Fimer PVS-50 [27] 	FIMER PVS-50 is a three-phase string inverter designed for commercial and industrial solar PV systems, supporting high DC input voltages and integrated communication via Modbus.	1
Solar park panels	10x JA Solar JAM72S30-550/MR 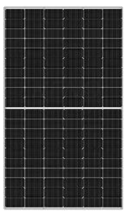	High-efficiency 550W JA Solar monocrystalline panels featuring half-cell PERC technology	10 (in pre-trial)

Sensors	Meteocontrol Si-I-420-T [28] 	Twenty-eight silicon irradiance sensors designed to accurately measure both solar radiation and cell temperature for precise photovoltaic system monitoring.	1
Gateway development	OrangePi 4A [29] 	OrangePi 4A, a compact ARM-based single-board computer equipped with a Rockchip RK3399 processor, 4 GB of RAM, and multiple interfaces. It will be used as the development platform for the solar park gateway prototype. It enables local testing of Modbus RTU/TCP communication with inverters, MQTT message routing, and edge-side telemetry buffering under real operating conditions.	1

2.3.2 Network infrastructure and functional components

The E3 use case leverages the B5G infrastructure described in detail in the section 2.1.3, using the same distributed 5G SA architecture connecting the Bucharest and Iasi 5G Labs with the Orange datacenter.

For E3, industrial-grade edge devices with 5G RedCap modems are deployed at photovoltaic plants to act as intelligent gateways for inverters, battery systems, and environmental sensors. These devices communicate via a dedicated URLLC slice, enabling control for inverter adjustments and grid balancing. The Bucharest 5G Lab integrates the 5G SA Core and orchestration systems, hosting part of the E3 use case software appliances in the associated edge computing facility.

The system allows for AI/ML-based forecasting, anomaly detection, and predictive control at the edge. Telemetry data is aggregated in the edge domain via APIs, where AI-driven models predict production patterns and optimize distributed solar energy assets.

Overall, E3 extends the B5G architectural principles validated in E1 — distributed telco edge nodes and slicing — to renewable energy management, demonstrating predictable performance, adaptive control, and cross-layer orchestration for large-scale solar energy systems (Figure 27 shows the dedicated E3 B5G architecture).

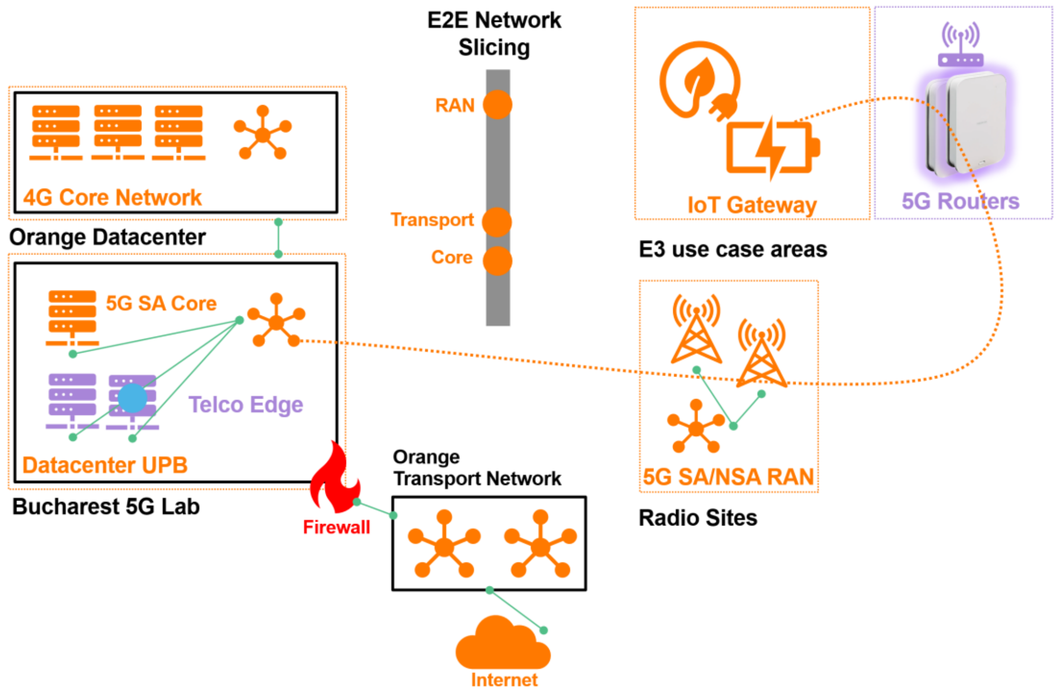

Figure 27. E3 use case network architecture.

2.3.3 Use case enablers

The E3 use case focuses on enabling real-time monitoring, forecasting, and control of PV installations through B5G/6G communication, edge intelligence, and cloud-based orchestration. The enablers selected for implementation combine communication resource management, network slicing, AI-aaS, IoT connectivity, and actuation mechanisms, forming the technical foundation for the full trial. Their deployment follows a phased approach, aligned with the project timeline and the progressive integration of hardware, communication infrastructure, and AI-based control logic. As presented in deliverable D3.1, E3 enablers are summarized in Table 12.

Table 12. List of enablers and sub-enablers for E3.

Category	Technology enablers	Sub-enablers
Communication enablers	Communication resource management	Network performance monitoring and control
Communication enablers	Communication resource management	Zero-touch network and service management
Communication enablers	Network slicing	
Application enablers and AI	AI-as-a-Service framework	AI-as-a-Service
Application enablers and AI	AI-as-a-Service framework	ML models catalogue
IoT and localization enablers	IoT Connectivity and Infrastructure	IoT Gateways
IoT and localization enablers	IoT Connectivity and Infrastructure	IoT Sensors and Devices
IoT and localization enablers	Data Ingestion and Telemetry	Data Collection from sensors and devices
IoT and localization enablers	IoT Contextual Awareness Systems	
IoT and localization enablers	Remote Control and Operation (Actuation)	Real-Time Actuation/Teleoperation

The *communication enablers* for network slicing and communication resource management are implemented by separating traffic classes and continuously steering E2E performance between field-level edge gateways, the B5G/6G transport, and the cloud orchestration platform. The solar sites rely on 5G RedCap devices embedded in industrial edge computers to transport inverter telemetry and environmental data with stable latency, while inverter reconfiguration or curtailment commands are delivered in real time. A dedicated network slice ensures deterministic behavior by isolating critical actuation commands from periodic telemetry and from edge-to-cloud AI model management traffic, avoiding cross-impact during load variations. Network performance monitoring complements this by measuring latency, jitter, packet loss, and throughput across the solar site, edge, and cloud, and by enabling policy-based adjustments when actuation timing or reliability constraints are at risk.

Application Enablers and AI: AI-aaS is integrated into the E3 use case (see Figure 28) to support on-demand execution of forecasting and optimization models. These models rely on telemetry streams, inverter performance metrics, and weather forecasting data to produce short- and medium-term predictions used for generation forecasting or real-time adjustment of inverter behavior. Depending on the latency and computational requirements, the AI workloads can run either at the edge or in the cloud. Short-term forecasting and anomaly detection are better suited to the edge gateway, while more complex, longer-horizon models are executed in the cloud environment.

Initial integration of AI-aaS, which includes data ingestion, linkage to the Open-Meteo weather API for inference-time weather forecasting, and the execution of basic forecasting algorithms, was implemented in March 2026. Further refinement will continue throughout 2027, forming part of the final operational configuration used in the field trials. The AI-aaS which supports model packaging, versioning and flexible deployment across both cloud and edge environments. With this combination of real-time telemetry streams, inverter performance metrics, and externally source weather-forecast data will enhance the accuracy of energy-production forecasts and enables more responsive operational control.

The ML catalogue provides structured storage, version control, and verification for all forecasting and optimisation models employed in E3. Currently validated models include Histogram-Based Gradient Boosting, Random Forest, and XGBoost. It ensures traceability throughout the update lifecycle and supports controlled deployments either toward the edge gateways or the cloud platform. During the trial, the catalogue enables efficient iteration, comparison, and replacement of models, thereby improving prediction accuracy and operational responsiveness.

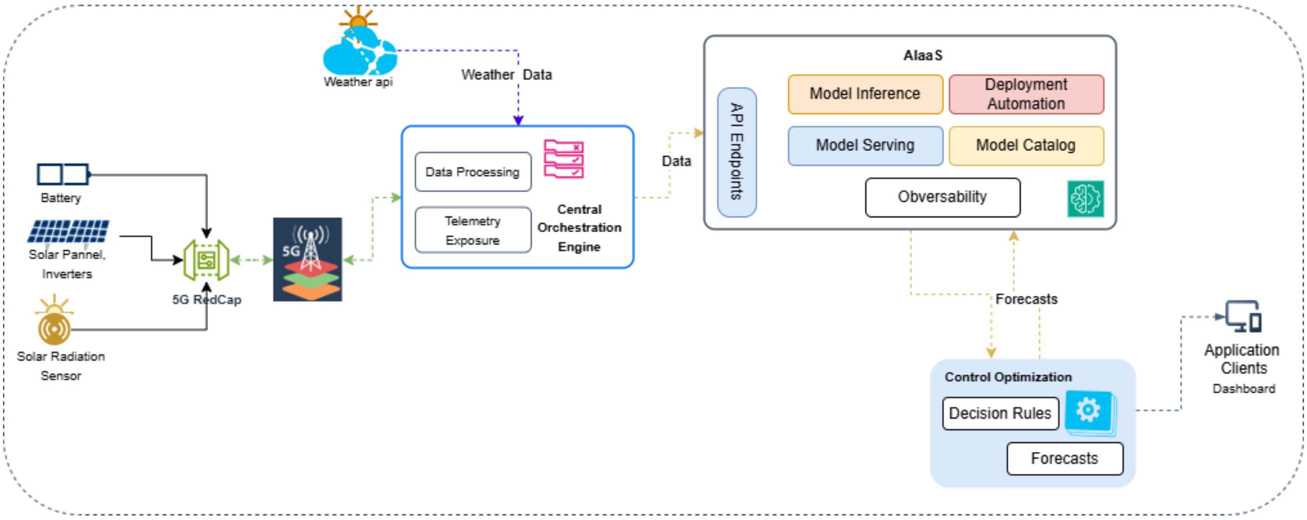

Figure 28. AI-aaS integration.

At the edge layer, 5G RedCap gateways perform real-time data ingestion, capturing high-frequency telemetry streams including inverter operational metrics, solar irradiance, and ambient/Module temperature. Incoming signals are normalized, validated, and timestamp-aligned to ensure schema consistency and temporal coherence. The processed data batches are then securely transmitted to the AI-aaS layer.

Within the AI-aaS platform, the forecasting model, hosted on the CAPG AI-aaS infrastructure, is deployed as a containerized microservice and versioned through the internal MLOps model registry. The platform manages inference serving, autoscaling, and observability. The control optimization service consumes the model's output through low-latency APIs to compute real-time inverter setpoint adjustments (e.g., reactive power, ramp-rate limits, curtailment thresholds).

Energy production predictions are continuously compared against actual telemetry to support automated validation, error monitoring, and drift detection. When drift thresholds or performance degradation metrics are met, the system triggers a retraining workflow, executing data refresh, feature re-computation, model training, evaluation, and approval gates. Successful models are automatically redeployed through the MLOps pipeline (see Figure 29).

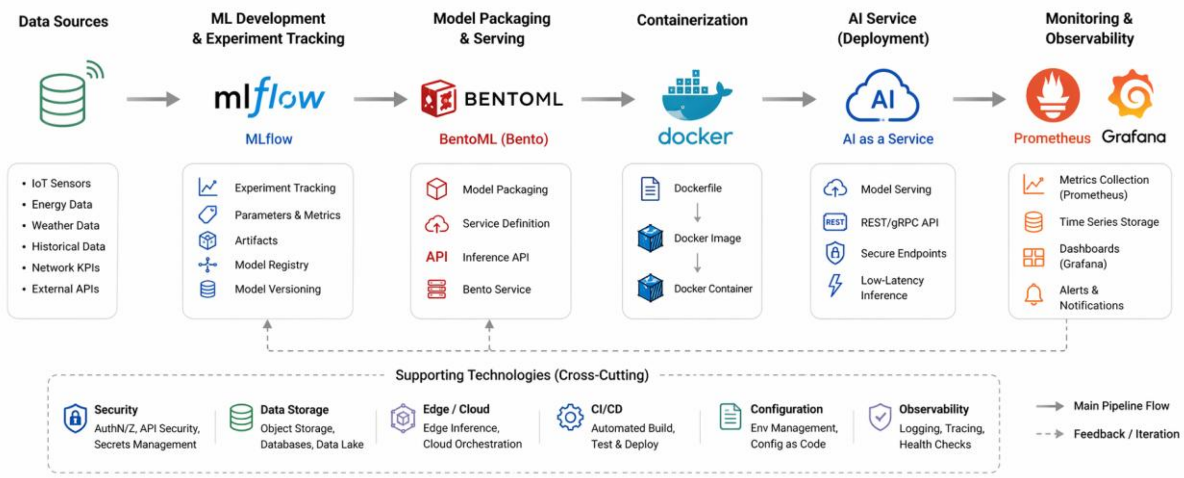

Figure 29. Forecasting pipeline.

This establishes a closed-loop continuous learning system, where models are systematically retrained, versioned, and rolled out to maintain optimal forecasting accuracy and control responsiveness.

IoT and Localization Enablers: are built around an industrial 5G RedCap gateway developed by SIM, which provides secure MQTT/TLS connectivity with slicing support between solar inverters, auxiliary sensors, and the communication infrastructure. The gateway aggregates and pre-processes inverter and environmental telemetry, supports remote configuration and firmware updates, and maintains local continuity during temporary connectivity losses, initially using rule-based logic and later enabling AI-driven command execution from the orchestration engine. To improve control accuracy under changing operating conditions, contextual awareness functions correlate site telemetry with external weather information, enabling the orchestration engine to adapt strategies and refine both predictions and real-time adjustments.

2.3.4 KPIs definition

Table 13 outlines the set of KPIs that will be measured during the final trial phase of the E3 use case, planned for mid-2027, in order to quantify the performance of the B5G-enabled smart solar energy management system deployed in a real solar park environment. The KPIs target network, application, and service-level behaviour required to support reliable monitoring, control, and forecasting of photovoltaic assets, including latency, throughput, availability, coverage, reliability, and connection density under operational conditions.

Table 13. E3 defined KPIs.

KPI name	Description/KPI definition	Scenario	KPI target
Network service latency (KPI_VA_3)	Time between the source node sends a data packet to when the destination node receives it	Solar park environment	10 ms
One-way application latency (KPI_VA_4)	Amount of time it takes at application level from the source to the destination application	Solar park environment	300 ms
DL User Experienced data rate (KPI_VA_10)	DL achievable data rate that is available to a device.	Solar park environment	50 Mbps
UL User Experienced data rate (KPI_VA_11)	UL achievable data rate that is available to a device.	Solar park environment	10 Mbps
Coverage (KPI_VA_22)	Geographic area where a network signal can be received and used by a device	Solar park environment	99 %
Network availability (KPI_VA_1)	Percentage of time the service can be delivered latency) within service space when requested	Solar park environment	>99.999 %
Network reliability (KPI_VA_2)	Period of time for which the service satisfies the required performance constraints (DL/UL capacity, E2E latency)	Solar park environment	>99.999 %
Connection density (KPI_VA_13)	Number of connected devices per square meter	Solar park environment	1 device/m2.
Inference speed (KPI_AI_1)	Latency incurred in the computation of an AI model in inference mode	CAPG testbed	<2 s
ML accuracy (KPI_AI_2)	Accuracy of ML classification model	CAPG testbed	<5%

2.3.5 KVIs definition

Table 14 summarizes the set of KVIs that will be evaluated during the final trial stage of the E3 use case, scheduled for Q2 2027, with the purpose of assessing how the proposed smart solar energy management solution is perceived in operational and business terms. During the trial execution, interested stakeholders, including operators, technology providers, and energy-sector representatives, will be invited to attend on-site and remote trial demonstrations. Following these activities, participants will be asked to complete structured questionnaires tailored to their roles, using standardized rating scales. The feedback collected through these questionnaires will be consolidated and analysed to derive the KVIs reported in the table, providing a direct link between the observed system behaviour during the trials and the perceived operational and economic impact of the E3 solution. In addition, these KVIs contribute to the assessment of sustainability by enabling efficient real-time monitoring, optimized energy production, and reduced operational losses, ultimately supporting greener energy generation and improved resource utilization within the solar park.

Table 14. E3 defined KVIs.

KVI name	Description/KVI definition	KVI target
Improvement in Energy System Responsiveness	Measures how operators and beneficiaries perceive the ability of the solar park to react in a timely and reliable manner to grid-related commands and operational changes.	At least 80% of respondents rate system responsiveness as *"good" or "very good"* (≥4 on a 5-point Likert scale).
Operator Confidence in Remote Energy Control	Assesses the level of trust operators have in remotely monitoring and controlling solar assets without on-site intervention.	An average score of ≥4/5 across reliability, safety, and usability questions, with ≥75% of operators expressing willingness to rely primarily on remote control for day-to-day operations
Perceived Increase in Renewable Energy Utilization	Measures how stakeholders perceive the system's ability to maximize usable solar energy production through forecasting and optimized control.	At least 70% of stakeholders report a *clear or significant* perceived improvement (≥4/5) in renewable energy utilization due to forecasting and optimized control, with <15% indicating no perceived benefit.
Reduction in Operational Costs	Measures operator perception of cost savings resulting from reduced manual interventions, fewer site visits, and optimized operations.	At least 65% of operators report a perceived operational cost reduction corresponding to ≥10–20% savings, based on reduced site visits, manual interventions, and operational overhead.
Business Scalability and Replicability	Captures stakeholder confidence that the solution can be economically scaled to additional solar sites or portfolios.	At least 75% of respondents rate scalability and replicability as *"good" or "very good"* (≥4/5), and ≥60% indicate confidence that the solution could be deployed to additional sites with limited customization.

2.3.6 Description of the pre-trial and trial setup

Table 15 presents the trial description details for E3.

Table 15. E3 pre-trial and trial setup.

Trial ID & Name	Trial 14/E3 – Solar energy monitoring, control and predictions using B5G/6G communications and edge-cloud
Trial Infrastructure / Venue	Solar park located near Bucharest, equipped with PV panels, energy meters, environmental sensors, and 5G/B5G connectivity from the trial network. Edge-compute node deployed at the site for local processing.
Trial Description	The trial demonstrates a smart energy management system for a solar park, using B5G connectivity and edge-cloud computing for real-time data collection, monitoring, analytics and control.

Technical Setup	**System Components**	• PV panels and inverters (Fimer PVS-50 string inverters with Modbus TCP interface) • Smart meters and IoT sensors • 5G SA/NSA RAN operating in the n78 band • Edge node (OrangePi 4A running the gateway application) • Cloud platform for analytics and visualization • API for control/monitoring dashboards • AI-aaS platform based on MLflow and BentoML
	System Configuration	• B5G network configured with dedicated slices for energy data traffic • Edge node configured for low-latency KPI extraction • Cloud dashboard with real-time streaming pipeline via MQTT • Device provisioning and remote management enabled via MEC APIs
Pre-Trial execution (initial phase, included in this deliverable, section 4.3)	**Pre-conditions**	• 5G network coverage active at the site. • Permission and access to monitor and control the solar park. • Deployed server application remotely, ready to be used in the pre-trial.
	Trial Steps	• Connect PV systems to the 5G/B5G network in CSOFT testbed. • Initiate real-time measurements from inverter. • Activate cloud analytics for data filtering and KPI extraction. • Stream processed data to cloud dashboards • Measure performance KPIs (latency, throughput, reliability) • Implement and calibrate the prediction model in dedicated CAPG testbed.
	Methodology	Incremental validation of the communication and data acquisition chain. The methodology focuses on establishing stable 5G/B5G connectivity, verifying IoT device onboarding, telemetry continuity, and basic localization/context awareness, without activating application-layer analytics or AI-based processing.
	Technical Enablers	• *Communication enablers*: 5G SA connectivity over n78 band • *Application enablers and AI:* implemented in dedicated CAPG testbed • *IoT enablers:* collect data and trigger commands to inverter
	Expected Outcomes	• Verification of end-to-end connectivity between PV systems, IoT gateways, edge node, and cloud endpoint • Validation of telemetry acquisition, transport stability, baseline latency, and inverter control • Validation of AI prediction model
Trial Execution (final phase, to be included in final deliverable D5.2)	**Pre-conditions**	• PV system operational • Sensors calibrated and synchronized • 5G/B5G network coverage active at the site • Edge–cloud integration fully validated • Monitoring dashboard linked to the edge pipeline
	Trial Steps	• Connect PV systems and sensors to the 5G/B5G network • Initiate real-time telemetry from meters and sensors • Activate edge analytics for data filtering and KPI extraction • Stream processed data to cloud dashboards

		• Integrate predictive analytics models (energy forecasting) in common testbed. • Validate performance KPIs (latency, throughput, reliability)
	Methodology	End-to-end validation of the solar park monitoring, forecasting, and control workflow under operational conditions.
	Complementary Measurements	• Power production variance • Forecasting accuracy
	Calculation Process	• Statistical aggregation of telemetry streams • Edge-to-cloud latency computation • KPI formulation for energy production forecasting accuracy
Technical Enablers	**Communication enablers**	• Full B5G network slicing (eMBB and URLLC), Network performance monitoring and control
	Application enablers and AI	• AI-aaS platform for energy forecasting and anomaly detection. • ML Model catalogue
	IoT and localization Enablers	• Industrial IoT gateways with 5G RedCap connectivity, synchronized PV and environmental sensors
Expected Outcomes	• Demonstration of stable real-time telemetry and remote actuation over B5G network slices • Quantified improvement in forecasting accuracy through AI-enabled edge and cloud processing • Verified reliability, scalability, and reproducibility of the E3 architecture for larger photovoltaic deployments	

3 Use case planning

3.1 Renewable Energy Communities planning

Table 16 details the high-level activities for E1.

