D5.1

D5.1 ANNEX – INDIVIDUAL PILOT REPORTS

D5.1 Annex: Individual Pilot Reports

Companion Annex to D5.1 — full individual pilot reports (descriptions, planning, execution, feedback, lessons and impact analysis)

Work package	WP5
Task	T5.1, T5.3, and T5.4
Due date	30/06/2026
Submission date	02/07/2026
Deliverable lead	NOVA (WP5 Leader)
Version	1.0
Authors	Anastasios Tzoumpas (NOVA), Ioannis Markopoulos (NOVA), Georgios P. Katsikas (UBI), Gregory Mygdakos (AGA), Christos Tranoris (PNET), Youngseok Choi (KUL), Hazana Halimi (KUL), Amizan Omar (UBRAD), Theodoros Rokkas (INC), Janani Kaushalya Hewage (UBRAD), Jaime Ruiz (NOK), María Crespo (UPV), Jesús Iglesias (TID), Emilio Garrido (TID), Dimitrios Brodimas (ENIK), Nikolaos Tzanis (ENIK), Elefterios Mylonas (UoP), Luis Rosa (ONE), Ulf Bodin (LTU).
Reviewers	Georgios P. Katsikas (UBI)
Abstract	This deliverable annex complements D5.1 with detailed pilot reports across the project's five validation clusters.
Keywords	D5.1, pilot reports; clusters.

Document Revision History

Version	Date	Description of change	List of contributors
V0.1	18/03/2026	Table of Contents	NOVA
V0.2	13/05/2026	Sections 1 and 2	NOVA, UBI, ONE, KUL, UBRAD
V0.3	19/06/2026	Sections 3 and 4, internal review version	NOVA, UBI, ONE, KUL, UBRAD, INC ALL CLUSTER PARTNERS
V0.4	26/06/2026	Addressing internal review comments	NOVA, UBI, ONE, KUL, UBRAD, ALL CLUSTER PARTNERS
V1.0	02/07/2026	Quality Check, Submission	NOVA, NETC

Grant Agreement No: 101189819 | **Topic:** HORIZON-CL4-2024-DATA-01-03
Call: HORIZON-CL4-2024-DATA-01 | **Type of action:** HORIZON-IA

DISCLAIMER

Co-funded by
the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

Co-funded by the European Union (COP-PILOT, 101189819). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

COPYRIGHT NOTICE

Project Co-funded by the European Commission in the Horizon Europe Programme		
Nature of the deliverable:	R	
Dissemination Level		
PU	Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page)	X
Classified R-UE/ EU-R	EU RESTRICTED under the Commission Decision No2015/ 444	
Classified C-UE/ EU-C	EU CONFIDENTIAL under the Commission Decision No2015/ 444	
Classified S-UE/ EU-S	EU SECRET under the Commission Decision No2015/ 444	

* R: Document, report (excluding the periodic and final reports)

DEM: Demonstrator, pilot, prototype, plan designs

DEC: Websites, patents filing, press & media actions, videos, etc.

DATA: Data sets, microdata, etc.

DMP: Data management plan

ETHICS: Deliverables related to ethics issues.

SECURITY: Deliverables related to security issues

OTHER: Software, technical diagram, algorithms, models, etc.

TABLE OF CONTENTS

TABLE OF CONTENTS 5

LIST OF FIGURES 8

LIST OF TABLES 9

ABBREVIATIONS 10

ANNEX A – INDIVIDUAL PILOT REPORTS 13

1.1 Cluster 1: Business Integration in Mining 13

- 1.1.1 Pilot 1: Dynamic/Auto Scaling of Compute Resources for Efficiency 13
  - 1.1.1.1 Summary 13
  - 1.1.1.2 Pilot description 13
  - 1.1.1.3 Objectives 14
  - 1.1.1.4 Pilot Relevance 14
  - 1.1.1.5 Roadmaps (UC and Cluster Level) 14
  - 1.1.1.6 Expected Outcomes 14
  - 1.1.1.7 Execution Steps 15
  - 1.1.1.8 Setup and validation tools 15
  - 1.1.1.9 Stakeholder Involvement 16
  - 1.1.1.10 Assessment Plan 16
  - 1.1.1.11 Pilot Execution and Results 17
  - 1.1.1.12 Showcasing 18
  - 1.1.1.13 Stakeholder Feedback and Impact 18
  - 1.1.1.14 Lessons learned 19
- 1.1.2 Pilot 2: Asset tracking and rock bolt monitoring 19
  - 1.1.2.1 Summary 19
  - 1.1.2.2 Pilot description 19
  - 1.1.2.3 Objectives 19
  - 1.1.2.4 Pilot Relevance 20
  - 1.1.2.5 Roadmaps (UC and Cluster Level) 20
  - 1.1.2.6 Expected Outcomes 20
  - 1.1.2.7 Execution Steps 21
  - 1.1.2.8 Setup and validation tools 21
  - 1.1.2.9 Stakeholder Involvement 22
  - 1.1.2.10 Assessment Plan 22
  - 1.1.2.11 Pilot Execution and Results 23
  - 1.1.2.12 Showcasing 24
  - 1.1.2.13 Stakeholder Feedback and Impact 24

1.1.2.14 Lessons learned 24
1.1.2.15 Showcasing 25
1.1.2.16 Stakeholder Feedback and Impact 27
1.1.2.17 Lessons learned 27
1.1.2.18 Showcasing 28
1.1.2.19 Stakeholder Feedback and Impact 28
1.1.2.20 Lessons learned 29
1.1.3 Pilot 3: Condition monitoring and predictive maintenance 29
1.1.3.1 Summary 29
1.1.3.2 Pilot description 29
1.1.3.3 Objectives 30
1.1.3.4 Pilot Relevance 31
1.1.3.5 Roadmaps (UC and Cluster Level) 31
1.1.3.6 Expected Outcomes 32
1.1.3.7 Execution Steps 33
1.1.3.8 Setup and validation tools 34
1.1.3.9 Stakeholder Involvement 36
1.1.3.10 Assessment Plan 36
1.1.3.11 Pilot Execution and Results 36
1.1.3.12 Showcasing 37
1.1.3.13 Stakeholder Feedback and Impact 37
1.1.3.14 Lessons learned 37
1.2 Cluster 2: Smart Sustainable IoT Solutions in Valencia 39
1.2.1 Pilot Scope and Planning 39
1.2.1.1 Pilot 1 (UC#2.1A) 39
1.2.1.2 Pilot 3 (UC#2.2) 42
1.2.2 Pilot Execution and Results 45
1.2.2.1 Pilot 1 (UC#2.1A) 45
1.2.2.2 Pilot 3 (UC#2.2) 46
1.3 Cluster 3A: AgriTech Transformation and Sustainability Initiative 49
1.3.1 Pilot Scope and Planning 49
1.3.2 Pilot Execution and Results 56
1.4 Cluster 3E: Edge Intelligence for Enhancing Grid Reliability 60
1.4.1 Pilot Scope and Planning 60
1.4.2 Pilot Execution and Results 63
1.5 Cluster 4: Smart Vineyards & Sustainable Winery Ecosystems 67
1.5.1 Pilot Scope and Planning 67
1.5.2 Pilot Execution and Results 68

1.5.2.1 Pilot 1: Vineyard Decision Support at Quinta da Pacheca 69
1.5.2.2 Pilot 2: Winery Production Monitoring at the Wine Next Generation Technical Days 69
1.5.2.3 Pilot 3: Transferability of the Vineyard Workflow at Quinta da Invejosa Poceirão 69
1.5.2.4 Pilot 4: Soil Sensors predictions and energy data for 5G coverage at Villalvaro in Soria, Spain........ 70
REFERENCES 77

LIST OF FIGURES

FIGURE 1: SEISMIC MONITORING VIEW SHOWING A LARGE EVENT FOLLOWED BY A SERIES OF AFTERSHOCKS, CREATING A TIME-LIMITED PEAK-LOAD SITUATION. 17

FIGURE 2: AUTO-SCALING DEMONSTRATION: COMPUTE RESOURCES (TOP), LOAD PRESENTED TO THE SYSTEM AS JOBS IN THE QUEUE (MIDDLE) AND JOBS COMPLETED (BOTTOM). 18

FIGURE 3: LARGE SCALE TEST OF SIMULATED SENSORS IN NUCLEUS, SHOWING A FICTIONAL MAP WITH 400 SIMULATED SENSORS. 24

FIGURE 4: DEMO OF PHYSICAL GEOFENCE INDICATING AN AREA IS SAFE. 26

FIGURE 5: DEMO OF PHYSICAL GEOFENCE, SHOWING WHEN A BOLT IS ELONGATED CLOSE TO MAXIMUM AND THUS THE INDICATION SHOWS THAT THE AREA IS UNSAFE. 26

FIGURE 6: DISPLAY OF HOW A TAG (ORANGE MARKER) IS SHOWN IN NUCLEUS MAP TOOL, AS USED IN INITIAL ASSET TRACKING DEMO. THE BLUE MARKER IS A GATEWAY AND THE GREEN MARKERS ARE ROCK BOLT SENSORS. 28

FIGURE 7: ROADMAP FOR THE HOSCH-PREDGE PILOT (PILOT 3 IN CLUSTER 1). 32

FIGURE 8: HOSCH PILOT SYSTEM ARCHITECTURE — SIMULATOR (RISE) → HOSCH CLOUD PLATFORM (GCP) → PREDGE CONVEYOR SOLUTION (PREDGE). 34

FIGURE 9: INFRASTRUCTURE SETUP. 35

FIGURE 10: CLUSTER 2 PILOT PLANNING. 39

FIGURE 11: INSTALLATION OF RADARMOVAVCID. 41

FIGURE 12: NES' SENSORS DEPLOYED INSIDE WASTE CONTAINERS. 44

FIGURE 13: WASTE INJECTION IN WASTE CONTAINER TO TEST CONTAINER FILL-UP. 46

FIGURE 14: ALERT NOTIFICATION SENT TO TELEGRAM BOT WITH OBJECTIVE CONTAINER INFORMATION. 47

FIGURE 15: AGROAPPS 360 FMIS (DRONE DATA INGESTION). 49

FIGURE 16: AGROAPPS 360 REGISTERING CULTIVATION ACTIVITIES PER PARCEL. 50

FIGURE 17: PLANT WEARABLE SENSORS DEPLOYED. 51

FIGURE 18: DRONE AND WEATHER STATIONS PROVIDED BY BARBA STATHIS. 51

FIGURE 19: UGV AND WEED DETECTION. 51

FIGURE 20: TELEMATIC DEVICES PREPARATION AND INSTALLATION ON FLEET. 52

FIGURE 21: OTE DEPLOYED 5GMPN COMPONENTS. 52

FIGURE 22: CLUSTER 3A FIELD DAY 1 - OPEN DEMO DAY. 53

FIGURE 23: CLUSTER 3A ARCHITECTURE. 54

FIGURE 24: CL3A DEMONSTRATION 1. 55

FIGURE 25: CLUSTER 4 PILOT PLANNING. 68

FIGURE 26: CLUSTER 4, PILOT 4 AT VILLÁLVARO, SORIA. 71

LIST OF TABLES

TABLE 1: PILOT 2 SCENARIO 2 KPIS RESULTS. 25
TABLE 2: PILOT 2 SCENARIO 2 KVIS RESULTS. 25
TABLE 3: PILOT 2 SCENARIO 3 KPIS RESULTS. 27
TABLE 4: PILOT 2 SCENARIO 3 KVIS RESULTS. 28
TABLE 5: ACTIVITIES. 32
TABLE 6: PILOT 3 KPIS. 33
TABLE 7: CLUSTER 3A SCENARIO VALIDATION STATUS ACROSS PILOTING PHASES. 58
TABLE 8: CLUSTER 4 USE-CASE OVERVIEW AND STAGE 1 VALIDATION STATUS. 67

ABBREVIATIONS

Abbreviation	Description
5G	Fifth Generation (mobile network)
5G MPN	5G Mobile Private Network
AI	Artificial Intelligence
API	Application Programming Interface
ATSI	AgriTech Transformation and Sustainability Initiative
BCS	Belt Conveyor System
BMP	Business Management Portal
CBM	Condition-Based Maintenance
CHP	Combined Heat and Power
CI/CD	Continuous Integration / Continuous Delivery
CL	Cluster
DEM	Demonstration (campaign)
DGS	Digital Ground Support
DM	Data Management
DMP	Data Management Plan
DO	Domain Orchestrator
DSO	Distribution System Operator
ESO	End-to-End Service Orchestrator
ETSI	European Telecommunications Standards Institute
EU	European Union
EV	Electric Vehicle
FMIS	Farm Management Information System
FR / NFR	Functional Requirement / Non-Functional Requirement
GCP	Google Cloud Platform
GDPR	General Data Protection Regulation
GKE	Google Kubernetes Engine
GPS	Global Positioning System

HypO	Hyper Orchestrator (ETSI HypO)
IM	Instant Messaging
IoT	Internet of Things
IP	Internet Protocol
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
KVI	Key Value Indicator
LED	Light-Emitting Diode
LoRaWAN	Long Range Wide Area Network
LTE	Long-Term Evolution
MQTT	Message Queuing Telemetry Transport
mTLS	Mutual Transport Layer Security
MVP	Minimum Viable Product
NB-IoT	Narrowband Internet of Things
NGSI-LD	Next Generation Service Interfaces – Linked Data
OC	Open Call
OEE	Overall Equipment Efficiency
OTA	Over-The-Air (updates)
PM	Predictive Maintenance
QR	Quick Response (code)
RBAC	Role-Based Access Control
RBM	Rock Bolt Monitoring
REST	Representational State Transfer
SIF	Secure Integration Fabric
SLA	Service Level Agreement
SME	Small and Medium-sized Enterprise
SMS	Short Message Service
TCP	Transmission Control Protocol
TMF	TM Forum
UAV	Unmanned Aerial Vehicle

UC	Use Case
UGV	Unmanned Ground Vehicle
UI	User Interface
UMDH	Urban Mobility Data Hub
VM	Virtual Machine
VPN	Virtual Private Network
WP	Work Package
YOLO	You Only Look Once (object detection model)

ANNEX A – INDIVIDUAL PILOT REPORTS

This annex accompanies deliverable D5.1 [1], "Outcomes of Initial Testing Cycle and Validation Results" and presents the individual pilot reports of the five COP-PILOT validation clusters. For each pilot, it provides the complete documentation of the initial testing cycle: the pilot description, objectives and relevance, the roadmaps at use-case and cluster level, the expected outcomes, the execution steps, the setup and validation tools, the stakeholder involvement, the assessment plan, the execution narrative with the supporting figures and showcasing material, the stakeholder feedback and lessons learned, and the detailed impact-assessment analysis. The main D5.1 report presents, per cluster, the consolidated pilot execution and results — including the KPI/KVI measurements as per D0.1 [2] and the initial impact assessment — while this annex provides the full underlying pilot-level documentation. The two documents are intended to be read together and constitute a single, coherent account of the Stage 1 piloting activity.

1.1 CLUSTER 1: BUSINESS INTEGRATION IN MINING

1.1.1 Pilot 1: Dynamic/Auto Scaling of Compute Resources for Efficiency

1.1.1.1 Summary

This pilot addresses dynamic/auto-scaling of compute resources for efficiency within Use Case UC1.1, IoT Mining Seismics (large-scale micro-seismic sensing, data collection and processing). Seismic processing load in a mine varies strongly between steady-state operation and peak situations, such as when a large seismic event initiates a series of aftershocks. The pilot integrates ColonyOS as a meta-orchestrator into RockSigma’s BEMIS seismic processing system and uses the ColonyOS blueprint and reconciliation mechanisms to automatically scale processing workers up and down based on the observed job queue.

A first version of the auto-scaling feature, together with a test case and demonstration scenario, was implemented and successfully demonstrated at the M13 project review, using a historical mine dataset replayed through a simulator to create a high-variability load. The pilot remains under development and integration towards a more sophisticated auto-scaling solution.

1.1.1.2 Pilot description

The pilot emulates a realistic mine seismic monitoring situation in which the processing load varies strongly over time, and compute capacity must follow actual demand:

- Actors: the BEMIS processing engine; the ColonyOS server, job queue, auto-scaling service and executors (workers); seismologists and mine operators relying on timely processing results.
- Context: underground mine seismic monitoring, where timely delivery of processed data is critically linked to the safety of workers and equipment.
- Key workflows: incoming waveform data generates processing jobs in the ColonyOS job queue; the auto-scaler monitors queue size and resource allocation; workers are added or removed through ColonyOS blueprints and reconciliation; jobs are completed and results delivered to end-users.
- Triggering events: intense seismic activity is observed by a dense seismic network, creating a burst of processing jobs (queue growth) that triggers scale-up, a short or empty queue triggers scale-down.

- Variations: a high-variability load scenario generated by replaying a historical mine dataset through a simulator, future variations towards more sophisticated scaling policies (e.g., refined detection algorithms and SLA-oriented thresholds).
- Integration point: the auto-scaler observes the ColonyOS job queue and actuates worker capacity via ColonyOS blueprints reconciled against the infrastructure.

1.1.1.3 Objectives

The key goals of the pilot are to:

- Implement auto-scaling of BEMIS compute resources using ColonyOS blueprint and reconciliation mechanisms as the core technology (KPI-VSS-CL1-1.4).
- Demonstrate scale-up in response to a growing job queue, and lossless scale-down in response to a short or empty queue.
- Use runtime observation of resource allocation and job queue size as a proxy for processing SLA compliance.
- Improve efficiency by letting compute capacity elastically follow actual demand, instead of static provisioning for peak load.

1.1.1.4 Pilot Relevance

This pilot contributes directly to the goals of Cluster 1 (Business Integration in Mining) and of UC1.1. In mine seismic monitoring, the processing load varies greatly between steady-state operation and time-limited peak situations, and timely delivery of processing results during peak load is critically linked to the safety of workers and equipment. Auto-scaling lets compute capacity elastically follow actual demand, improving efficiency compared to static peak provisioning while safeguarding SLA compliance. The pilot thereby supports the cluster-level value indicators on improved real-time decision-making capabilities in critical situations (KVI-VSS-CL1-1.1) and higher resolution seismic monitoring (KVI-VSS-CL1-1.3).

1.1.1.5 Roadmaps (UC and Cluster Level)

The activities of this pilot are:

- Design and implementation of the auto-scaling feature (first version), test case and demonstration scenario (2025-10-01 to 2026-02-28).
    - Status: completed; used for the M13 project review.
- Development towards a more sophisticated auto-scaling solution during the intermediate evaluation phase.
- Integration and validation in the full UC 1.1 system demonstration by M34.

Design and development activities are carried out in the RockSigma local development environment, testing and validation in a cloud-hosted pilot environment, with final system demonstration and validation planned in the RISE test/demo environment during the final evaluation phase.

1.1.1.6 Expected Outcomes

The following KPIs will be measured, as defined in D0.1:

- KPI-VSS-CL1-1.4 Compute dynamic/auto-scaling for SLA compliance
    - baseline at M1: ‘No’; target at M36: ‘Yes’

The related value indicators are

- KVI-VSS-CL1-1.1 Improved real-time decision-making capabilities in critical situations, and
- KVI-VSS-CL1-1.3 Higher resolution seismic monitoring

Technical behaviours to be observed: scaling up of compute resources in response to a growing queue of jobs and scaling down without loss in response to a short or empty queue, under a high-variability load scenario.

1.1.1.7 Execution Steps

The step-by-step pilot procedure is:

- System setup and initialization: deployment of BEMIS with the integrated ColonyOS version in the cloud-hosted pilot environment, configuration of the auto-scaler (blueprints and scaling thresholds), preparation of the historical mine dataset in the simulator environment.
- Replay of the historical dataset through a simulator to create a high-variability load scenario, including a burst corresponding to a large seismic event followed by aftershocks.
- Data collection: time series of allocated compute resources, job queue length and completed jobs, together with screenshots and recordings used for the demonstration.
- Variations: steady-state load versus time-limited peak load. Future variations will address more sophisticated scaling policies.
- Wrap-up and teardown: verification of lossless scale-down, archiving of demonstration data, and release of pilot resources.

1.1.1.8 Setup and validation tools

This section describes the pilot environment and the tools used to set up, operate and validate the pilot.

Infrastructure setup

During the initial evaluation phase, the pilot has been operated in the RockSigma cloud-hosted pilot environment. This environment hosts the BEMIS processing engine together with the ColonyOS server, and the simulator used to replay recorded datasets and generate load scenarios; it is used for validation runs with recorded sensor data.

In later phases, the pilot will be integrated into the RISE test/demo environment with an edge-to-cloud architecture, providing a COP-PILOT pre-production infrastructure for final testing, validation and demonstration (timeline to be defined).

Technical Setup

The main service is RockSigmas BEMIS Compute Engine, a distributed seismic processing and monitoring system for large-scale seismic networks in mining. BEMIS ingests seismograms from underground sensor arrays, processes them into a high-density seismic event catalogue, and exposes the results to end-users through a REST-API, viewable through RockSigmas web-based visualization and data interrogation application, the BEMIS Quake Monitor. Within COP-PILOT, BEMIS integrates ColonyOS as a meta-orchestration layer: processing workloads are submitted as function specifications to a shared job queue, from which executors pull tasks matching their capabilities. Autoscaling is performed through ColonyOS using its Blueprint and Reconciliation features.

Technical Components Involved

The following systems, platforms and components participate in the pilot:

- Applications: the RockSigma BEMIS Compute Engine (seismic processing pipeline) and the BEMIS Quake Monitor (web-based monitoring and visualization tool).
- Communication and compute infrastructure: cloud infrastructure on the RockSigma cloud. In later phases, the RISE edge-to-cloud test/demo environment.
- Middleware/platforms: ColonyOS acting as job orchestrator / meta-orchestrator, including its blueprint and reconciliation mechanisms, with executors running the processing workloads.
- Monitoring and analytics tools: ColonyOS job queue and executor state monitoring (queue length, running processes, executor liveness and capability profiles), BEMIS API based monitor metrics (n.o events, n.o seismograms, rates, pipelines), both collected through Prometheus and visualized in Grafana. BEMIS Quake Monitor (BEMIS native event visualization in a web-based 3D environment)
- Integration points and data flows: vendor sensor data is converted to a generic data format, orchestrated through ColonyOS into the BEMIS compute engine (executors), stored in the event database, and exposed via API to the visualization layer.

Measurement and validation Tools

KPI measurement follows the COP-PILOT unified validation and evaluation methodology (D0.1). The KPIs of this pilot are measured through P22 (ROC) self-assessment, with status reviewed at each evaluation milestone (initial at M17, intermediate at M25, final at M33).

Validation is based on runtime observation of allocated compute resources and the ColonyOS job queue: the KPI is achieved when the system scales up in response to a growing job queue and scales down without loss when the queue is short or empty. The growth trend and size of the job queue are interpreted as a proxy for processing SLA compliance. Time series of allocated resources, queue length and completed jobs are collected from the demonstration environment.

DMP and Ethics

The pilot uses recorded seismic waveform data obtained from customer mine sites and is stored within the access-controlled RockSigma pilot environment and used solely for validation of the pilot objectives. No personal data is processed in the pilot.

1.1.1.9 Stakeholder Involvement

The following stakeholders participate in the pilot:

- End-users: mining companies operating seismic monitoring systems. In the initial phase, end-user involvement is through recorded datasets from a customer mine site; direct end-user participation is planned for later demonstration phases.
- Cluster 1 partners.

1.1.1.10 Assessment Plan

The pilot-specific KPIs are evaluated through P22 (ROC) self-assessment against the baselines and targets defined in D0.1, with status reviewed at the initial (M17), intermediate (M25) and final (M33) evaluation milestones and reported in the corresponding project deliverables and in this report.

RockSigma analyses the results internally; cluster-level consolidation and reporting are coordinated by the Cluster 1 lead (LTU).

1.1.1.11 Pilot Execution and Results

This section reports on the execution of the pilot during the initial evaluation phase, centred on the auto-scaling demonstration prepared for and shown at the M13 project review.

Scenario 1: Auto-scaling under high-variability load

This pilot was the focus of the demonstration at the M13 project review. While successfully demonstrated, the pilot remains under development and integration towards a more sophisticated auto-scaling solution. At the basic level, the pilot has been validated by observing up- and down-scaling of compute resources using a historical dataset from a mine site, replayed through a simulator environment to create a high-variability load scenario.