Table 16. E1 planning.

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q3 2025-done
Devices	Acquisition of hardware end devices (sensors, redcap devices, etc.)	Q2 2025-done
Infrastructure & Application Design	Design of B5G infrastructure and applications for the Bucharest site	Q2-Q4 2025-done
Preliminary Integration including enablers	Initial integration of all subsystems to verify interoperability	Q1 2026-done
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q2 2026-done
Extension to multiple building configuration	Extend the use solution to simulated multiple buildings	Q2-Q4 2026
Final Application Development	Complete application implementation, optimize performance and reliability	Q2-Q4 2026
AI/ML integration	Integrate AI/ML in Romanian testbed	Q3-Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyze trial results	Q2-Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2-Q4 2027

No deviations occurred during the first 17 months of the project.

3.2 Robotized offshore wind turbine blade inspection and maintenance use case planning

The high-level planning activities for the E2 use case are provided in Table 17.

Table 17. E2 planning.

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q1 2025-done

Hardware acquisition	Acquisition of 5G small cell, 5G modems, sensors and localization modules	Q2-Q4 2025-done
Infrastructure & Application Design	Design of B5G infrastructure and applications for the Dutch site	Q2-Q4 2025-done
Evaluation of enablers	Initial tests of the localization enabler	Q1 2026-done
5G connectivity tests	Perform connectivity tests in the lab and collect the relevant KPIs	Q2 2026-done
Drone acquisition	Acquisition of drones. The first (smaller) drone has been acquired. The larger drone is to be acquired.	Q2-Q4 2026
Final Application Development	Complete application implementation, optimize performance and reliability	Q2-Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyze trial results	Q2-Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2-Q4 2027

No planning deviations occurred during the first 17 months of the project. While the drone to be used has to be changed, the acquisition of the new drone should fall within the same timeline as initially planned.

3.3 Solar energy monitoring, control and predictions using B5G/6G communications and edge-cloud planning

Table 18 illustrates the high-level activities planning for E3.

Table 18. E3 planning.

Phase	Details	Schedule
UC Definition	Finalize the target architecture of the E3 use cases	Q3 2025-done
Devices	Acquisition and production of hardware end devices (sensors, redcap devices, etc.)	Q2 2025-done
Infrastructure & Application Design	Design of B5G infrastructure and applications for the Bucharest site	Q2-Q4 2025-done
Preliminary Integration including enablers	Initial integration of all subsystems to verify interoperability	Q1 2026-done
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios in Simtel's lab	Q2 2026-done
Devices evolution	Create and deploy the dedicated IoT Gateway device that will replace the prototype, at a solar park.	Q3 2026
Final Application Development	Complete application implementation, optimize performance and reliability	Q2-Q4 2026

AI/ML integration	Integrate AI/ML in Romanian testbed	Q3-Q4 2026
Full Trial Integration	Integrate full architecture, network, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring, KPIs and KVIs measurements	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyze trial results	Q2-Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2-Q4 2027

No deviations occurred during the first 17 months of the project.

4 Preliminary integration and results

This section presents the preliminary activities related to integration, testing, and early dissemination events for each use case, carried out up to M17. It provides an overview of the initial implementation progress, highlights key milestones achieved, and outlines the validation steps undertaken to ensure readiness for subsequent phases of development and large-scale experimentation.

4.1 Renewable Energy Communities

4.1.1 Integration status

For the E1 use case, preliminary integration activities for the first pre-trial planned at the end of Q1 2026 focused on preparing the B5G infrastructure and aligning it with the SB-EMS application components to ensure readiness for execution in April 2026.

Figure 30 illustrates the Private 5G SA infrastructure deployed at the Politehnica University Campus in Bucharest, comprising two outdoor 5G NR antennas and an indoor 5G SA installation, which formed the RAN backbone validated for the E1 pre-trial.

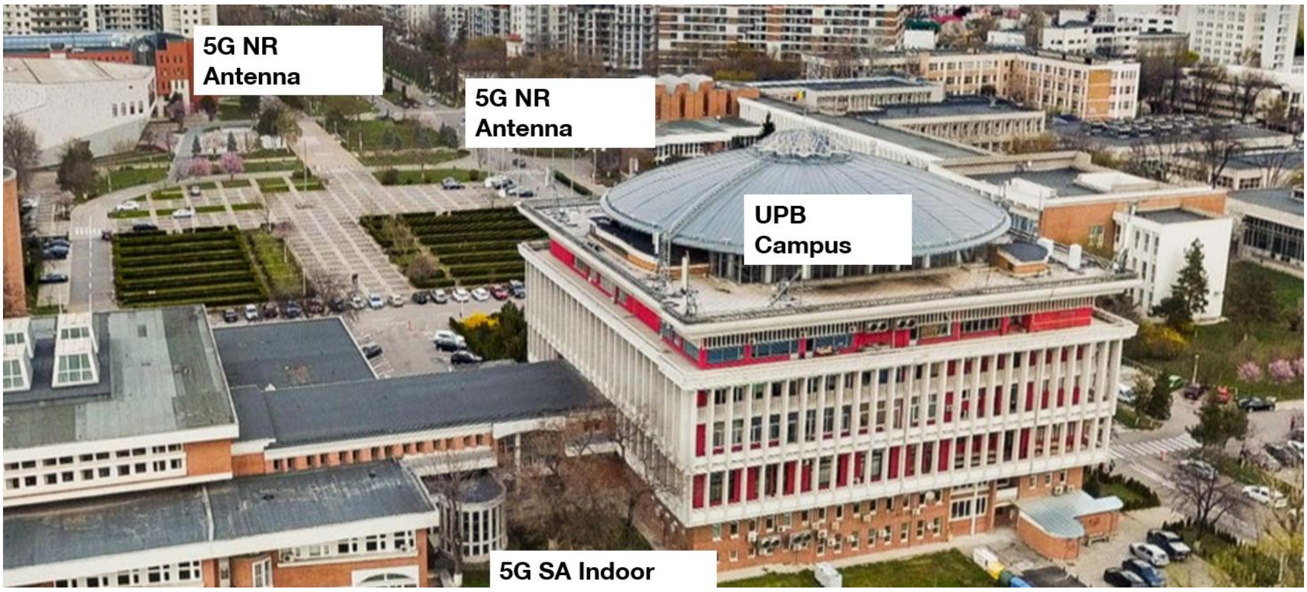

Figure 30. Private 5G SA infrastructure deployed at the UPB Campus in Bucharest.

Figure 31 shows the outdoor 5G radio equipment deployed at the Smart Campus site, including the base station hardware and antenna installations that were verified as part of the upgrade path toward SA operation in preparation for the April 2026 pre-trial.

Figure 31. Outdoor 5G radio equipment.

At network level, the 5G SA core deployed in the Bucharest 5G Lab was validated in its distributed configuration, with central control plane functions (AMF, SMF, PCF, AUSF, UDM) running in the virtualized cloud environment and User Plane Functions instantiated both centrally and at the Iasi edge location.

The orchestration layer, based on OpenStack and Kubernetes, was tested before the pre-trial for lifecycle management of core and edge components.

Edge and cloud computing resources are hosted in the datacentre at the ORO 5G lab. VMs running Kubernetes clusters are used to represent both edge nodes at different buildings (some of them emulated) and cloud resources for centralized processing at the REC level. As shown in Figure , the system deployment for each building includes:

- 1 VM (represented in blue) to host the platform components, i.e., the service and resource orchestrator, the CL Governance, the MLOps platform, the monitoring platform, and the IoT platform;
- 1 VM (represented in orange) to host a Kubernetes worker node where the SB-EMS application components are instantiated.

An additional Virtual Machine (VM), in green, will be used in the second phase of the project to host the REC-level components.

The first version of the software for the single building scenario has been installed in the ORO testbed and it is accessible via VPN from Nextworks' lab for testing and troubleshooting.

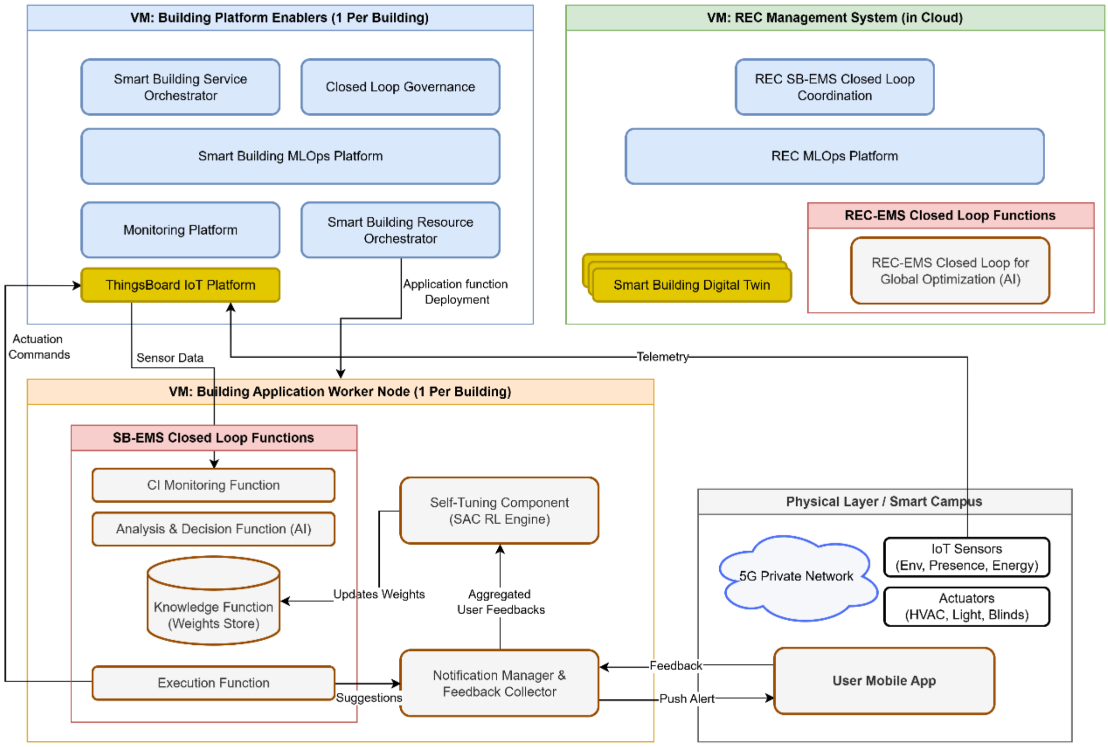

Figure 32. E1 – Setup in ORO 5G lab in Bucharest.

At edge and cloud level, virtual machines were instantiated in the ORO datacenter to emulate building-level edge nodes and REC-level cloud resources. Kubernetes clusters were deployed and configured to host the SB-EMS components. For each building instance, one VM hosts the platform components (Service Orchestrator, Resource Orchestrator, Closed Loop Governance, MLOps, Monitoring Platform and IoT platform), while a second VM acts as Kubernetes worker node for SB-EMS application functions. VPN access was enabled to allow remote testing and troubleshooting from Nextworks premises.

The IoT platform (ThingsBoard Community Edition) was installed (as depicted in Figure 33) and integrated with the MQTT-based sensors available in the Bucharest 5G Lab and Iasi campus. REST APIs and RPC interfaces were tested to ensure correct data ingestion and actuation flows. Integration between the IoT platform and the Monitoring Platform was validated, including data aggregation, filtering and storage according to the unified data model used by SB-EMS.

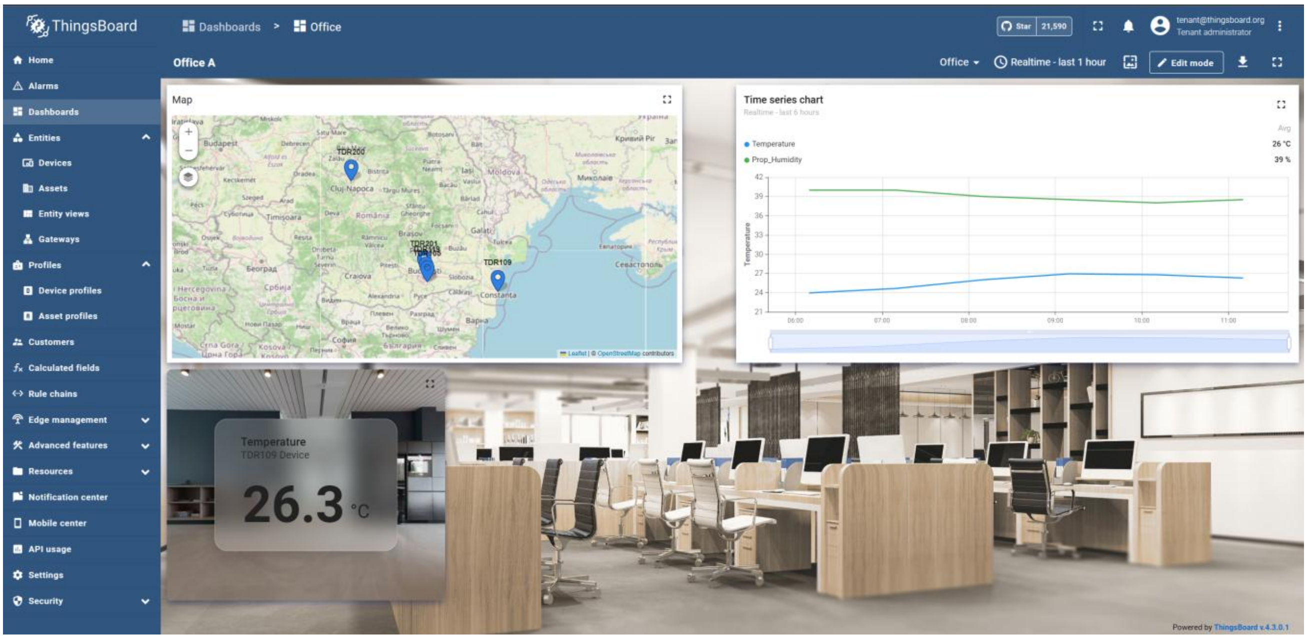

Figure 33. ThingsBoard IoT platform installed in the ORO lab.

Building upon the completed infrastructure setup, the full SB-EMS baseline has been deployed and successfully tested within the laboratory environment. The ThingsBoard IoT platform is fully operational, with all device entities from the ORO lab provisioned to enable continuous telemetry ingestion and real-time visualization through dedicated dashboards. Complementing this, the Monitoring Platform has been instantiated, featuring validated plugins that automate telemetry retrieval from ThingsBoard and ensure data persistence in InfluxDB.

Furthermore, the initial version of the core SB-EMS application services and functional components has been deployed and verified. This includes the SB-EMS closed loop functions, which actively process environment data and power consumption to compute optimal building configurations that balance energy savings and user comfort. The Human-In-The-Loop (HITL) framework, comprising the Notification Manager, User Mobile App (depicted in Figure 34 and), and the reinforcement learning-based Self-Tuning system, has also been integrated to collect and process occupant feedback for dynamic policy updates.

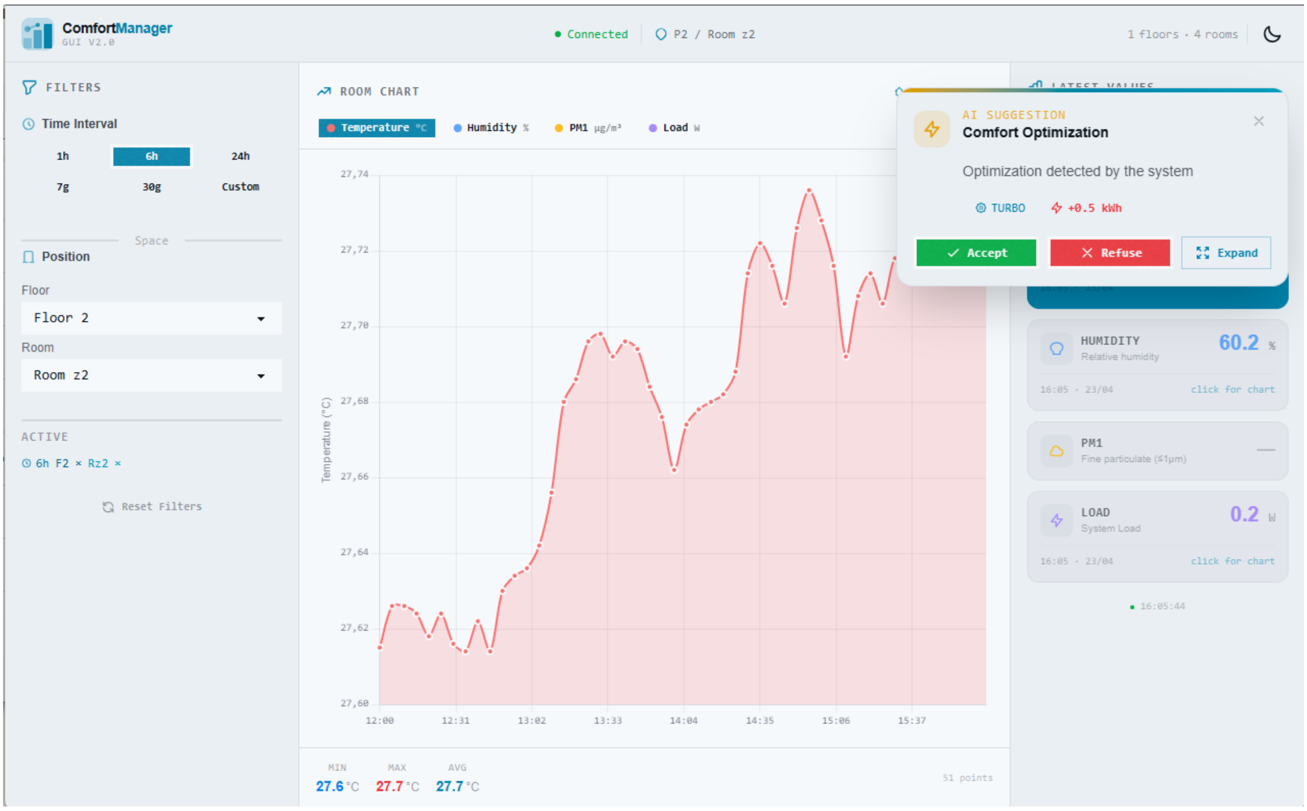

Figure 34. Comfort manager graphical user interface.

Figure 35. Example of notification sent by the SB-EMS notification manager to the User App.

By the end of Q1 2026, the B5G infrastructure, edge computing environment, IoT integration and basic intent-based management were functionally integrated and validated in lab conditions. This baseline ensured that the April 2026 pre-trial was focused on end-to-end validation of SB-EMS data collection, prediction and control workflows over the prepared B5G network.

Integration of AI-based classification services

In parallel with the ORO testbed integration, a preliminary version of the intent processing chain was integrated in the CAPG testbed, allowing high-level intents related to latency or reliability to be translated into slice selection and QoS configuration actions. Closed-loop monitoring mechanisms were connected to network telemetry sources (RAN and core KPIs) and to application-level metrics. Initial tests confirmed that deviations from expected KPIs can trigger alarms and basic corrective actions through the orchestration layer.

The ML catalogue and MLOps platform were deployed in initial form, enabling registration and deployment of ML models for inference at the edge. For the pre-trial, focus was placed on validating model onboarding, containerized deployment and interaction with the Closed Loop Analytics function, without yet activating full self-tuning or federated learning features.

Building on the initial deployment of the ML catalogue and MLOps platform, CAPG contributed to the E1 pre-trial activities by implementing an AI-based traffic profile classification component for CPE/UE telemetry. This component extends the monitoring capabilities of the platform by enabling the classification of short-term traffic windows into operational profiles, supporting predictive monitoring and service-awareness functions within the E1 use case.

The implemented solution follows an AI-aaS architecture, combining historical telemetry ingestion, fixed-window aggregation and containerized model serving. The objective of this integration is to provide real-time classification capabilities that can be leveraged by Closed Loop Analytics and intent-based management functions.

The solution is organized in a modular architecture composed of three main layers:

- Data and Processing Layer, responsible for ingesting historical CPE telemetry and preparing fixed 60-second windows through data cleaning, temporal ordering and feature engineering.
- MLOps and Training Layer, where multiple classification models were evaluated, including baseline approaches, Logistic Regression, Histogram Gradient Boosting and Random Forest. MLflow is used for experiment tracking, model versioning and artifact management.
- Serving and Inference Layer, where trained models are deployed as classification services using BentoML endpoints, enabling both scripted validation and real-time inference requests.

A monitoring stack based on Prometheus and Grafana is used to collect system metrics and ensure observability of the deployed AI services.

The architecture of the implemented AI classification pipeline is illustrated in Figure 36.

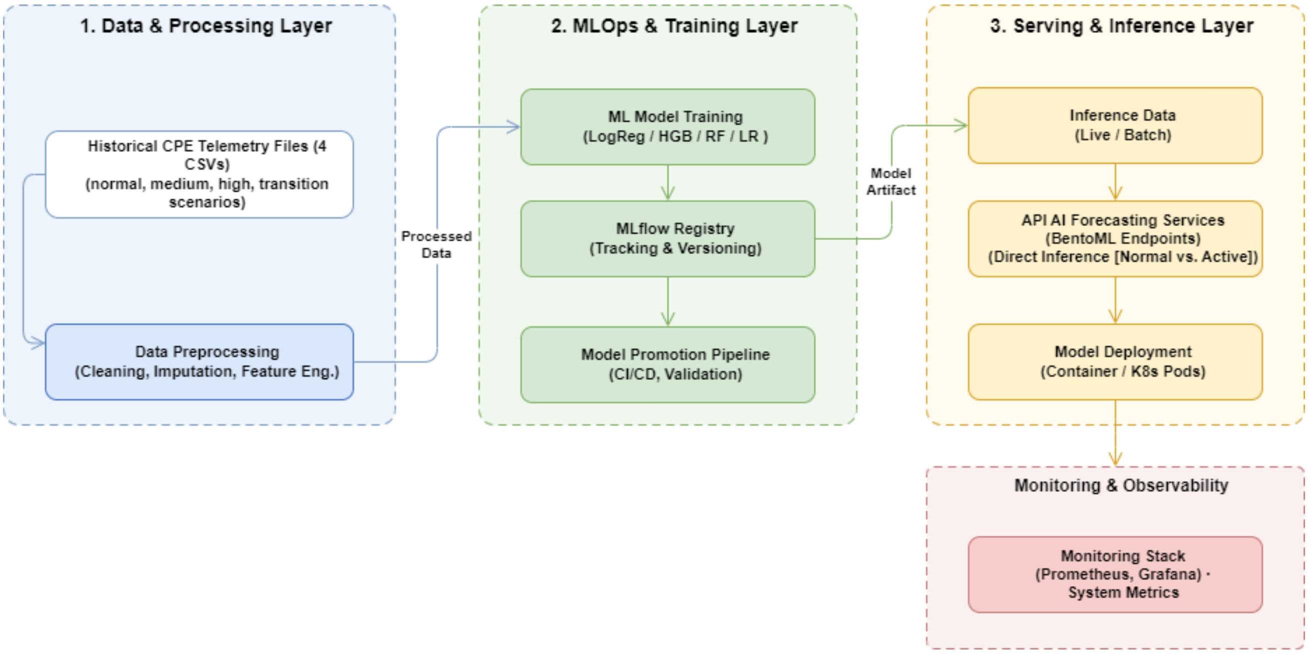

Figure 36. AI-based CPE traffic profile classification pipeline implemented for the E1 use case.

4.1.2 Preliminary test results and initial demo

The pre-trial activities for the E1 use case confirmed that the SB-EMS has been successfully deployed and functionally validated within the ORO 5G Lab environment. The laboratory integration verified the E2E operational chain, encompassing the B5G infrastructure, edge computing Kubernetes clusters, and the ThingsBoard IoT platform for telemetry ingestion and device command execution. Key functional components, including the occupancy-centered control closed loop and the human-in-the-loop feedback mechanisms, were tested under lab conditions to establish a stable baseline. This preliminary validation ensures that all core subsystems ready for end-to-end performance trials over the prepared 5G private network. The E1 configuration for pre-trial execution is presented in Table 19.

Table 19. E1 pre-trial execution.

E1.1 / Lab01	Trial 12/E1 – Renewable Energy Communities
Facility / Site	Orange 5G Lab in Bucharest; CAPG LAB E1.1/Lab01
Objective / Description	Joint optimization of energy consumption and user comfort in single smart buildings and REC scenarios by automatically finding the optimal compromise between these two objectives
Executed By	NXW, ORO, CAPG
Devices / Components Involved	Environmental sensors (temperature, humidity, air quality)
5G Infrastructure	5G SA private network with indoor and outdoor coverage at the UPB Campus, utilizing network slicing (URLLC and mIoT profiles);
Targeted KPIs	Service provisioning latency, MLOps pipeline automation ratio, Occupancy prediction accuracy
Measurement Tools	Derived from system logs of the service orchestrator tool and from MLOps operational dashboard
Ethics / Beta Users	No GDPR data collected
Technical Enablers	Compute-as-a-Service enablers: Intelligent service and resource orchestration in extreme edge/edge/cloud compute continuum Deployed at CAPG LAB E1.1/Lab01: IBN Framework; Agentic AI; AI-aaS

To evaluate the efficiency of the orchestration layer within the E1 architecture, a series of experimental runs were conducted to measure the Service Provisioning Latency. This KPI defines the time elapsed from the initial service request issued via the Service Orchestrator REST APIs to the point where the SB-EMS application functions are fully instantiated and operational in the emulated edge environment.

The SB-EMS service was instantiated ten times under consistent laboratory conditions at the Orange 5G Lab. These runs covered the complete deployment workflow, including container provisioning on Kubernetes worker nodes, the deployment of ML models for inference through the MLOps platform, and the activation of monitoring jobs. The preliminary results, summarized in Table 20, demonstrate a high degree of deployment stability. Across the ten iterations, the system achieved an average provisioning time of 15 seconds with a standard deviation of 2 seconds, significantly outperforming the project's target threshold of <30 seconds.

Table 20. E1 Service provisioning latency measurements

Run ID	Provisioning Time (s)	Target KPI	Status
Run 01	14	<30s	Pass
Run 02	17	<30s	Pass
Run 03	15	<30s	Pass
Run 04	12	<30s	Pass
Run 05	18	<30s	Pass
Run 06	16	<30s	Pass
Run 07	13	<30s	Pass
Run 08	15	<30s	Pass
Run 09	17	<30s	Pass
Run 10	13	<30s	Pass
Average	15	<30s	Pass

These results confirm that the orchestration mechanisms, including the integration between the Service Orchestrator, Resource Orchestrator, and the Kubernetes-based edge nodes, are capable of meeting the low-latency requirements necessary for agile energy management services.

In addition to deployment latency, the automation of the machine learning lifecycle was evaluated to verify the MLOps pipeline automation ratio. Since the SB-EMS closed-loop functions rely entirely on AI models managed through the MLOps platform for tasks such as room occupancy and energy consumption forecasting, full automation of model lifecycle tasks is critical. During the preliminary integration tests, the system achieved an automation ratio of 100%, exceeding the target of >90%. This was evidenced by the platform's ability to automatically handle model onboarding, containerized deployment to the edge nodes, and the triggering of inference services without manual intervention.

Furthermore, the occupancy prediction accuracy was assessed using reports generated by MLflow, which is utilized within the architecture for experiment tracking, model versioning, and lifecycle management. Based on four distinct experiments conducted during the preliminary integration phase, the system consistently demonstrated high performance in identifying room occupancy states. The models achieved a measured occupancy prediction accuracy of 96%, successfully surpassing the initial project target of >90%. This high level of accuracy ensures that the SB-EMS can make reliable decisions regarding comfort and energy optimization based on actual building usage patterns. Figure 37 shows the metrics of one of the four runs of the occupancy model training.

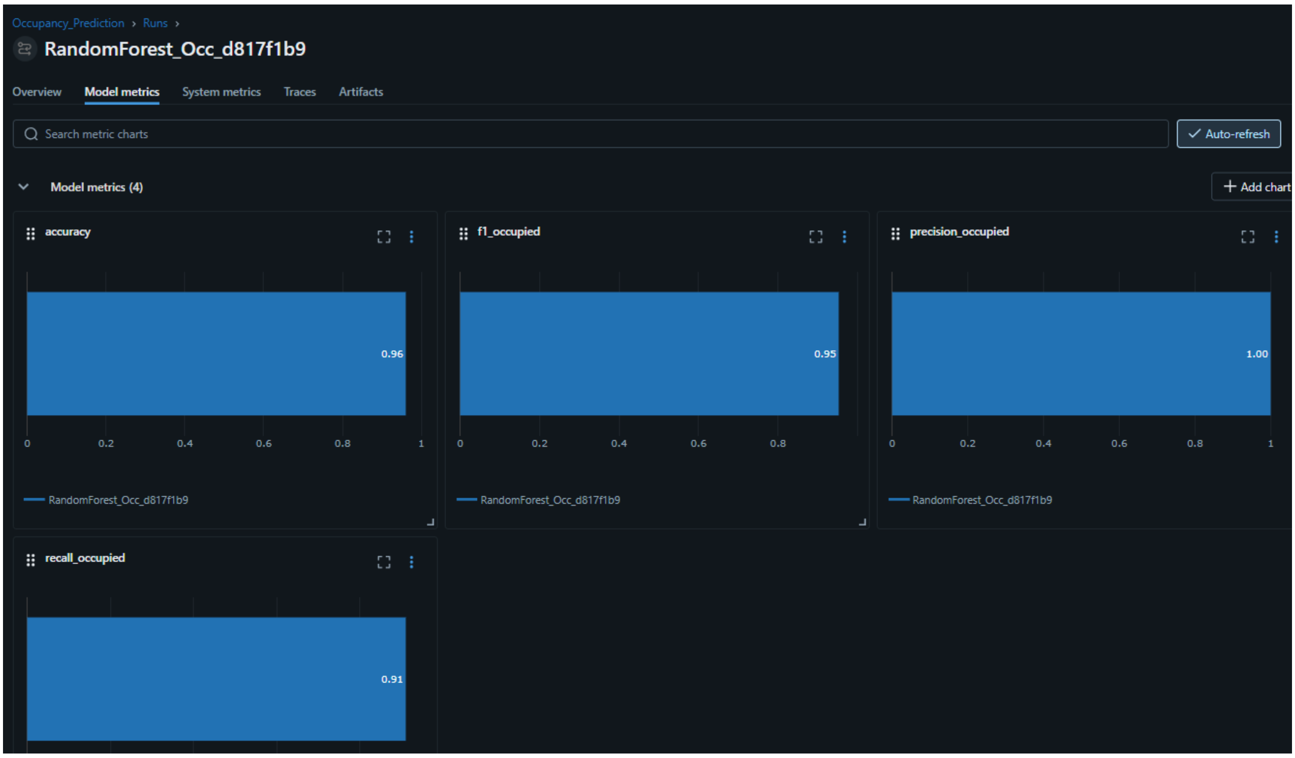

Figure 37. Metrics of one of the 4 runs of the Occupancy Predictor Model.

Initial experiments were conducted using historical operational data provided offline by the Orange 5G Lab for the E1 use case. The dataset contains approximately 8,700 time-series samples across multiple traffic scenarios, including normal, medium, high, and transition conditions. After fixed-window aggregation using 60-second windows, the dataset was transformed into structured samples for classification and split into training, validation, and testing subsets to ensure robust evaluation of model generalization. Feature engineering was applied to capture traffic dynamics within each window, including temporal aggregation statistics such as mean, standard deviation, percentiles, and trends, along with volume and throughput descriptors for both uplink and downlink traffic. In addition, derived features were introduced to characterize traffic asymmetry and burstiness, enabling a more comprehensive representation of network behavior. Several ML models were evaluated for binary traffic profile classification. The results presented in Table 21 indicate that the Logistic Regression model provides the best generalization performance, achieving a test macro F1-score of 0.93 and a balanced accuracy of 0.95.

Table 21. Performance of the evaluated classification models for the 60-second window configuration.

Model	Train Macro F1	Val Macro F1	Test Macro F1	Test Balanced Accuracy	Test Accuracy
Logistic Regression	1	1	0.930	0.951	0.935
Baseline	1	1	0.911	0.937	0.916
HistGradientBoosting	1	1	0.911	0.937	0.916
Random Forest	1	1	0.911	0.937	0.916

To further analyze the classification performance, Table 22 presents the confusion matrix of the deployed model on the test dataset.

Table 22. Confusion matrix of the deployed classification model on the test dataset.

True label	Predicted normal	Predicted active
normal	36	0
active	7	64

The results show that all normal traffic windows are correctly classified, while only a small number of active samples are misclassified. This confirms the robustness of the model in distinguishing between normal and active traffic conditions. The evaluated models also demonstrate stable behaviour in transition scenarios, correctly identifying both the initial normal state and the final active state, which is essential for enabling reliable context-aware decision-making. These results demonstrate the feasibility of integrating AI-based traffic profile classification into the E1 monitoring platform, supporting predictive analytics and service-aware orchestration within the B5G infrastructure.

In addition to the model evaluation results, the AI-based classification service was deployed and monitored using an observability stack based on Prometheus and Grafana. Figure 38 presents the runtime monitoring dashboard of the AI-aaS classification component.

Figure 38. AI-aaS observability dashboard for the E1 classification service.

The monitoring results confirm that the deployed classification service operates in a stable manner, maintaining consistent availability and low resource utilization, supporting its integration within the E1 monitoring and control platform.

These initial integration results validate the deployment of AI-based classification services within the E1 architecture and provide the foundation for their use in closed-loop control and intent-driven network management during the upcoming pre-trial activities.

Table 23 presents the E1 KPIs measurements during the E1 pre-trial.

Table 23. E1 KPIs measurements.

Test Case ID	Target KPIs / Requirements	KPIs measurements	Evaluation / Notes
E1.1_Lab01_01	KPI_AI_1: < 2 s KPI_AI_2: > 0.9	KPI_AI_1: 0.02 s KPI_AI_2: 0.935	Inference time shorter than specified Accuracy above minimum.
E1.1_Lab01_02	KPI_VA_5: 30s	KPI_VA_5: 15s	Average of 10 runs
	KPI_VA_3: 20 ms	10 ms	
	KPI_VA_1:99.999%	100%	
	KPI_VA_11:10 Mbps	100 Mbps	
E1.1_Lab01_03	KPI_AI_2: 90%	KPI_AI_2: 100%	
E1.1_Lab01_04	KPI_AI_2: 90%	KPI_AI_2: 96%	

Test Case ID: E1.1_Lab01_01

The AI-related KPIs were successfully achieved. The measured inference time of 0.02 s remained well below the target threshold of 2 s, while the model accuracy of 0.935 exceeded the minimum required value of 0.9.

Test Case ID: E1.1_Lab01_02

- Objective: Validate the SB-EMS service provisioning latency using the integrated orchestration platform.
- Target KPI: Service provisioning latency < 30 seconds.
- KPI Measurement: 15 seconds (average of 10 runs).

Conclusions after the tests: The tests confirm that the orchestration mechanisms, including the integration between the Service Orchestrator, Resource Orchestrator, and the Kubernetes-based edge nodes, are capable of meeting the low-latency requirements necessary for agile energy management services. The consistent performance across multiple runs indicates a robust automated lifecycle management process.

Test Case ID: E1.1_Lab01_03

- Objective: Validate the degree of automation in the machine learning lifecycle, specifically for model onboarding and containerized deployment to edge nodes.
- Target KPI: MLOps pipeline automation ratio > 90%.
- KPI Measurement: 100% automation ratio.

Conclusions after the tests: The MLOps platform demonstrated full capability to handle model lifecycle tasks without manual intervention. This validates that the system can autonomously manage the AI models required for continuous closed-loop operation.

Next steps: Implement and integrate full self-tuning and re-training logic based on user feedback.

Test Case ID: E1.1_Lab01_04

- Objective: Evaluate the performance of the AI models in identifying room occupancy states to ensure reliable comfort and energy optimization.
- Target KPI: Occupancy prediction accuracy > 90%.
- KPI Measurement: 96% accuracy (measured via MLflow across 4 conducted experiments).

Conclusions after the tests: The high measured accuracy of 96% confirms that the models can effectively capture building usage patterns. This ensures the SB-EMS can make trustworthy decisions regarding room settings.

Looking forward, the evolution towards 6G is essential to support ultra-scalable, AI-native, and intent-driven energy management systems, enabling real-time distributed intelligence, tighter integration of edge–cloud resources, and seamless operation across large-scale multi-building and Renewable Energy Community scenarios.

Next steps: Transition from laboratory data to live telemetry ingestion from building sensors to validate performance under operational conditions. The next phase will focus on the integration of the AI-based classification and intent-driven control components into the ORO 5G Lab testbed, enabling a complete end-to-end operational workflow that combines real-time monitoring, traffic classification and closed-loop decision-making over the B5G infrastructure.

This will be complemented by the transition from offline data processing and validation in CAPG's internal environment to real-time execution within the Orange testbed, where live telemetry ingestion, AI-aaS deployment, and interaction with the IBN framework will be validated under operational conditions.

Early dissemination and demos.

The first outputs from the E1 use case integration and testing activities will be presented during the EuCNC 2026 in Malaga in June 2026. Building on these initial results, further dissemination activities will be planned starting in September 2026, ensuring alignment with the broader timeline of the Romanian cluster and close synchronization with related use cases, including E1. These coordinated efforts aim to maximize visibility, foster cross-use-case validation, and support the consolidation of results across the cluster.

4.2 Robotized offshore wind turbine blade inspection and maintenance

In this section, the integration status and preliminary test results for the E2 use case are described.

4.2.1 Integration status

For the E2 use case, the integration work started with deploying the B5G network infrastructure at the TNO lab in The Hague. In March 2026, the Amarisoft Callbox Classic was successfully installed and configured, as shown in Figure .

Figure 39. Amarisoft Callbox deployed at TNO lab.

The callbox acts as a 5G SA small cell. It is configured to operate at 3925 MHz according to the obtained spectrum license, with a carrier bandwidth of 50 MHz.

The Quectel 5G modem is placed on the opposite side of the lab (as shown in Figure), with a distance of approximately 8 m from the Amarisoft gNB. Preliminary connectivity tests have been performed to ensure that the 5G modem can connect to the Amarisoft gNB. The results of these tests are presented in Section 4.2.2.

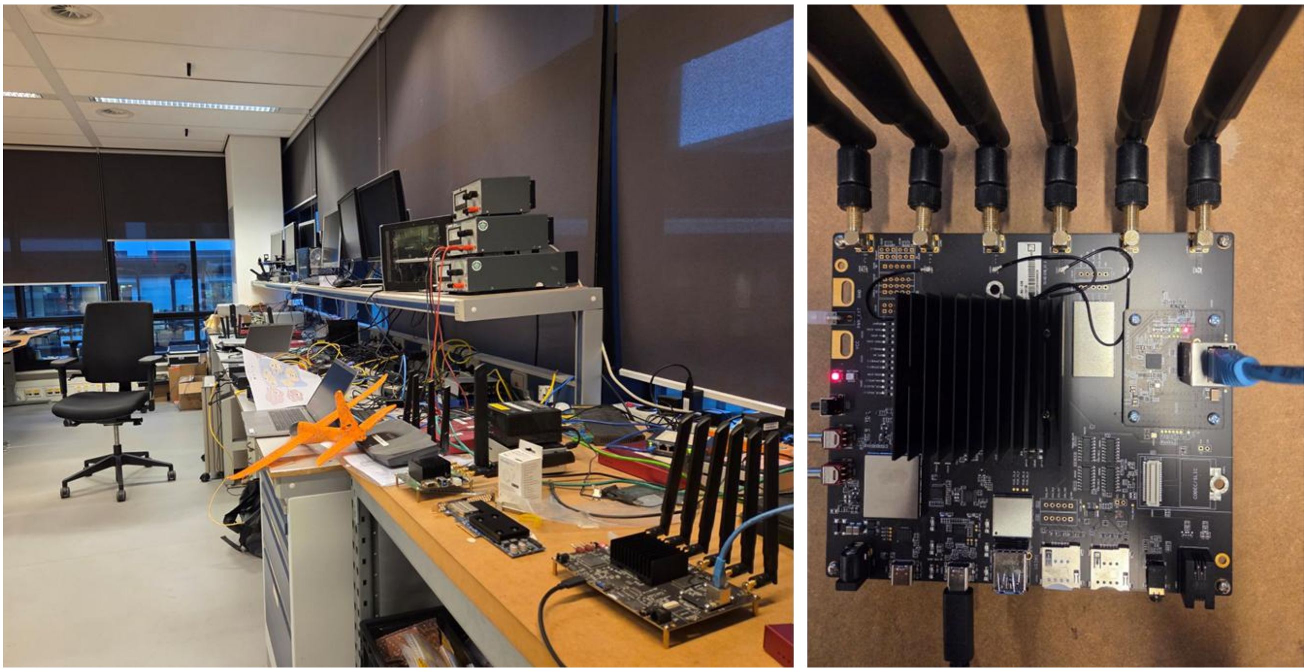

Figure 40. Quectel 5G modem deployed at TNO lab.

In terms of the E2 technology enablers, the current focus is on the GNSS and RTK positioning. As described in Section 2.2.4, we have decided to use a second GNSS/RTK module as a base station in order to provide the RTK correction data to the GNSS/RTK module mounted on the drone. The development steps and integration results are presented in Section 4.2.2.

4.2.2 Preliminary test results and initial demos

After the successful installation of the B5G small cell at the TNO lab in March 2026, the uplink data rate and latency KPIs measurements were performed. The service availability KPI will be tested in the next phase.

Table 24 illustrates the E2 pre-trial configuration.

Table 24. E2 pre-trial configuration.

E2.1/Lab01	E2 – Robotized offshore wind turbine blade inspection and maintenance
Facility / Site	TNO lab in The Hague, The Netherlands
Objective / Description	Validation of B5G connectivity and collecting preliminary KPIs such as uplink data rate and end-to-end latency. Initial assessment of the GNSS and RTK positioning accuracy.
Executed By	TNO
Devices / Components Involved	Quectel 5G modems and u-blox GNSS and RTK modules
5G Infrastructure	Amarisoft Callbox Classic acting as a 5G SA small cell
Targeted KPIs	Uplink data rate, end-to-end latency and positioning accuracy.
Measurement Tools	Ping, iPerf and localization software
Ethics / Beta Users	Not applicable
Technical Enablers	GNSS and RTK positioning

In March 2026, the uplink data rate and latency values were measured in the lab and the results are presented in Table 25. It can be seen that the uplink data rate is below the target value of 50 Mbps. This is mainly due to the configured downlink and uplink frame structure of DDDDDDDSUU, which is downlink heavy.

New tests were performed in April 2026 with a new frame structure of DDDSUUUUUU which has three times as many uplink slots compared to the previous configuration. As shown in Table 25, the uplink data rate is now around the target value. It is important to note that the measured KPIs can vary depending on the environment and the distance between the UE and gNB.

Table 25. E2 KPIs pre-trial measurements.

Test Case ID	Target KPIs / Requirements	KPIs measurements	Evaluation / Notes
E2.1_Lab01_01 (TDD frame structure DDDDDDDSUU)	KPI_VA_11: Uplink data rate: > 50 Mbps KPI_VA_3: One-way end-to-end latency: < 100 ms	KPI_VA_11: 14.5 Mbps KPI_VA_3: 10.3 ms	Uplink data rate is below target. Latency fulfils the target.
		KPI_VA_11: 12.4 Mbps KPI_VA_3: 10.9 ms	Uplink data rate is below target. Latency fulfils the target.
		KPI_VA_11: 14.0 Mbps KPI_VA_3: 9.9 ms	Uplink data rate is below target. Latency fulfils the target.
		KPI_VA_11: 17.6 Mbps KPI_VA_3: 8.9 ms	Uplink data rate is below target. Latency fulfils the target.
		KPI_VA_11: 16.6 Mbps KPI_VA_3: 7.9 ms	Uplink data rate is below target. Latency fulfils the target.