The figure below shows a typical case where auto-scaling creates end-user value: a large seismic event initiates a series of aftershocks, creating a time-limited peak-load situation for the processing software. In a scenario like this, timely delivery of processed data is of the essence, as time is critically linked to the safety of workers and equipment.

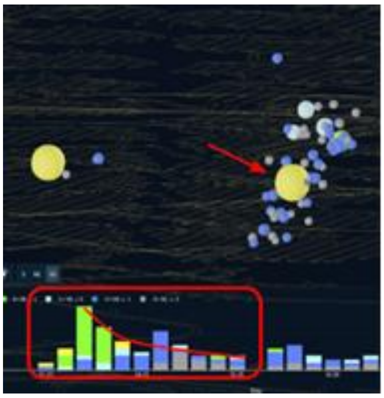

Figure 1: Seismic monitoring view showing a large event followed by a series of aftershocks, creating a time-limited peak-load situation.

The next figure is from the demonstration of this pilot. The top graph shows how compute resources are scaled up and down, the middle graph shows the load presented to the system (measured as jobs in the queue), and the bottom graph shows jobs completed.

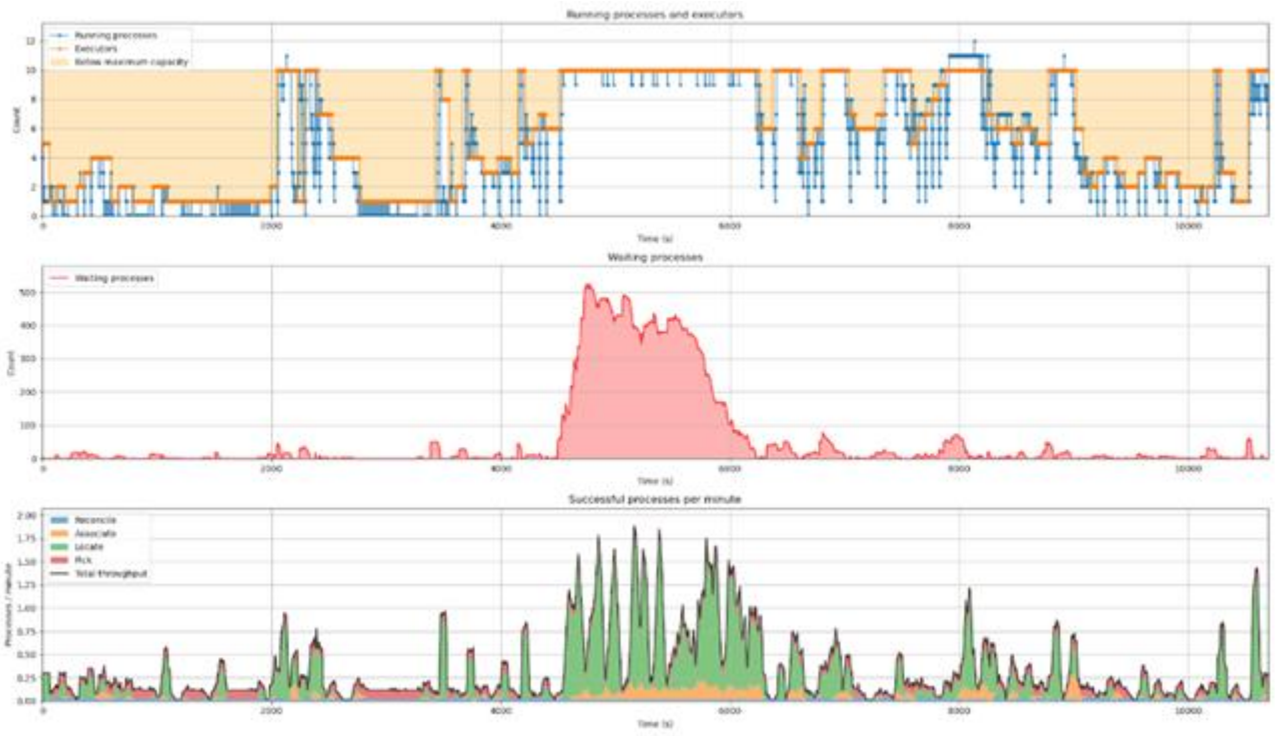

Figure 2: Auto-scaling demonstration: compute resources (top), load presented to the system as jobs in the queue (middle) and jobs completed (bottom).

KPI and KVI validation summary

The tables below summarize the KPI and KVI status for this pilot at the current stage, with definitions, baselines and targets as per the COP-PILOT unified validation and evaluation methodology (D0.1).

Open Data availability

The customer datasets used in this pilot are commercially sensitive and are not published as open data. Where feasible, anonymized or synthetic datasets and aggregated results may be made available in line with the COP-PILOT Data Management Plan.

1.1.1.12 Showcasing

The pilot was demonstrated at the COP-PILOT M13 project review, where the first version of the auto-scaling feature, its test case and the demonstration scenario were presented.

1.1.1.13 Stakeholder Feedback and Impact

Discussions with industry stakeholders highlighted challenges that reinforce the relevance of the demonstrated approach. Several stakeholders noted that infrastructure is often dimensioned to meet demanding service-level requirements, resulting in significant resource consumption and operational costs. In addition, stakeholders expressed interest in expanding the number of monitored assets, sensors, and data streams, but identified the associated increase in computing requirements and costs as a limiting factor.

These observations support the underlying rationale of the pilot, demonstrating the need for more flexible and efficient resource management mechanisms that can dynamically adapt computing capacity to actual workload demands while maintaining required performance levels.

1.1.1.14 Lessons learned

Key takeaways from UC1.1 and this pilot so far:

- Using the ColonyOS as a job orchestrator makes it easier to build and scale a microservices-based system
- Starting with basic auto-scaling algorithms and iterating towards more sophisticated ones and reusing an existing orchestration platform (ColonyOS) rather than building new infrastructure, reduces technical risk.

1.1.2 Pilot 2: Asset tracking and rock bolt monitoring

1.1.2.1 Summary

Cluster1-Pilot 2 is meant to develop, integrate and demonstrate ThingWave's asset tracking and rock bolt monitoring featuring ColonyOS deployment orchestration using OpenSlice. The pilot is divided into two distinct parts, one continuously running part and several smaller distinct scenarios to showcase specific parts to stakeholders during a short timeframe.

1.1.2.2 Pilot description

The pilot aims to emulate the real world scenario of ThingWave's sensors installed in underground mines, and the use of those sensors to enable underground asset tracking. The sensors themselves monitor movement inside the rock mass in order to determine the status of the ground support, using ThingWaves Digital Ground Support (DGS) system. The key actors in the pilot are ThingWave, providing their IoT sensors and platform, which are connected with ColonyOS and Arrowhead, RISE, who provides the emulated testbed, and LTU, who develops and supports ColonyOS and Arrowhead.

For the continuous pilot, the key workflow is an incremental development of the testbed, with more physical demo devices and simulated devices being added into the setup. This simulates day-to-day operations as well as allows for development and testing of new software features and improved processing.

For the distinct scenarios, the workflow is a little different. The main workflow is based on distinct testing of new features, with a higher level of documentation of that feature. The results can then be added into the continuous pilot in order to simulate and test the feature during regular, long-term operation. These scenarios often include specific triggers from the sensor in order to simulate certain events that take place in mining operations. Examples include Rock Bolt Monitoring (RBM) sensor being elongated past a threshold or an asset tracking tag entering an area.

1.1.2.3 Objectives

The pilot has three key goals. It aims to demonstrate the new features developed by ThingWave within the project. This includes underground asset tracking, underground LED status indication, geofences, the new software, Nexus, for ThingWave's gateways as well as integration with Arrowhead and ColonyOS to connect with the COP-PILOT platform. The second goal is validation of improvements to ThingWave's overall solution. This includes improvements to the IoT platform Nucleus, improvement to the existing sensor solutions and scalability improvements for the installation and integration process. Finally, the pilot aims to validate the KPIs and KVIs described in D0.1 regarding UC1.2.

1.1.2.4 Pilot Relevance

The pilot's main contribution to Cluster 1's overall goal is the development and integration of Arrowhead services. This ensures the possibility for the Cluster 1 platform to handle significant data flows to and from different platforms that have Arrowhead integration. In order for this verification, improvements to Nucleus are required as well as the development of a new gateway software, Nexus. To verify sufficient flexibility in the system, it will transmit data from a newly developed seismic sensor, a newly developed asset location data and information provided by existing smart rock bolt sensors.

1.1.2.5 Roadmaps (UC and Cluster Level)

The timeline for the pilot activities are tied to the three COP-PILOT cycles, with the first phase being December 2025 to May 2026, the second phase from August 2026 to January 2027 and the third phase from April 2027 to September 2027. During the first phase the key activities are internal development to enable future integration and development of asset tracking technologies, which includes geofencing development. The asset tracking and geofencing developed will be shown in smaller scenarios, while software development is tested continuously. The asset tracking scenario will be demonstrated in December of 2025 while the geofence scenario will be demonstrated at NorthStar day in March of 2026. A key software development is enabling support for Arrowhead services, as well as development of a new Arrowhead data forwarder. Finally, there is continuous work in order to enable mass deployment of IoT enabled rock bolts. Pilots are vertically integrated with the COP-PILOT project by supporting data streams starting with sensors and actuators in the form of rock bolt sensors (RBM-VX, RBM-XR, vibration sensors, SmartTags) and optical alert systems (LED strips) as well as Instant Messaging (IM) alarms and SMS. Data from sensor devices are forwarded in real-time to ThingWave's IoT platform and forwarded to Eclipse Arrowhead for further COP-PILOT integration. These application systems are to be auto deployed by OpenSlice and Kubernetes.

During the second phase, the software developed in cycle one will help integrate the continuous testing into the RISE testbed and couple UC1.2 to UC1.4. There will also be a smaller demonstration of a new seismic sensor shown in October 2026. During this cycle there will be continued improvements of the system scalability and integration with Arrowhead and ColonyOS. The main focus of the cycle is, beside deploying into UC1.4, improving solutions from the first phase with the help of feedback received from stakeholders.

The final phase of the project, the focus will be on improving and making the integration into Arrowhead and ColonyOS more robust. This will use lessons learned from phase 2. During this phase, autonomous alarm generation will be implemented with support for several ThingWave sensors to generate alarms. There will also be validation of all solutions. The demonstration for this phase will be a collaboration with the rest of Cluster 1 and can be approximated to September of 2027. The final key activity during this phase is the investigation of suitable business models and go to market strategies.

1.1.2.6 Expected Outcomes

The KPIs ranging from KPI-VSS-CL1-2.1 to KPI-VSS-CL1-2.12 are all expected to be measured, either during the ongoing pilot or during smaller scale scenario testing. The measuring of the KPIs result in validation of the KVIs.

No specific stakeholder feedback is expected, but it is rather expected that the feedback received from stakeholders is most relevant to the part of the platform that stakeholder uses. Thus, the feedback can range from hardware installation to network stability or to required improvements of critical software features. What is expected is that each stakeholder will give feedback on the part of the platform that is most critical to them and thus helps remove bottlenecks for scaling ThingWave's EvoMining platform and enabling asset tracking.

1.1.2.7 Execution Steps

The execution of the pilot is generally the same regardless of scenario, with two key differences between the continuous testing and smaller testing scenarios. The differences are that the continuous testing is not torn down before the end of the project and that the smaller scenarios have a distinct scenario script testing smaller, single features instead of running all implemented features.

The system is set up by starting a Nucleus server to run all COP-PILOT testing, which is only done twice. Once to set-up an internal server run by ThingWave, and once to set up the same type of server at RISEs premises. After server initialization a gateway is set-up to connect to the correct instance of Nucleus. This gateway can then connect to any ThingWave sensors relevant for the scenario that is to be executed. Scenarios can be, for example, the tracking of a tag mounted to a specific asset, simulation of a large number (400 to 1000+) sensors or demonstration of physical geofencing equipment. The chain of sensor, gateway and Nucleus is responsible for data collection. In case of a smaller scenario, the test can be temporarily torn down after displaying it, by shutting down the gateway and sensor used. It is easy to restart the equipment by powering up the devices, which helps when integrating into the continuous testing.

When the final solution is implemented and running, tests will be performed to demonstrate how sudden and dangerous rock mass displacement events will trigger LEDs underground to show warning colours to staff near the hazardous area. The entire vertical data path will be powered with existing IoT solutions in collaboration with COP-PILOT technologies, all executing on Cluster 1 data centre by RISE and with COP-Pilot integration in collaboration with LTU.

1.1.2.8 Setup and validation tools

This section outlines the pilot setup and required measurement tools for Cluster1-Pilot 2, with a focus on the continuous pilot deployed at RISE. There is also a description for the already conducted scenarios.

Infrastructure setup

This pilot's infrastructure consists of both hardware and software systems. Hardware in the form of rock bolt sensors, SmartTags, gateways, and servers. The software is a complex distributed platform with ThingWave Nucleus as its core, with additional software components such as Eclipse Arrowhead systems, ColonyOS applications, cluster processing management (Docker and Kubernetes), etc. Then there are multiple Arrowhead applications that interact with data streams provided by Nucleus. The core processing and storage infrastructure is shared between RISE's datacentre and ThingWave's own virtualization infrastructure.

Technical Setup

Test beds include rock bolt sensors, SmartTags and connected LED strips to interact with the physical world. Then the software previously mentioned runs the entire pilot applications. The RISE datacentre also manages data sharing with the upper COP-PILOT infrastructure and orchestration tools.

Technical Components Involved

Below is a list all systems/platforms/devices that participate in the scenarios:

- *Applications: (e.g., vertical-specific software)*

ThingWave Nucleus with Eclipse Arrowhead support, Arrowhead DataForwarder, Nexus for data acquisition,

- *Communication infrastructure: (e.g., 5G/6G nodes, edge/cloud)*

  5G and LTE networks, Edge and Cloud servers, data clusters,

- *Middleware/platforms*

  Arrowhead Core systems, ColonyOS, storage clusters

Measurement and validation Tools

The measurement and validation tools required for Cluster1-Pilot 2 that covers UC1.2 are mainly timing equipment in order to measure time elapsed during certain tasks. This equipment covers KPI-VSS-CL1-2.1, KPI-VSS-CL1-2.2, KPI-VSS-CL1-2.6, KPI-VSS-CL1-2.7, KPI-VSS-CL1-2.8 and KPI-VSS-CL1-2.9. There is also a need for counters, that can help count packets, system restarts and number of sensors. These types of counters are already built into the system, and it covers KPI-VSS-CL1-2.3, KPI-VSS-CL1-2.8 and KPI-VSS-CL1-2.9. For KPI-VSS-CL1-2.5, KPI-VSS-CL1-2.10 and KPI-VSS-CL1-2.11, access to ThingWave's Nucleus platform is required in order to measure them. For KPI-VSS-CL1-2.4, a confirmation from partner companies is required. Finally, there is a need for demo-sensor and simulated sensors, in order to populate the platform with controllable data.

DMP and Ethics

Data that does not originate from ThingWave demo devices and simulation is kept strictly confidential. Otherwise, there are no DMP or ethics issues that need to be adhered to.

1.1.2.9 Stakeholder Involvement

Participates who interact with the scenarios:

- *End-users (e.g., operators, citizens)*

  Mining staff, industry experts,

- *External observers (e.g., city officials, industry reps)*

  Cluster 1 members' network

The two scenarios have additional specific stakeholders;

- Continuous: RISE technical team, LTU team and TW staff, other Cluster 1 members.
- North Star: Telia, Ericsson, LTU, Boliden

  Installations at stakeholders: Geotechnical personnel, …

1.1.2.10 Assessment Plan

All scenarios are to be evaluated by ThingWave personnel. This includes using the measurement tools and methods outlined under the headline "Measurement and validation Tools" above. The analysis and reporting of the results are done by ThingWave personnel as well.

The exception is any scenario involving an installation at a stakeholder, as those scenarios will be evaluated by the external stakeholder AND by ThingWave personnel.

1.1.2.11 Pilot Execution and Results

This section describes the different scenarios of the pilot execution. Scenario 1 describes the continuous testing and will thus be updated with each iteration of the pilot report. The scenarios that follow concern smaller, distinct scenarios used for stakeholder validation or testing of new technologies.

Scenario 1 – Continuous Testing

The current scenario specific results are:

- Determined current installation time for rock bolts (via stakeholder)
- Measured large scale KPIs (2.2, 2.3, 2.5, 2.8, 2.11)
- Continuously track system uptime on simulated system
- Received feedback from stakeholder, both on platform and sensor scalability issues
- Used to improve functionality and performance
- Have platform ready for partner integration
- Discovered bugs in the system from having the system running
- Test COP-PILOT integration technologies

The specific results achieved during the continuous testing can be divided into two parts, one that relies on stakeholders involved and one that comes from the internal work during the testing. The internal work of running the test and by using a large set of sensors have given measurable results for the following KPIs: KPI-VSS-CL1-2.2, KPI-VSS-CL1-2.3, KPI-VSS-CL1-2.5, KPI-VSS-CL1-2.8 and KPI-VSS-CL1-2.11. The results are shown in Table 3-4 of the main D5.1 report. The running test has also helped very specifically with KPI-VSS-CL1-2.9, as a continuous test can measure the systems uptime in a controlled manner. There have also been several smaller bugs found and fixed from having ThingWave's Nucleus and Nexus systems running, as well as feature improvements such as a better data-plot tool.

There are also results that rely on stakeholders involved with ThingWave. One specific example is the opportunity to measure the installation time for smart rock bolts in a real setting. That involvement has also resulted in feedback both on the platform and on the sensors. This feedback has been used for improvements that are necessary in order to scale the business.

KPI and KVI validation summary

The KPI and KVI connected to the scenario are displayed in Table 3-4 and Table 3-5 of the main D5.1 report. The scenario helped get initial values for all KPIs, while the KVI cannot be evaluated at this time in the pilot.

Open Data availability

Data suitable for public disclosure will be made available: https://github.com/thingwave

1.1.2.12 Showcasing

The continuous nature of the scenario gives a different type of showcasing, as stakeholders, presences and attendees are not relevant measurement tools. The results from the continuous testing instead serve a different type of showcasing ability, as results from the scenario are used as presentation material or in other dissemination. One example is in Figure 3, showing the large scale test conducted to measure KPIs.

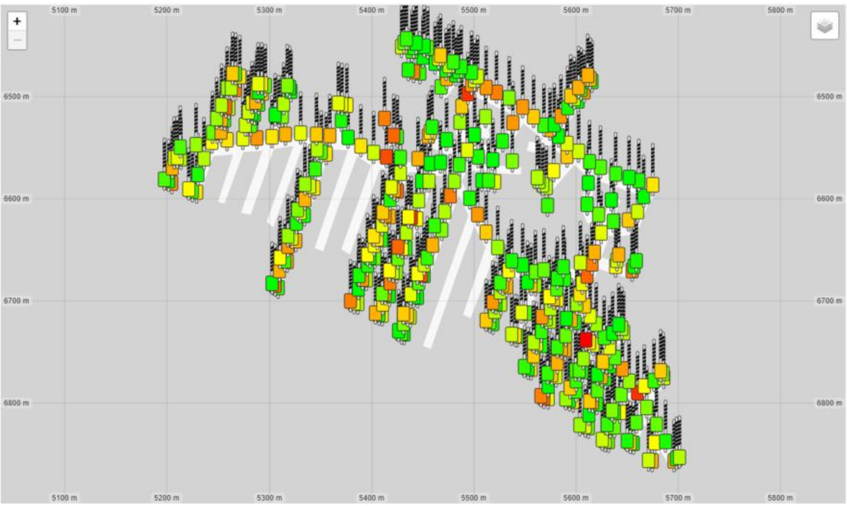

Figure 3: Large scale test of simulated sensors in Nucleus, showing a fictional map with 400 simulated sensors.

ThingWave staff have also demonstrated and discussed our technologies as major events, such as IMARC, EACG, and Euro Mine Expo.

1.1.2.13 Stakeholder Feedback and Impact

The continuous aspect of this pilot due to deployment at ThingWave's headquarters and deployment at RISE means that there is no set number of attendees and stakeholders. Within the continuous testing, installation at stakeholder sites is a part, but due to confidentiality, as well as varying number of attendees, no set number of attendees can be given.

The impact of the continuous testing has been largely positive, with a large amount of information regarding the current scalability of the system having been extracted. The continuous testing has also given the opportunity to discover smaller bugs in the system that only show up during longer test runs.

1.1.2.14 Lessons learned

Key takeaways are:

- Improved hardware needed for large scale installation
- Need dockers for SW deployment in order to visualize the required number of bolts.
- Numerous small platform functionality upgrades for better user experience

Scenario 2 – Geofence demonstration at NorthStar day

The scenario involved a display of a brand new geofence function using ThingWave technologies. During the NorthStar day in Luleå (11-12/3 - 2026) a physical LED-strip was used as a geofence equipment that can mark if an area underground is safe or not. This was demonstrated by pulling a demo-bolt, visualizing to the crowd how the rock bolt was tensioned close to its breaking point. That pull was used to generate an event in the system which changed to LEDs to shine bright red.

KPI and KVI validation summary

The KPI and KVI connected to the scenario are displayed below in Table 1 and Table 2. The scenario helped get initial values for all KPIs, while the KVI cannot be evaluated at this time in the pilot.

Sub-scenario	KPI name	Objective	Result
NorthStar day	**KPI-VSS-CL1-2.6:** Alarm detection latency	Get initial evaluation of alarm latency	30s
NorthStar day	**KPI-VSS-CL1-2.12:** Data processing distribution	Data processing tested	Manual at cloud

Table 1: Pilot 2 Scenario 2 KPIs results.

Sub-scenario	KVI name	Objective	Result
na	**KVI-VSS-CL1-2.6:** Timely hazard alerting for underground workers	na	na

Table 2: Pilot 2 Scenario 2 KVIs results.

Open Data availability

No permanent sensor data was saved. KPI measurements were noted and saved.

1.1.2.15 Showcasing

The event consisted of two parts, one demonstration in Mjölkuddsberget in Sweden and a presentation of ThingWave's technologies. The event was hosted by Telia, and attended by partners such as LTU, Ericsson, Boliden, Epiroc and more. The most important demonstration of technologies was in Mjölkuddsberget, with results shown in Figure 4 and Figure 5.

Figure 4: Demo of physical geofence indicating an area is safe.

The plan is to continue testing inside Mjölkuddsberget, if permission can be obtained. Several mines have also offered to assist with real world underground testing.

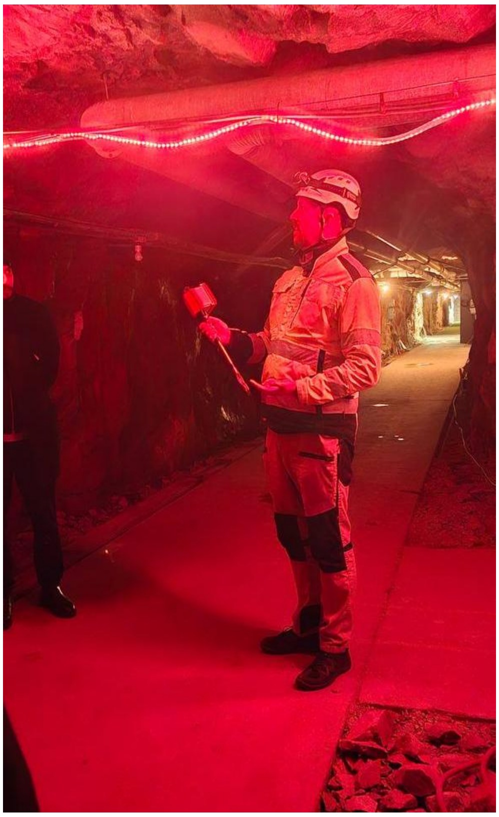

Figure 5: Demo of physical geofence, showing when a bolt is elongated close to maximum and thus the indication shows that the area is unsafe.

1.1.2.16 Stakeholder Feedback and Impact

The pilot was held at NorthStar day in Luleå, an event arranged by Telia. Key stakeholders included Telia, Ericsson, Boliden, LTU, Epiroc, and RISE. Other stakeholders partaking were Karlstad Universitet, Region Stockholm, SAAB, Traton, Lindholmen Science Park, Nobina, Westermo, Rejlers, Svevia

Pilot has also been presented to America based companies in Utah, Nevada, Australia, and Canada.