E2.1_Lab01_02 (TDD frame structure DDDSUUUUUU)	KPI_VA_11: Uplink data rate: > 50 Mbps KPI_VA_3: One-way end-to-end latency: < 100 ms	KPI_VA_11: 51.0 Mbps KPI_VA_3: 10.9 ms	Both uplink data rate and latency fulfil the target values.
		KPI_VA_11: 51.4 Mbps KPI_VA_3: 12.9 ms	Both uplink data rate and latency fulfil the target values.
		KPI_VA_11: 51.9 Mbps KPI_VA_3: 9.9 ms	Both uplink data rate and latency fulfil the target values.
		KPI_VA_11: 53.4 Mbps KPI_VA_3: 8.9 ms	Both uplink data rate and latency fulfil the target values.
		KPI_VA_11: 55.6 Mbps KPI_VA_3: 7.9 ms	Both uplink data rate and latency fulfil the target values.
E2.1_Lab01_03 (GNSS/RTK)	KPI_VA_24: Positioning accuracy: < 10 cm	KPI_VA_24: 1.72 cm where horizontal and vertical accuracy are 1.4 cm and 1.0 cm respectively.	Positioning accuracy fulfils the target.

The KPI measurements showed that uplink throughput and latency targets can be achieved with appropriate configuration. These results validate the feasibility of supporting real-time drone-based inspection and data transmission in controlled environments. However, the transition towards 6G is essential to enable large-scale offshore operations, providing enhanced uplink capacity for high-resolution sensor and video data and integrated positioning capabilities required for precise and safe drone navigation in complex maritime environments.

On the GNSS/RTK positioning enabler part, the initial implementation has resulted in an overall accuracy of 1.72 cm. The work has been performed following a step-wise approach where each step is built on its previous step, starting from the most basic configuration and progressively moving towards a fully autonomous, wireless and remotely monitored GNSS/RTK system.

The first step was to validate the fundamental RTK link using the simplest possible configuration. Both u-blox EVK-X20P modules were connected via a USB-UART (Universal Asynchronous Receiver-Transmitter) cable to a laptop running the u-center-2 software, which is the official u-blox GNSS evaluation software. This configuration is depicted in Figure .

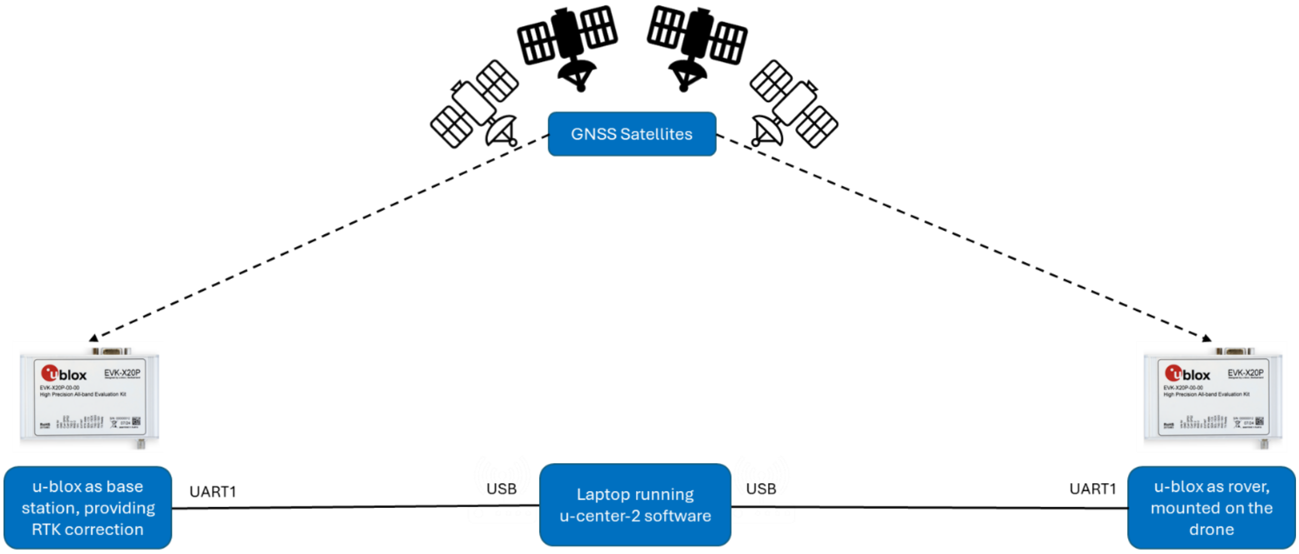

Figure 41. Initial configuration with the u-blox modules connected by cable to a laptop.

In this configuration, the base station and the rover were configured as a NTRIP server and an NTRIP client respectively, with the RTK correction stream being relayed locally via the laptop. NTRIP is an application-level protocol that supports the streaming of GNSS data over the Internet. Furthermore, both u-blox modules were connected to GNSS antennas positioned outdoor to ensure optimal GNSS signal reception.

The test results are shown in Figure and Figure , respectively. Figure shows the u-center-2 Graphical User Interface (GUI) for the base station. On the "Data sources" panel in the top left section, the GUI indicates the base station is configured as an NTRIP server. On the "Satellite Position View" and the "Satellite Signal View" panels in the middle section, the available satellites and the received signal quality are shown. On the "Map View" panel in the upper right section, the position of the base station is given. On the "Data View" panel in the lower right section, the fix mode indicated a time fix as well as the number of satellites used which is 17 out of the 35 available satellites. The shown 3D and 2D accuracy values were rather low, as the screenshot was taken immediately after the setup. The base station usually needs about 10 – 15 minutes to reach a better accuracy.

In Figure , the similar GUI for the rover is shown. On the "Data View" panel, it can be seen that the module had an RTK fix with a 3D position accuracy of 1.72 cm and a 2D accuracy of 1.40 cm. These outcomes confirmed that the hardware configuration and the RTK correction protocol were working correctly.

Figure 42. u-center-2 GUI for the base station configured as an NTRIP server.

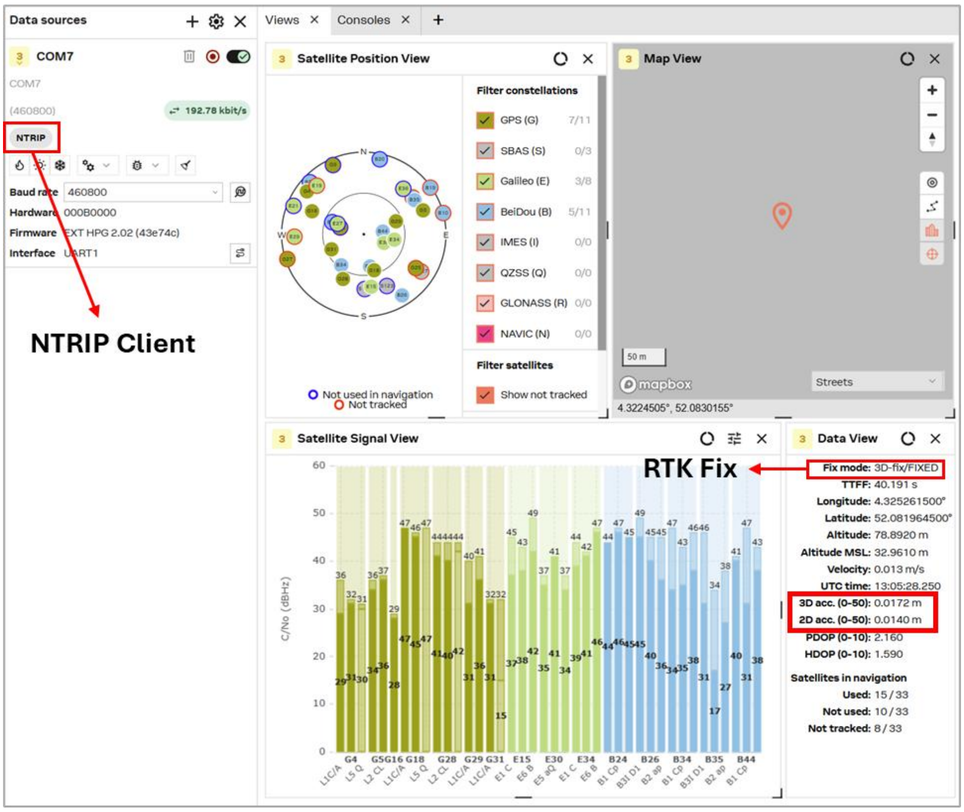

Figure 43. u-center-2 GUI for the rover configured as an NTRIP client.

The second step was to replace the laptop with Raspberry Pi units as shown in Figure Each Pi unit was connected to the EVK-X20P module via UART, using the Future Technology Devices International (FTDI)

USB-to-UART bridge on the EVK board. The two Pis were linked via an Ethernet cable, and both continued running u-center-2 software under Linux for initial verification.

A key technical challenge at this stage was loading the serial driver for the u-blox vendor IDs (0x1546:0x050C and 0x1546:0x050D), which were not recognised by the Linux *ftdi_sio* driver by default. The EVK slide switch had to be set to the Inter-Integrated Circuit (I2C) position for the UART ports to enumerate correctly. The driver issue was resolved with a persistent *udev* rule that loads *ftdi_sio* and injects the vendor ID pair on device insertion. Finally, both devices could communicate seamlessly and the RTK link was functioning as expected.

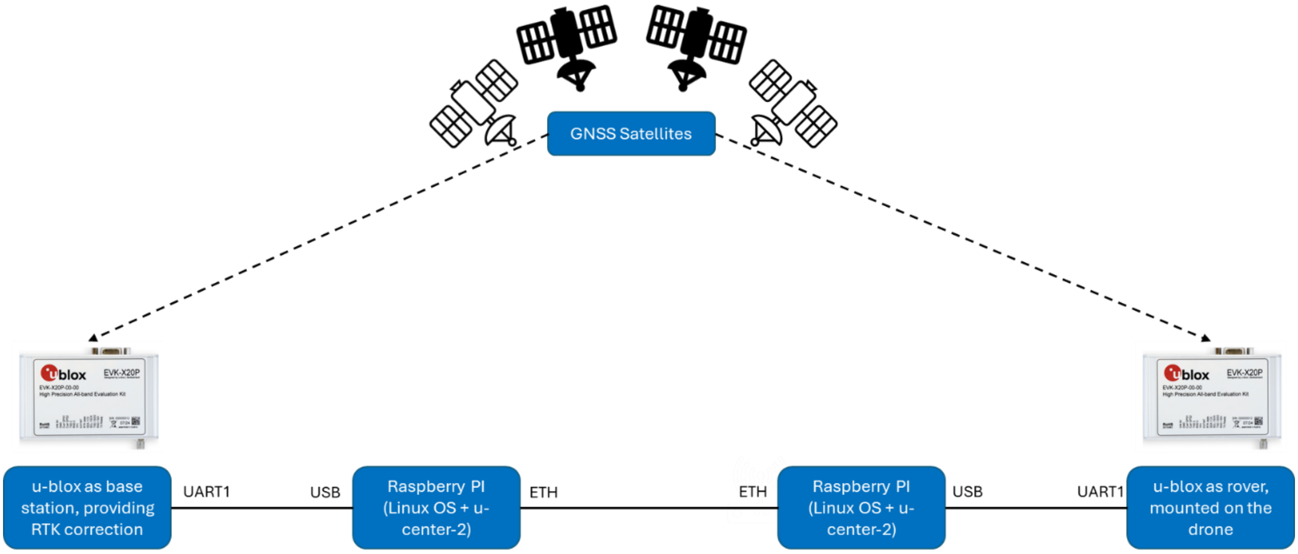

Figure 44. The laptop in the initial configuration is replaced by two Raspberry Pi units.

The third step involved replacing the wired Ethernet link with Wi-Fi as shown in Figure . In addition, the u-center-2 software on the Pi units was replaced entirely with TNO's custom Python-based code with the same functionalities. Since the u-center-2 GUI only works on Windows operating system, we have developed a web-based monitoring dashboard that connects to the API servers on both Pi units. The API server runs as a background thread on both Pi units and exposes a status endpoint that returns the current state as a JSON object containing among others the position, RTK fix state, satellite data, NTRIP connection status, and Pi system metrics. The dashboard application, which runs on a laptop, polls these status endpoints at regular intervals and renders the data without requiring any GNSS software.

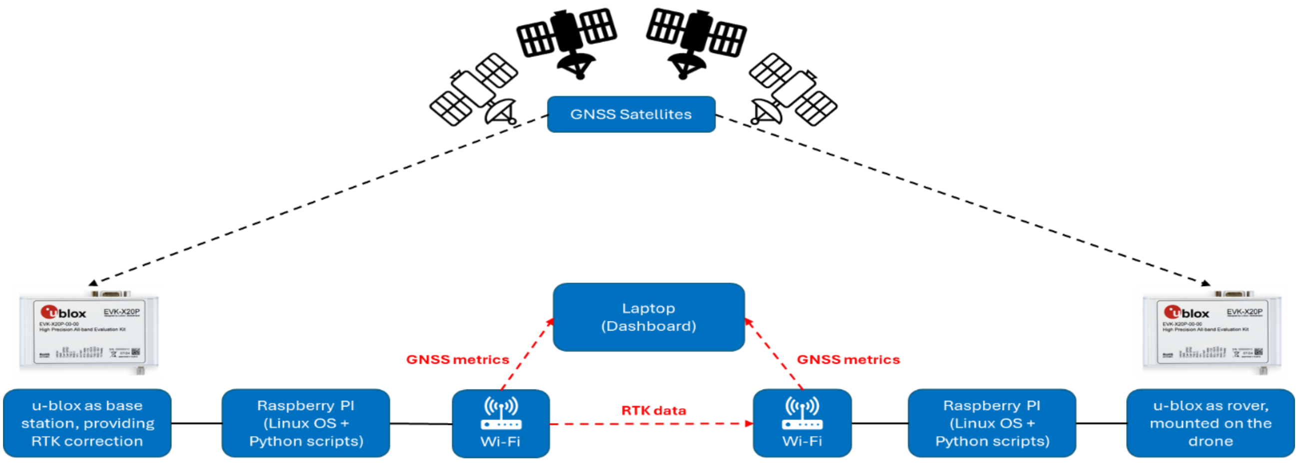

Figure 45. Raspberry PI units are connected via Wi-Fi and providing GNSS metrics to the dashboard.

The dashboard provided a unified view of the GNSS/RTK system status in real time. For visibility purposes, the dashboard is split into multiple panels as shown in Figure 46 to Figure 49. The first panel (Figure 46) displays the positions of the base station and the rover. The second panel (Figure 47) illustrates the satellites visible to the base station along with their signal quality, while the third panel (Figure 50) provides the corresponding view for the rover. Finally, the rover-specific metrics are presented in the fourth panel (Figure 49) which shows the horizontal (hAcc) and vertical (vAcc) position accuracy of 1.4 cm and 1.0 cm are shown, respectively.

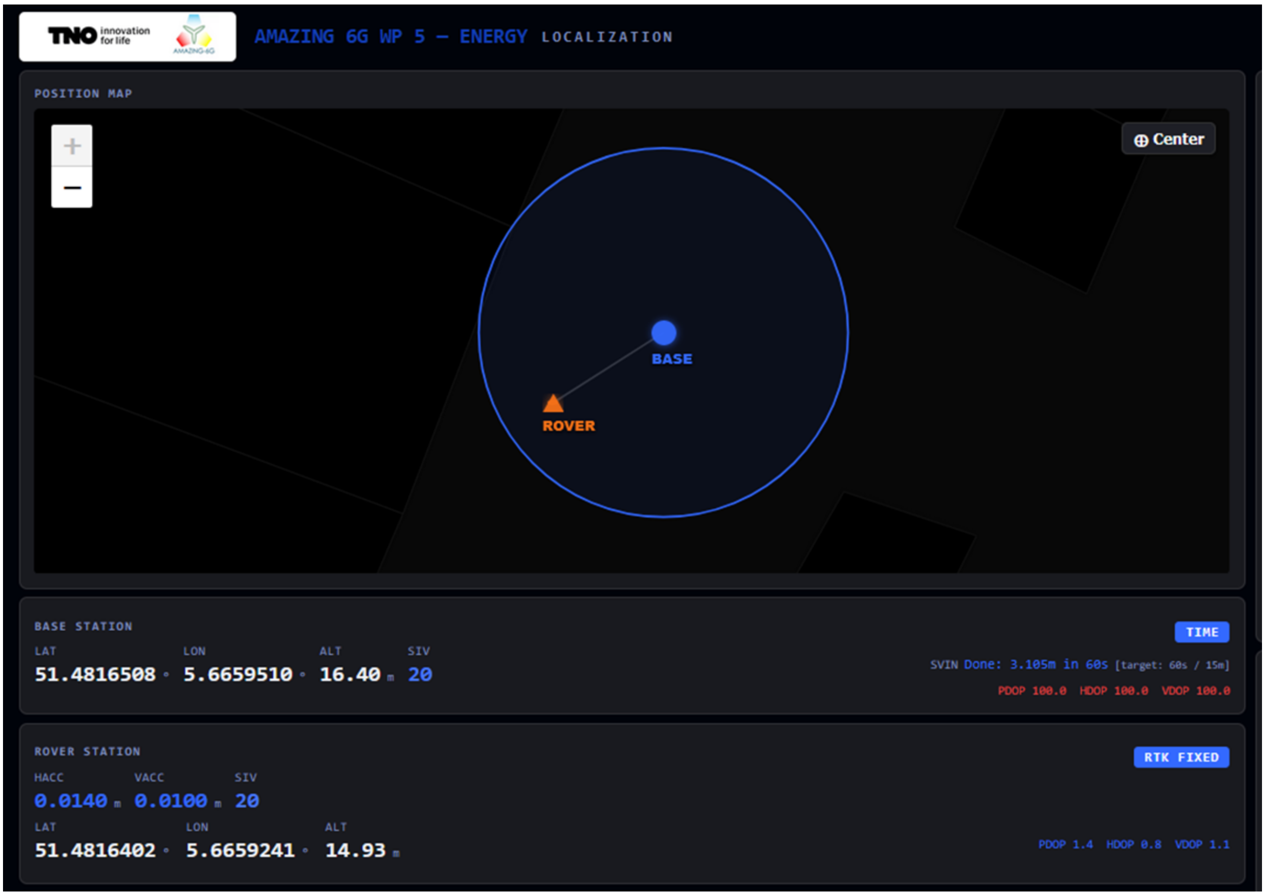

Figure 46. Dashboard panel showing the position of the base station and rover.

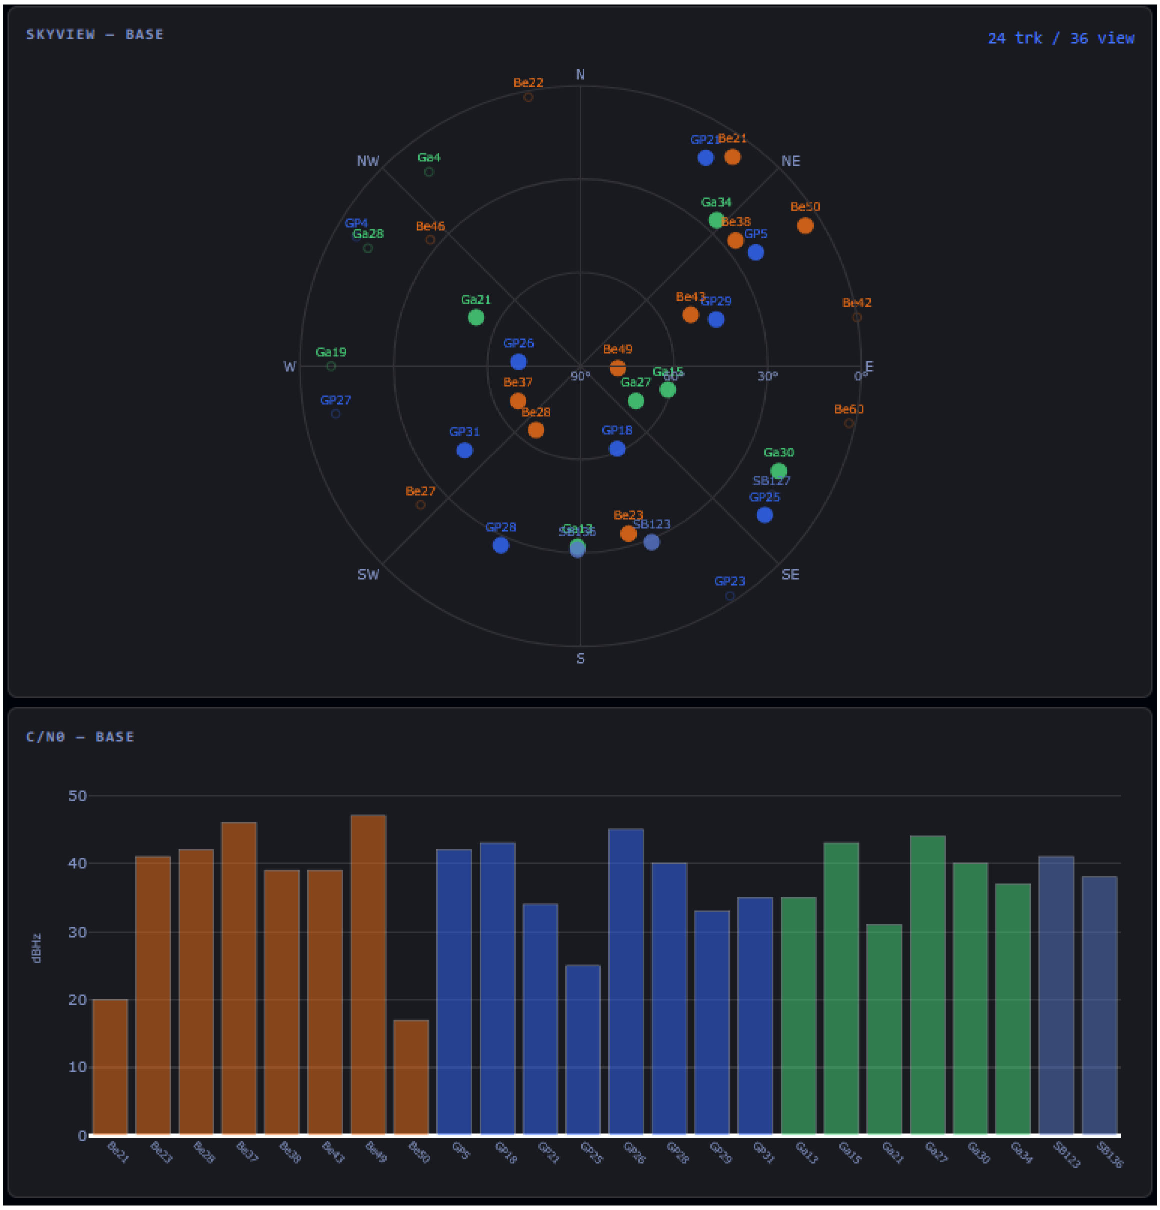

Figure 47. Dashboard panel showing the satellites visible to the base station and their signal quality.