1.1.2.17 Lessons learned

Key takeaways:

- Interest for the technologies
- LEDs powerful tool in dark areas to increase awareness of hazards
- Automated smart bolt installation capabilities vital to achieve good coverage for asset tracking

Scenario 3 – Initial asset tracking demo

The specific scenario tested was to accurately track a tag between two gateways mesh networks. This was done by connecting a tag to a gateway and seeing that it shows up on Nucleus. The tag was then physically carried away from the office to disconnect it from the gateway. When the tag was shown as disconnected, and thus not in the same area, the original gateway was disconnected from its power source. A new gateway, with a “different” location, was then plugged in before the tag was physically carried back to the office. The tag connected to the new gateway, which means it was detected in a new area of the mine.

KPI and KVI validation summary

The KPI and KVI connected to the scenario are displayed below in Table 3 and Table 4. The scenario helped get initial values for all KPIs, while the KVI cannot be evaluated at this time in the pilot.

Sub-scenario	KPI name	Objective	Result
Initial asset tracking demo	**KPI-VSS-CL1-2.7:** Position update rate	Get initial measurement value	Automated positional update every 10 minutes
Initial asset tracking demo	**KPI-VSS-CL1-2.10:** Capability integration	Get initial measurement value	Manually integrated position

Table 3: Pilot 2 Scenario 3 KPIs results.

Sub-scenario	KVI name	Objective	Result
na	**KVI-VSS-CL1-2.3:** Cost efficiency of lost asset localisation	na	na

na	**KVI-VSS-CL1-2.5:** Improved situational awareness of worker location	na	na

Table 4: Pilot 2 Scenario 3 KVIs results.

Open Data availability

Data suitable for public disclosure will be made available on ThingWave's GitHub account: https://github.com/thingwave

1.1.2.18 Showcasing

The initial demo was showcased as a video used for the M13 review within the project and it was uploaded to the project website and YouTube. The scenario did not have stakeholders physically present for the initial use of the demo, only ThingWave personnel. Therefore, the material was captured on film in order to create a video for dissemination purposes.

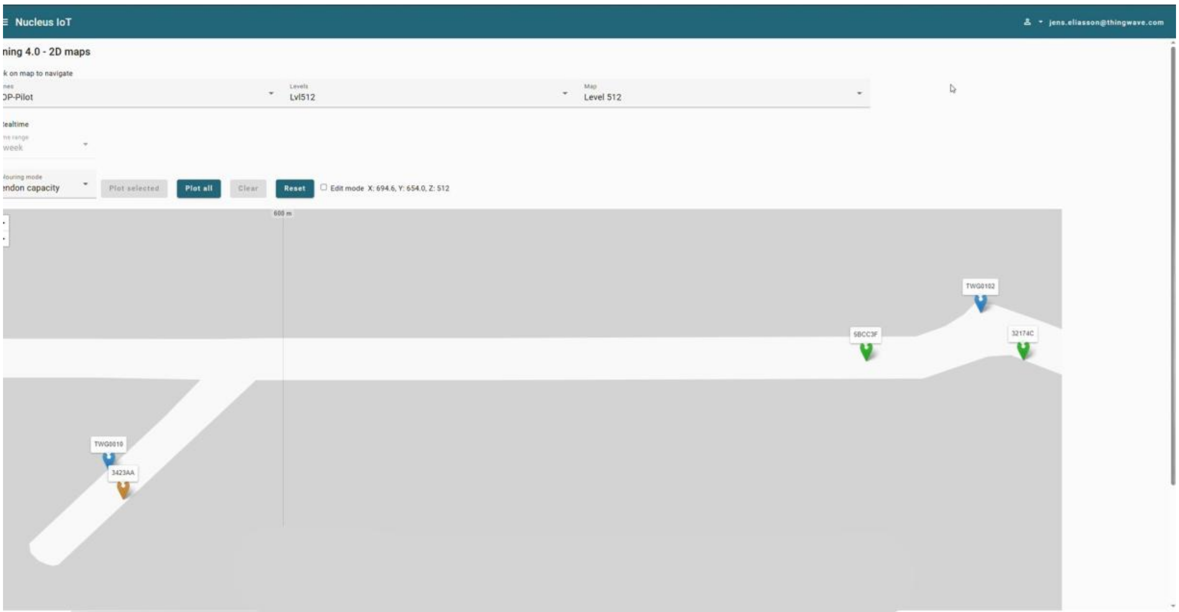

nes: OP-Pilot
Levels: Lvl512
Map: Level 512
Realtime
ne range: week
ouring mode: endon capacity
☐ Edit mode X: 694.6, Y: 654.0, Z: 512

Figure 6: Display of how a tag (orange marker) is shown in Nucleus map tool, as used in initial asset tracking demo. The blue marker is a gateway and the green markers are rock bolt sensors.

1.1.2.19 Stakeholder Feedback and Impact

This pilot has been presented at several industry expos and fairs. The feedback has been very positive! We have also discussed the planned work with other major enterprises in the mining industry. Companies such as IBM and Ericsson are closely looking at what we are working on with great interest.

We have also presented our ideas and demos to business partners and other SMEs in the mining industry and have also received positive feedback from them.

Co-funded by
the European Union

1.1.2.20 Lessons learned

We have Key takeaways:

- Initial demo showed that the technologies is feasible
- Stakeholder feedback shows the need for the proposed solution
- Additional applications can be based in the developed technologies for:
- Reduced CO2 emissions by smarter ventilation control
- Increased safety
- Reduced time waste due to better situational awareness

1.1.3 Pilot 3: Condition monitoring and predictive maintenance

1.1.3.1 Summary

In industrial and mining operations, belt conveyor systems (BCS) form the backbone of raw material transport. Any downtime or underutilisation of these assets has an immediate effect on overall equipment efficiency (OEE).

To address the maintenance challenges of BCS, the pilot integrates IoT-based condition monitoring with cloud analytics across two complementary platforms:

HOSCH - provides the industrial sensing hardware and web platform for continuous conveyor belt monitoring

Predge Conveyor Solutions - delivers the analytics and predictive maintenance capabilities, built on data streamed from the HOSCH platform.

The pilot aims to demonstrate the following outcomes enabled by COP-PILOT technologies:

- Efficient and dependable conveyor belt condition monitoring at scale
- Realistic industrial data flow within the COP-PILOT pre-production environment
- Scalable device connectivity — demonstrating reliable pipeline performance under high device concurrency at the RISE pre-production environment
- Reduced commissioning time and increased automation for onboarding new customers and data sources
- Improved data acquisition at higher sampling rates enabled by the RISE hardware test bed
- Seamless integration between the HOSCH platform and Predge Conveyor Solution for end-to-end data flow

1.1.3.2 Pilot description

Due to the strict requirements of the industrial monitoring context, the pilot is set up in the pre-production environment at RISE in Luleå, Sweden. Device simulation runs at RISE; all processing, storage, and application components are managed by HOSCH on GKE and the Predge Conveyor Solution by Predge.

Actors:

- Development engineers (HOSCH and Predge)

Context:

- The pilot is executed in an industrial mining and raw material handling context, emulated at the RISE pre-production environment in Luleå, Sweden, where device simulators replicate real industrial hardware and publish telemetry via MQTT with mTLS authentication to the HOSCH VerneMQ broker on GKE.
- Customer and BCS onboarding is performed as part of the pilot to measure commissioning time and degree of automation (KPI-3.1, KPI-3.2)
- The RISE hardware test bed is used to evaluate improved sensing hardware capable of higher data acquisition rates, ranging from 100 Hz baseline up to 500 Hz (KPI-3.3)
- The pipeline is tested under full 50-device concurrency to validate scalability of data acquisition, storage, and processing (KPI-3.4)
- Processed data forwarded from the HOSCH platform to Predge Conveyor Solution for analytics, with results written back to the HOSCH platform for viewing via the HOSCH platform interface

Key Workflows:

- Deployment and integration of new sensing equipment and data sources
- Device onboarding: certificate provisioning, registration in HOSCH platform, MQTT connection validation
- Data ingestion: MQTT → VerneMQ (RISE → HOSCH GCP) → worker → InfluxDB
- Data forwarding: Raw Voltage data sent from HOSCH platform to Predge Conveyor Solution via the Predge API

Triggering Events:

- Customer or operator request to integrate a new data source
- Simulated voltage data published from RISE-hosted device simulators

Variations:

- High device concurrency: all 50 devices publishing simultaneously at 100 Hz baseline
- Burst load: devices increasing publish rate for a 5-minute window
- Device reconnection: selective disconnect/reconnect to validate session recovery
- The integration and commissioning process is timed and monitored to capture KPI-VSS-CL1-3.1 and KPI-VSS-CL1-3.2

1.1.3.3 Objectives

- Demonstrate increased efficiency in onboarding and maintaining new customers, and data sources (KPI-VSS-CL1-3.1, 3.2)

- Validate the end-to-end data pipeline from RISE device simulators through MQTT ingestion, InfluxDB storage, Predge Conveyor Solution analytics, and visualisation via the HOSCH platform interface
- Demonstrate increased data acquisition throughput with higher sampling rates (KPI-VSS-CL1-3.3)
- Demonstrate increased reliability of data streams when scaling to a larger number of concurrently connected devices (KPI-VSS-CL1-3.4)
- Confirm VerneMQ mTLS authentication and topic-level authorisation at 50-device concurrency

1.1.3.4 Pilot Relevance

The pilot demonstrates how data-driven CBM and PM, combining edge, IoT, and cloud infrastructure, can improve OEE in industrial operations. The use case is established around BCS and their maintenance in a simulated mining-adjacent context.

While current industry practices can realise such solutions, there are shortcomings in scalability and reliability. Using the HOSCH platform and Predge Conveyor Solution — it is expected to address these shortcomings and improve efficiency, operability, and flexibility.

Expected Innovations:

- Efficient onboarding of new data sources for an existing customer and BCS — measured via KPI-VSS-CL1-3.1 and 3.2
- Efficient onboarding of new customers and BCS
- Improved data acquisition at higher sampling rates enabled by the RISE hardware test bed — KPI-VSS-CL1-3.3
- Scalability of the pipeline to many concurrently connected devices — KPI-VSS-CL1-3.4

1.1.3.5 Roadmaps (UC and Cluster Level)

The roadmap below aligns pilot activities with both the HOSCH use-case milestones and the COP-PILOT Cluster 1 schedule. Activities in Phases P1 and P2 correspond to the Domain #3 (Predge/HOSCH) task rows visible in the cluster Gantt chart.

Cycle	Planned Dates	Location	Key Activities
CP-1 (P1)	Q1–Q3 2026 (Jan–Aug)	RISE (simulators) + GKE / remote (HOSCH)	Simulation environment development; API implementation to replicate IoT device data transfer; MVP demonstration of the platform; commissioning time baseline (KPI-3.1); internal KPI capture; scalability stress test (KPI-3.4).
CP-2 (P2)	Q3–Q4 2026 (Aug–Dec)	RISE (simulators) + GKE / remote (HOSCH)	Hardware test-bed deployment; testing alternative hardware setups; deployment of hardware, data transmission and

			orchestration concept; burst load testing (KPI-3.3); commissioning automation improvements (KPI-3.2).
CP-3 (P3)	Q1–Q4 2027 (Feb–Mar)	RISE (simulators) + On-site / Partner (TBC)	Optimisation of hardware setup and data flow protocols; deployment of optimised hardware and cloud components; final testing and validation including UI, alarm and notification functions; full end-to-end demo with external stakeholders; final KPI reporting and lessons-learned documentation.

Table 5: Activities.

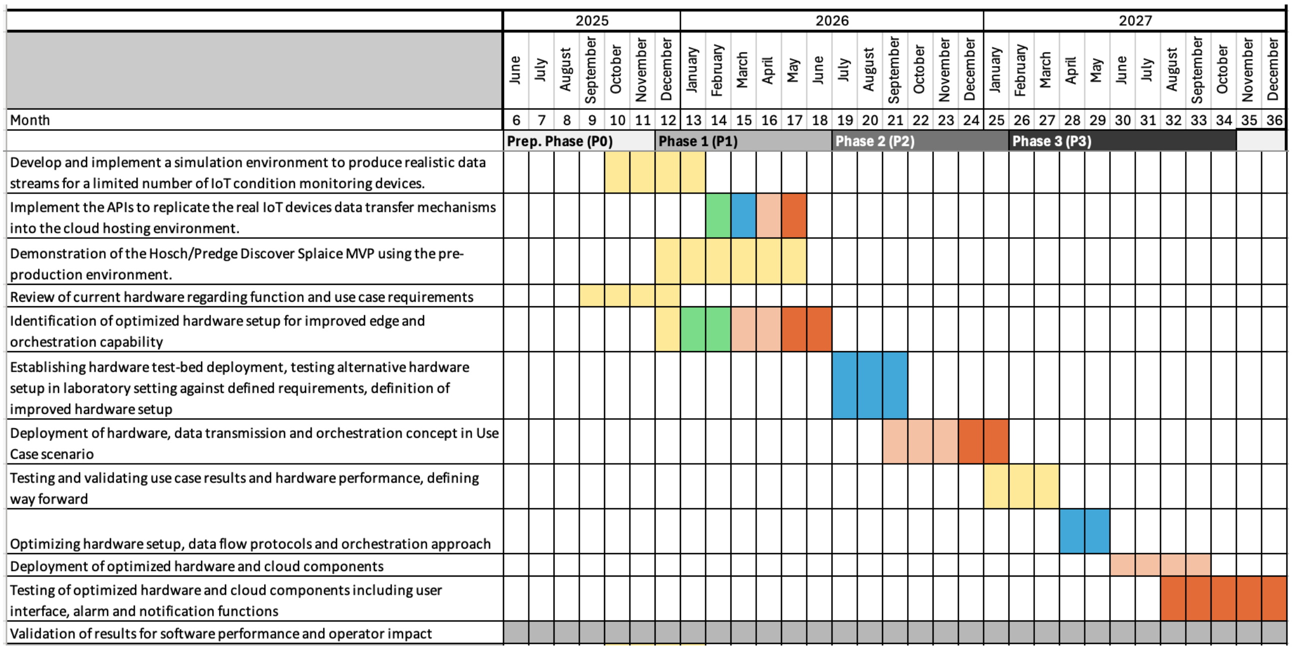

	2025							2026												2027											
	June	July	August	September	October	November	December	January	February	March	April	May	June	July	August	September	October	November	December	January	February	March	April	May	June	July	August	September	October	November	December
Month	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
	Prep. Phase (P0)						Phase 1 (P1)							Phase 2 (P2)							Phase 3 (P3)										
Develop and implement a simulation environment to produce realistic data streams for a limited number of IoT condition monitoring devices.					🟨	🟨	🟨	🟨																							
Implement the APIs to replicate the real IoT devices data transfer mechanisms into the cloud hosting environment.									🟩	🟦	🟧	🟥																			
Demonstration of the Hosch/Predge Discover Splaice MVP using the pre-production environment.							🟨	🟨	🟨	🟨	🟨	🟨																			
Review of current hardware regarding function and use case requirements				🟨	🟨	🟨	🟨																								
Identification of optimized hardware setup for improved edge and orchestration capability							🟨	🟩	🟩	🟧	🟧	🟥	🟥																		
Establishing hardware test-bed deployment, testing alternative hardware setup in laboratory setting against defined requirements, definition of improved hardware setup														🟦	🟦	🟦															
Deployment of hardware, data transmission and orchestration concept in Use Case scenario																🟧	🟧	🟧	🟥	🟥											
Testing and validating use case results and hardware performance, defining way forward																				🟨	🟨	🟨									
Optimizing hardware setup, data flow protocols and orchestration approach																							🟦	🟦							
Deployment of optimized hardware and cloud components																									🟧	🟧	🟧	🟧			
Testing of optimized hardware and cloud components including user interface, alarm and notification functions																											🟥	🟥	🟥	🟥	🟥
Validation of results for software performance and operator impact	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒	▒

Figure 7: Roadmap for the HOSCH-Predge pilot (Pilot 3 in Cluster 1).

Design & Development activity	🟨
Integration activity	🟩
Deployment/Provisioning activity	🟦
Testing activity	🟧
Validation activity	🟥

1.1.3.6 Expected Outcomes

The following KPIs are expected to be measured during the pilot's different phases. Not all KPIs will be measured during all phases.

KPI Name	Measured Value	Description
KPI-VSS-CL1-3.1	Commissioning time	Commissioning time for data pipelines — from initial configuration to first validated data flow, covering device registration, broker setup, and backend connectivity.
KPI-VSS-CL1-3.2	Commissioning automation	Degree of automation in onboarding: (a) new data sources / devices, (b) new customers. Measured as the percentage of onboarding steps completed without manual intervention.
KPI-VSS-CL1-3.3	Data sample rate	Maximum sustained sampling rate (Hz) at which the hardware and pipeline can acquire and process data without loss, measured per device and aggregated across the 50-device test bed.
KPI-VSS-CL1-3.4	Scalability to large number of devices	Number of devices that can be concurrently managed by the complete data acquisition, storage, and processing pipeline without degradation of throughput, latency, or data integrity.

Table 6: Pilot 3 KPIs.

1.1.3.7 Execution Steps

Case 1 – Customer / BCS / Data Source Onboarding

- Start measuring the time
- Setting up the customer and/or BCS in the HOSCH platform
- Setting up available data sources at Predge Conveyor Solution
- Stop measuring the time
- Validate the uncompromised arrival of the data

Case 2 – Scalability and Performance

- System setup and initialisation: all 50 devices active and streaming at RISE
- Scenario script execution: ramp publish rate from 100 Hz to 500 Hz across all devices
- Data collection: measure throughput, end-to-end latency, and message loss
- Fault injection or variations: device reconnection, sustained burst load
- Wrap-up and teardown: confirm clean state

1.1.3.8 Setup and validation tools

The HOSCH platform architecture for the pilot is shown in Figure 8. The system spans three distinct environments: device simulators at RISE as RISE Test Bed, the HOSCH platform on GCP/GKE, and the Predge Conveyor Solution hosted at Predge. Data flows from RISE-hosted Device Simulator through the HOSCH platform (via MQTTs, InfluxDB) and onward to the Predge Conveyor Solution via the Predge API. The computed analytics are written back to the HOSCH InfluxDB, the data is displayed via the HOSCH platform interface.

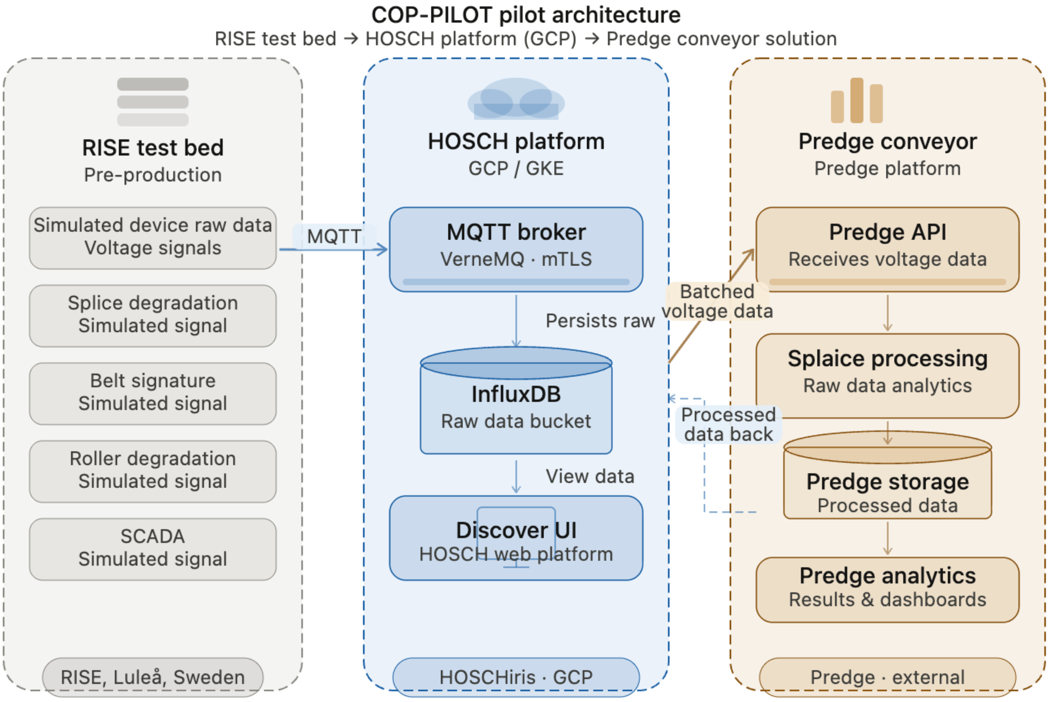

Figure 8: HOSCH PILOT system architecture — Simulator (RISE) → HOSCH cloud platform (GCP) → Predge Conveyor Solution (Predge).

Infrastructure setup

The device data is simulated at the RISE test bed and streamed to the HOSCH platform on GCP, where it is stored. The data is then forwarded to Predge for further analysis and made available for viewing through the HOSCH platform interface.

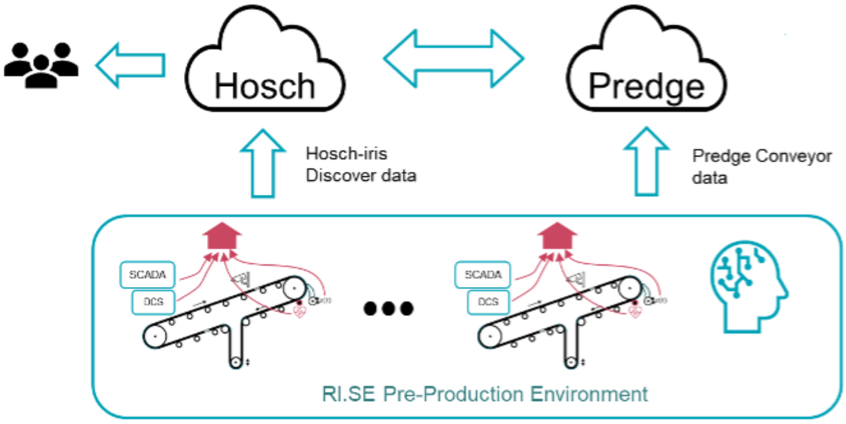

Figure 9: Infrastructure Setup.

Technical Setup

The BCS for specific customers are deployed in the pre-production environment at RISE, with all involved data sources simulated as containerised components. The deployment of data sources at RISE is not in scope of the KPI measurement. The onboarding of data sources occurs in the HOSCH platform where a digital twin of the BCS is also established.

For KPI assessment, only data streams are aggregated, and engineering hours are timed.

Technical Components Involved

- Applications: HOSCH platform, Predge Conveyor Solution, RISE pre-production environment
- Communication infrastructure: Simulated edge devices at RISE streaming over MQTT to HOSCH platform (GCP)
- Middleware/platforms: VerneMQ (MQTT broker), InfluxDB (data storage)
- Monitoring or analytics tools: Python, Grafana, Loki
- Integration points and data flows: Device simulators at RISE → HOSCH platform (GCP) via MQTT → InfluxDB storage → Predge Conveyor Solution via Predge API

Measurement and validation Tools

- Time-taking: engineering hours and commissioning steps logged manually — primary KPI-3.1 and 3.2 measurement
- Data stream comparison: published message count from RISE simulator logs vs. stored measurement count in InfluxDB — primary data loss validation
- Prometheus + Grafana: real-time dashboards for broker connection count, ingestion rate, worker CPU/memory; snapshots exported at end of each scenario
- HOSCH platform Interface: visual confirmation of data arrival and correct display of sensor streams and conveyor asset status — end-to-end pipeline validation
- Predge Conveyor Solution: validation that processed analytics results are correctly received, computed, and written back to the HOSCH platform

DMP and Ethics

Not applicable, as only synthetic data is used, which can be freely distributed and used.

1.1.3.9 Stakeholder Involvement

It is important to note that different stakeholders from the list below will be involved at different stages of the pilot.

An overview of the roles that are relevant is as follows:

- End-users: Platform operators and maintenance engineers using the HOSCH platform interface for data monitoring and BCS onboarding
- External observers: Maintenance managers or digitalisation officers from industrial operations
- Internal technical teams: Developers at HOSCH managing cluster provisioning, pipeline operation, and the platform. Developers at Predge managing the simulator and the Conveyor Solution

1.1.3.10 Assessment Plan

Assessments are conducted within the validation activities of the pilot, see Figure 7. Namely these instances occur starting at M17 of the project.

1.1.3.11 Pilot Execution and Results

The pilot scenarios were defined and integration tests conducted in the lead-up to the initial evaluation phase. This chapter reports on the execution of the scalability scenario carried out during the initial evaluation at M17.

Scenario 1 – 50 Device Concurrency – stress test

The scalability scenario was executed at the RISE pre-production environment with 50 simulated field devices streaming voltage data to the HOSCH platform on GKE. The scenario aimed to validate the end-to-end data pipeline under full 50-device concurrency, covering data acquisition, ingestion via MQTT, storage in InfluxDB, and forwarding to the Predge Conveyor Solution.

The scenario was not successfully completed at this evaluation phase due to technical issues encountered during execution. Data loss was observed under full 50-device concurrency, preventing the successful validation of KPI-VSS-CL1-3.4 at M17. The identified issues are under investigation, and mitigation measures are being implemented ahead of the intermediate evaluation at M25.

KPI and KVI validation summary

Open Data availability

Not applicable as the mining sector is a critical infrastructure that will not share data.

1.1.3.12 Showcasing

No showcasing events have been conducted at this stage of the pilot. The current evaluation phase is based on a simulated pre-production environment at RISE and is not yet suitable for external demonstration.

1.1.3.13 Stakeholder Feedback and Impact

Stakeholder discussions confirmed the relevance of several key concepts addressed within Cluster 1. Feedback indicated that secure on-premises integration is often perceived as more relevant than direct cloud connectivity, particularly in industrial control environments where data security and operational requirements are critical. Several stakeholders expressed concerns regarding the use of cloud connections for operational data due to security considerations.

At the same time, edge device orchestration and management capabilities, such as over-the-air updates, were identified as having significant operational value. Stakeholders also highlighted that the practical applicability of edge-to-cloud solutions depends on available bandwidth and data volumes, especially in underground mining environments.

Overall, the feedback supports the Cluster 1 focus on edge computing, secure data handling, and distributed orchestration, while emphasizing the importance of flexible deployment models that can accommodate both cloud-connected and on-premises industrial infrastructures.

1.1.3.14 Lessons learned

- RISE simulator and HOSCH platform performed reliably — the pre-production environment at RISE and the HOSCH data ingestion pipeline demonstrated stable performance under full 50-device concurrency, confirming the robustness of the device simulation and cloud infrastructure
- Data loss detection needs early validation — testing the data loss detection scheme with mock data prior to full scenario execution would improve confidence in the measurement methodology for KPI-3.4

Initial Impact Assessment – Cluster 1

Cluster 1 focuses on the validation of edge-enabled solutions for digital mining environments, including distributed seismic monitoring, edge-to-cloud deployment capabilities, asset localisation, and automated tunnel inspection. As highlighted in Section 3.0, the cluster exhibits a predominantly business-oriented value profile, with 80% of identified KVIs associated with business value outcomes, compared to 15% societal and 5% sustainability-related outcomes. Table 3-8 of the main D5.1 report presents a selected summary of the principal KPI-KVI relationships and initial validation findings available at M17, based on the indicators defined in D0.1. Given the breadth of validation activities undertaken within the cluster, the following analysis also considers supporting evidence reported across the wider Cluster 1 pilot activities, including operational results relating to deployment efficiency, monitoring capabilities, hazard alerting, and worker situational awareness. This broader perspective provides a more comprehensive understanding of the emerging business, societal, and sustainability value pathways being established through the initial validation cycle.

The strongest evidence currently relates to the **business dimension**, where several quantitative and qualitative indicators demonstrate encouraging progress towards improved operational resilience, scalability, and deployment capability. Initial results include a 90% state reconciliation success rate in test scenarios against a ≥95% target — clear progress, not yet achieved — testbed availability exceeding 98%, successful demonstrations of automated workload placement and reassignment, and improved observability and auditability capabilities. Additional evidence from UC1.2 includes reductions in rock bolt installation and configuration time, increased sensor infrastructure scale, improved data reception performance, and automated positional updates. While many of these results have been obtained within controlled environments and therefore do not yet represent validated operational impact, they provide meaningful proxy evidence that the platform is progressing towards its intended business value outcomes.

The **societal dimension** remains at an early stage of validation but already demonstrates several promising value pathways. In addition to the worker safety and critical decision-making KVIs presented in Table 3-8, UC1.2 provides evidence related to improved situational awareness and hazard alerting capabilities. Automated positional updates have replaced manual worker-location communication processes, while alarm detection and SMS alerting have been implemented with a measured latency of approximately 30 seconds. Although these results do not yet demonstrate realised safety outcomes or improved emergency response performance, they represent important enabling capabilities that support the cluster's longer-term societal objectives. Similarly, while distributed edge-to-cloud processing for critical seismic decision support has not yet been fully validated, initial autoscaling and workload management demonstrations indicate progress towards the required technical foundations.

Sustainability outcomes are also in the early stages of validation. At present, the available evidence should be interpreted as establishing the value pathway rather than quantifying final environmental impact. For example, the transition from manual inspection activities towards automated monitoring and reporting through Nucleus creates a credible basis for reducing inspection-related travel and associated emissions in future operational deployments. Additional sustainability-relevant value pathways include extended functional lifetime of tracked assets, improved infrastructure efficiency, and reduced cloud dependency through edge-enabled processing. While further modelling work will be required to quantify these benefits, the current results provide encouraging proxy evidence that supports the sustainability logic embedded within the Cluster 1 evaluation framework.

1.2 CLUSTER 2: SMART SUSTAINABLE IOT SOLUTIONS IN VALENCIA

1.2.1 Pilot Scope and Planning

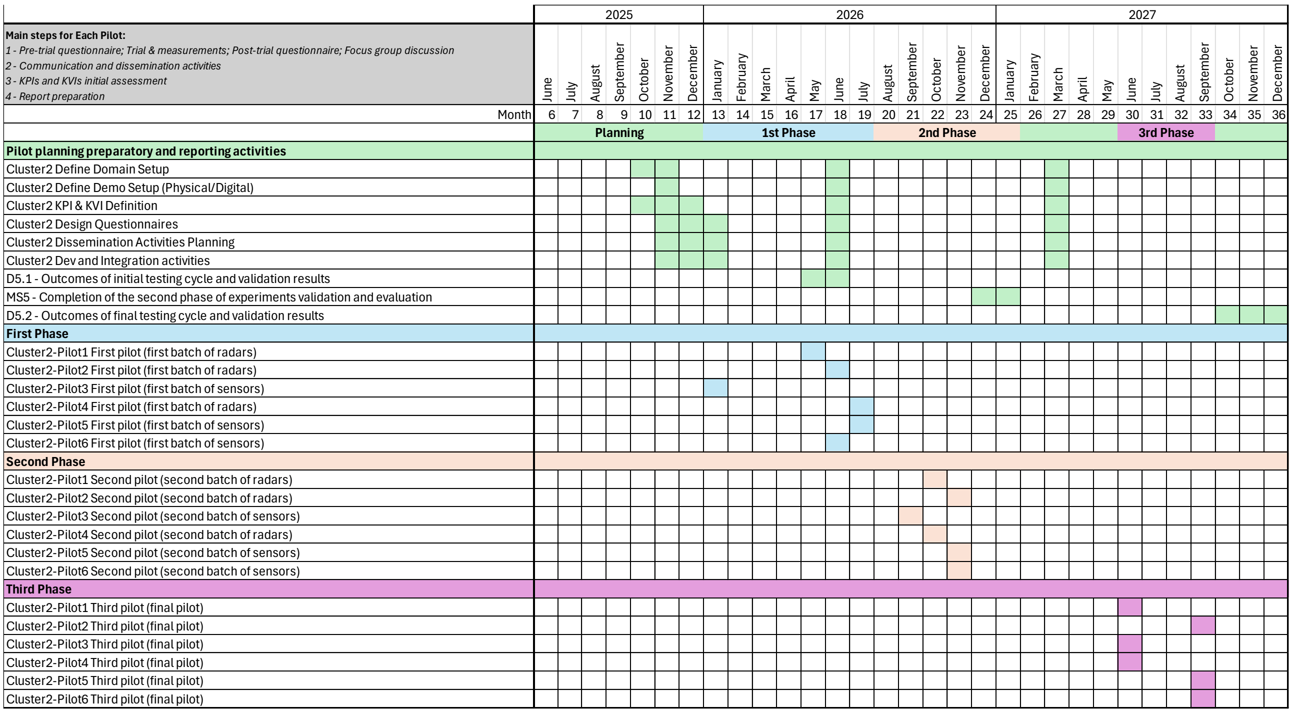

	2025							2026												2027											
Main steps for Each Pilot: *1 - Pre-trial questionnaire; Trial & measurements; Post-trial questionnaire; Focus group discussion* *2 - Communication and dissemination activities* *3 - KPIs and KVIs initial assessment* *4 - Report preparation*	June	July	August	September	October	November	December	January	February	March	April	May	June	July	August	September	October	November	December	January	February	March	April	May	June	July	August	September	October	November	December
Month	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
	Planning							**1st Phase**							**2nd Phase**										**3rd Phase**						
Pilot planning preparatory and reporting activities																															
Cluster2 Define Domain Setup					🟩	🟩							🟩									🟩									
Cluster2 Define Demo Setup (Physical/Digital)						🟩							🟩									🟩									
Cluster2 KPI & KVI Definition					🟩	🟩	🟩						🟩									🟩									
Cluster2 Design Questionnaires						🟩	🟩	🟩					🟩									🟩									
Cluster2 Dissemination Activities Planning						🟩	🟩	🟩					🟩									🟩									
Cluster2 Dev and Integration activities						🟩	🟩	🟩					🟩									🟩									
D5.1 - Outcomes of initial testing cycle and validation results												🟩	🟩																		
MS5 - Completion of the second phase of experiments validation and evaluation																			🟩	🟩											
D5.2 - Outcomes of final testing cycle and validation results																													🟩	🟩	🟩
First Phase																															
Cluster2-Pilot1 First pilot (first batch of radars)												🟦																			
Cluster2-Pilot2 First pilot (first batch of radars)													🟦																		
Cluster2-Pilot3 First pilot (first batch of sensors)								🟦																							
Cluster2-Pilot4 First pilot (first batch of radars)														🟦																	
Cluster2-Pilot5 First pilot (first batch of sensors)														🟦																	
Cluster2-Pilot6 First pilot (first batch of sensors)													🟦																		
Second Phase																															
Cluster2-Pilot1 Second pilot (second batch of radars)																	🟧														
Cluster2-Pilot2 Second pilot (second batch of radars)																		🟧													
Cluster2-Pilot3 Second pilot (second batch of sensors)																🟧															
Cluster2-Pilot4 Second pilot (second batch of radars)																	🟧														
Cluster2-Pilot5 Second pilot (second batch of sensors)																		🟧													
Cluster2-Pilot6 Second pilot (second batch of sensors)																		🟧													
Third Phase																															
Cluster2-Pilot1 Third pilot (final pilot)																									🟪						
Cluster2-Pilot2 Third pilot (final pilot)																												🟪			
Cluster2-Pilot3 Third pilot (final pilot)																									🟪						
Cluster2-Pilot4 Third pilot (final pilot)																									🟪						
Cluster2-Pilot5 Third pilot (final pilot)																												🟪			
Cluster2-Pilot6 Third pilot (final pilot)																												🟪			

Figure 10: Cluster 2 Pilot Planning.

1.2.1.1 Pilot 1 (UC#2.1A)

Pilot Description:

The pilot emulated a modern, data-driven traffic management ecosystem operating under real-world municipal infrastructure constraints. The operational context consisted of a live urban mobility and planning smart city testbed setup at key entry points and critical routes within the city of Valencia, Spain. On the hardware and edge device tier, the pilot deployed five 5G-connected radar units, equipped with dedicated solar panels for standalone power to guarantee grid-independent execution. These units were deployed alongside pre-existing municipal traffic cameras and supplementary third-party open-data mobility sensors. On the software and systems tier, the architecture integrated the FIVE IoT Platform, the local UPV Orion Context Broker, and the Telefónica Smart City Platform. The primary users interacting with these platforms included Valencia city traffic management authorities, urban mobility operators, and transportation policy planners.

The core workflow operated as a continuous real-time data loop. The 5GRadar units continuously scanned active traffic lanes to detect passing vehicles. These telemetry payloads were transmitted over high-speed 5G networks directly to the central data management platform. Specialized algorithms processed the incoming traffic patterns in real time to estimate transit flow status and isolate traffic congestion points based on collected statistics and trends. Simultaneously, third-party open-data sensor streams were ingested through the overarching COP-PILOT infrastructure, enabling the coordinated deployment of adaptive mobility applications and dynamic rerouting rules. This process was driven by continuous real-time data payloads pushed from the 5GRadar edge installations or external sensors during standard traffic flow or peak travel hours.

Objectives:

The key goals of this pilot focused on validation, demonstration, and systemic improvement across both technical and operational domains. Technically, the pilot aimed to validate the physical field rollout, network link reliability, and hardware uptime of the 5G-connected radar units running autonomously on independent solar power arrays. It also aimed to demonstrate multi-sensor interoperability by successfully combining data from the 5GRadars and third-party open-data sensors. Furthermore, technical teams introduced and validated efficient traffic state classification algorithms that combined the inputs of the 5GRadars and available open data. The final operational objective was to alleviate municipal transit pain points by supplying authorities with automated trend logging and statistics to reduce traffic congestion, enhance road safety, inform urban planning, secure environmental benefits, and maximize public transport efficiency.

Roadmap:

The practical deployment and validation procedure moved through five core operational steps.

1. Addressed site selection and feasibility assessment, where technical leads determined the desired locations for monitoring and assessed the physical installation feasibility at each site. The specific number of radars utilized at each location was determined based on street characteristics and traffic conditions, after which units were deployed and validated in the field.

2. Operated in parallel with location scouting, during which FIVE worked on adapting the 5GRadar's edge firmware to meet the exact structural needs of the project's central data management platform while evaluating whether additional metrics should be included in the telemetry payloads.

3. Executed real-time data collection and processing, with radar units continuously scanning the target areas. The collected data was transmitted via 5G networks to the central urban mobility platform where advanced AI algorithms processed the data in real time to identify traffic patterns and active congestion points.

4. Focused on the deployment of mobility applications integrating third-party sensors, where open data from external sensors was integrated through the COP-PILOT platform to expand the reach and utility of the applications.

5. Established continuous monitoring and optimization, allowing the urban mobility platform to continuously refine its underlying models based on historical and real-time data, provide automated trend reports highlighting inefficiencies to authorities, and push ongoing system updates to ensure radar units adapt to evolving urban challenges.

Setup & Validation:

The underlying infrastructure setup was hosted entirely within the on-premise UPV Domain private testbed infrastructure. The computational infrastructure was structured across virtualized machine environments. The internal core framework was divided into three primary virtual machines: the Orion Context Broker VM for real-time message brokering, the SIF VM containing secure Cloud Ziti network identities, and the OpenSlice orchestration VM. This underlying compute capability was managed via a production Kubernetes cluster where a master control plane node dynamically scheduled and coordinated software execution threads over general cluster worker nodes. External gateway connections routed through a public IP infrastructure using port redirection mappings to safely bridge incoming device telemetry streams to internal nodes and expose configuration interfaces to web operators.

Figure 11: Installation of radarMovAvCid.

The technical setup on the hardware edge tier comprised five 5G-connected radar units mounted onto municipal traffic pillars and coupled with dedicated standalone solar panels to achieve full grid independence and operational power autonomy. The active worker environment contained dedicated virtual processing nodes provided by UPV, including an Application and Complex Event Processing worker node developed by Telefónica, and a YOLO-based Verificator worker node developed by Fivecomm. Visual tracking and historical data management were handled by Telefónica's Smart City Stack worker node to allow management teams to access live monitoring dashboards.

The structural ecosystem integrated specialized software applications, communication components, and middleware platforms to handle data loops from the radars to the operator client.

- Applications: Used real-time AI traffic pattern processing software and Telefónica's localized Smart City Stack to manage data visualization layers and historical traffic logs.

- Communication Infrastructure: Employed high-speed cellular network layers via Telefónica's Public 5G Network to maintain stable, real-time data links from the deployed radars.

- Middleware/Platforms: The core data middleware was a local Orion Context Broker for immediate context ingestion, which automatically synchronized data entities with a central FIWARE Scorpio Context Broker acting as the NGSI-LD compliant smart city data backbone.

- Monitoring or Analytics Tools: Deployed advanced AI algorithms for real-time traffic density estimation, lane detection accuracy tracking, and traffic jam pattern recognition.

Regarding integration points and data flows, telemetry parameters captured vehicle variables at the edge and transmitted them over 5G. Payloads were formatted using standard NGSI-LD JSON schemas into the UPV's local Orion Context Broker and synchronized with Scorpio. Through persistent Scorpio subscriptions, data was forwarded to FIWARE Draco, which coordinated structured historical data persistence into a PostgreSQL storage layer. Curated datasets were then safely published through an extended CKAN Open Data portal using RESTful APIs to facilitate secure third-party consumption and academic reuse.

Measurement and validation tools featured automated payload parsing utilities inside the Orion Context Broker to check data structure integrity and ingestion frequencies to evaluate system uptime and packet delivery logs. For accuracy cross-referencing, technical teams utilized a specialized YOLO-based verification tool to cross-examine the radar's vehicle counting precision and classification performance against baseline traffic camera video feeds.

Data Management Plan:

From a data management and ethics perspective, all telemetry data processed, routed, and stored during this trial strictly complied with GDPR and the project's overarching Data Management Plan. Ingested datasets consisted entirely of context-free, machine-generated variables, including timestamped vehicle type classifications for anonymous count totals, velocity markers, lane identification indices, device temperatures, and voltage logs. The technical workflow completely avoided the collection or tracking of personal data, license plates, facial images, or individual driver characteristics, successfully guaranteeing a completely anonymous testing environment.

Stakeholder Involvement:

Internal technical teams featured several distinct responsibilities: FIVE handled radar hardware provisioning; ETRA managed physical roadside deployment, and solar panel integration; Telefónica provided the public 5G networks and Smart City Stack; and internal technical teams from UPV handled virtual machine network routing, data integration layers, and orchestrator API bridges. On the external tier, the primary end-users and observers consisted of the Valencia City Traffic Management Authorities and municipal transportation planners, who actively supervised deployment compliance, monitored live traffic dashboards, and extracted automated long-term statistics to shape future urban logistics and transit regulations.

1.2.1.2 Pilot 3 (UC#2.2)

Pilot Description:

The pilot emulated a modern, data-driven waste management ecosystem operating within a smart city campus environment. The operational context consisted of a realistic, controlled smart city testbed established at the UPV Vera Campus in Valencia, Spain. On the hardware and edge device tier, the pilot deployment plan involved non-invasive fill-level IoT hardware sensors, specifically the NM-EM400-MUD model, integrated directly inside upgraded campus waste bins. On the software and systems tier, the architecture seamlessly integrated NES' IoT Platform, the local UPV Orion Context Broker, the Telefónica Smart City Platform, and the ETSI OpenSlice Service Orchestrator. The primary users interacting with these operational layers included UPV Environmental Unit operators and collection dispatchers who required dynamic tracking capabilities.

The core workflow operated as an automated, event-driven loop. Non-invasive edge sensors continuously tracked volume capacities inside general waste containers and pushed this data through local broker layers to the central city platform. When a dynamically configured volume threshold was breached, system's background monitoring logic autonomously signalled the OpenSlice orchestrator layer. OpenSlice instantly generated and deployed an ephemeral notification application pod inside the Kubernetes cluster infrastructure without any manual administrative oversight. This short-lived application established an outbound interface connection to deliver an instantaneous, geotagged operational routing alert straight to field technicians via a dedicated

Telegram messaging bot. To enable rapid data logging and fast live loop testing during early scenarios, a distinct configuration variation was introduced by compressing the data reporting windows and transmission frequencies of the physical hardware sensors.

Objectives:

The core goals established for this pilot centered on comprehensive validation and operational optimization across multiple technical and infrastructural domains. Technically, the pilot aimed to validate the physical field rollout and overall data integrity of the non-invasive fill-level IoT sensors deployed inside campus bins. It aimed to demonstrate cross-system interoperability by seamlessly routing payloads between NES' platform acting as the IoT agent, the on-premise Orion Context Broker, and Telefónica's central Smart City Platform. Furthermore, technical teams prioritized validating zero-touch edge provisioning by leveraging ETSI OpenSlice to automatically deploy, manage the lifecycle of, and terminate ephemeral services based on real-time IoT events. From an operational perspective, the primary objective was to replace traditional, static collection routines with a dynamic, demand-driven route model triggered by exact traced asset metrics. This logistical shift was designed to optimize truck deployments, reduce fuel consumption, and lower greenhouse gas footprints, directly supporting UPV's strategic target of achieving complete carbon neutrality by 2030.

Setup & Validation:

The underlying infrastructure setup was hosted entirely within the on-premise UPV Domain private testbed infrastructure. The computational environment was structured across isolated virtualized machine environments to segregate platform functions and maintain secure data flows. The internal core framework was divided into three primary virtual machines: the Orion Context Broker VM for real-time message brokering, the SIF VM containing secure Cloud Ziti network identities, and the OpenSlice orchestration VM. This underlying compute capability was managed via a production Kubernetes cluster where a master control plane node dynamically scheduled software execution threads over general cluster worker nodes. External gateway connections routed through a public IP infrastructure using port redirection mappings to safely bridge incoming device telemetry streams to internal nodes and expose configuration interfaces to web operators.

The technical setup on the hardware edge tier comprised the specialized NES non-invasive fill-level IoT hardware sensors integrated within upgraded campus waste bins. The active worker environment contained dedicated virtual processing nodes provided by UPV, including an Application and Complex Event Processing worker node developed by Telefónica, and a localized AI Agent worker node developed by Telefónica. Visual tracking and historical data management were handled by Telefónica's Smart City Stack worker node to allow management teams to access live monitoring dashboards.

Figure 12: NES' sensors deployed inside waste containers.

The structural ecosystem integrated specialized software applications, communication components, and middleware platforms to handle data loops from the bin to the operator client.

- Applications: Utilized NES utility for initial hardware log management and calibration, and Telefónica's localized Smart City Stack to manage data visualization layers and historical container tracking.
- Communication Infrastructure: Employed an external commercial LTE and NB-IoT cellular network layer managed by Telefónica to maintain low-power data links from the containers, while high-speed internal virtual networks managed routing between local cluster virtual machines.
- Middleware/Platforms: Core data middleware utilized the FIWARE Orion Context Broker as the standardized NGSI-LD translation layer, while ETSI OpenSlice functioned as the primary service orchestrator to handle system-to-system automation.
- Monitoring or Analytics Tools: Deployed native Kubernetes command-line tools to audit cluster states, while operational notification loops used the external Telegram Bot API framework.

Regarding integration points and data flows, telemetry parameters sampled fill levels at the physical container and transmitted them over the NB-IoT network to NES, which ensured data quality and adapted the payloads to a smart data model before passing them via an automated webhook to the public proxy endpoint, routing it directly into the internal OrionCB VM. The structured JSON updates were then forwarded to the Smart City Stack to refresh map interfaces. Upon a threshold breach, the monitoring logic fired an API command to the OpenSlice VM, which automatically instructed the Kubernetes master node to instantiate an ephemeral notification pod on a cluster worker node. This temporary pod executed an outbound REST call to the Telegram Bot API, delivering a geo-routed alert message straight to the field operator's mobile device detailing exact coordinates, fill capacity, sensor distance, device temperature, battery margins, and tilt orientations.

Measurement and validation tools featured automated payload parsing utilities inside the Orion Context Broker to check data structure integrity and ingestion frequencies to evaluate system uptime and protocol compatibility metrics. The Kubernetes command-line interface (*kubectl*) served as the primary diagnostic tool for technical teams to inspect active pod states, parse container logs, and calculate the exact latency of zero-touch service provisioning following an event trigger. End-to-end telemetry receipt was verified using mobile interface logs from the Telegram alert channel to check delivery success, while qualitative assessment metrics were collected using digital operator feedback forms to measure user satisfaction rates.