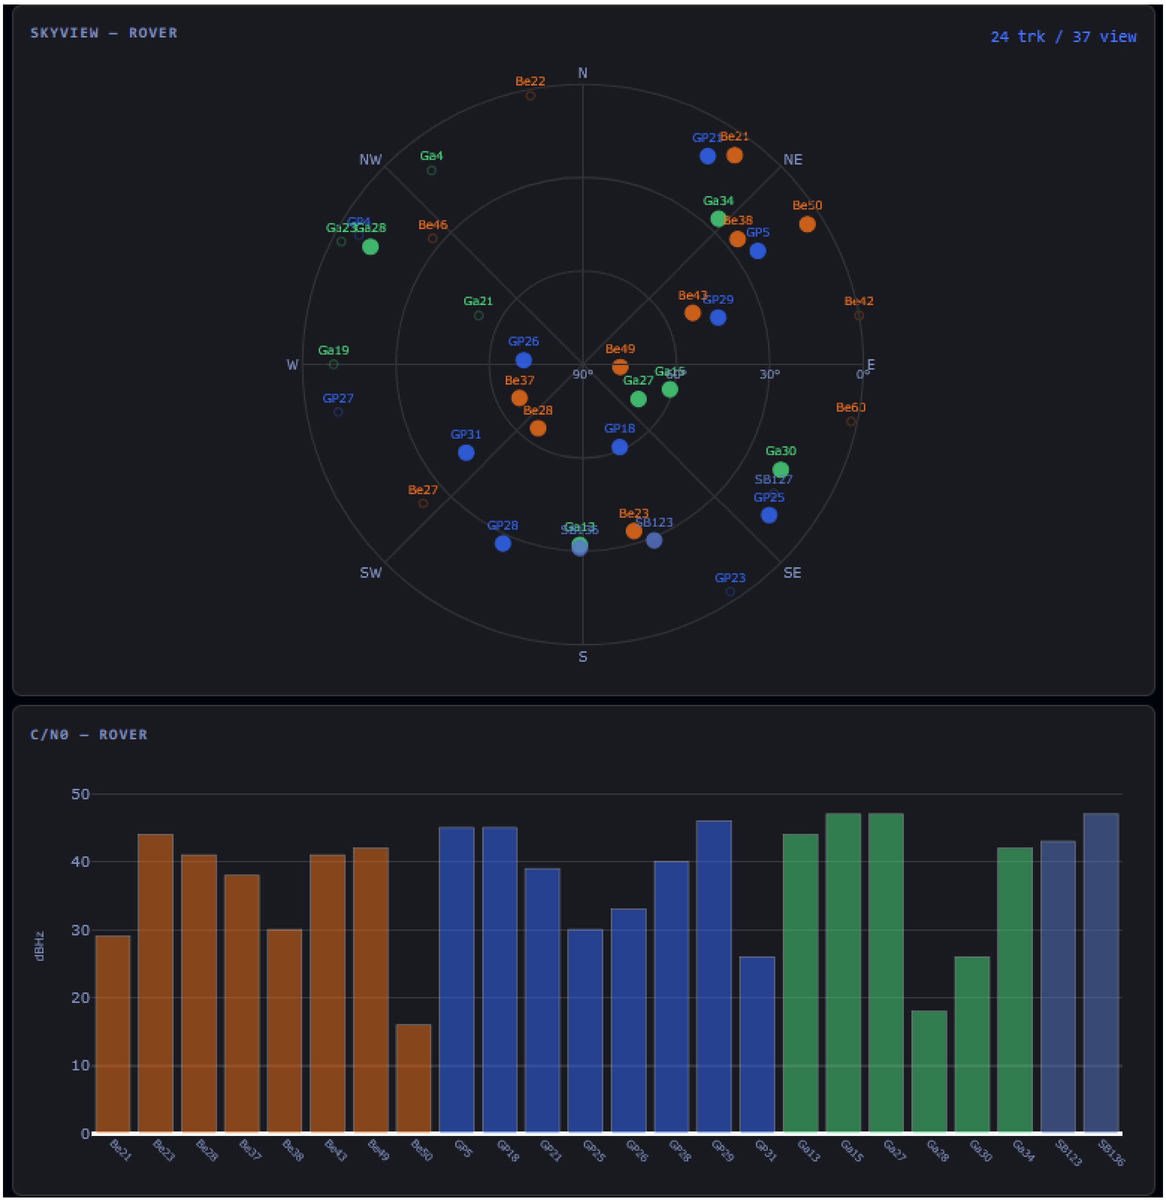

Figure 48. Dashboard panel showing the satellites visible to the rover and their signal quality.

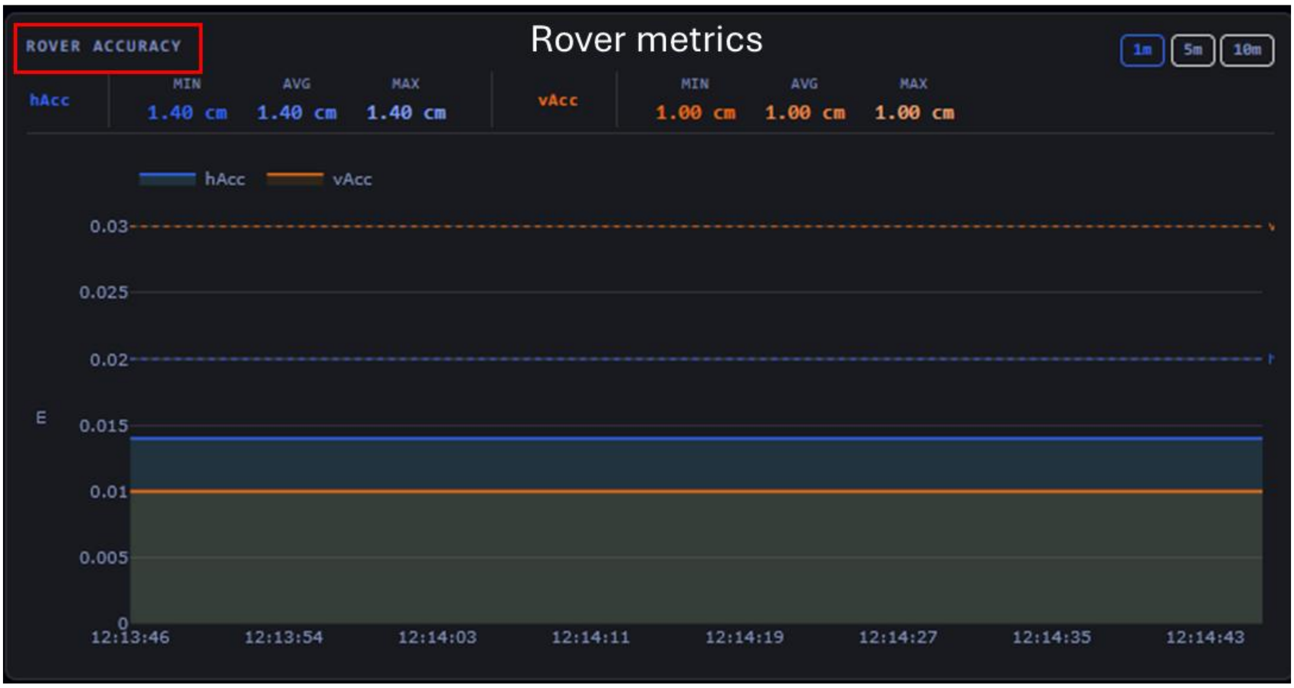

Figure 49. Dashboard panel showing the rover position accuracy.

Overall, the conducted integration work confirmed that combining GNSS positioning with RTK correction services significantly improves positioning accuracy. The test results demonstrated that centimetre-level accuracy was consistently achieved.

Having completed the lab tests, the first field test for the localisation enabler using drone was implemented in end of May 2026 at the Zephyros lab in Vlissingen and full results will be presented in D5.2 (M33). Drone inspections require high uplink data rate to support the transfer of sensor data. However, the current 5G systems are mainly optimised for downlink data rate hence 6G networks supporting gigabits of uplink data rate are required especially when a drone fleet is considered.

Next steps

In the next phase, the GNSS and RTK localisation architecture will be extended to include a cloud server that acts as an internet-accessible NTRIP relay, enabling the rover to receive corrections via cellular network rather than local Wi-Fi network. This is required when flying the drone Beyond Visual Line-of-Sight (BVLOS) where a shared local Wi-Fi network is not available. As depicted in Figure 50, the base station will transmit the RTK correction stream to the server over the internet. The rover will connect to the server as an NTRIP client via a cellular connection. A VPN or secure tunnel can be used to protect the correction stream.

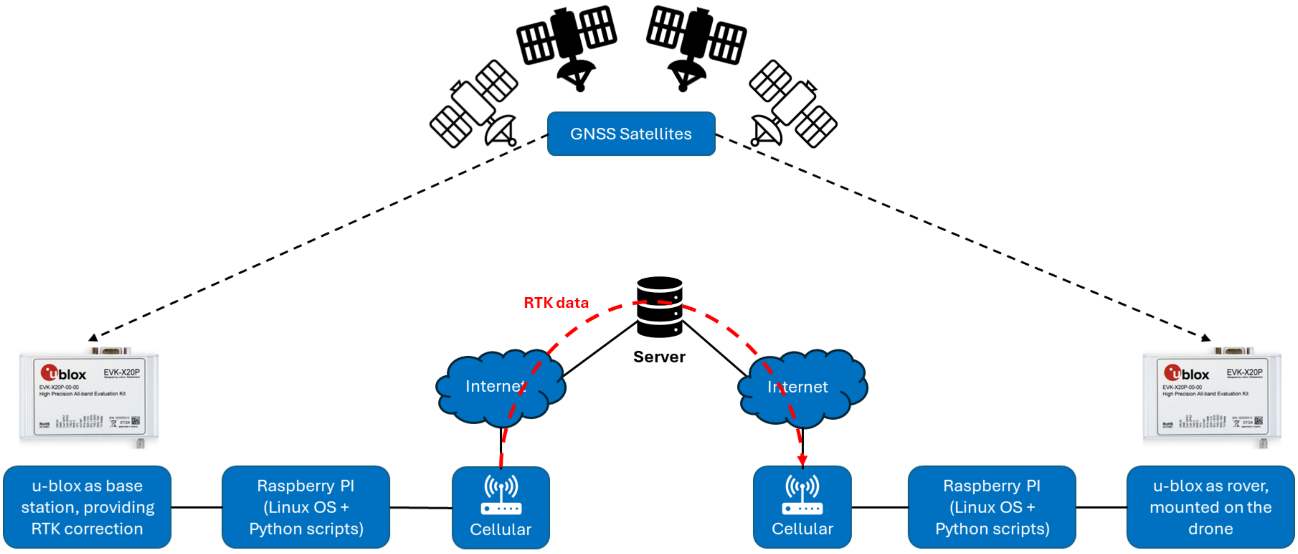

Figure 50. Final architecture with internet-connected base station and rover.

The work on other E2 technology enablers such as network slicing, public-private network integration and CAMARA APIs for quality on demand will continue in the next phase. The integration of all the hardware components (i.e. 5G modem, GNSS/RTK module and ultrasound or similar sensor) into a drone will be completed before the final trial.

The final trial in H1 2027 will demonstrate a B5G-enabled inspection drone capable of performing sensor measurements on a turbine blade. The measurement data will be transmitted to a DT system via the B5G network for analysis, where the results can be used for helping decision making. The development of DT system will be started from H2 2026.

Early dissemination and demos

As part of the dissemination efforts, TNO has presented the E2 use case at the Fieldlab Zephyros AIRTuB-ROMI annual event in the Netherlands (as shown in Figure 51). The event took place on 28 November 2025 in Vlissingen. The two-day annual event[2] showcased advancements in automated offshore wind turbine maintenance, highlighting the move toward unmanned operations. Featuring 14 partners from the AIRTuB-ROMI project, the event demonstrated drones, crawlers, and sensors designed to enable real-time, autonomous inspection and maintenance of wind turbine blades. Figure 52 provides an impression of this event.

2 https://www.worldclassmaintenance.com/fieldlab-zephyros-airtub-romi-tweedaags-jaarevent-offshore-onderhoud-wordt-slim-en-onbemand/

Figure 51. TNO presented the E2 use case at the Fieldlab Zephyros annual event in November 2025.

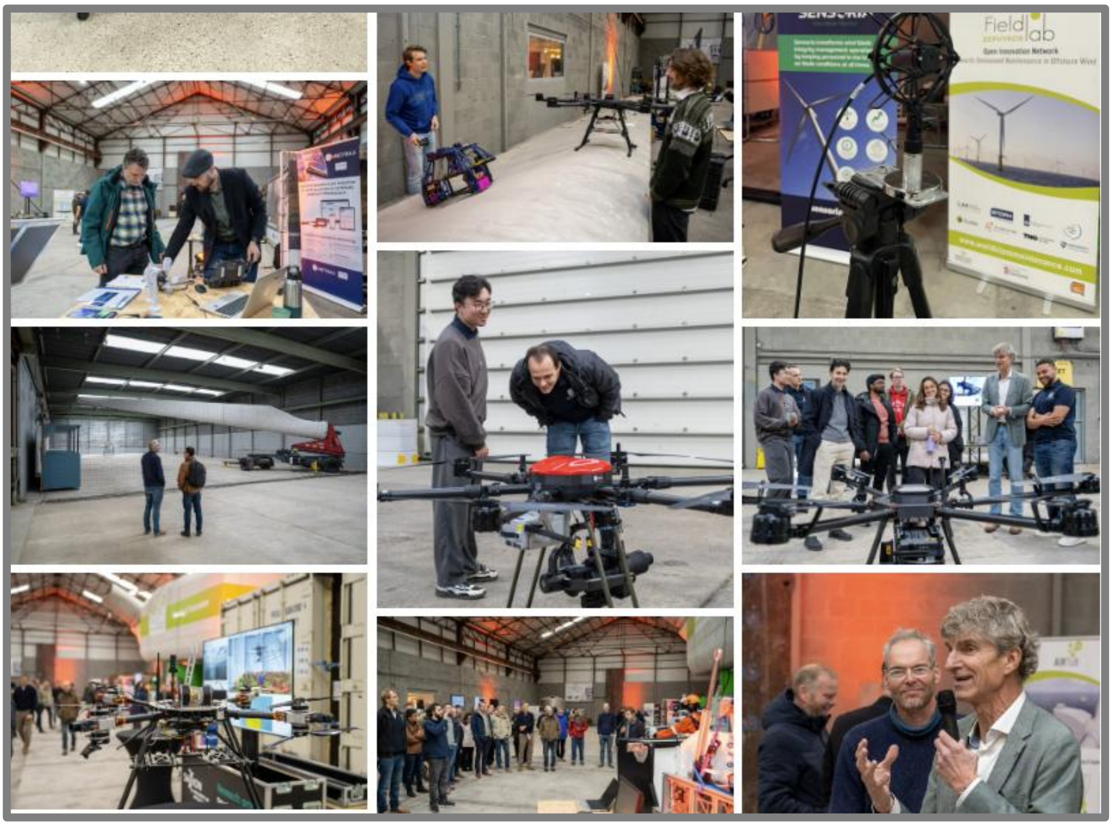

Figure 52. An impression of the Fieldlab Zephyros AIRTuB-ROMI annual event in November 2025.

Looking ahead, the first E2 results will also be presented during the EuCNC 2026 in Malaga.

4.3 Solar energy monitoring, control and predictions using B5G/6G communications and edge-cloud

4.3.1 Integration status

The integration activities performed during January-March 2026 focused on the progressive realization of an E2E system connecting PV field equipment, 5G-enabled IoT gateways, B5G network infrastructure, and the cloud-based energy management platform. The work followed an incremental approach, starting from hardware-level validation and extending toward full system interoperability across the field-network-cloud architecture.

In the initial phase, efforts concentrated on validating the core communication capabilities between the gateway and the inverter. A prototype gateway was implemented using an Orange Pi platform interfaced with the inverter via an RS-485 HAT module to enable Modbus RTU communication with a Fimer PVS-50 inverter. Read and write operations were successfully executed, allowing access to inverter registers and confirming the ability to retrieve telemetry (power, voltage, current, temperature) and send control commands (e.g., power limitation). This phase established the functional baseline for field-device interaction. Figure 53 illustrates the laboratory setup used during this phase, highlighting the physical connection between the gateway and the inverter.

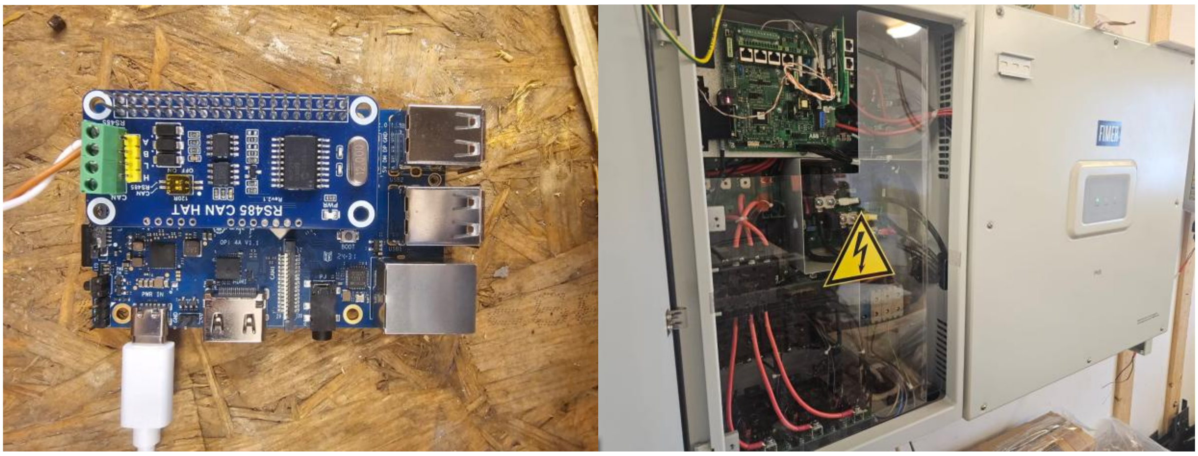

Figure 53. Orange Pi acting as the IoT Gateway and the inverter used for tests in the testbed.

Following this validation, integration activities progressed toward enabling wide-area connectivity via 5G. A RedCap modem compliant with 3GPP Release 17 was introduced using a Quectel development board. The modem was configured in ECM mode and connected via USB to the gateway, exposing a standard IP interface. This allowed the gateway to establish connectivity with the B5G network and exchange data with the cloud platform. Figure 54 presents the gateway prototype with the integrated 5G modem, demonstrating the transition from a locally connected system to a fully networked IoT device capable of remote communication over B5G infrastructure.

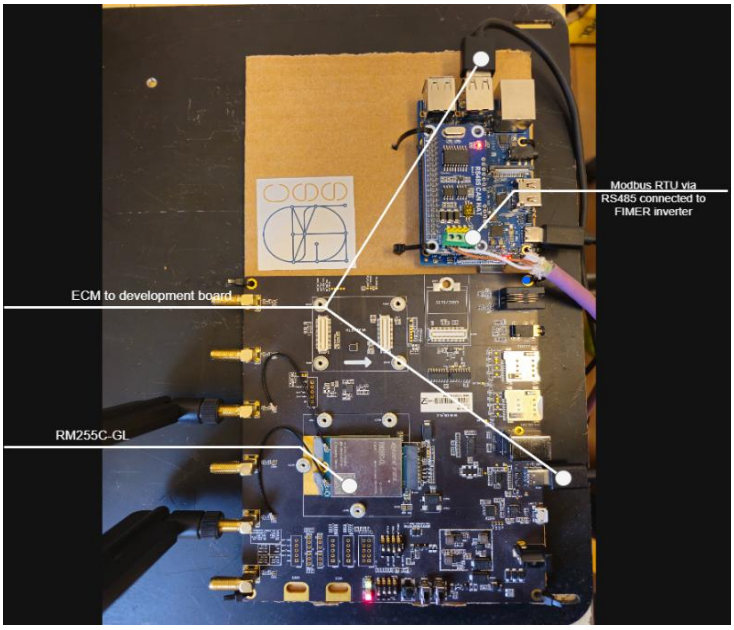

Figure 54. Gateway prototype device using the 5G RedCap Quectel modem.

Subsequently, the hardware design was refined through the development of a compact integration solution based on an M.2 modem form factor, using a custom USB-to-M.2 adapter. This optimization reduced system complexity and improved suitability for deployment in operational environments.

At network level, the integration leveraged the B5G infrastructure provided by ORO, operating in the n78 band (3.5 GHz). The gateway successfully attached to the B5G network and established stable IP connectivity through the user plane. Initial configurations aligned with the target architecture supporting network slicing, separating telemetry flows and control traffic. While full URLLC validation is scheduled for the final field trials, baseline tests confirmed stable connectivity, low latency, and reliable bidirectional communication.