Data Management Plan:

From a data management and ethics perspective, all telemetry data processed, routed, and stored during this trial strictly complied with GDPR and the project's overarching Data Management Plan. Ingested datasets consisted strictly of non-personal, machine-generated variables, including volumetric capacity percentages, raw sensor distances, internal hardware battery voltages, container tilt orientations, and precise geographic coordinates. The technical workflow completely avoided the collection, tracking, or storage of volunteer profiles or personal user characteristics, successfully guaranteeing a completely anonymous and non-invasive testing environment throughout the entire execution.

Stakeholder Involvement:

The pilot coordinated structural contributions from three core stakeholder groups executing a unified evaluation matrix. Internal technical teams from NES, UPV, and Telefónica handled hardware configurations, virtual machine network routing, and orchestrator API bridges to ensure technical baseline alignment. On the operational and user tier, campus facilities interacted with the system by monitoring active container entities to schedule dynamic, physical asset collections based on real-time needs. Finally, external observers representing the UPV Environmental Unit supervised active demonstrations, verified deployment compliance, and evaluated overall operational utility to shape future campus sustainability models.

1.2.2 Pilot Execution and Results

1.2.2.1 Pilot 1 (UC#2.1A)

This scenario evaluated the continuous live execution, edge data collection, and integration of five 5G-connected radar units co-located with existing municipal traffic camera structures. Due to physical roadway infrastructure limits and placement constraints, a best-effort deployment layout was executed across five junctions featuring non-ideal geometric configurations. This included adapting a single radar node (*radMovSerreria*) to track a non-standard, wide three-lane span instead of a dual-lane layout and installing the remaining four edge devices on narrow inferior street passages characterized by distinct structural slopes.

Showcasing:

The live system execution, vehicle streams, and data integrations were formally demonstrated at the Urban Mobility Data Hub (UMDH) event on May 26, 2026, at the ADEIT facility in the City of Valencia to a broad group of regional public and private innovation stakeholders.

Observed System Behaviour:

During live tracking sequences, the physical 5GRadar installations successfully maintained continuous edge detection and transmission operations 24/7 across all hours of the day. The underlying 5G network topology remained fully stable, sustaining real-time data links without any data loss, communication dropouts, or forced fault injections. Despite non-perpendicular grade conditions, the tracking systems cleanly logged vehicle profiles, speed variations, lane alignments, and directional distributions. Traffic movement patterns, localized density spikes, and slow-moving congestion strings over the sloped passages were clearly displayed on the central smart city monitoring dashboard to supply real-time operational statistics.

KPIs & KVIs:

Sub-scenario	KPI name	Objective	Initial Measurement (M17)
Real-Time Traffic Monitoring	KPI-VSS-CL2-1. Traffic State Classification Accuracy	Improvement in urban traffic control and traffic flow inside the target zones	Not achieved. Current accuracy reached 80% against the 95% target.
Real-Time Traffic Monitoring	KPI-VSS-CL2-6. System uptime	Measures the reliability of the system to guarantee continuous operation with minimal downtime. Captured as a percentage (%).	Achieved. Maintained complete network availability and stable connections across topology without data loss or downtime during execution.
Real-Time Traffic Monitoring	KPI-VSS-CL2-14. Traffic-jam Detection False Positive Rate	Measures the percentages of traffic jam notifications obtained from 5GRadars that are false positive. Captured as a percentage (%)	Not achieved. Current false positive rate stands at 20% against the 5% target.

1.2.2.2 Pilot 3 (UC#2.2)

This scenario evaluated an end-to-end automated waste tracking loop by monitoring real-world accumulation behaviour inside a specific target container (*urn:ngsi-ld:Waste Container:67-49-e4-49-70-22-00-28*). Database logs initially established a stable resting baseline, showing that the container was empty at 42% capacity. Firstly, the background tracking service was initialized and dynamically provisioned through the OpenSlice catalog. Field operators simulated an accumulation surge by physically throwing multiple large waste bags directly into the container.

Figure 13: Waste injection in waste container to test container fill-up.

The integrated hardware registered the volume change, updating its internal status to 61.5% capacity (rendered as 62% on the localized map interface), successfully breaching the custom-configured 50% alert threshold. The active monitoring service identified the violation and sent a zero-touch programmatic call to the orchestrator layer. OpenSlice automatically moved the service order to an acknowledged state and instantly instantiated a short-lived Waste Collection Alert Service application pod on a Kubernetes worker node without manual administrative intervention. This ephemeral pod executed an outbound REST call to the Telegram Bot API, delivering an instantaneous operational dispatch warning straight to the field operator's mobile device. The transmitted text layout detailed critical parameters, including 61.5% fill capacity, a 578 mm sensor distance, a 70% remaining battery margin, and normal tilt orientation. Crucially, it embedded a direct geo-routing hyperlink using the container's exact physical coordinates to guide maintenance trucks.

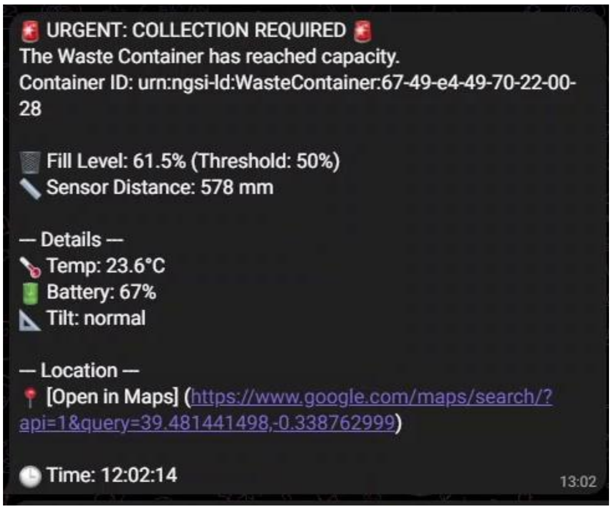

```
🚨 URGENT: COLLECTION REQUIRED 🚨
The Waste Container has reached capacity.
Container ID: urn:ngsi-ld:WasteContainer:67-49-e4-49-70-22-00-
28

🗑 Fill Level: 61.5% (Threshold: 50%)
📏 Sensor Distance: 578 mm

— Details —
🌡 Temp: 23.6°C
🔋 Battery: 67%
📐 Tilt: normal

— Location —
📍 [Open in Maps] (https://www.google.com/maps/search/?
api=1&query=39.481441498,-0.338762999)

🕒 Time: 12:02:14
```

Figure 14: Alert notification sent to Telegram Bot with objective container information.

Upon verifying message delivery, the temporary alert pod executed its cleanup routing scripts and automatically vanished from the cluster namespace to prevent computing resource bloat, while the background monitoring thread remained online.

Showcasing:

The technical workflow was formally demonstrated at an on-site event inside the private UPV campus testbed infrastructure. Campus logistics coordinators and environmental sustainability officers reviewed the live execution run, covering real-time container mapping, OpenSlice portal lifecycle ordering, physical bag dumping, pod generation logs, and final Telegram alert verification. Key partners from Nespra, UPV, and technical leads from the university's environmental management branch attended.

Observed System Behaviour:

During the live trial run, the system demonstrated highly responsive event-driven behaviour. The hardware edge sensors instantly captured volumetric capacity shifts and dispatched telemetry packages without data loss. The orchestration middleware handled lifecycle state transitions smoothly, deploying containerized pods automatically on the Kubernetes cluster within expected provisioning latency bounds. The communication infrastructure securely routed outbound alerts to the external API, while the cluster control plane effectively executed cleanup routines to release computing cores back to the common pool once the message loop closed.

KPIs & KVIs:

Sub-scenario	KPI name	Objective	Result
Automated Event-Driven Waste Alerting	KPI-VSS-CL2-6. System uptime	Measures the reliability of the system to guarantee continuous operation with minimal downtime. Captured as a percentage (%).	Achieved. Maintained complete network availability and stable connections across topology without data loss or downtime during execution.
Automated Event-Driven Waste Alerting	KPI-VSS-CL2-5. Waste Threshold Alerts Notified	Percentage of waste container capacity exceedances that successfully trigger and deliver an automated alert to operators via the COP-PILOT Multi-Modal Notification Service.	Achieved. Initial iterations successfully delivered 100% of triggered alerts. The underlying workflow is currently undergoing architectural redesign and process modifications.

Initial Impact Assessment – Cluster 2

Cluster 2 focuses on the validation of edge-enabled solutions for intelligent urban and environmental management through the Real-Time Traffic Monitoring and Automated Event-Driven Waste Alerting pilots. As highlighted in Section 3.0, the cluster exhibits a sustainability-oriented value profile, reflecting its objective of improving operational efficiency, environmental responsiveness, and data-driven decision-making within urban and industrial environments. Tables 3-10 and 3-11 of the main D5.1 report present the principal KPI-KVI relationships and initial validation findings from the first validation cycle (M17).

The business dimension is primarily represented through the KVIs of Validation Performance and Market Readiness. The results indicate that both pilot scenarios have successfully established stable operational environments, with full system availability and successful deployment of the required sensing and communication infrastructure. This is significant because reliability and operational continuity are prerequisites for wider adoption and future commercialisation. While market readiness is currently assessed at 50%, the findings suggest that Cluster 2 has progressed beyond proof-of-concept activities and is entering a phase where the focus shifts from demonstrating technical feasibility towards demonstrating operational value. The results therefore indicate a positive trajectory towards broader deployment, although continued optimisation and stakeholder engagement will be required before full operational maturity is achieved.

The societal dimension is represented through the KVI of User Satisfaction. A notable finding is the difference observed between the two pilot scenarios. The Real-Time Traffic Monitoring pilot achieved 100% user satisfaction, suggesting that stakeholders perceive clear operational benefits from improved visibility of traffic conditions and enhanced situational awareness. In contrast, the Automated Event-Driven Waste Alerting pilot recorded lower satisfaction despite strong technical performance and successful alert delivery. This highlights an important insight emerging from the initial assessment: value realisation is influenced not only by technical functionality but also by how effectively solutions integrate into existing operational processes and stakeholder workflows. The findings therefore suggest that user adoption and organisational integration may become increasingly important determinants of impact as the pilots mature.

The sustainability dimension is represented through KVIs relating to Data Accuracy and operational effectiveness. Across both pilots, the results demonstrate that continuous monitoring, automated event detection, and real-time notification capabilities are functioning as intended. Although several performance indicators have not yet achieved their final targets, the pilots have established the operational foundations required to support more responsive traffic management and waste monitoring processes. From an impact perspective, these capabilities provide important value pathways towards improved resource utilisation, reduced operational inefficiencies, and more sustainable management practices. While the wider environmental benefits cannot yet be quantified at M17, the current findings indicate that the necessary technological and operational mechanisms are now in place to support future sustainability outcomes.

1.3 CLUSTER 3A: AGRITECH TRANSFORMATION AND SUSTAINABILITY INITIATIVE

1.3.1 Pilot Scope and Planning

The pilot Planning of Cluster 3A was structured as a progressive validation roadmap, moving from individual component and infrastructure validation through to integrated, multi-use-case demonstrations conducted under operational field conditions. The primary objective of Stage 1 was to establish that the four ATSI use cases, precision agriculture and crop monitoring, AI-enabled AgriRobotics, secure data management and traceability, and smart logistics optimisation, can operate as a coherent, platform-orchestrated system within a real contract farming environment. A secondary objective is to collect the first evidence base for the cluster KPIs, covering treatment frequency reduction, traceability coverage, logistics optimisation, and platform-level security and automation indicators. Stage 1 spanned the first operational spinach cultivation season at the Kilkis testbed, conducted across 3 pilot and 3 control parcels within the Barba Stathis contract farming network.

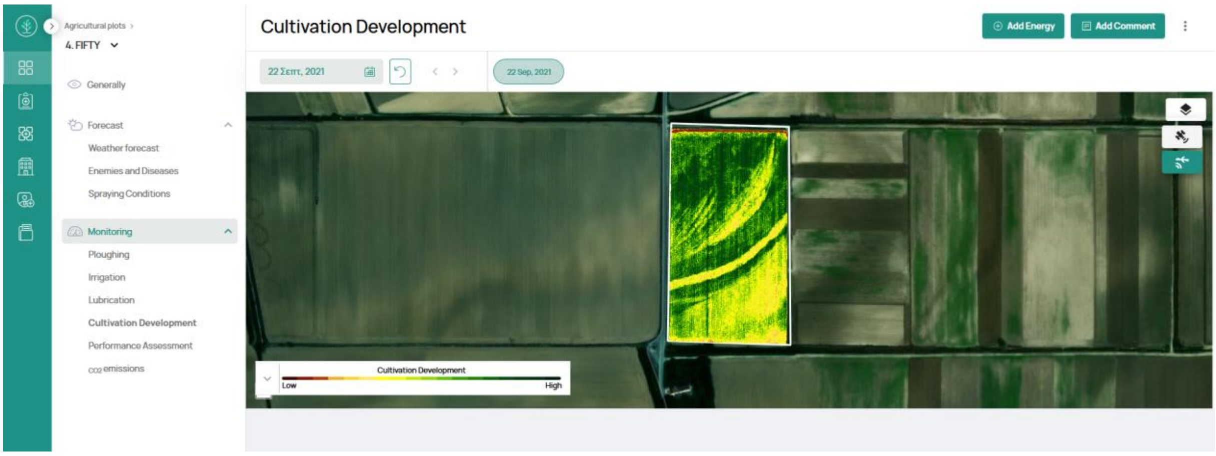

Figure 15: AgroApps 360 FMIS (drone data ingestion).

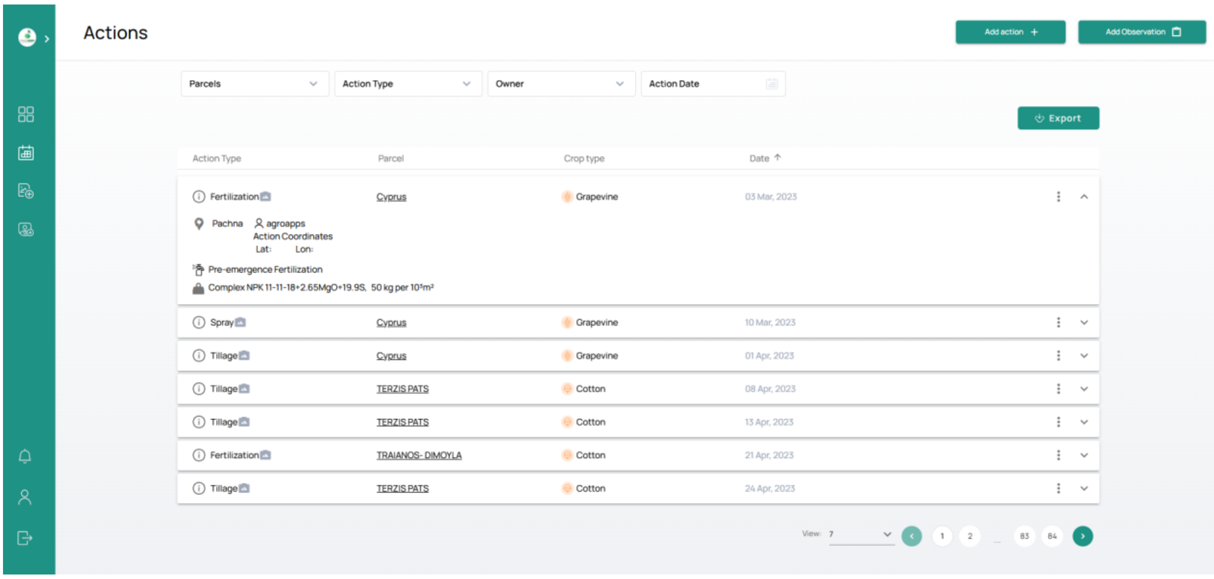

Action Type	Parcel	Crop type	Date ↑
Fertilization	Cyprus	Grapevine	03 Mar, 2023
Spray	Cyprus	Grapevine	10 Mar, 2023
Tillage	Cyprus	Grapevine	01 Apr, 2023
Tillage	TERZIS PATS	Cotton	08 Apr, 2023
Tillage	TERZIS PATS	Cotton	13 Apr, 2023
Fertilization	TRAIANOS- DIMOYLA	Cotton	21 Apr, 2023
Tillage	TERZIS PATS	Cotton	24 Apr, 2023

Pachna agroapps
Action Coordinates
Lat: Lon:
Pre-emergence Fertilization
Complex NPK 11-11-18+2.65MgO+19.9S, 50 kg per $10^3 m^2$

Figure 16: AgroApps 360 registering cultivation activities per parcel.

Figure 17: Plant Wearable sensors deployed.

Figure 18: Drone and Weather Stations provided by Barba Stathis.

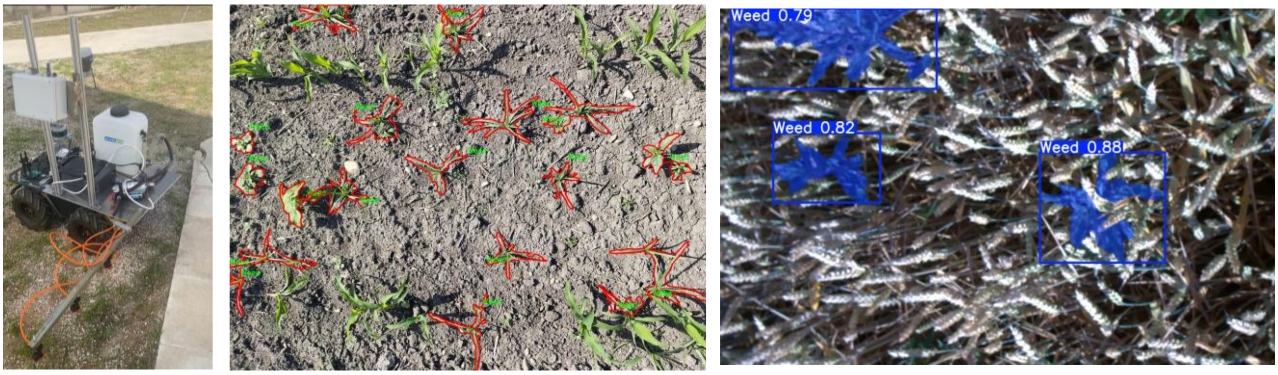

Figure 19: UGV and weed detection.

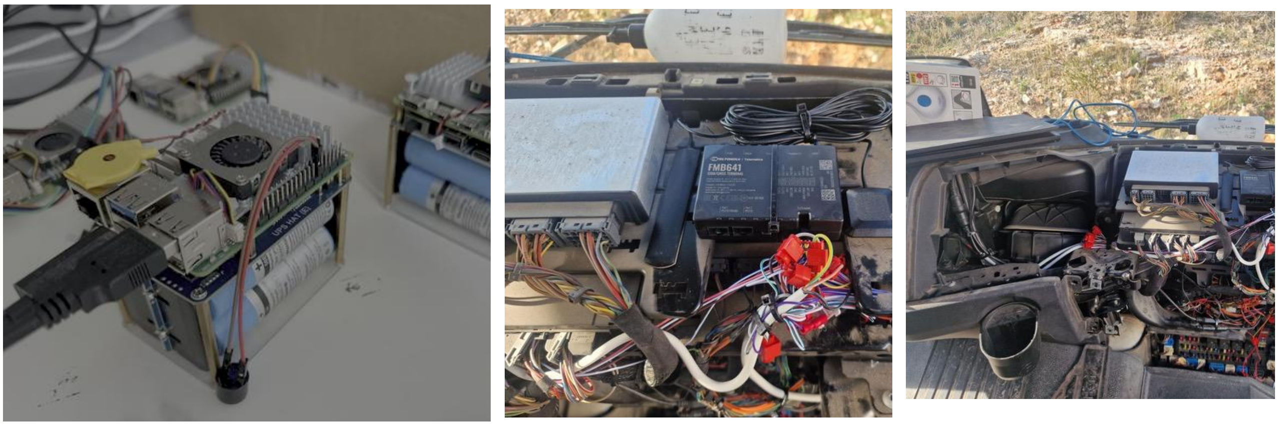

Figure 20: Telematic devices preparation and installation on fleet.

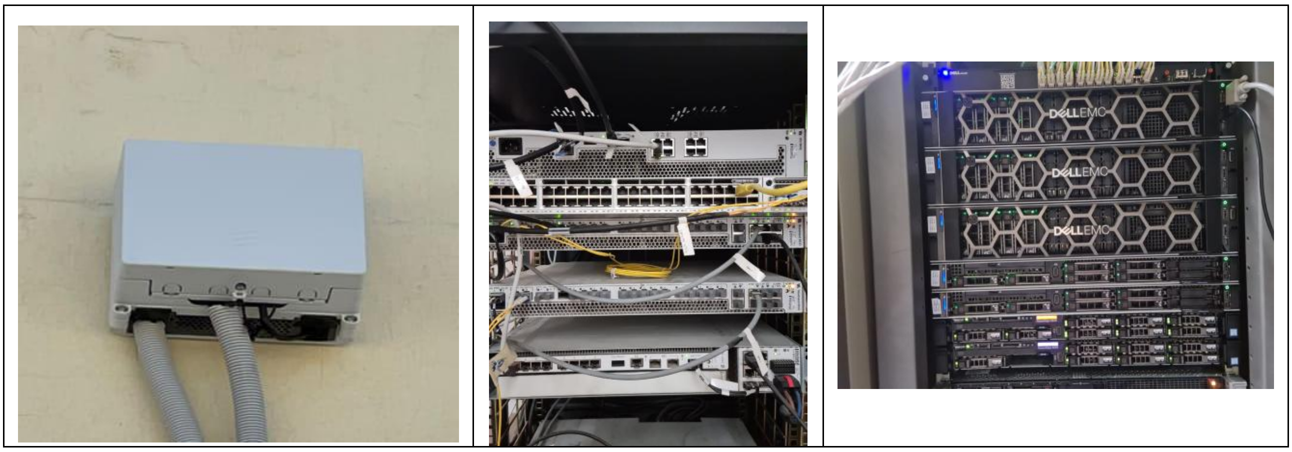

Figure 21: OTE deployed 5GMPN components.

The pilot roadmap of Stage 1 was organised into three sequential phases. The first phase covered infrastructure preparation and integration setup: the OTE private 5G Mobile Private Network was commissioned at the Kilkis testbed, Domain Orchestration (DO) and Data Management (DM) instances were deployed, all domain services were integrated with the COP-PILOT Secure Integration Fabric (SIF), and individual component validation was carried out across all partner systems prior to integrated operation.

The second phase, designated DEM #1, involved live integrated field operations: real-time data ingestion from field sensors and UAV surveys, AI-driven crop monitoring and weed infestation detection, UGV dispatch and targeted intervention, blockchain-attested activity recording, and PowerFleet-coordinated harvest logistics. Field Day 1, a public demonstration of the complete integrated workflow, was conducted during this phase and successfully demonstrated the end-to-end operational pipeline.

coppilot

Co-funded by the European Union

Aerial mapping

Inspect & Execute

Monitoring & Auditing

Human in the loop

Precision Application

Cluster 3A team

Cluster 3A team

FIELD DAY

CLUSTER 3A
AGRITECH TRANSFORMATION AND SUSTAINABILITY INITIATIVE (ATSI)

AgroApps

AGRICULTURAL UNIVERSITY OF ATHENS

Barba Stathis

Figure 22: Cluster 3A Field Day 1 - Open demo day.

The third phase, designated DEM #2, will validate the learning effects and performance improvements achieved between the first and second demonstration campaigns and will consolidate the evidence base for KPI measurement ahead of the M18 review. Field Days 2 and 3 and Demo Day 1, a full-ecosystem public demonstration, will be conducted within this phase.

Stage 1 exercised a broad set of COP-PILOT platform capabilities across all cluster domains. The Secure Integration Fabric was configured to establish encrypted, authenticated connectivity across the five fixed sites and the moving-vehicle domain, enabling secure data exchange between field assets at Kilkis, application infrastructure in Thessaloniki and Athens, and the Barba Stathis Sindos processing facility. UBITECH's ServOrch and InfraOrch components were deployed to provide service and infrastructure orchestration, binding the distributed partner services into a coherent managed environment. Domain Orchestration and Data Management instances were deployed per site, and the FIWARE NGSI-LD context management layer was established as the standardised IoT data federation interface through which field sensor streams, UAV detection outputs, and FMIS records are exchanged across domains. A custom integration layer was developed and deployed to bridge the Hyperledger Fabric permissioned blockchain with the FIWARE NGSI-LD infrastructure, an architectural component specific to Cluster 3A that extends the COP-PILOT interoperability framework to support blockchain-based supply chain traceability.

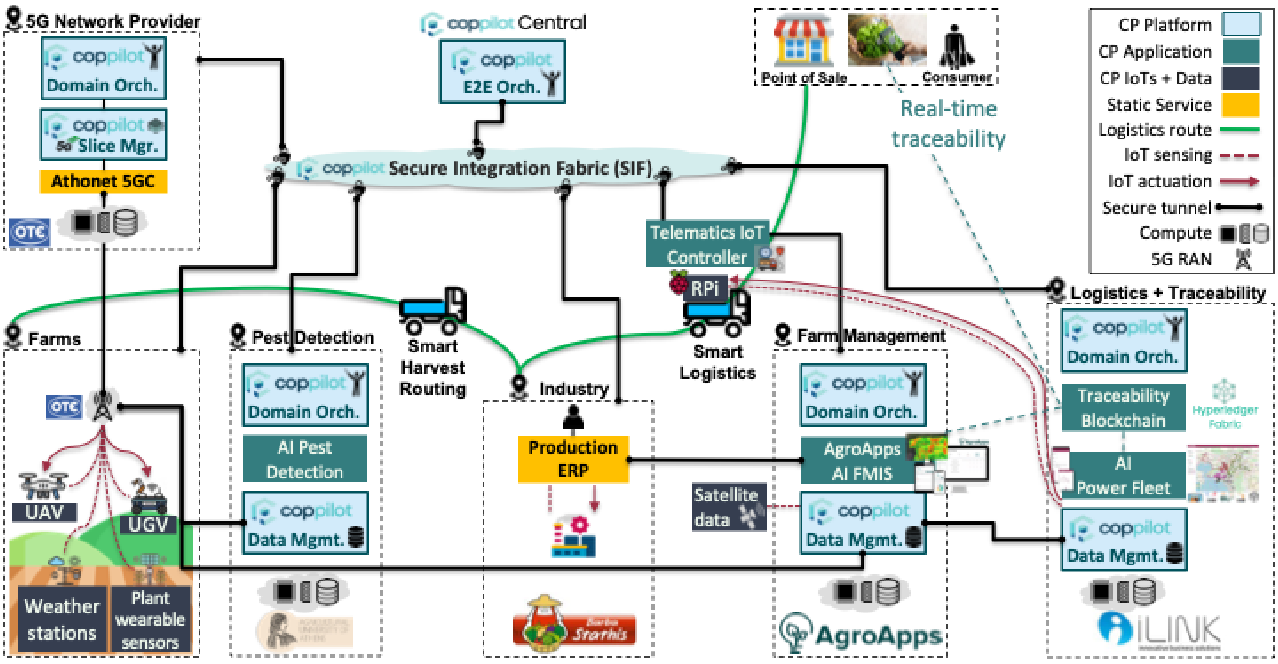

Figure 23: Cluster 3A architecture.

The cluster infrastructure was deployed across five fixed sites and 2 moving-vehicle domains. At the Kilkis farm testbed, the field-level sensing and actuation environment was established comprising the OTE private 5G MPN as primary connectivity, drone for canopy surveys and weed detection, an AUA-developed Unmanned Ground Vehicle for ground-level inspection and targeted intervention, TOR plant wearable sensors for continuous antinutrient proxy monitoring deployed across spinach parcels, and IoT weather stations providing localised environmental data. Satellite imagery was integrated at the data management layer to complement UAV survey data with wide-area crop status and growth stage information. AgroApps 360 was deployed on Kubernetes infrastructure in Thessaloniki and operates as the cluster FMIS hub; AUA's AI pest detection and weed identification inference stack was deployed on Kubernetes in Athens; and iLink's Hyperledger Fabric blockchain node and PowerFleet logistics platform were deployed from Athens. The Barba Stathis Sindos processing facility in Thessaloniki was configured to receive blockchain-attested lot records and PowerFleet logistics routing outputs. OTE's 5G core and radio access network was deployed across the Kilkis and Athens sites, with public 5G coverage providing the resilience failsafe layer.

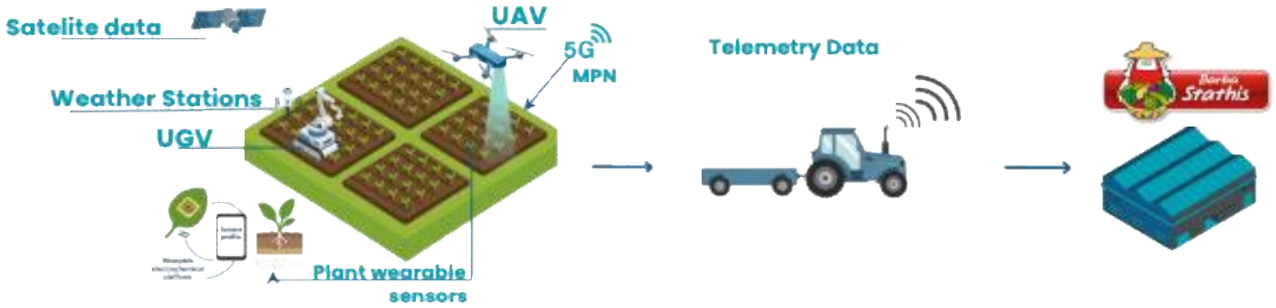

Figure 24: CL3A demonstration 1.

From a technical setup perspective, Stage 1 required the validated integration of seven partner systems across five geographic domains. The primary integration dependencies addressed were: the AUA AI inference pipeline configured to receive UAV imagery via the FIWARE data management layer and return structured detection outputs to AgroApps 360; the UGV communication and dispatch interface validated over the OTE 5G MPN; the TOR wearable sensor data stream integrated with the COP-PILOT DM instance at the Kilkis edge; the Hyperledger Fabric to FIWARE NGSI-LD bridge developed and tested to enable cultivation events and inspection records to be simultaneously available in the IoT data federation layer and the blockchain ledger; and the PowerFleet integration configured to receive crop readiness signals from AgroApps 360 and return dispatch plan confirmations. The multi-tenancy and service isolation requirements of the COP-PILOT platform were exercised across all sites, with particular emphasis on the Kilkis-to-Athens data path for AI inference and the Athens-to-Sindos supply chain data exchange.

Measurement and validation in Stage 1 follows the framework defined in D0.1, COP-PILOT's Unified Methodology for Project Validation and Evaluation. Initial KPI measurement is currently underway as part of DEM #1 operations; consolidated KPI evidence will be collected and reported at the conclusion of DEM #2. Stakeholder feedback is collected through structured surveys and interviews conducted during Field Day events and the Demo Day events, covering perceived usability, relevance, and adoption intent across farmer, agronomist, food processor, and logistics operator stakeholder groups. All pilot activities are documented using the COP-PILOT Pilot Report Template, with individual pilot reports.

The stakeholder involvement plan for CL3A is structured around two primary groups. The operational stakeholders: Barba Stathis agronomists, contract farmers, and logistics coordinators, all are engaged throughout the pilot period as active users of the deployed platform components: agronomists register parcels and interact with AgroApps 360 advisory outputs; farmers record cultivation activities in real time; and logistics coordinators interact with PowerFleet dispatch plans. Their day-to-day engagement provides continuous usability and relevance feedback outside of formal demonstration events. The validation stakeholders, including representatives from agricultural advisory organisations, food certification bodies, retailers, and consumer-facing platforms, are engaged through the Field Day and Demo Day events, where the integrated system is and will be further presented in operational conditions and structured feedback is collected. The involvement of Barba Stathis as both an operational and validation stakeholder is central to CL3A pilot, as the organisation simultaneously provides the testbed environment and constitutes the primary commercial target for the cluster's integrated service offering.

Data management and ethical considerations for Cluster 3A are addressed at both platform design and operational levels. AgroApps 360 was designed on privacy-by-design principles in compliance with GDPR: cultivation records, parcel location data, and farmer activity logs are processed under defined data processing agreements that establish the respective roles of the farmer, Barba Stathis as contract farming operator, and AgroApps as platform provider. Farmers retain ownership of their farm-level data, with access rights structured to ensure that aggregated network-level outputs are available to Barba Stathis agronomists while individual farm records remain under farmer-controlled permissions. The EU Data Act (Regulation 2023/2854) is directly relevant to the IoT sensor and robotics components of the cluster; the platform architecture was designed to ensure that machine-generated data from UAV operations, UGV field activities, and field IoT sensors is handled within the data sharing and portability framework the Data Act establishes. The UGV's mandatory human confirmation requirement before any spray intervention is issued reflects both the current regulatory framework for autonomous agricultural operations and the human-in-the-loop design principle applicable to AI-assisted advisory systems under the EU AI Act. All data management activities are conducted in accordance with the COP-PILOT project Data Management Plan.

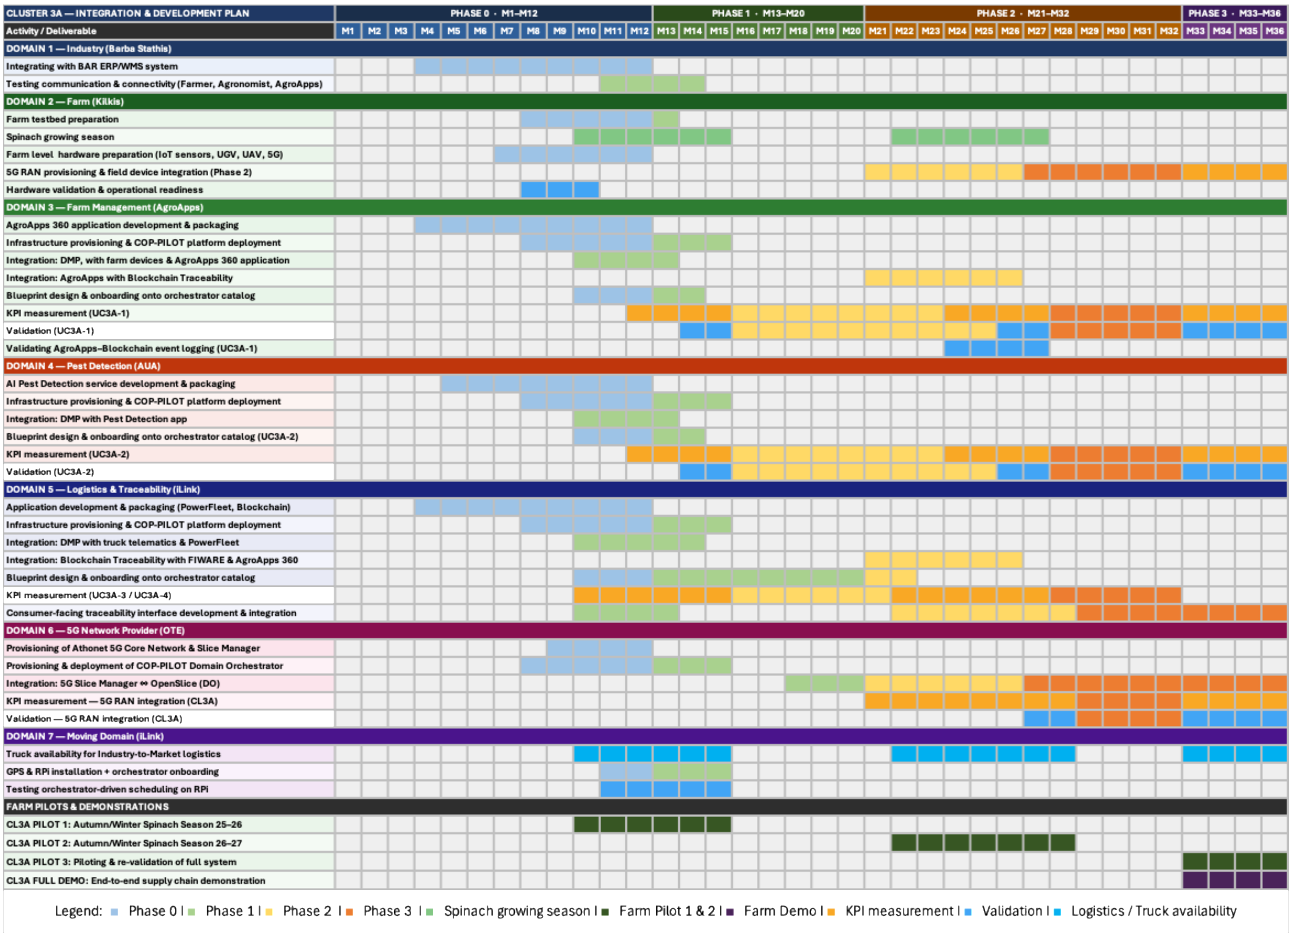
CLUSTER 3A — INTEGRATION & DEVELOPMENT PLAN

PHASE 0 · M1–M12 | PHASE 1 · M13–M20 | PHASE 2 · M21–M32 | PHASE 3 · M33–M36

Activity / Deliverable | M1 … M36

DOMAIN 1 — Industry (Barba Stathis)
- Integrating with BAR ERP/WMS system
- Testing communication & connectivity (Farmer, Agronomist, AgroApps)

DOMAIN 2 — Farm (Kilkis)
- Farm testbed preparation
- Spinach growing season
- Farm level hardware preparation (IoT sensors, UGV, UAV, 5G)
- 5G RAN provisioning & field device integration (Phase 2)
- Hardware validation & operational readiness

DOMAIN 3 — Farm Management (AgroApps)
- AgroApps 360 application development & packaging
- Infrastructure provisioning & COP-PILOT platform deployment
- Integration: DMP, with farm devices & AgroApps 360 application
- Integration: AgroApps with Blockchain Traceability
- Blueprint design & onboarding onto orchestrator catalog
- KPI measurement (UC3A-1)
- Validation (UC3A-1)
- Validating AgroApps–Blockchain event logging (UC3A-1)

DOMAIN 4 — Pest Detection (AUA)
- AI Pest Detection service development & packaging
- Infrastructure provisioning & COP-PILOT platform deployment
- Integration: DMP with Pest Detection app
- Blueprint design & onboarding onto orchestrator catalog (UC3A-2)
- KPI measurement (UC3A-2)
- Validation (UC3A-2)

DOMAIN 5 — Logistics & Traceability (iLink)
- Application development & packaging (PowerFleet, Blockchain)
- Infrastructure provisioning & COP-PILOT platform deployment
- Integration: DMP with truck telematics & PowerFleet
- Integration: Blockchain Traceability with FIWARE & AgroApps 360
- Blueprint design & onboarding onto orchestrator catalog
- KPI measurement (UC3A-3 / UC3A-4)
- Consumer-facing traceability interface development & integration

DOMAIN 6 — 5G Network Provider (OTE)
- Provisioning of Athonet 5G Core Network & Slice Manager
- Provisioning & deployment of COP-PILOT Domain Orchestrator
- Integration: 5G Slice Manager ⇔ OpenSlice (DO)
- KPI measurement — 5G RAN integration (CL3A)
- Validation — 5G RAN integration (CL3A)

DOMAIN 7 — Moving Domain (iLink)
- Truck availability for Industry-to-Market logistics
- GPS & RPi installation + orchestrator onboarding
- Testing orchestrator-driven scheduling on RPi

FARM PILOTS & DEMONSTRATIONS
- CL3A PILOT 1: Autumn/Winter Spinach Season 25–26
- CL3A PILOT 2: Autumn/Winter Spinach Season 26–27
- CL3A PILOT 3: Piloting & re-validation of full system
- CL3A FULL DEMO: End-to-end supply chain demonstration

Legend: ■ Phase 0 | ■ Phase 1 | ■ Phase 2 | ■ Phase 3 | ■ Spinach growing season | ■ Farm Pilot 1 & 2 | ■ Farm Demo | ■ KPI measurement | ■ Validation | ■ Logistics / Truck availability

1.3.2 Pilot Execution and Results

To deliver on Cluster 3A's objectives, reducing agrochemical inputs through precision crop monitoring, enabling targeted robotic intervention, establishing digital supply chain traceability, and optimising harvest and distribution logistics, pilot execution was organised around four validation scenarios, one per use case. Each scenario covers a defined segment of the agricultural value chain, from parcel registration and field sensing through to consumer-accessible product traceability. The following subsections describe each scenario, report on what has been validated to date, and set out what will be addressed in Phase 2 and Phase 3. All activities are conducted on the spinach pilots at the Kilkis testbed within the Barba Stathis contract farming network.

Scenario 1: Integrated Field Monitoring and Crop Management Cycle

The first scenario addresses Use case 3A1 objectives and covers the complete crop monitoring and management cycle across a spinach growing season. Starting from parcel registration and activation of satellite and weather data collection, it progresses through plant wearable sensor deployment for antinutrient monitoring, scheduled UAV surveys for weed detection and crop health assessment, data fusion in AgroApps 360, generation of zone-specific recommendations, and execution of targeted interventions. All events are logged in AgroApps 360 and recorded on the blockchain, with the scenario closing on a crop readiness notification that triggers harvest and logistics planning.

Parcel registration was completed with BAR agronomists and contracted farmers. TOR wearable sensors were deployed across the Kilkis pilot parcels and data collection is currently underway. UAV survey campaigns were conducted, and the full data path from UAV imagery to AgroApps 360 advisory output was validated. Farmer activity logging is operational, and the first complete spinach cultivation cycle was recorded digitally across the BAR network. In Phase 2, antinutrient modelling outputs will be integrated into the FMIS recommendation layer, and the automated harvest readiness notification to Barba Stathis will be fully validated alongside second-season KPI measurement. Phase 3 will demonstrate a fully optimised, multi-season advisory cycle with consolidated performance evidence.

Scenario 2: UAV-UGV Weed Detection and Precision Spraying

The second scenario addresses Use case 3A2 objectives and covers the weed management workflow from aerial detection to targeted ground intervention. UAV imagery is processed by AI models to classify weed, crop, and soil zones and assess weed cover percentage per field area. When infestation exceeds the predefined threshold, a UGV mission is dispatched to the affected locations. The UGV performs an on-site assessment using its onboard cameras and edge AI, transmits its findings for agronomist confirmation, and executes precision spot spraying on confirmed weed patches. All activity is logged in the AgroApps 360 and recorded on the blockchain.

This scenario was fully demonstrated at Field Day 1. All nine steps were executed in live field conditions at the Kilkis testbed: UAV detection, AI weed segmentation, UGV dispatch, ground-level assessment, agronomist confirmation, and targeted spot spraying were completed and logged, confirming that the UAV-to-AI-to-UGV pipeline operates as a coherent, platform-orchestrated workflow. In Phase 2, the focus will be on validating the learning effects between campaigns, including improvements in detection accuracy, reduced false positive rate, and extended coverage across a second growing season. Phase 3 will produce consolidated cross-season performance evidence and demonstrate the workflow at scale.

Scenario 3: Batch Traceability with SLA-triggered Orchestration

The third scenario addresses Use case 3A3 objectives and establishes the blockchain-based traceability layer covering the full production and logistics lifecycle. Each harvested batch is registered on the Hyperledger Fabric ledger with a cryptographic signature, timestamp, and identifier. Subsequent events, packaging, transport telemetry, and chain-of-custody handovers, are all recorded as immutable ledger entries. SLA deviations detected during transport trigger orchestrated corrective actions via PowerFleet. The scenario closes with a consumer-accessible QR interface exposing the complete lot history from farm parcel to retail shelf.

The blockchain node was deployed and the FIWARE NGSI-LD integration layer was validated during Phase 1. Batch registration and packaging metadata logging are operational and are continuously exercised in DEM #1; transport telemetry integration and SLA-triggered orchestration are currently being validated. In Phase 2, the consumer QR traceability interface will be fully validated and publicly demonstrated, completing the field-to-gate chain. Phase 3 will demonstrate the end-to-end consumer-facing traceability chain from spinach parcel to retail point of sale at full operational scale.

Scenario 4: Smart Logistics and Supply Chain Optimisation

The forth scenario addresses Use case 3A4 objectives and covers the complete logistics chain from harvest trigger to consumer traceability access. When crop readiness thresholds are met, AgroApps 360 emits an event to PowerFleet, which computes an optimised dispatch plan for the Barba Stathis harvester fleet. Inbound transport is monitored via vehicle IoT sensors feeding temperature, humidity, and GPS data to PowerFleet and the blockchain. The batch is processed and packaged at Sindos with a QR-linked blockchain record. PowerFleet after accepting incoming orders and fullfillemnt dates, orchestrates outbound delivery to wholesalers and retailers, optimising routes for freshness, distance, and service levels. Consumers scan QR codes at the point of sale to access the full lot traceability record.

Harvest dispatch triggering, PowerFleet plan generation, inbound routing coordination, and processing facility integration were validated during Phase 1 and DEM #1; onboard transport telemetry (temperature/GPS) is scheduled for the Phase 2 outbound segment. Inbound logistics coordination with the Barba Stathis harvester fleet was confirmed as operational. In Phase 2, outbound delivery optimisation to wholesalers and retailers, and the consumer QR traceability interface, will be validated and demonstrated. Phase 3 will deliver the fully integrated logistics scenario with real-time freshness-aware rerouting across the complete supply chain, from field to consumer.

Table below summarises the validation status of each blueprint scenario across the three piloting phases.

Scenario	Phase 1 / DEM #1 (Stage 1 — validated)	Phase 2 / DEM #2 (Stage 2 — to be validated)	Phase 3 (Stages 2–3 — to be validated)
Scenario 1 (UC3A.1)	Parcel registration, wearable deployment, UAV campaigns, AI crop monitoring, full first-season activity recording	Antinutrient modelling in FMIS, harvest readiness automation, second-season KPI measurement	Fully optimised multi-season advisory cycle, consolidated performance evidence
Scenario 2 (UC3A.2)	All nine steps demonstrated at Field Day 1: UAV detection, AI segmentation, UGV dispatch, ground assessment, agronomist confirmation, precision spot spraying	Detection accuracy improvement, reduced false positive rate, second-season coverage	Cross-season performance evidence, full-scale demonstration
Scenario 3 (UC3A.3)	Blockchain node deployed, FIWARE bridge validated, batch registration and packaging logging operational; transport telemetry and SLA orchestration ongoing	Consumer QR traceability interface validated and publicly demonstrated	End-to-end consumer transparency from spinach parcel to retail shelf
Scenario 4 (UC3A.4)	Harvest dispatch triggered, PowerFleet plans generated, inbound dispatch and routing coordinated, processing facility integration confirmed	Outbound delivery to wholesalers and retailers, consumer QR scan at point of sale	Freshness-aware dynamic rerouting, full supply chain demonstration field to consumer

Table 7: Cluster 3A scenario validation status across piloting phases.

Initial Impact Assessment – Cluster 3A

Cluster 3A focuses on the validation of digital agriculture and AgriFood supply chain solutions aimed at improving farming efficiency, crop protection, environmental sustainability, and supply chain transparency. As highlighted in Section 3.0, the cluster exhibits a predominantly sustainability-oriented value profile, with approximately 60% of identified KVIs associated with sustainability outcomes, compared to 27% business and 13% societal outcomes. Table 3-12 of the main D5.1 report presents the principal KPI-KVI relationships and initial validation findings from the first validation cycle (M17).

A key theme emerging from the assessment is the development of an integrated digital agriculture ecosystem that combines AI-driven decision support, autonomous field operations, logistics monitoring, and supply chain traceability. Unlike traditional agricultural improvement initiatives that focus on isolated interventions, Cluster 3A seeks to connect multiple stages of the agricultural value chain into a unified data-driven ecosystem. The M17 findings indicate that many of the enabling technologies required to support this vision are already operational and generating data, providing a strong foundation for future impact measurement.

The **business** dimension is represented through the KVI of Agricultural Input Cost Savings for Contract Farmers. While quantitative evidence relating to yield improvement and resource optimisation is not yet available due to the need for full growing-season comparisons, the successful deployment of the monitoring and optimisation infrastructure represents an important milestone. Initial system deployment has been validated, creating the conditions required for future assessment of input efficiency, productivity improvements, and operational cost savings. The findings therefore suggest that the cluster has progressed beyond technology development and is now positioned to evaluate how digital farming practices translate into measurable economic/ business value for agricultural stakeholders.