At cloud level, the CSOFT platform was integrated through secure APIs, enabling device onboarding, telemetry ingestion, and real-time visualization. A continuous data pipeline was established between the gateway and the cloud environment, supporting streaming of inverter measurements. The dashboard provides a centralized view of gateway status, including connectivity and activity indicators. During integration, this interface was used to confirm successful device registration and persistent connectivity over the B5G network, as can be seen in Figure 55.

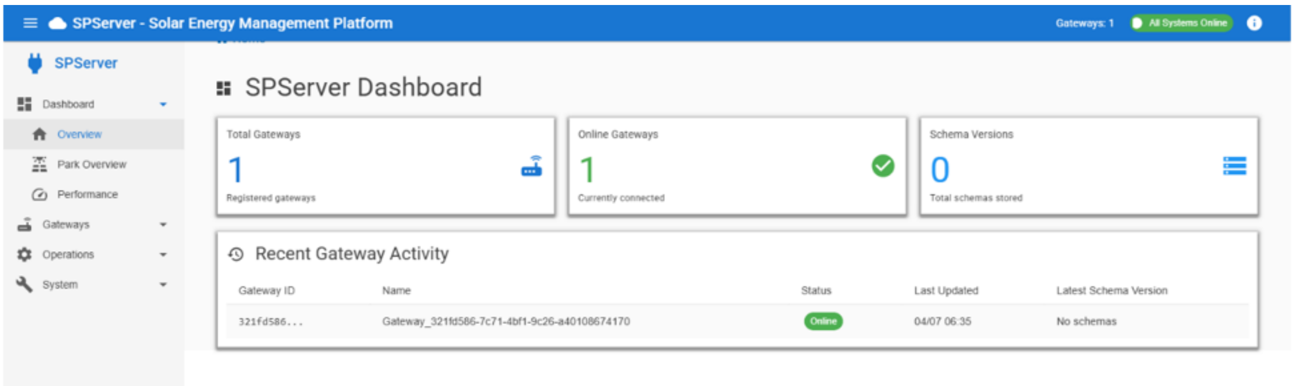

Figure 55. Monitoring dashboard – fleet of deployed IoT gateways.

The platform also provides aggregated and device-level visualization capabilities, enabling both system-wide monitoring and granular inspection of individual assets. The visualization layer is built on top of the real-time ingestion pipeline and supports dynamic querying, historical data storage, and correlation between multiple telemetry streams. Data is normalized at ingestion stage and mapped to a unified schema, allowing consistent representation across dashboards and APIs.

Figure 56 presents a consolidated view of the PV system performance at park level. It includes daily, monthly, and cumulative energy production metrics computed from inverter-level telemetry, as well as derived indicators such as hourly production profiles and relative contribution of each device. The Production Share visualization enables quick identification of load distribution across inverters, while the hourly generation chart reflects the temporal variability of solar output and confirms correct alignment with expected irradiation patterns. The Device Inventory section provides status information (online/offline, operational state, applied power limits), ensuring traceability between aggregated values and individual devices. During integration, this dashboard was used to verify that aggregated energy values were consistent with raw measurements retrieved via Modbus, confirming correct data ingestion, aggregation logic, and time-series alignment in the cloud layer.

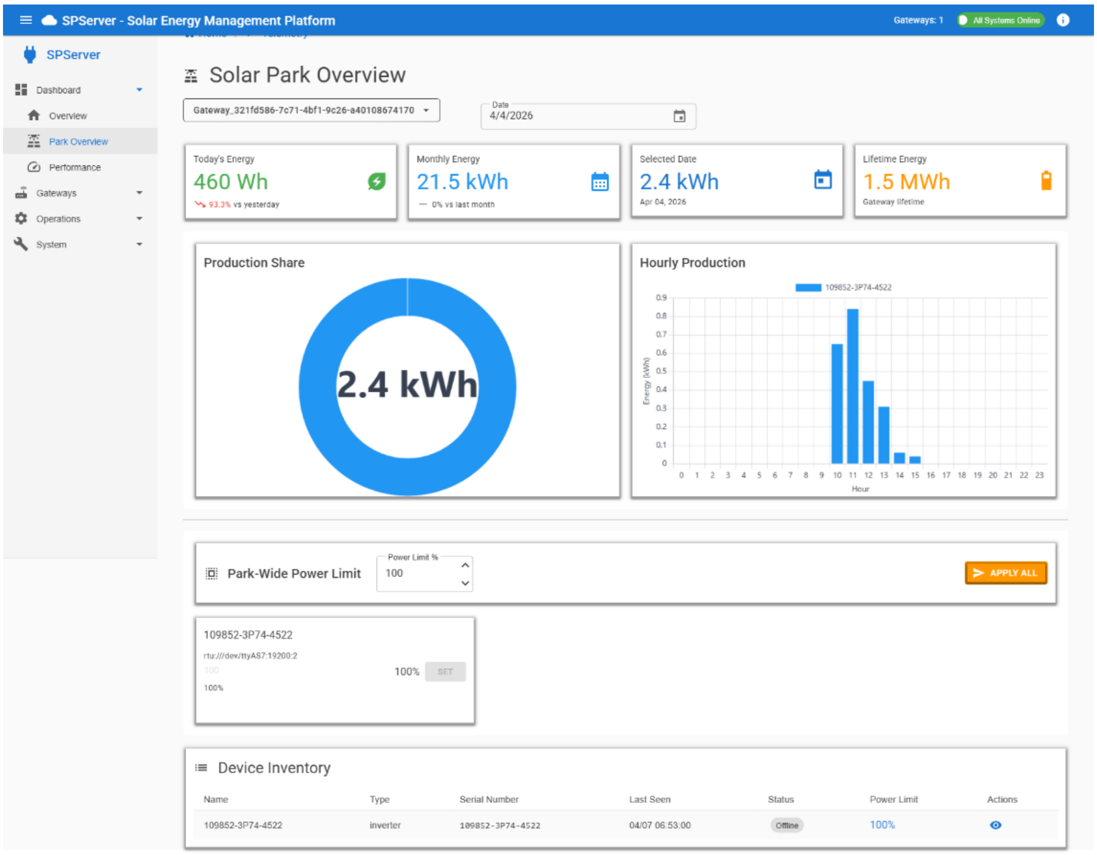

Figure 56. Park-level energy production dashboard.

A key aspect of the integration was the validation of the E2E control loop, ensuring that actuation commands can be reliably propagated from the cloud to the field devices and that their effects are accurately reflected in telemetry. Commands generated at the application layer (e.g., inverter power limitation) are transmitted via secure APIs to the gateway, forwarded over the B5G network, and executed locally through Modbus write operations. The gateway implements command acknowledgment and basic retry mechanisms to ensure delivery reliability. Upon execution, the inverter updates its operational state, which is subsequently captured through the telemetry pipeline and reflected in the dashboards with minimal delay. This closed-loop behavior confirms synchronization between control and monitoring paths and validates the system's capability to support real-time operational scenarios such as grid-driven power adjustments.

In addition to the solar park monitoring components described above, CAPG contributed to the E3 pre-trial activities by implementing an AI-based solar production forecasting component. This component is designed to extend the monitoring capabilities of the platform by enabling predictive analytics for photovoltaic energy generation. These pre-trial activities were implemented in the CAPG testbed between January and April 2026 and were not yet integrated into the Romanian testbed. This integration is in scope of final trial scheduled in H1 2027.

The forecasting service is integrated within an AI-aaS architecture, enabling the processing of historical inverter telemetry together with environmental data sources. The objective is to support predictive monitoring functionalities. The implemented solution follows a modular architecture composed of three main layers:

- *Data and Processing Layer:* responsible for ingesting historical inverter telemetry and environmental information obtained through weather APIs, and for preparing the data through cleaning, imputation and feature engineering processes.
- *MLOps and Training Layer:* where machine learning models are trained and managed. Multiple regression models were evaluated for solar production forecasting, including XGBoost, Histogram Gradient Boosting and Random Forest. Model training experiments are tracked using MLflow, which enables experiment tracking, model versioning and artifact management.
- *Serving and Inference Layer:* where trained models are deployed as forecasting services using BentoML endpoints, enabling both batch and live prediction requests. The deployed services can be integrated with the monitoring platform and deployed in container-based environments.

A monitoring stack based on Prometheus and Grafana was used to collect system metrics and ensure observability of the deployed AI services.

The architecture of the implemented AI forecasting pipeline is illustrated in Figure 57.

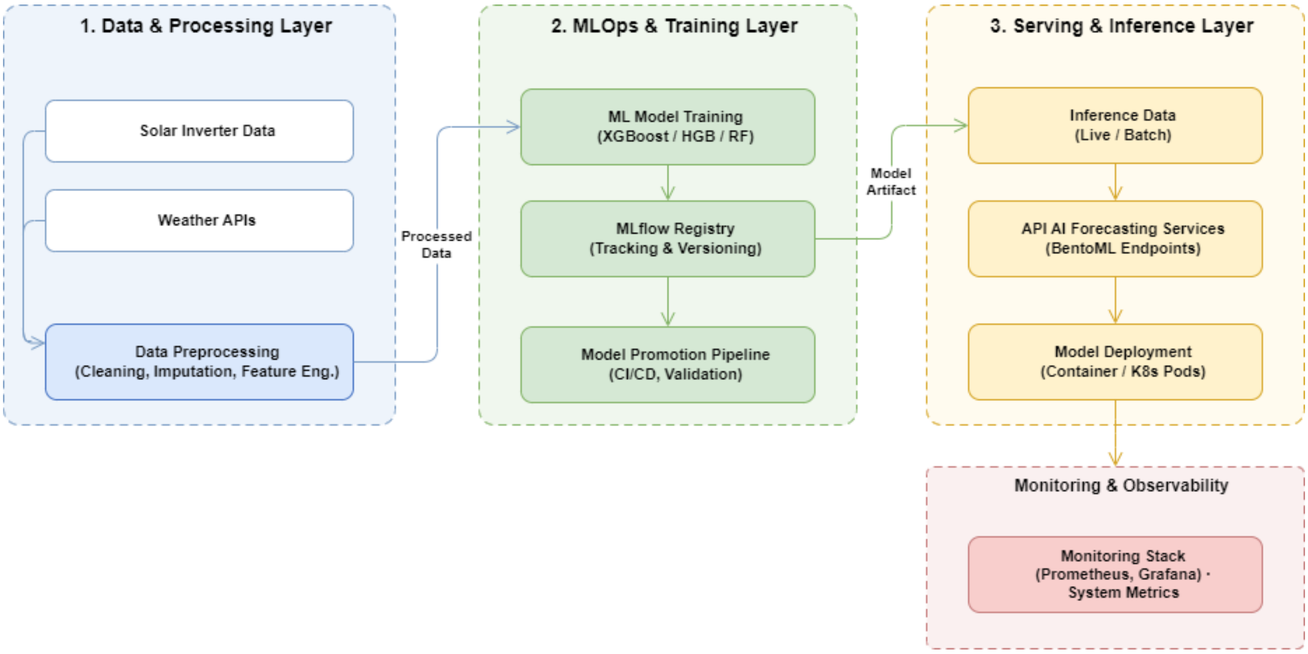

Figure 57. AI-based solar forecasting pipeline implemented for the E3 use case.

4.3.2 Preliminary test results and initial demos

Following the completion of the integration activities at the end of March 2026, the system reached a stable operational baseline, enabling the transition toward pre-trial validation. At that stage, all key building blocks were functionally interconnected: Modbus-based communication with the inverter was reliable, the B5G connectivity via the RedCap modem was operational, the telemetry pipeline between field and cloud was continuously active, and the cloud-to-edge control path had been successfully verified.

The pre-trial activities were carried out between 3-6 April 2026 at the CSOFT office, in Bucharest. A testbed environment was configured to replicate a simplified photovoltaic installation consisting of one inverter connected to a string of 10 solar panels, all interfaced with the cloud platform through the 5G-enabled gateway. This setup allowed controlled validation of the system while preserving the essential characteristics of a real deployment. Table 26 illustrates the E3 pre-trial configuration.

Table 26. E3 pre-trial execution.

E3.1 /Lab 01	E3 – Solar energy monitoring, control and predictions using B5G/6G communications and edge-cloud
Facility / Site	CSOFT office testbed, Bucharest; setup including 1 inverter and 10 solar panels connected via 5G gateway
Objective / Description	Validation of end-to-end integration, including telemetry acquisition, 5G communication, cloud ingestion, visualization, and remote control
Executed By	CSOFT, SIMTEL, CAPG, ORO
Devices / Components Involved	PV panels, inverter (SunSpec Modbus), Orange Pi gateway, 5G RedCap modem, 5G SA network (n78), cloud platform
5G Infrastructure	5G SA network with RedCap connectivity, private APN, core functions (AMF, SMF, UPF), initial slicing configuration
Targeted KPIs	UL throughput, DL throughput, Network availability, Coverage, Connection density
Measurement Tools	iPerf, Modbus validation, cloud dashboards, timestamp correlation
Ethics / Beta Users	Project partners, no ethical and personal data impact
Technical Enablers	IoT gateway, Modbus communication, 5G RedCap, cloud ingestion, remote control mechanisms

The testing phase focused on verifying the full operational chain, from data acquisition at field level to visualization and control at cloud level. The approach followed an incremental validation logic, where each segment of the system-data collection, communication, ingestion, visualization, and actuation-was exercised both independently and as part of the end-to-end workflow.

From the beginning of the tests, the gateway established and maintained persistent connectivity with the cloud platform over the 5G SA network. Throughout the entire testing window, the communication session remained stable, with no observed interruptions or reconnections. Telemetry data originating from the inverter (see Figure 58) was continuously transmitted and successfully ingested into the cloud environment. These measurements were updated at regular intervals, ensuring near real-time visibility of system behaviour.

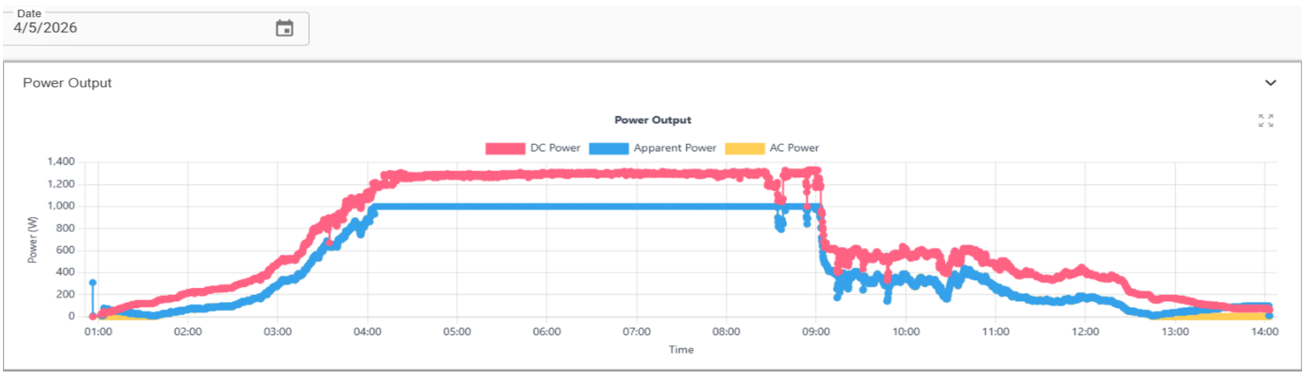

Figure 58. Inverter data.

The visualization layer accurately reflected the incoming data. At system level, the monitoring dashboard consistently indicated the gateway as online and active, confirming correct device registration and connectivity status. At the same time, the production dashboard displayed energy values aligned with the expected generation profile of the solar panels under the given environmental conditions. This confirmed that the aggregation logic and time-series processing in the cloud platform were functioning correctly.

At device level, detailed telemetry was accessible through the inverter interface, where all parameters were presented with proper units, scaling, and timestamps. The correctness of these values confirmed that Modbus registers were properly decoded and mapped within the application layer. Historical visualization further demonstrated continuity of the data stream, allowing observation of trends and short-term variations in production and electrical parameters.

A central part of the pre-trial validation was the execution of control operations. Commands generated in the cloud platform were transmitted through the B5G network to the gateway and forwarded to the inverter using Modbus write operations. The system demonstrated consistent behavior during these tests, with commands being executed reliably and without errors.

A representative example of this process is illustrated in Figure 59. When a command was issued to limit the inverter output to 0%, the effect was immediately visible in the telemetry data. The AC power dropped to zero, reflecting the enforced limitation on energy injection into the grid, while DC-side values remained above zero due to the physical generation of the solar panels. This behavior confirms both the correctness of the control mechanism and the expected physical response of the system.

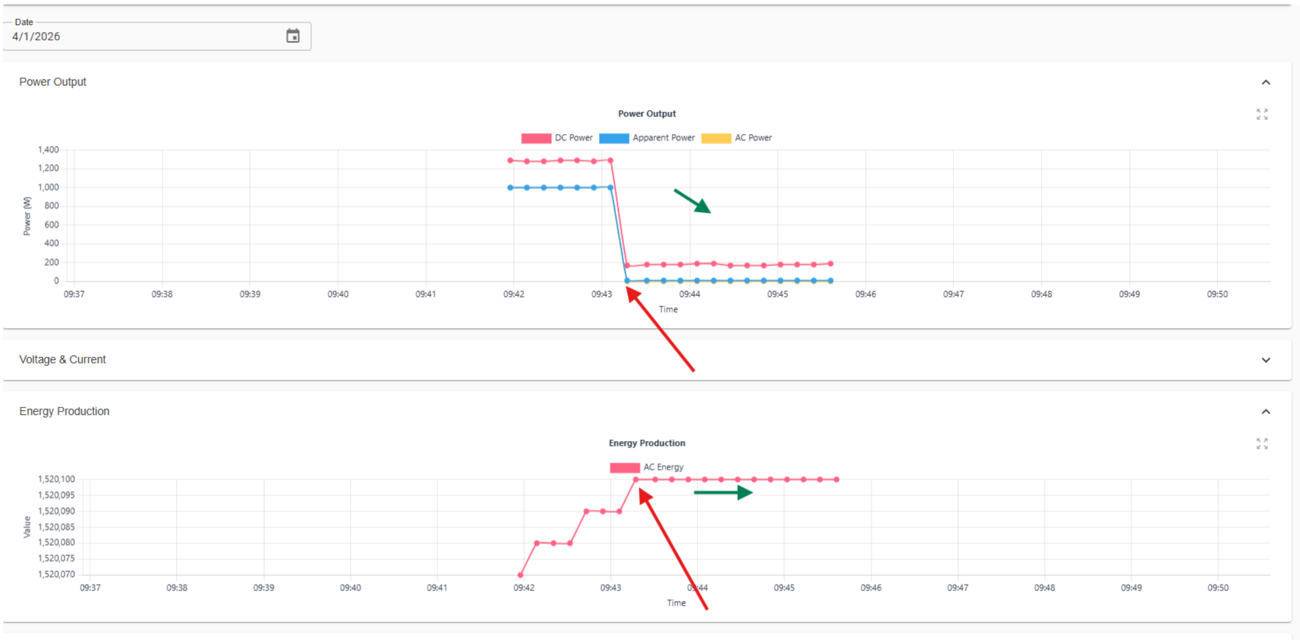

Figure 59. Inverter power limitation/increase effect.

For the CAPG separate implementation of AI-aaS, initial experiments were conducted using historical operational data collected from the photovoltaic installation used in the E3 use case. The dataset contains approximately 105,000 time-series samples and 14 variables, covering around one year of solar energy production measurements.

Feature engineering techniques were applied to capture temporal patterns and production dynamics, including temporal features (hour, day of week and month), lag variables representing recent production history, and rolling statistics describing short-term production trends.

Several ML models were evaluated for short-term solar energy forecasting (see Table 27).

Table 27. Performance of the selected models for different prediction horizons.

Prediction Horizon	Model	Train RMSE	Train R^2	Test RMSE	Test R^2
5 min	XGBoost	0.925	0.969	0.970	0.966
10 min	HistGradientBoosting	1.202	0.948	1.332	0.935
60 min	HistGradientBoosting	1.636	0.906	2.331	0.807

Figure 60 illustrates the comparison between the actual and predicted solar energy production values for the 5-minute forecasting horizon on the test dataset using the XGBoost model. The results show that the model is able to accurately capture the temporal dynamics of PV energy production, including peak generation periods and daily production patterns.

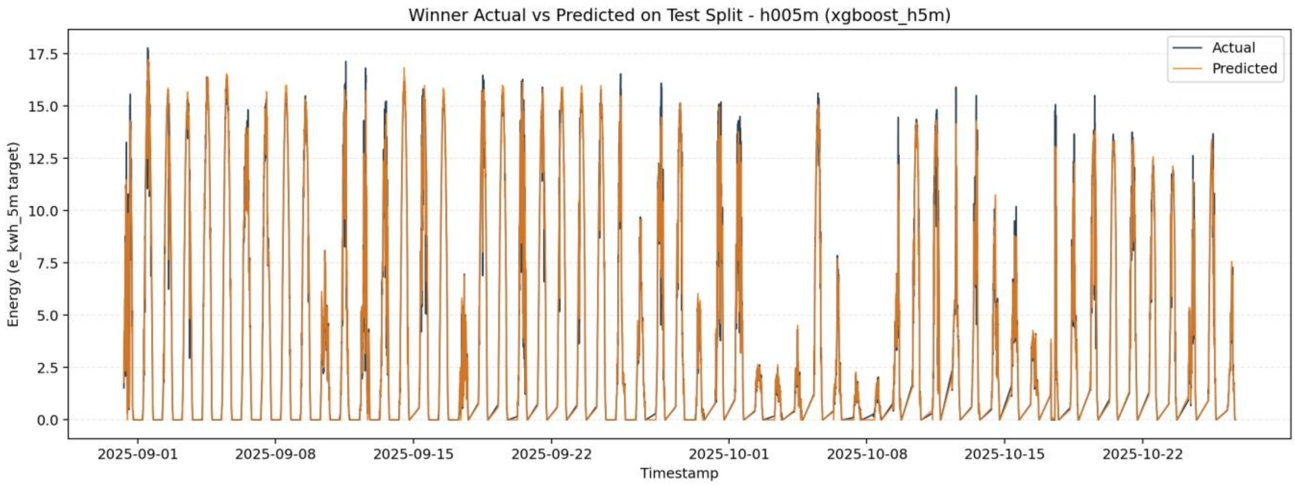

Figure 60. Actual and predicted solar energy production values.