The **societal** dimension is reflected through the KVI of *Shift from Reactive to Predictive Crop Protection*. The deployment of the AgroApps360 AI alert system, UAV-based crop monitoring activities, agronomist scouting records, and initial UGV field validation collectively indicate a transition towards more proactive and data-driven crop management practices. Rather than relying solely on manual observation and reactive interventions, the pilot is establishing mechanisms that enable earlier detection of crop issues and more informed decision-making. Although the full effectiveness of these capabilities will require validation across complete growing cycles, the current results demonstrate meaningful progress towards the cluster's objective of supporting predictive agricultural management.

The **sustainability** dimension represents the strongest aspect of the Cluster 3A value framework and extends beyond environmental performance alone. The findings suggest that sustainability is being addressed at multiple levels of the agricultural ecosystem. At farm level, resource optimisation capabilities have been established through AI-supported monitoring and autonomous field operations. At supply chain level, telematics systems are actively collecting logistics data, while blockchain-enabled traceability mechanisms are already supporting pilot delivery cycles. Furthermore, the integration of multiple platforms and data domains across the COP-PILOT ecosystem demonstrates progress towards a more transparent and trustworthy AgriFood supply chain. Although quantitative evidence relating to resource savings, spoilage reduction, product freshness, and CO_2 reduction remains dependent on full-season data aggregation, the operational foundations required to support these outcomes are now largely in place.

1.4 CLUSTER 3E: EDGE INTELLIGENCE FOR ENHANCING GRID RELIABILITY

1.4.1 Pilot Scope and Planning

For each pilot executed during Stage 1, describe the scope, objectives, and relevance to the project goals. Identify the use case(s) addressed, the COP-PILOT platform features exercised, and the business/technical goals. Follow the structure from the Pilot Report Template (Sections 1.1–1.4). Describe the planning timeline, infrastructure setup (testbed details, COP-PILOT components deployed), technical setup, technical components involved, measurement and validation tools, stakeholder involvement plan, and DMP/ethics considerations. Follow the Pilot Report Template (Sections 2.1–2.6).

ID	UC#	Date	Partners	Stakeholder Type	Stakeholders examples
First Phase					
Cluster3E-Pilot-1	UC3	M14 – Feb 26	BPO, ENIC	Biogas Plant Owners, RES Owners	AegeanTerra, Quest Energy
Cluster3E-Pilot-2	UC2	M16 – Apr 26	PPC, ENIC	EV Charger Owners, Asset owners	Ares2t, E-GAP
Second Phase					
Cluster3E-Pilot-3	UC1	M20 – Sep 26	UOP, PNET, ENIC	DSOs, Energy Market Operators, Aggregators	HEDNO, IPTO, HENEX
Cluster3E-Pilot-4	Integrated Pilot	M23 – Nov 26	UOP, PPC, BPO, PNET, ENIC	All of the above	All of the above
Third Phase					
Cluster3E-Pilot-5	Final Integrated Pilot	M33 – Sep 27	UOP, PPC, BPO, PNET, ENIC	All of the above, OC partners	All of the above

Cluster3E-Pilot1 focused exclusively on UC#3E.3 and was executed on 12 March 2026 at the BPO biogas electricity co-production plant in Preveza, Western Greece. The pilot demonstrated end-to-end data flow from IoT sensors installed at the plant through the COP-PILOT infrastructure to a real-time operator dashboard, establishing the foundational data integration layer for the biogas domain.

The pilot is deployed at an industrial biogas facility managing approximately 100,000 tonnes of organic waste annually, producing biogas through a two-stage anaerobic digestion process fuelling two CHP units with a combined nominal capacity of 2 MW. Two instruments are connected via Modbus TCP: an AwiFLEX Gas Analyser monitoring CH_4, CO_2, and H_2S across primary and secondary tanks, and a Memosens CPS11E pH sensor monitoring the digester. Data flows from the sensors through an Ewon Flexy 205 industrial gateway to a Jetson Orin Nano edge server, where the FIWARE context broker handles ingestion and harmonisation before the application dashboard presents live and historical data to the plant operator. The main actors are BPO (plant owner, sensor deployment and demo coordination) and ENIC (technology provider, full platform and application deployment).

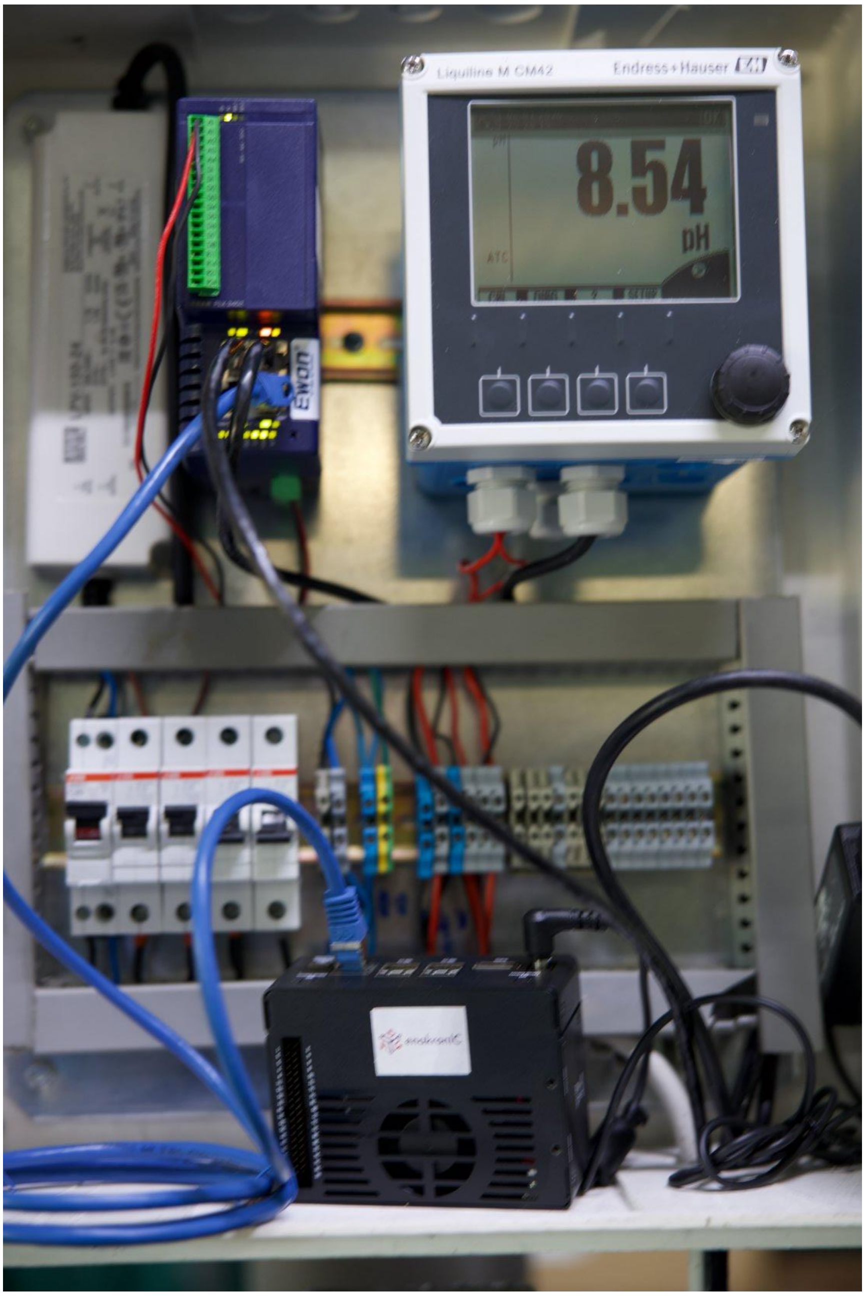

Technical objectives include demonstrating real-time data flow from IoT sensors to the operator dashboard (addressing FR3.1, FR3.3, FR3.6, FR3.7, NFR3.6), validating FIWARE-based data ingestion and harmonisation from Modbus TCP sensors, validating OpenSlice-based automated service deployment and SIF integration, and demonstrating both live monitoring and historical data storage and retrieval. Business objectives focus on demonstrating the value of continuous automated monitoring for biogas process management, and establishing the digital foundation for future development of anaerobic digestion forecasting and predictive maintenance services.

This pilot directly contributes to Cluster 3E's strategic ambition by transforming a manually monitored biogas plant into a smart, data-driven asset. Prior to the project, the plant relied on intermittent manual sampling, leaving it vulnerable to undetected anomalies and reactive maintenance. The pilot establishes the digital integration layer needed to evolve the facility towards predictive maintenance and electricity generation forecasting, making it a more reliable and intelligent contributor to the regional renewable energy mix.

1.4.2 Pilot Execution and Results

The scenario validated the full data pipeline from IoT sensors installed at the Preveza biogas plant through the COP-PILOT infrastructure to an operator-facing dashboard. Real-time and historical data from the AwiFLEX gas analyser (CH_4, CO_2, H_2S across primary and secondary tanks) and the Memosens CPS11E pH sensor were ingested via Modbus TCP, harmonised through the FIWARE context broker, and visualised on the dashboard application deployed via OpenSlice orchestration.

Execution followed a structured sequence: system setup and connectivity verification preceded the demonstration, including validation of VPN connectivity and certificate validity between the edge and central servers and SIF connectivity to the Athens infrastructure. The demonstration itself combined a BPO-led plant tour explaining the anaerobic digestion process and CHP operations with a live dashboard walkthrough, covering both real-time sensor data and historical data retrieval. A Q&A and stakeholder discussion session concluded the event. Both live monitoring and historical data retrieval were demonstrated successfully in the presence of 12 attendees including plant operators, external industry observers, and the internal technical team.

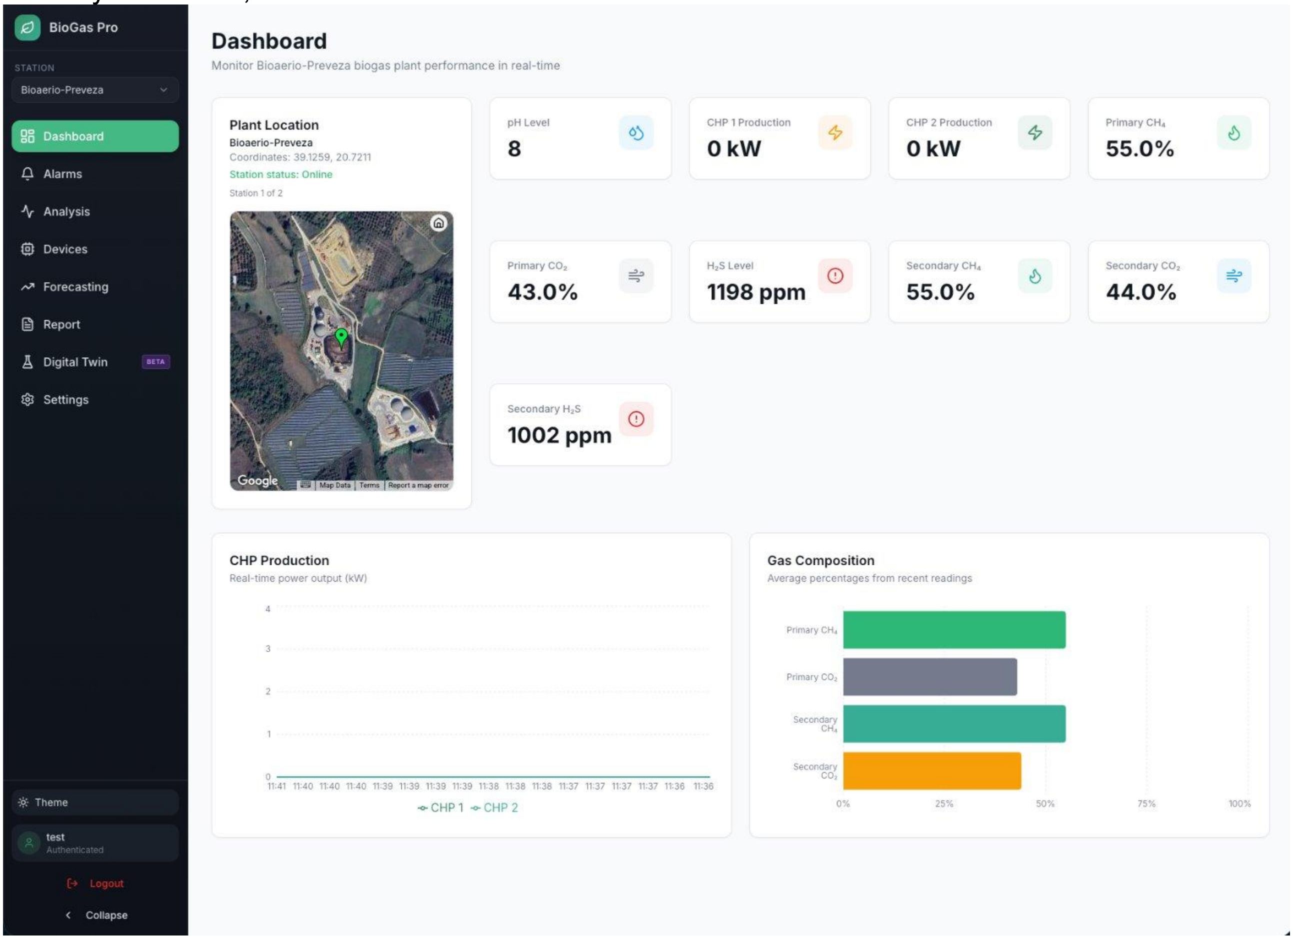

Dashboard

Monitor Bioaerio-Preveza biogas plant performance in real-time

Plant Location

Bioaerio-Preveza

Coordinates: 39.1259, 20.7211

Station status: Online

Station 1 of 2

pH Level 8

CHP 1 Production 0 kW

CHP 2 Production 0 kW

Primary CH_4 55.0%

Primary CO_2 43.0%

H_2S Level 1198 ppm

Secondary CH_4 55.0%

Secondary CO_2 44.0%

Secondary H_2S 1002 ppm

CHP Production

Real-time power output (kW)

Gas Composition

Average percentages from recent readings

The demonstration took place on 12 March 2026 at the biogas plant in Preveza and was attended by 12 people: 3 biogas plant owners and 2 engineers from their respective plants, 1 environmentalist, approximately 5 workers of the host plant, and 1 engineer from the construction company responsible for building and maintaining the plant. The event combined a physical plant tour with a live technical demonstration, giving attendees both operational context and a hands-on view of the COP-PILOT-enabled monitoring application.

Attendees responded positively to the demonstrated capabilities. The real-time dashboard and the ability to visualise both live and historical process data were well received, particularly by plant owners and technical staff. Two concrete extension proposals emerged from the discussion: integration with an energy aggregator to enable the plant to receive and respond to demand response signals, and incorporation of environmental curtailment monitoring to track and manage methane venting during peak generation hours.

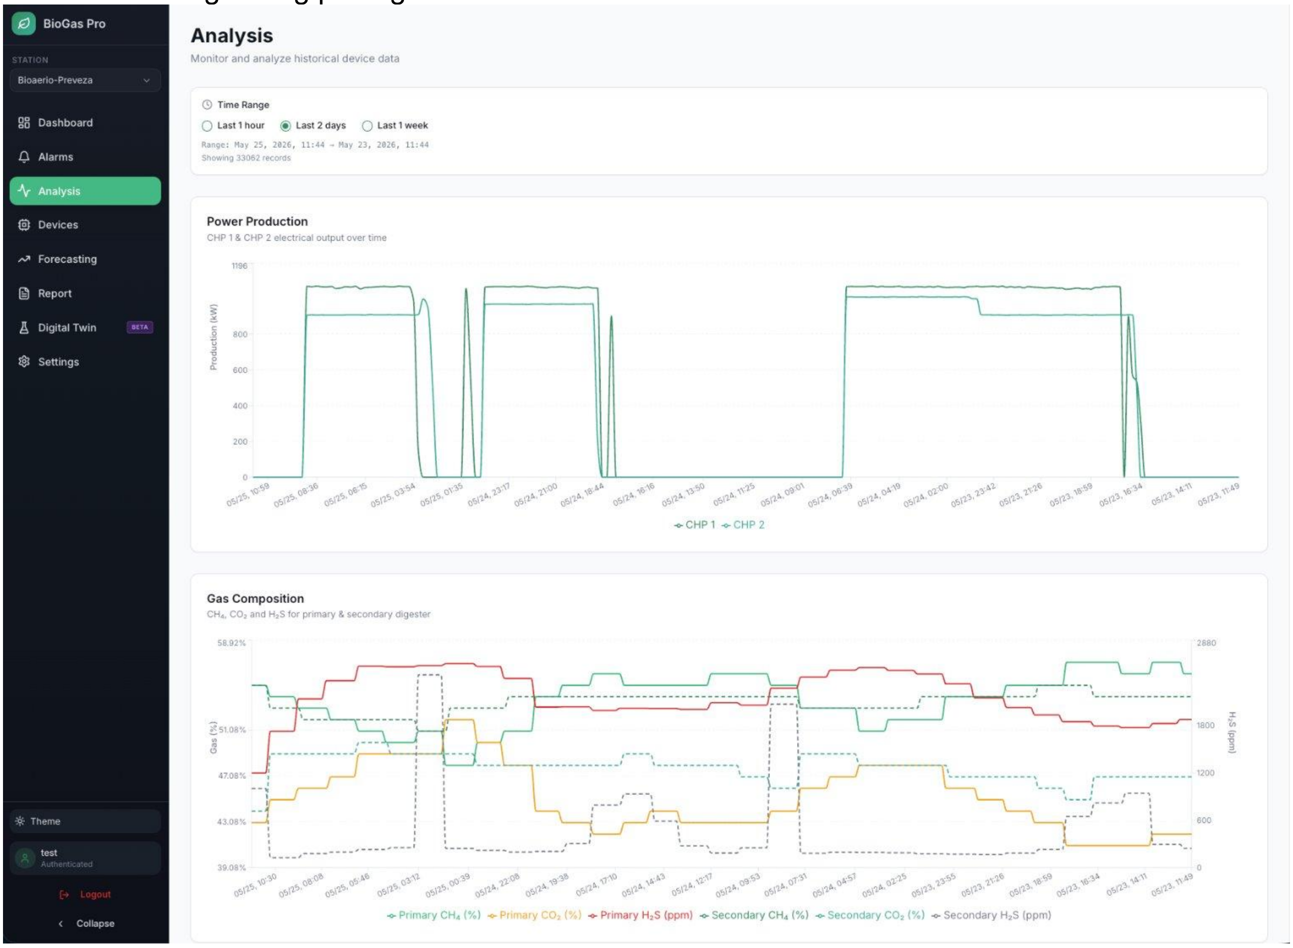

Range: May 25, 2026, 11:44 → May 23, 2026, 11:44
Showing 33062 records

Prior to the demonstration, expired VPN certificates disrupted connectivity between the edge and central servers. A fix was applied on-site and the demo proceeded without further issues. This highlighted the need for a pre-demo readiness checklist covering connectivity verification, certificate validity, and overall service health. Additionally, stakeholder feedback — particularly around demand response integration and curtailment monitoring — underlined the value of engaging domain experts and plant operators early, as their input provides operationally grounded directions for future application development.

Initial Impact Assessment – Cluster 3E

Cluster 3E focuses on creating value across the energy ecosystem through improved operational flexibility, renewable energy integration, and sustainable energy management. As presented in Section 3.0, the cluster exhibits a relatively balanced value profile, with sustainability representing the largest proportion of identified KVIs (44%), followed by business (33%) and societal outcomes (22%). This distribution reflects the pilot's objective of delivering operational, environmental, and community benefits through edge-enabled energy solutions that support the transition towards more resilient and sustainable energy systems.

Table 3-13 of the main D5.1 report presents the principal Key Value Indicators (KVIs), supporting Key Performance Indicators (KPIs), and initial validation findings for Cluster 3E. The results demonstrate that the pilot's value proposition extends beyond technical performance and encompasses broader objectives relating to energy flexibility, renewable energy integration, community benefits, workforce development, and long-term sustainability. Collectively, the findings suggest that the cluster is beginning to establish the operational foundations required to support future energy transition objectives while creating value across business, societal, and sustainability dimensions.

Business Dimension – Enabling Flexible and Adaptive Energy Systems

The business dimension is represented through the KVIs of Edge-Based Flexibility Orchestration and Renewable Integration at Distribution Level. Initial validation results demonstrate a flexibility activation time of approximately 600 ms and flexibility estimation accuracy of approximately 75%, indicating encouraging capability for near real-time orchestration of distributed energy resources. In addition, renewable energy integration activities reported renewable energy source intermittency of approximately 4% and grid overload levels of approximately 1%, suggesting that the pilot is beginning to establish the operational capabilities required to support increasing levels of distributed renewable generation while maintaining network stability.

A key finding emerging from the initial assessment is the pilot's progress towards enabling more responsive and adaptive energy system operations. The ability to activate flexibility services within sub-second timescales (600 ms, progressing towards the 300 ms KPI target) demonstrates the potential of edge-enabled orchestration to support rapid operational decision-making and improved management of distributed energy assets. Similarly, the achieved flexibility estimation accuracy provides encouraging evidence that the platform can support the situational awareness and forecasting capabilities required for more efficient utilisation of available energy resources.

The results also indicate early progress towards the cluster's renewable integration objectives. While these measurements should be viewed as initial operational indicators rather than final impact outcomes, they demonstrate that the pilot is beginning to address key challenges associated with renewable energy variability and grid management. From a value perspective, these capabilities are important because they create the enabling conditions required for higher levels of renewable energy penetration, improved system resilience, and more efficient operation of decentralised energy networks.

Societal Dimension – Early Evidence of Adoption and User-Centred Value Pathways

The societal dimension focuses on improving reliable access to EV charging and clean energy services, alongside supporting the digital upskilling of energy sector workers. Initial results indicate approximately 5% progress towards improving reliable access to EV charging and clean energy services, while workforce upskilling activities remain at an early stage of deployment. Given that both value areas are closely linked to user adoption, operational practices, and stakeholder engagement, future assessment cycles will incorporate qualitative stakeholder interviews to complement quantitative measurements and provide a richer understanding of the societal benefits generated by the pilot.

The initial findings suggest that the societal value pathways are beginning to emerge as the pilot transitions from technology deployment towards operational adoption. In the case of Digital Upskilling of Energy Sector Workers, the current baseline result provides an important reference point against which future improvements can be assessed. As adoption of COP-PILOT tools and dashboards increases, indicators such as reductions in manual monitoring effort and increased utilisation of digital decision-support capabilities are expected to provide valuable insights into workforce transformation and skills development within the energy sector.

Similarly, the reported improvement in Reliable Access to EV Charging and Clean Energy provides encouraging early evidence that the pilot is progressing towards its intended societal objectives. As validation activities continue and adoption increases, additional operational evidence will help provide a more comprehensive understanding of how the platform contributes to improved accessibility, reliability, and user experience. Together, these findings suggest that Cluster 3E is beginning to establish the foundations for wider societal benefits through enhanced access to digital energy services, improved operational visibility, and increased stakeholder engagement within the evolving energy ecosystem.

Sustainability Dimension – Emerging Evidence of Decarbonisation and Resource Optimisation

The sustainability dimension is represented through KVIs relating to CO_2 optimisation of charging energy, renewable energy integration and grid decarbonisation, biogas equipment lifetime extension, and local energy autonomy. As summarised in Table 3-13, initial results demonstrate EV charging demand forecasting accuracy of approximately 70%, together with early improvements of around 2% in biogas plant profitability and reductions in unplanned downtime. While several sustainability outcomes remain in the early stages of validation, these findings provide encouraging evidence that predictive and data-driven energy management approaches can contribute to both operational efficiency and longer-term decarbonisation objectives.

A key theme emerging from the initial assessment is the role of predictive and optimisation capabilities as enablers of future sustainability outcomes. The EV charging demand forecasting capability demonstrates how data-driven approaches can support more efficient energy management and charging optimisation, while the renewable energy integration activities establish important foundations for future decarbonisation benefits. Although many of the anticipated environmental impacts will require longer operational periods to fully materialise, the current results indicate that the pilot is beginning to establish the operational conditions necessary to support more sustainable energy system management.