The obtained results indicate that the forecasting models achieve high accuracy for short-term prediction horizons, particularly for the 5-minute interval, where the model achieved a test R^2 of 0.966.

The similarity between training and testing performance suggests that the models generalize well and no significant overfitting is observed. As expected, prediction accuracy decreases for longer forecasting horizons due to the increased uncertainty associated with weather conditions and solar radiation variability.

These initial results demonstrate the feasibility of integrating AI-based forecasting capabilities into the solar monitoring platform, enabling predictive analytics for photovoltaic energy generation within the E3 use case.

In addition to the model evaluation results, the AI-aaS deployment was monitored through an observability stack based on Prometheus and Grafana. Figure 61 presents an example of the runtime monitoring dashboard, including service availability, request rate, latency, and resource utilization metrics for the deployed forecasting services.

Figure 61. AI-aaS observability dashboard.

Following the integration and tests done in CSOFT environment, Table 28 presents the KPIs measurements from April 2026 pre-trial session.

Table 28. E3 KPIs measurements.

Test Case ID	Target KPIs / Requirements	KPIs measurements	Evaluation / Notes
E3.1_Lab01_01	KPI_VA_3 Network service latency 10 ms KPI_VA_4: One-way application latency: 300 ms KPI_VA_10: DL User Experienced data rate 50 Mbps KPI_VA_11: UL User Experienced data rate 10 Mbps KPI_VA_22: Coverage 99 % KPI_VA_2: Network availability 99.999 % KPI_AI_1: Inference speed < 2 s KPI_AI_2: ML accuracy < 5%	KPI_VA_3: 20 ms KPI_VA_4: 250 ms KPI_VA_10: 28 Mbps KPI_VA_11: 21 Mbps KPI_VA_22: 100 % KPI_VA_2: 100 % KPI_VA_4: complaint KPI_AI_1: 0.5 s KPI_AI_2:4.66%;	Indoor 5G coverage limitations impacted latency and throughput; functionality validated Inference time shorter than specified Forecasting error below the specified threshold.

The measured KPIs confirm overall functional performance of the system; however, some deviations from target values-particularly for latency and user-experienced data rates-are influenced by suboptimal indoor 5G radio conditions in the laboratory environment. The CSOFT office setup does not provide full 5G coverage optimization (limited indoor signal penetration), which resulted in increased network latency and reduced effective throughput compared to target values.

Despite these constraints, the one-way application latency remained within acceptable bounds for the tested use case, confirming that the end-to-end control loop is operational. Coverage and availability were effectively maintained during the test period, with no observed service interruptions. Connection density requirements are considered compliant given the small-scale setup.

From a communication perspective, the system maintained stable 5G connectivity throughout the entire duration of the tests. No packet loss or session instability was observed at application level, however indoor coverage 5G limitations impacted latency and throughput. These results confirm that the

RedCap-based gateway can sustain continuous data exchange and control signalling over the B5G infrastructure.

The system also demonstrated stability under continuous operation. Over the multi-day testing period, no crashes, communication failures, or performance degradation were detected. This indicates that the integration between hardware, network, and cloud components is robust and suitable for extended operation.

The AI forecasting module met both performance targets: the average inference time was 0.5 s over ten runs, remaining below the 2-second threshold, while the normalized forecasting error was 4.66%, remaining below the 5% target.

Overall, the pre-trial results confirm that the E2E architecture is fully functional. The telemetry acquisition chain operates reliably, the cloud platform correctly ingests and visualizes data, and the control loop enables responsive and deterministic actuation of field devices. The controlled setup using a single inverter and 10 solar panels proved sufficient to validate the core functionalities of the system and to de-risk the transition toward full-scale deployment.

Looking ahead, the evolution towards 6G is necessary to support large-scale solar park deployments by providing improved uplink capacity, deterministic low-latency communication, and native AI integration, enabling real-time optimization, high-density device connectivity, and more accurate distributed energy forecasting and control across geographically dispersed photovoltaic assets.

Next steps

The next phase will focus on the full integration of the AI-based forecasting and intent-driven control components into the Romanian testbed, enabling a complete end-to-end operational workflow that combines real-time monitoring, predictive analytics, and automated actuation over optimized B5G network slices. This will be complemented by the transition from laboratory conditions to the operational solar park environment, where outdoor deployment, enhanced 5G coverage, and URLLC configurations are expected to validate target KPIs and system scalability.

In parallel, the next version of the gateway device will be a dedicated industrial grade IoT Gateway, that had the previous capabilities, but will also be smaller, more efficient in terms of size and costs of production, and will be further tested during the integration tests and deployments to other solar parks.

The final trial scheduled in H1 2027 will consolidate these advancements by demonstrating a fully integrated, and scalable energy management solution under real conditions, including multi-device coordination, improved forecasting accuracy, and closed-loop control performance, ultimately confirming the readiness of the E3 architecture for large-scale photovoltaic deployments.

Early dissemination and demos

The demo was conducted in Bucharest, in March 2026, and builds upon the existing 5G SA infrastructure and experimentation ecosystem developed in Romania, including facilities such as the Orange 5G Lab, which supports testing of advanced 5G use cases and industrial applications.

The early demo also provides initial validation of the performance characteristics of the system. While conducted in a controlled indoor environment, the setup confirmed stable connectivity, continuous data transmission, and responsive control execution. These results align with the expected capabilities of 5G RedCap and edge-cloud architectures for industrial IoT applications. Video demonstration of the system[3] and Figure 62 illustration.

3 https://youtu.be/qkITO3rn07g?si=uOxa6hp_oXNfEyvw

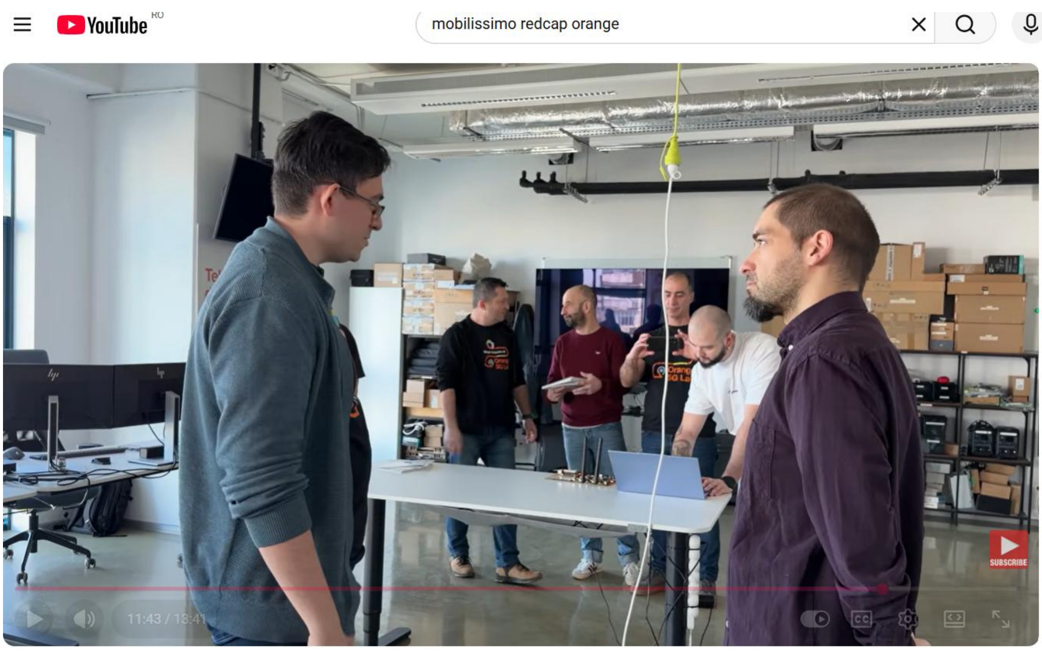

Figure 62. Early demo at Orange 5G laboratory in March 2026.

In addition to the technical validation, the demo contributes to the broader Romanian 5G innovation landscape by demonstrating a concrete vertical use case in the energy sector. It complements ongoing national efforts to expand 5G capabilities and apply them in domains such as smart infrastructure, environmental monitoring, and digital transformation.

The first results of the E3 implementation were also showcased during PoliFest 2026 event (see Figure 63), organized at Politehnica University of Bucharest. The event is a large-scale technology and education fair that brings together academia, industry, and students, focusing on innovation, engineering solutions, and emerging technologies, providing a platform to demonstrate practical implementations and engage with potential stakeholders and future specialists.

Figure 63. Early demo and dissemination during PoliFest event, Bucharest, March 2026.

Further insights into the implementation and demonstration can be found in the following material: LinkedIn article on 5G RedCap deployment and relevance in Romania[4].

Overall, the early demos demonstrate that the integration of 5G connectivity, edge computing, and cloud-based intelligence can enable reliable, real-time monitoring and control of distributed solar energy assets, forming a foundation for more advanced functionalities such as predictive analytics and automated grid response in the next phases. As next steps, the first E3 results will also be presented during the EuCNC 2026 in Malaga in June 2026.

4 https://www.linkedin.com/pulse/5g-redcap-o-nou%C4%83-premier%C4%83-pentru-re%C8%9Belele-de-telecomunica%C8%9Bii-eficiente-wdn3f?utm_source=share&utm_medium=member_android&utm_campaign=share_via

5 Conclusion and next steps

This deliverable, D5.1, documents the intermediate engineering outputs for WP5 Energy Domain, based on work performed from M3 to M17. The main conclusion is that the architectural foundations and key enabling functions required for E2E energy-domain operation were successfully instantiated and validated in controlled environments, and that next developments are now well scoped for the next integration and trial phases in WP7. The reported progress demonstrates measurable readiness for system-level demonstration, while also providing early evidence of where performance or operational coverage must be improved to meet the planned target KPIs in more representative conditions.

In particular, D5.1 concludes that the E2E chains required across field connectivity, edge processing, cloud orchestration, and application-level control/analytics are operational at the baseline level. Preliminary AI integration readiness was validated through onboarding and inference workflows designed for edge or edge-cloud execution, alongside early classification/forecasting pipeline evaluations that quantify generalization performance and data-to-decision readiness. Communication and IoT integration conclusions confirm that secure device connectivity and telemetry/control data paths function correctly and that control actions propagate through the system and are reflected back in telemetry, with measured service continuity. Where target KPI deviations were observed, such as latency or user-experienced data-rate differences under lab radio constraints, the deliverable identifies the contributing factors and provides direct technical direction for remediation (e.g., radio environment representativeness, configuration changes, and slice/throughput tuning). The next-step scope therefore concentrates on closing the remaining workflow gaps, activating full closed-loop behaviors and model lifecycle automation, and validating performance under the trial deployment constraints and traffic mixes planned for WP7. Main conclusions are presented below for each energy use case.

Renewable Energy Communities: By the end of the initial validation phase, Q1 2026, the E1 use case demonstrated a stable E2E deployment and measurable performance against the defined KPI framework for B5G-enabled smart building energy management. The SB-EMS platform was successfully integrated and validated over a private 5G SA infrastructure with edge-based user plane deployment and Kubernetes-managed edge/cloud resources, ensuring full support for telemetry ingestion, AI inference, and closed-loop control. The pre-trial results confirm that the system is capable of meeting stringent performance targets. Service provisioning latency was consistently measured at an average of 15 seconds, well below the <30 s KPI target, indicating efficient orchestration across service and resource layers. E2E latency for actuation commands reached 10 ms, outperforming the <20 ms requirement, while reliability of actuation execution achieved 100%, exceeding the 99.999% target. Uplink throughput reached 120 Mbps in the tested configuration, significantly above the >10 Mbps requirement, supporting scalability for dense IoT deployments. On the AI side, inference latency was measured at 0.01 seconds, far below the <2 s threshold, and classification accuracy reached 0.935–0.96 across evaluated models, satisfying the >0.9 KPI target and confirming robustness for operational deployment. In addition, full automation of the MLOps lifecycle (100%) and high occupancy prediction accuracy (96%) validate the readiness of the system for autonomous closed-loop operation. These results demonstrate that the E1 architecture can support reliable, low-latency, and AI-driven energy management services within indoor REC environments, providing a solid baseline for scaling towards real-time operation with live data and for extending validation to additional KPIs such as sensor data latency, energy efficiency of edge nodes, and large-scale connection density in the next trial phases.

In terms of next steps, the final integration and validation plan will continue during 2026 and first half of 2027 and will concentrate on E2E execution of SB-EMS monitoring → prediction → decision → actuation/suggestion workflows over the prepared sliced private network. The enablers to be further activated include full closed-loop self-tuning/re-training logic, expanded AI-aaS model catalogue coverage for additional predictive/anomaly services, and scaling of closed-loop orchestration with tighter coupling to network telemetry. For REC scope, the subsequent phase will introduce additional cloud-side REC components and Federated Learning management, enabling multi-building collaborative learning while maintaining data confidentiality through federated updates. In context of IBN framework, it will be migrated to ORO testbed, integrated with the live 5G UDR APIs for slice provisioning and subscriber binding, enabling E2E validation of the intent-driven control chain under operational telemetry conditions. In the context of massive deployment, 6G will be needed to support AI-native, distributed, and federated intelligence at scale, enabling real-time coordination across multiple buildings and energy communities with ultra-low latency and seamless edge–cloud continuum integration.

Robotized Offshore Wind Turbine Blade Inspection and Maintenance: For the E2 pre-trial baseline, the project implemented and validated key infrastructure and positioning capabilities required for the integrated drone-to-DT inspection trial. At network level, the B5G standalone environment was deployed in the lab using Amarisoft Callbox as gNB, Open5GS as the 5G core, and an edge UPF instance conceptually aligned with the intended offshore replication model. Network slicing design was prepared to isolate drone/inspection traffic from background traffic, targeting consistent connectivity characteristics for real-time data transfer to the edge DT system.

Quantitative pre-trial results showed that both the uplink data rate and latency requirements were met. The uplink data rate measurements were done with two different TDD frame configurations. Under an uplink heavy TDD frame configuration, the uplink data rate target of 50 Mbps was met. The measured one-way E2E latency values (i.e. between 8 - 12 ms) were consistently below the < 100 ms target. Looking forward, 6G will be required to enable fully autonomous offshore inspection missions by providing integrated sensing and communication, higher uplink capacity for high-resolution data streaming, and ultra-reliable real-time control in dynamic and remote maritime environments.

For the localization enabler, the implemented GNSS/RTK positioning achieved centimetre-level accuracy. Across the multi-step integration approach (including RTK correction link operation and headless execution with a web dashboard), the rover positioning accuracy reached approximately 1.40 cm and ~1.0 cm for horizontal and vertical accuracy respectively, hence supporting the spatial referencing needs for blade inspection. Next steps are to extend the RTK correction delivery to an internet-accessible NTRIP relay via cellular network in order to support drone operation beyond visual line-of-sight, continue network slicing/public-private validation with the finalized KPI checks, and complete full drone integration so that the H1 2027 trial can stream sensor data to the DT system and facilitate onsite decision-making within an inspection window.

Solar Energy Monitoring, Control and Predictions Using B5G/6G Communications and Edge-Cloud: By the end of April 2026, E3 implemented and validated a functional solar monitoring and control architecture spanning field devices, B5G connectivity, and the cloud-based orchestration and visualization layer. Field integration was established by implementing an edge gateway prototype (Orange Pi-based) that performed Modbus read/write operations against a Fimer PVS-50 inverter, enabling reliable telemetry retrieval (power, voltage, current, temperature/operating state) and execution of remote-control commands (e.g., inverter power limitation) with telemetry reflecting the commanded operating mode. Connectivity was then integrated using a 5G SA setup with RedCap modem-based gateways to achieve stable IP connectivity to the cloud platform and continuous telemetry streaming.

In the April 2026 pre-trial, the system's KPI measurements confirmed overall functional performance and service continuity, while highlighting indoor radio limitations for certain performance metrics. Specifically, network service latency was measured at 20 ms versus a target of 10 ms, and one-way application latency was 250 ms versus a target of 300 ms. For data-rate KPIs, DL user experienced data rate was 28 Mbps versus a target of 50 Mbps, and UL user experienced data rate was reported as 21 Mbps versus a target of 10 Mbps. Service-level KPIs showed coverage = 100% (target 99%) and network availability = 100% (target 99.999%). The connection density KPI was assessed as complaint for the small-scale testbed configuration. These deviations were attributed to the indoor 5G limitations in the lab environment, while the E2E monitoring and control functionality was validated successfully.

Next steps focus on completing the full operational E3 workflow in the solar park environment by integrating the AI forecasting and intent/command optimization components into the Romanian testbed for the final trial in H1 2027. In parallel, the communication enablers (including network slicing for separating time-critical control from bulk telemetry) and the zero-touch orchestration/assurance mechanisms will be implemented under outdoor radio conditions, where the intended KPI targets-especially for latency and user experienced throughput-can be validated under more representative deployment constraints. Going to field massive rollout, 6G will be essential to support large-scale photovoltaic deployments by enabling massive device connectivity, improved uplink performance, and AI-driven real-time optimization across distributed energy assets with deterministic latency and reliability.

As next steps, the results for D5.1 will be used in WP7 to accelerate transition to system-level demonstrations by reusing the validated interface bindings and baseline platform deployments as starting points. Specifically, WP7 will use the evidence collected on component readiness, early KPI measurements, and identified limiting factors to prioritize what must be re-tested versus what is already stable, thereby shortening the time needed to reach final trial readiness. The performance measurement methodology refined in D5.1 will also be applied in WP7 to ensure consistent KPI collection across lab and field-representative setups, enabling direct comparison of improvements and targeted verification of remediation actions. Communication and dissemination activities derived from these trial preparations will continue until the end of the project as first results are already available. All energy-domain use case implementations will continue with enablers completion and will conclude with field trials in the first part of 2027, building on the validated interface bindings and baseline platform deployments established in D5.1. Final integrations and tests including KPIs and KVIs measurements will be reported in the final WP5 deliverable, D5.2.

6 References

[1] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.

[2] A. Adadi and M. Berrada, “Peeking inside the black-box: A survey on explainable artificial intelligence (XAI),” IEEE Access, vol. 6, pp. 52138–52160, 2018.

[3] McMahan et al., “Communication-efficient learning of deep networks from decentralized data,” in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017.

[4] ETSI, “Zero-touch network and Service Management (ZSM); Reference Architecture,” ETSI GS ZSM 002, 2019.

[5] Quectel Wireless Solutions Co., Ltd., *“RM500Q-AE 5G Sub-6 GHz M.2 Module – Specification,”* Version 1.0, Aug. 2020.

[6] MobileControl, *“eW-B19/20 Energy and I/O Monitoring Module – Technical Datasheet,”* MobileControl GmbH, 2022.

[7] OMEGA Engineering, *“Pt100 Platinum RTD Temperature Sensors – Technical Reference and Datasheet,”* OMEGA Engineering Inc., Norwalk, CT, USA, 2021.

[8] Environmental Monitor A3, *“Environmental Monitor A3 – Air Quality Monitoring Device Datasheet,”* A3 Environmental Monitoring System, 2022.

[9] 3GPP, "System architecture for the 5G System (5GS)," TS 23.501, v18.0.0, Dec. 2023.

[10] 6G Smart Networks and Services Industry Association (6G-IA), "European Vision for the 6G Network Ecosystem," SNS JU Architecture Working Group White Paper, v5.0, Jun. 2024.

[11] VITAL-5G Project, "Vertical Innovations in Transport and Logistics over 5G: Final Architecture and Open Online Platform Specifications," EU Horizon 2020 Project Deliverable D2.3, May 2023. https://www.vital5g.eu/

[12] AIRTuB-ROMI Project. https://www.worldclassmaintenance.com/sub-project/airtub-romi/

[13] Dassault Systemes, “Abaqus Finite Element Analysis software,”. https://discover.3ds.com/test-and-validate-your-product-faster-abaqus

[14] Sonatest,”Ultrasound roller probe,”. https://sonatest.com/resources/application-notes/roller-probe-inspection

[15] Quectel Wireless Solutions, “Quectel RM500Q-GL 5G Sub-6 GHz M.2 Module,” Version 1.3. https://www.quectel.com/download/quectel_rm50xq_series_hardware_design_v1-3/

[16] DJI Enterprise, “DJI Matrice M600,”. https://www.dji.com/nl/support/product/matrice600

[17] DJI Enterprise, “DJI Matrice 4 Series,”. https://enterprise.dji.com/matrice-4-series

[18] u-blox,”EVK-X20P multi-band GNSS receiver,”. https://www.u-blox.com/en/product/evk-x20p

[19] Amarisoft, “AMARI Callbox Classic,”. https://www.amarisoft.com/test-and-measurement/device-testing/device-products/amari-callbox-classic

[20] Open5GS, “Open-source implementation of 5G Core”. https://open5gs.org/

[21] PointPerfect Flex, “PPP-RTK GNSS correction data service”. https://www.u-blox.com/en/product/pointperfectflex

[22] S. N. K. Veedu et al., "Toward Smaller and Lower-Cost 5G Devices with Longer Battery Life: An Overview of 3GPP Release 17 RedCap," IEEE Communications Standards Magazine, vol. 6, no. 3, pp. 84–90, Sep. 2022.

[23] X. Lin, "3GPP Evolution from 5G to 6G: A 10-Year Retrospective," arXiv, 2024.

[24] S. Tamboli, M. Rawale, R. Thoraiet, and S. Agashe, "Implementation of Modbus RTU and Modbus TCP communication using Siemens S7-1200 PLC for batch process," in 2015 International Conference on Smart Technologies and Management for Computing, Communication, Controls, Energy and Materials (ICSTM), 2015, pp. 258–263.

[25] SunSpec Alliance, "SunSpec Information Model Specification," SunSpec Alliance, San Jose, CA, USA, Standard 701, 2020.

[26] Quectel Wireless Solutions, "RM500Q-AE 5G Sub-6 GHz M.2 Module Datasheet," V1.1, Oct. 2020.

[27] FIMER SpA, "PVS-50/60-TL String Inverter Datasheet," Vimercate, Italy, 2021.

[28] Meteocontrol GmbH, "Si Series Silicon Irradiance Sensors - Si-I-420-T Datasheet," Augsburg, Germany, 2022.

[29] Shenzhen Xunlong Software CO., Limited, "Orange Pi 4A User Manual," V1.0, 2023.

Annex A: Energy 1 SB-EMS and REC-EMS Workflows

This annex provides a detailed technical overview of the workflows and architectural interactions required to implement energy optimization and user comfort management within the single-building and renewable energy community scenarios.

The workflows for the single building scenario are shown in the following figures, while the implementation of the REC scenario workflows is planned for the second half of the project and it will be documented in the next WP5 deliverable D5.2. Figure Figure shows the sequence diagram for the provisioning and configuration of the SB-EMS application in the edge environment of each building, triggered by the service orchestrator REST APIs (based on TM Forum API for service catalogue, service ordering, and service management - TMF633, TMF640, TMF641). The workflow assumes that the IoT platform is already up and running, with sensors and actuator already registered with the IoT platform. Moreover, the network slices for the sensing and commands traffic are already established in the private network, the former with mIoT profile and the latter with URLLC profile.

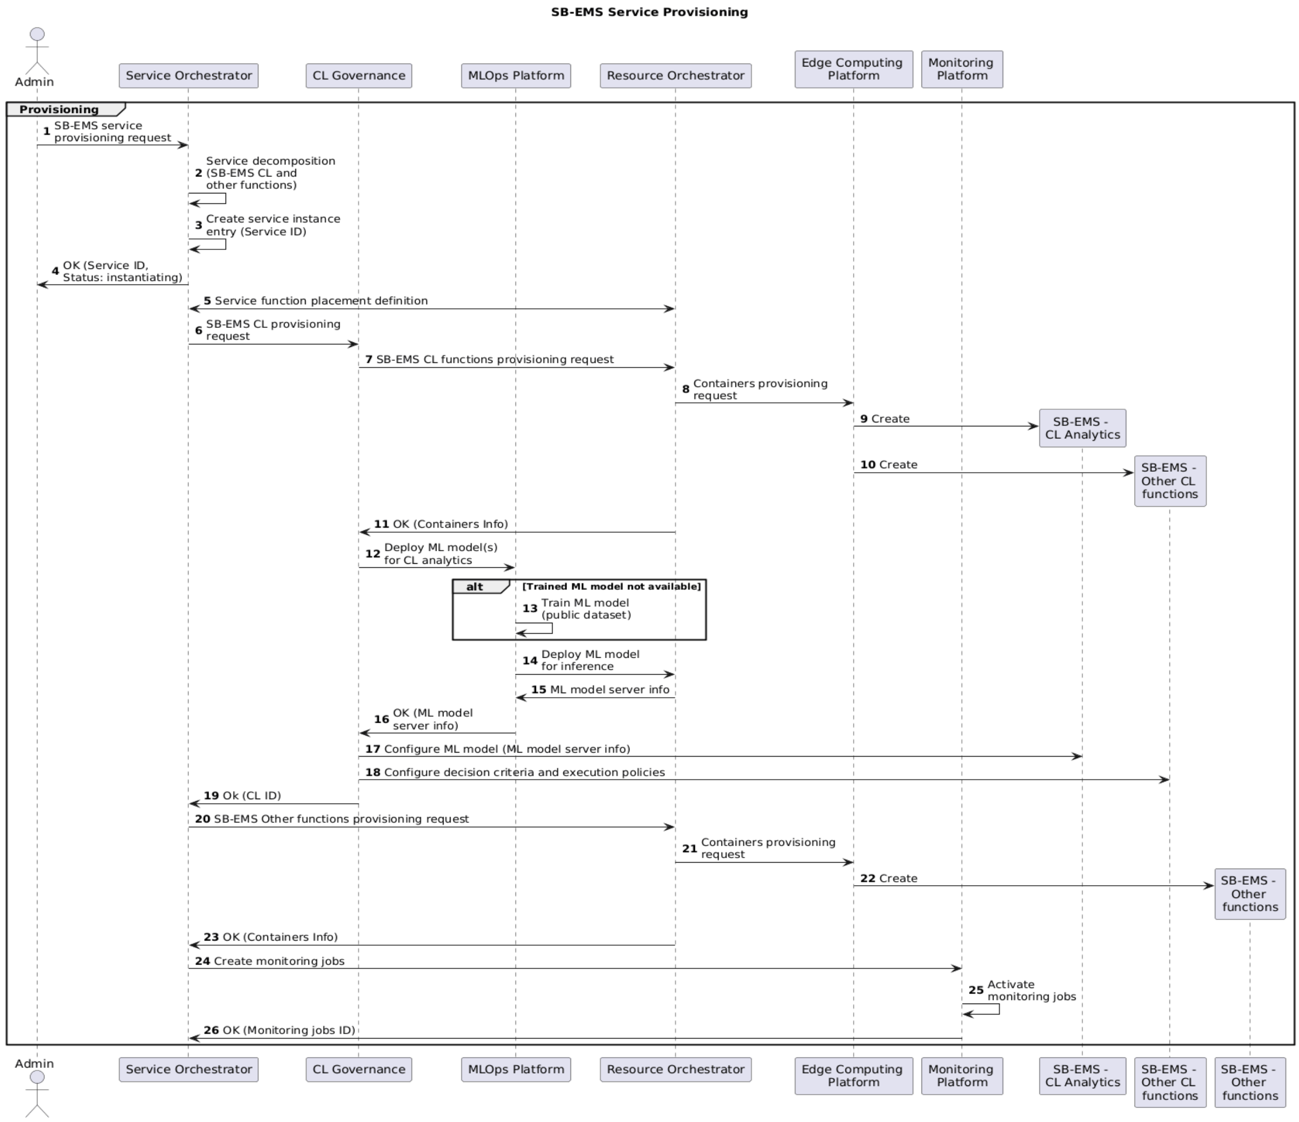

Figure 64. E1 - Workflow for SB-EMS service provisioning.

The workflow starts with the service provisioning request issued by the system administrator (step 1), with a reference to the service descriptor defined in the catalogue.

The Service Orchestrator identifies the components to be deployed on the basis of the service descriptor, creates a new entry in the record and returns the unique ID of the service (steps 2-4), which can be used later on by the administrator to monitor its status and perform operations at runtime. Interacting with the Resource Orchestrator and making use of the Resource Allocation Engine, the placement of the various service functions is elaborated (step 5) and the provisioning process starts with the instantiation of the closed loop functions.

The related request is issued to the CL Governance (step 6), which coordinates the provisioning of the CL functions interacting with the Resource Orchestrator that creates the various containers in the edge nodes (steps 7-11). Since the CL Analysis function relies on ML models, the CL Governance interacts with the MLOps to deploy the models for the inference tasks (step 12). If suitable trained models are not yet available in the ML Catalogue, a new training task is started, potentially using a public dataset (step 13). When the trained model(s) become available they are deployed in the edge node of the building interacting with the Resource Orchestrator (steps 14-15) and their reference is returned to the CL Governance (step 16). The following phase involves the configuration of the CL functions, where the CL Governance configures e.g., the ML models on the CL Analytics function, the parameters for the decision criteria in the CL Decision function, and the policies for automatic vs. supervised building settings at the CL Execution function (steps 17-18).

After the successful creation of the CL (step 19), the Service Orchestrator triggers the instantiation of the other functions of the SB-EMS application through the Resource Orchestrator (steps 20-23). As a final phase, the monitoring jobs related to network performance, usage of computing resources, and collection of IoT sensors data from the IoT platform are activated (step 24-26). As result, the SB-EMS is up and running at the building, ready to process the IoT data and execute the actions to guarantee the continuous optimization of user comfort and energy consumption at runtime, as described in the following workflow in Figure .

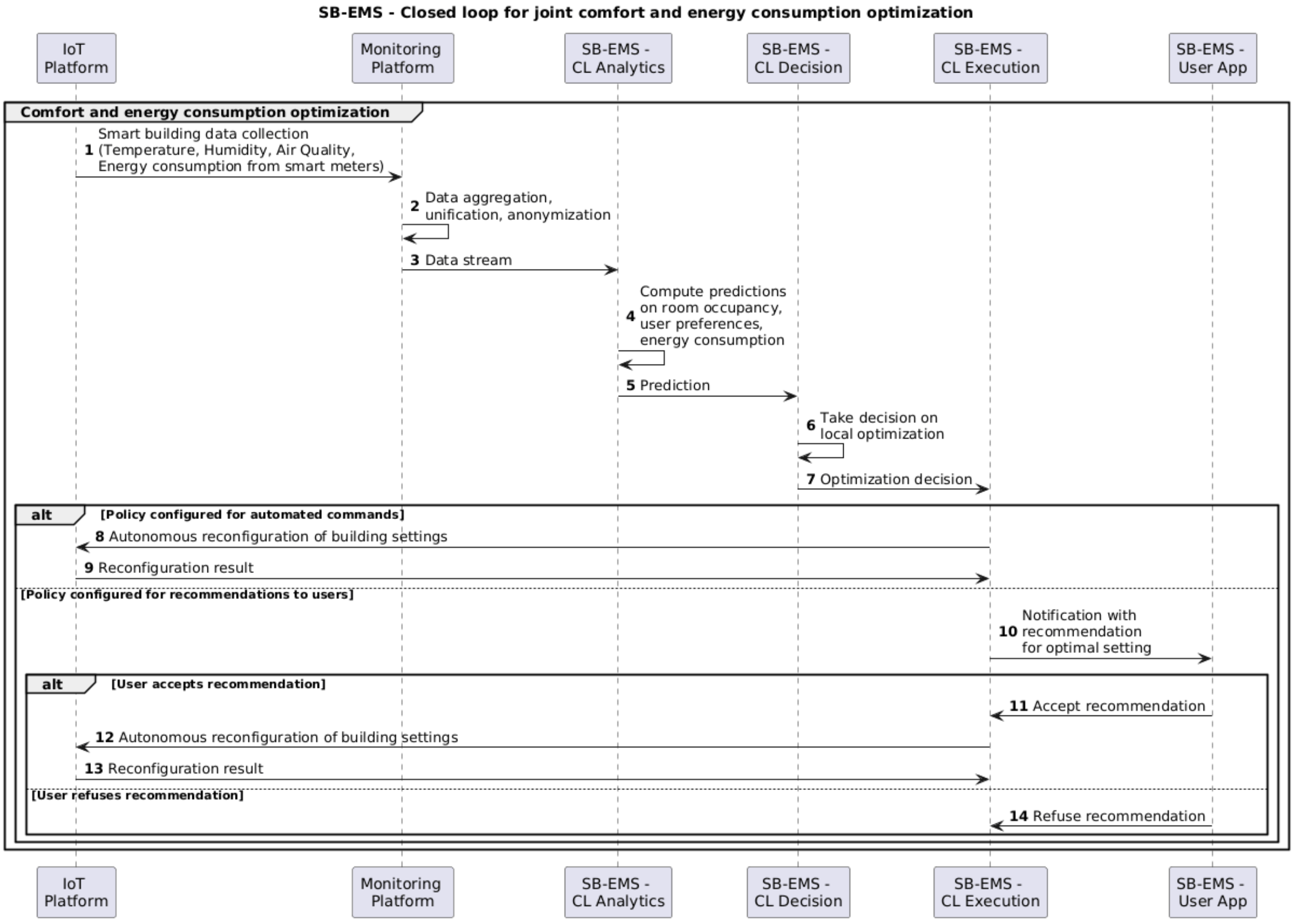

Figure 65. E1 - Workflow for SB-EMS comfort and energy consumption optimization.

The relevant data from the IoT platform are transferred to the Monitoring Platform, where they are filtered, pre-processed, and aggregated following a unified data model. If needed, data can also be anonymized (steps 1-2). The processed data feed the ML algorithms at the CL Analytics function, where predictions on room occupancy, requirements on comfort conditions, appliance workloads, and energy consumption in the near future are elaborated, feeding the following CL Decision stage (steps 3-6). The optimization decision with the suggested smart building reconfiguration is then notified to the CL Execution function (step 7). If the system is configured to automatically execute the commands, the Execution function interacts with the Actuation API of the IoT platform to reconfigure the smart building settings, e.g., at the HVAC system (steps 8-9). Otherwise, a notification is sent to the user mobile application (step 10) and the user can decide to either accept (steps 11-13) or refuse the suggestion (step 14). In the former case, commands are issued towards the IoT platform for reconfiguration.

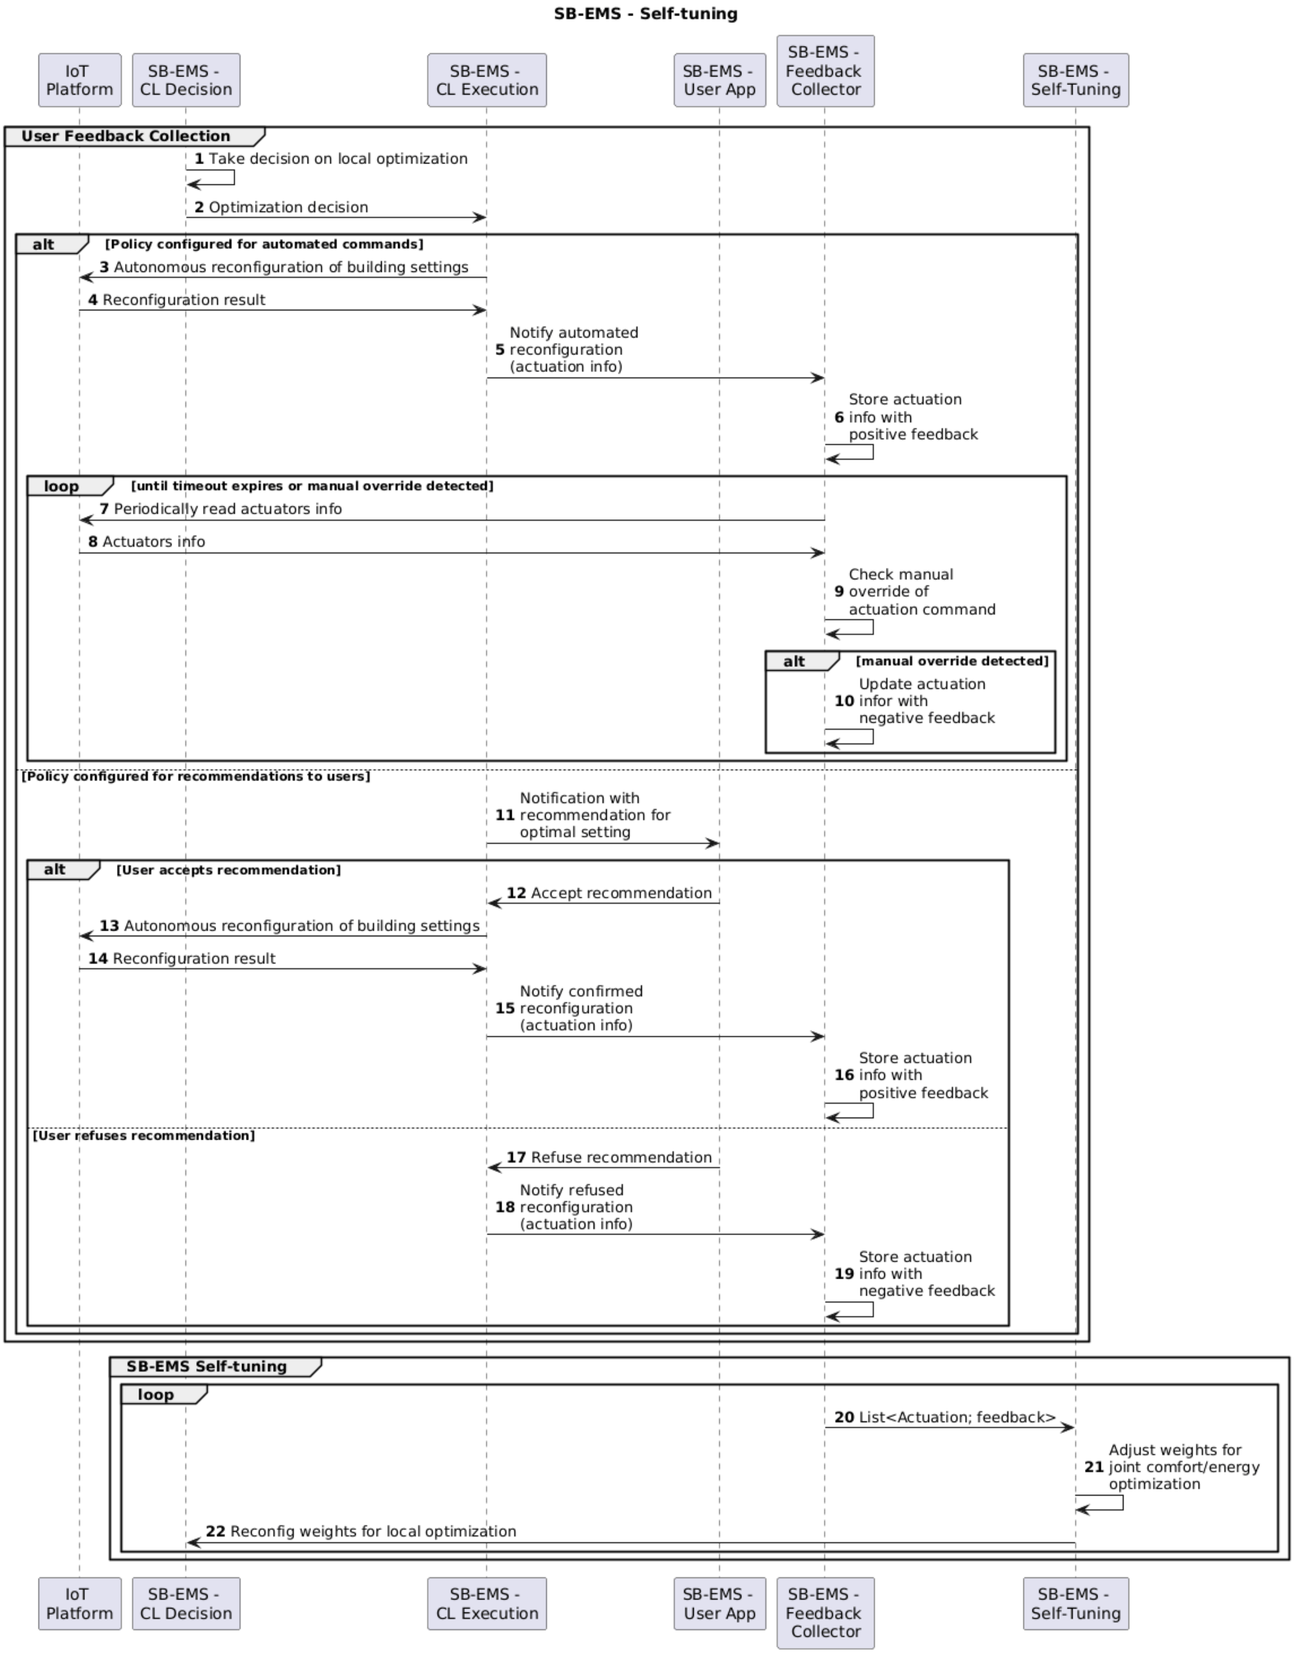

Figure 66. SB-EMS Workflow-self-tuning-balance between comfort and energy consumption optimization.

Figure shows the workflow for self-tuning the decision policies on the basis of the implicit or explicit feedback provided by the user on the actions or suggestions issued by the SB-EMS. This part of the system will be implemented in the second half of the second year of the project and the related implementation will be documented in the following deliverable D5.2. The first part of the workflow shows the procedures to collect the feedback from the user on a decision taken by the CL (steps 1-2). This procedure follows two different patterns, depending on the policy configured at the CL Execution function, i.e. if the system is configured for automated smart building reconfigurations (steps 3-10) or for sending suggestions to the user (steps 11-19).

In the first case, the re-configuration command is immediately issued by the CL Execution function to the IoT platform (step 3-4) and the reconfiguration action is notified to the SB-EMS Feedback Collector component (step 5) where it is stored assuming a positive feedback from the user (step 6). In order to evaluate if the user is actually happy with the reconfiguration, the Feedback Collector monitors periodically the status of the associated actuators for a given period of time (steps 7-8) to detect if the user has manually changed the configuration set by the SB-EMS system (step 9). If this the case, the entry in the local record is updated indicating a negative feedback from the user (step 10).

In case of SB-EMS configured to send suggestions to the user, the notification is sent to user's mobile application (step 11). If the user accepts the suggestion (step 12), the CL Execution function reconfigures the smart building interacting with the IoT platform (steps 13-14) and it notifies the SB-EMS Feedback Collector about the action confirmed by the user (step 15-16). If the suggestion is refused (step 17), the Feedback Collector is informed about the proposed action and the refusal of the user (steps 18-19).

The second part of the workflow (steps 20-22) shows how the system uses the information stored at the Feedback Collector for self-tuning. The list of reconfiguration actions issued and accepted or refused by the user is periodically sent to the SB-EMS Self-Tuning component (step 20). This information is used to understand if the user is happy about the weights that are used in the decision phase to balance the mix of comfort perceived by the user and the level of energy consumed by the building (step 21) and, if needed, they are updated on the CL Decision function (step 22).

The SB-EMS generate a diverse and dynamic set of network demands. Real-time sensor telemetry, actuation commands, federated learning, and ML model inference each require different levels of latency, reliability, and throughput. The Intent-Based Network (IBN) framework addresses this directly by operating as the network intelligence layer beneath the SB-EMS. The IBN detects the corresponding shifts in traffic patterns from CPE data and RAN metrics, and automatically provisions the appropriate network slice interacting with the B5G/6G network and compute resources without requiring operator intervention.