Similarly, the Biogas Equipment Lifetime Extension KVI provides encouraging early evidence that predictive maintenance and operational optimisation can contribute to improved asset performance and resource utilisation. The observed improvements in profitability and reductions in unplanned downtime suggest that value is already beginning to emerge at plant level. Together with ongoing activities relating to local processing, distributed energy management, and digital infrastructure deployment, these findings indicate that Cluster 3E is establishing important foundations for future sustainability benefits associated with decarbonisation, energy resilience, and long-term asset optimisation.

In addition, the Local Energy Autonomy and EU Digital Sovereignty KVI represents an important strategic sustainability objective for the cluster. While its impacts are not yet fully reflected through quantitative indicators, ongoing activities relating to local processing, distributed energy management, and digital infrastructure deployment provide a foundation for future evaluation. As implementation progresses, a combination of quantitative metrics and qualitative stakeholder feedback may help capture the broader value associated with increased energy autonomy, resilience, and digital sovereignty within the energy ecosystem.

1.5 CLUSTER 4: SMART VINEYARDS & SUSTAINABLE WINERY ECOSYSTEMS

1.5.1 Pilot Scope and Planning

The Cluster 4 pilot plan is structured as a progressive validation roadmap for the smart-vineyard and sustainable-winery ecosystem. The roadmap moves from preparatory planning and initial use-case testing towards increasingly integrated demonstrations across the wine value chain. Cluster 4 is organised around four complementary use cases, depicted in Table 8.

Use case	Scope	Main workflow and capabilities
UC4.1	Circular IoT asset lifecycle management	Tracks sensor and IoT-device health, maintenance, replacement, reuse, recycling and traceability, while also supporting worker health monitoring through wearable or connected health sensors, to enable safer operations and more circular asset-management processes.
UC4.2	Smart vineyard monitoring and water-use efficiency	Combines soil-moisture sensing, drone observations, edge processing, FIWARE-based exchange, and vineyard-wide decision-support analytics and visualisation.
UC4.3	Sustainable winery production optimisation	Integrates winery data for real-time production monitoring, OEE analysis, and anomaly detection.
UC4.4	Rural connectivity and energy-aware infrastructure	Addresses resilient rural connectivity, network-status monitoring and energy-aware operation of off-grid or remote vineyard infrastructure.

Table 8: Cluster 4 use-case overview and Stage 1 validation status.

UC4.1 addresses circular lifecycle management for IoT assets; UC4.2 addresses smart-vineyard monitoring and water-use efficiency; UC4.3 addresses sustainable winery production monitoring and optimisation; and UC4.4 addresses rural connectivity and energy-aware infrastructure management.

Figure 25 represents the Gantt chart of the Cluster 4 validation activities planning, across the project timeline.

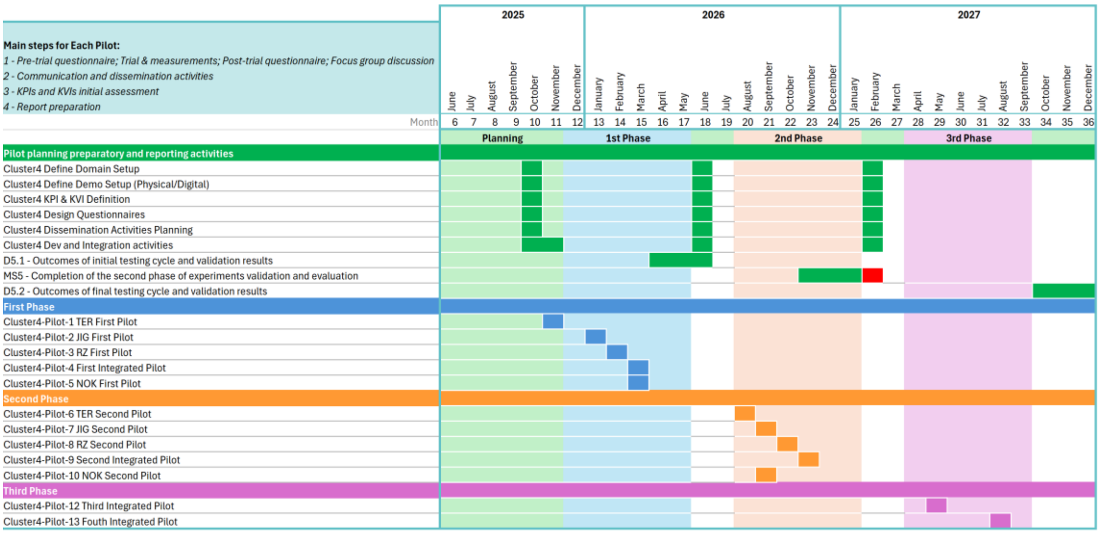

Figure 25: Cluster 4 Pilot Planning.

The first phase plan begins with use-case-level pilots before attempting the integrated cluster demonstrator. This allows the partners to validate the principal application chains, identify onboarding and observability gaps, and collect initial evidence from stakeholders on performance. This also includes validating the COP-PILOT platform capabilities required by the cluster, including edge-to-cloud deployment and orchestration, secure inter-domain connectivity, and FIWARE-based data exchange. Therefore, it includes pilots linked to UC4.1, UC4.2, UC4.3, and UC4.4, covering IoT sensor lifecycle management, smart vineyard irrigation, winery production monitoring, and rural connectivity and energy-aware vineyard infrastructure, as well as an integrated pilot focusing on the overall Cluster 4 workflow in collaboration with all the use cases.

The second phase is planned to deepen validation of each cluster branch and increase cross-domain integration. This phase includes further Terraview, JIG, RedZinc and Nokia pilots, as well as a integrated Cluster 4 pilot. The emphasis shifts from initial functionality towards scalability, interoperability, orchestration, secure data exchange, and more mature operational relevance across the participating domains. The third phase focuses on fully integrated Cluster 4 demonstrations. These pilots are intended to bring together the vineyard, winery, connectivity, energy, and IoT lifecycle-management dimensions into common scenarios.

The objective is to validate the cluster as a connected wine-value-chain ecosystem, where vineyard intelligence can support winery planning, infrastructure and energy data can improve the reliability of field operations, and lifecycle information from sensors and devices can support circular asset-management processes.

1.5.2 Pilot Execution and Results

During the first testing phase, Cluster 4 executed three piloting activities covering the vineyard and winery branches of the wine value chain, as summarised in Table 3-14 of the main D5.1 report.

Pilot 1 and Pilot 3 addressed UC4.2, focusing on soil-moisture analytics, edge/cloud processing, and vineyard decision support. Pilot 2 addressed UC4.3, focusing on winery data ingestion, real-time OEE monitoring, and dashboard visualisation. Finally, Pilo 4 focused on vineyard data ingestion, sensor, and energy use data. Together, these pilots provided initial validation of both the use-case workflows and the COP-PILOT platform mechanisms supporting them. The demonstrations exercised data ingestion, edge/cloud processing, FIWARE-based data exchange, OpenSlice/Maestro-supported deployment and orchestration, and secure interconnection through SIF. Subsections below summarise the main execution outcomes, while detailed technical setups, screenshots, survey results, and interview material are provided in the corresponding annexes.

1.5.2.1 Pilot 1: Vineyard Decision Support at Quinta da Pacheca

The first Cluster 4 pilot was conducted on 18 November 2025 at Quinta da Pacheca, Portugal, during the Euro Tech Day forum. The validation scope was the data orchestration with FIWARE integration, edge/cloud deployment and water-efficiency analytics. The pilot demonstrated how point soil-moisture measurements and drone-derived observations could be processed locally and combined with cloud analytics to create broader vineyard-management information. Thus, the workflow consisted of the acquisition and preprocessing of vineyard observations at the edge; transfer harmonised data to the OneSource domain through FIWARE-based interfaces; deploy the required software through Kubernetes/Helm and OpenSlice/Maestro-supported mechanisms; generate vineyard-wide AquaView soil-moisture outputs; and obtain initial evidence on usability, trust, sustainability relevance and adoption conditions.

The event reached approximately 100 participants, including vineyard owners and managers, agronomists and viticulture practitioners, agricultural technology providers, drone operators, researchers, and representatives of the wider innovation and regulatory ecosystem. Stakeholder results, extracted through eleven completed surveys and three targeted interviews, record 90.9% positive feedback regarding ease of understanding and use, 81.8% trust in data handling, and unanimous agreement among respondents that data-driven irrigation planning contributes to sustainable practices.

1.5.2.2 Pilot 2: Winery Production Monitoring at the Wine Next Generation Technical Days

The second Cluster 4 pilot was conducted on 19 March 2026 during the Wine Next Generation Technical Days in Vilafranca del Penedès, Spain. The validation scope was the integration of the UC4.3 winery application with the COP-PILOT platform, covering with FIWARE integration, and edge/cloud deployment. The demonstration focused on machine status, production orders, cumulative OEE indicators, anomaly visibility, and user and configuration management functions. Thus, the workflow consisted of generating data through an IoT source; ingesting and harmonising the information through the FIWARE Orion Context Broker; processing the incoming data in the JIG backend; storing the results in the application database; presenting machine status, production orders and cumulative OEE indicators through the real-time dashboard.

The event brought together approximately 60 participants, including winery owners and managers, production technicians, technology providers, senior management representatives, and members of Spanish wine-sector associations. Sixteen survey responses and three targeted interviews were collected. The feedback showed strong interest in the use of real-time monitoring to identify inefficiencies, reduce production waste, limit manual intervention, and improve operational visibility. In parallel, evidence supplied by Bodegas Franco-Españolas reported a reduction in production waste from 6.1% to 2.9%, a decrease in annual production downtime from 312 hours to 89 hours, and estimated energy savings of 47.3 MWh per year. These results support the operational relevance of UC4.3, while also indicating that future uptake will depend on progressive integration with existing winery systems, user training, and the digital maturity of each production environment.

1.5.2.3 Pilot 3: Transferability of the Vineyard Workflow at Quinta da Invejosa Poceirão

The third Cluster 4 pilot was conducted on 23 April 2026 at Quinta da Invejosa, Poceirão, Portugal, as part of an event organised with AVIPE. Its main objective was to validate how the COP-PILOT platform could enable interoperability across different vineyard environments by connecting heterogeneous field devices, communication technologies, edge resources, and analytics services within a common workflow. Two edge setups were used, with soil-moisture sensors connected through LoRaWAN and Wi-Fi, while multispectral drone imagery was processed locally, and the resulting data was exchanged through FIWARE-based interfaces. The COP-PILOT platform supported the harmonisation and transfer of these heterogeneous data streams to the AquaView analytics and visualisation services, demonstrating that the UC4.2 workflow could be replicated at a second site without requiring a fully bespoke integration.

The event involved approximately 40 participants, mainly viticulturists, winemakers and technical practitioners. 34 surveys and 4 interviews were collected to assess stakeholders' perceptions of the vineyard decision-support solution. The feedback highlighted the potential value of data-driven irrigation, reduced unnecessary pumping and more precise use of agricultural inputs.

1.5.2.4 Pilot 4: Soil Sensors predictions and energy data for 5G coverage at Villalvaro in Soria, Spain

The fourth Cluster 4 pilot was conducted on 16 and 17 June 2026 at Villalvaro and Matanza de Soria, at the Nokie 5G Green-powered site. Initial validation activities for Use Case UCE-4.4 were carried out at Nokia’s solar-powered site serving the Matanza and Villálvaro area in the province of Soria, Spain. These activities represented an important first step in the practical assessment of the use case under real operational conditions. The selected site combines 5G connectivity, agricultural sensing infrastructure, and renewable energy generation, providing a suitable environment for evaluating the integration of advanced digital technologies for precision viticulture and energy-efficient network operation.

As part of these initial trials, soil temperature and soil moisture data collected by sensors provided by Nokia and TerraView were integrated with a FIWARE Onesource instance. FIWARE was used to support the ingestion, management, and exchange of contextual information prior to data transfer to the COP-PILOT platform. This integration made it possible to combine measurements originating from different sensor systems within a common data-processing environment. It also demonstrated the interoperability of the components involved in the use case and their ability to provide agricultural information in a structured and reusable manner.

The integrated sensor data were subsequently used to implement an artificial intelligence algorithm. The purpose of this algorithm is to predict future soil temperature and moisture values based on measurements collected by deployed sensors. In addition to forecasting the behaviour of these environmental variables, the algorithm estimates the probability of disease affecting vineyards around the 5G site. This functionality is intended to support earlier and better-informed decision-making by local winegrowers, allowing them to anticipate potentially harmful conditions and to optimise the use of agricultural resources.

Predicting soil conditions and disease risk is especially relevant for precision viticulture. Continuous monitoring can provide winegrowers with more detailed information than occasional manual inspections, while predictive models can identify trends before they become evident in the field. In this context, the UCE-4.4 solution aims to demonstrate how 5G connectivity, Internet of Things sensors, data platforms, and artificial intelligence can be integrated to improve vineyard monitoring and enable more preventive, data-driven agricultural practices.

An important component of the validation activities was the direct participation of winegrowers from the surrounding area. The technologies deployed in the use case were presented to them, including the soil sensors, the 5G communication infrastructure, the data integration process, and the artificial intelligence functionalities. The participants were also informed about the expected benefits of the solution, such as improved awareness of soil conditions, earlier identification of disease risks, better planning of agricultural interventions, and potentially more efficient use of water, energy, and other resources.

Figure 26: Cluster 4, Pilot 4 at Villálvaro, Soria.

The engagement with local winegrowers also provided an opportunity to collect feedback from potential end users. Their observations are essential for assessing whether the proposed technologies address actual agricultural needs and whether the information generated by the platform can be presented in a useful and understandable manner. On 17 June 2026, the participating winegrowers completed a set of questionnaires and evaluation forms prepared for the project. The collected responses are currently being processed and analysed to calculate the relevant project KPIs for user acceptance, perceived usefulness, expected benefits, and overall suitability of the proposed solution.

In parallel with the agricultural and user-engagement activities, a series of 5G network performance tests was conducted during the two-day validation period. These tests were performed in the area covered by the deployed 5G infrastructure to which the agricultural sensors are connected. The objective was to characterise the network conditions available to the use case and assess whether the communication system can support the reliable transmission of sensor information and the operation of the associated digital services.

The network measurements generated an additional set of technical KPIs for use in evaluating the COP-PILOT project. These indicators provide information about the performance of the 5G network and its capacity to support connected sensing and agricultural applications in a rural environment. The resulting data are currently being ingested into the project platform, where they will be combined with the agricultural sensor measurements, application-level information, and user-evaluation results.

The trials also included integrating data on the site’s energy consumption and solar energy production. Since the Nokia site is powered by solar energy, these measurements are particularly relevant to the use case's energy-efficiency objectives. The available data enable analysis of the relationship among energy generation, network operation, connected devices, and the services executed at or through the 5G site.

Although the initial activities focused primarily on sensor integration, agricultural predictions, user engagement, and network validation, the energy data will support additional functionalities planned for future phases of UCE-4.4. These new features will address the optimisation of energy consumption and the more efficient use of locally generated renewable energy. For example, future developments may consider the availability of solar energy when scheduling computing, communication, or data-processing activities, while maintaining the performance and reliability required by the agricultural applications.

Overall, the initial validation activities successfully demonstrated the integration of heterogeneous soil sensors, FIWARE-based data management, the COP-PILOT platform, artificial intelligence models, 5G connectivity, and renewable-energy monitoring within a single operational scenario. The participation of local winegrowers also ensured that the technological evaluation was complemented by feedback from the solution's intended users.

The resulting KPIs will contribute to the technical and impact assessment of UCE-4.4 and will guide the implementation of the next set of functionalities, particularly those related to energy optimization and the long-term application of the solution in precision viticulture.

Initial Impact Assessment – Cluster 4

Cluster 4 focuses on the validation of digital agriculture solutions designed to improve operational efficiency, sustainability, and stakeholder engagement through advanced monitoring, data-driven decision support, and resource optimisation. As highlighted in Section 3.0, the cluster exhibits a well-balanced value profile, with sustainability representing the largest proportion of identified KVIs while maintaining strong business and societal value objectives. This balance reflects the pilot's ambition to demonstrate how digital technologies can simultaneously improve economic performance, environmental sustainability, and user acceptance within agricultural environments. Table 3-15 of the main D5.1 report presents the principal KPI-KVI relationships and initial validation findings from the first validation cycle (M17).

A key theme emerging from the assessment is the successful translation of digital capabilities into measurable operational value. Cluster 4 has already generated tangible evidence across multiple value dimensions. The results suggest that the pilot is progressing beyond technology deployment towards demonstrating practical benefits for agricultural stakeholders, while simultaneously supporting longer-term sustainability objectives.

The business dimension is represented through the KVIs of Digital Transformation MES and Operational Cost Savings. The results demonstrate encouraging progress towards both objectives. Scalability readiness has reached 50%, with supporting documentation already prepared for deployment and reuse in additional contexts, indicating that the solution is moving beyond pilot-specific implementation towards broader applicability. More notably, operational cost savings have reached 28.5%, approaching the ≥30% target and providing strong evidence that digitalisation and process optimisation can generate meaningful economic benefits. This finding is particularly significant, as it suggests that the solution is already delivering measurable efficiency improvements while laying the foundation for future scalability and broader adoption.

The societal dimension focuses on User Satisfaction Improvement and Ease of Use. The pilot achieved a user satisfaction score of 90.9%, indicating high stakeholder acceptance and confidence in the deployed solution. This is an important finding because long-term impact depends not only on technical performance but also on the user's willingness to adopt and integrate new digital tools into existing practices. The Stakeholder Inclusivity Score further demonstrates positive progress towards accessibility and engagement objectives, suggesting that the pilot is successfully supporting participation across relevant user groups. Collectively, these findings indicate that the solution is not only technically viable but also aligned with stakeholder expectations and operational needs.

The sustainability dimension represents the strongest component of the Cluster 4 value framework. The results demonstrate substantial progress across multiple sustainability-related value pathways. Energy efficiency gains of 53.65% at KVI level, exceeding the 40% target, provide encouraging evidence that the solution can contribute to more resource-efficient agricultural operations. Achieving a 100% Soil Health Monitoring Score confirms that the monitoring infrastructure required to support long-term environmental management is fully operational. In addition, the reduction of ground infrastructure requirements demonstrates how digital technologies can support more efficient deployment models while maintaining monitoring capabilities. Taken together, these findings suggest that the pilot is successfully establishing the foundations for more sustainable and data-driven agricultural practices.

Cluster	Key challenges encountered	Solutions applied	Lessons learned	Feedback to WP3/WP4
CL1 — Mining (LTU)	Restricted access to operational mines, data and field events; a data-loss issue under full 50-device concurrency (Pilot 3), identified for resolution ahead of the M25 re-evaluation; hardware limits for testing at scale.	Historical datasets replayed via simulator; large-scale simulated sensor populations (400+); emulated installations at RISE pre-production environment; conceptual/needs-oriented stakeholder validation at M18.	ColonyOS as job orchestrator simplifies building and scaling microservices-based systems; starting with basic auto-scaling and iterating reduces risk; improved hardware and automated smart-bolt installation needed for large-scale deployment; LEDs are a powerful awareness tool in dark environments; data-loss detection schemes must be validated early with mock data; continuous testing exposes defects invisible to short tests.	WP3-DM: fix concurrency-related data loss ahead of M25; WP3-DO: evolve autoscaling policies (SLA-oriented thresholds); WP4: complete UC1.1–1.3 migration into the RISE edge-to-cloud environment; WP3/WP4: support on-premises deployment models and OTA edge management per stakeholder feedback.

CL2 — Valencia (UPV)	Non-ideal physical deployment geometry (wide lane spans, sloped passages); heterogeneous protocols and external/legacy data sources.	Best-effort radar placement with algorithmic adaptation; NGSI-LD harmonisation via Orion/Scorpio; zero-touch OpenSlice provisioning of ephemeral alert services; YOLO-based verification against camera feeds.	24/7 stable edge detection and 5G transmission is achievable even under constrained installations; event-driven ephemeral services with automatic cleanup prevent resource bloat and remove manual operations; threshold-based orchestration loops (fill-level → OpenSlice → Telegram alert) work end-to-end.	WP3-DO/ESO: generalise the validated zero-touch ephemeral service pattern; WP4: containerise and orchestrator-deploy the remaining services (traffic-jam detection, alert/notification integration).
CL3A — AgriTech (AGA)	Integration of seven partner systems across five geographic domains; bridging permissioned blockchain with NGSI-LD; regulatory constraints on autonomous interventions.	Custom Hyperledger Fabric ↔ FIWARE NGSI-LD bridge; mandatory human (agronomist) confirmation before spray actions; private 5G MPN with public 5G failsafe; per-site DO/DM instances bound through SIF.	The full nine-step UAV-to-AI-to-UGV pipeline operates as a coherent, platform-orchestrated workflow in live field conditions; human-in-the-loop design satisfies both regulatory and trust requirements; the blockchain–FIWARE decoupled data-plane pattern (with PowerFleet as authorised bridge actor) is viable.	WP3-DM: incorporate the Fabric–NGSI-LD bridge pattern into the platform interoperability assets; WP4: support DEM #2 cross-season validation (detection accuracy, reduced false positives, consumer QR chain).
CL3E — Energy (UOP)	Expired VPN certificates disrupted edge-central connectivity	On-site certificate fix; FIWARE-based ingestion/harmonisation of Modbus TCP streams; OpenSlice-	A pre-demo readiness checklist (connectivity, certificate	WP3-SIF: automated certificate lifecycle monitoring/renewal and standardised readiness checks; WP4: prepare the

	before the demonstration; transformation of a manually monitored industrial plant into a data-driven asset.	automated dashboard deployment; structured demo combining plant tour and live walkthrough.	validity, service health) is essential; early engagement of plant operators and domain experts yields operationally grounded development directions (demand-response integration, methane curtailment monitoring).	cross-domain federation path towards Pilots 3–5.
CL4 — Wine (ONE)	Heterogeneous field devices and connectivity (LoRaWAN, Wi-Fi) across different vineyard environments; varying digital maturity of wineries.	FIWARE-based harmonisation and edge/cloud processing; Kubernetes/Helm and OpenSlice/Maestro-supported deployment; redeployment of the UC4.2 workflow at a second site; structured surveys and interviews at each event.	The platform enables workflow transferability across sites without bespoke integration; strong stakeholder trust and usability results (90.9% ease of use, 81.8% trust); quantified operational value at the winery (waste 6.1%→2.9%, downtime 312h→89h, 47.3 MWh/yr savings); uptake depends on integration with existing winery systems, user training and digital maturity.	WP3/WP4: consolidate the redeployment/onboarding path as a repeatable capability; support the UC4.1 and UC4.4 branches entering validation in the next phases, towards the integrated cluster workflow.
CL4 — Vineyards (Nokia, Terraview, Onesource)	Contact with local wineyards owners provides relevant feedback that we should use to focalize the next	Initial test was conducted first and presented later. Now that we have feedback, we know how relevant is to consider the particularity of the	It is important to recognize that pre-phylloxera vineyards, some of which are more than 100 or even 200	Recommendation to WP3 and WP4: It is essential to define and plan a dedicated deployment architecture combining physical sensors and virtual sensors for the

	OpenCalls and dissemination of information	vineyards in the area. People were very interested in the virtual stations presented in this initial pilot, asking for having some permanent data available during the project livetime.	years old, require a differentiated approach. Monitoring methods, management practices, and evaluation criteria should be adapted to the specific characteristics of these historic vines. Giving proper attention to their unique needs is essential for local growers and communities, for whom these vineyards represent not only an agricultural asset but also an important part of the region's cultural and historical heritage.	vineyard owners participating in the project. Their direct involvement will be crucial for validating the real-world benefits, usability, and practical value of the proposed technology under local operating conditions. We also recommend developing a specific plan to document and take into account the age and distinctive characteristics of the local vineyards, including the presence of very old vines, and to identify the most appropriate measures for improving their performance. The proposed approach should be adapted to the needs and resources of the area's small vineyard owners, who collectively manage approximately 60 hectares distributed across around 1,000 small vineyard plots. This highly fragmented ownership structure should be considered when designing the sensing architecture, defining deployment costs, selecting services, and planning future scalability.

REFERENCES

[1] COP-PILOT D5.1 “Outcomes of initial testing cycle and validation results”, Jun. 2026, Available: Link

[2] COP-PILOT, D0.1 “COP-PILOT’s Unified Methodology for Project Validation and Evaluation: Definition and Evaluation of Project-wide (WP1-WP7) Key Performance and Value Indicators”, v1.0, June 2026, Available: Link

Co-funded by
the European Union