



Deliverable D5.3

BeGREEN Use Case Demonstration and KPI Assessment

September 2025



Co-funded by
the European Union



Contractual Date of Delivery:	June 30, 2025
Actual Date of Delivery:	September 15, 2025
Editor(s):	Keith Briggs (BT)
Author(s)/Contributor(s):	Josep Xavier Salvat, Jose A. Ayala Romero (NEC) German Castellanos, Simon Pryor (ACC) Vladica Sark, Jesús Gutiérrez (IHP) Guillermo Bielsa (TSA) Miguel Catalan-Cid, Esteban Municio, David Reiss (i2CAT) Joss Armstrong, Jimmy O'Meara, Joseph McNamara (LMI) Israel Koffman, Baruch Globen (RunEL) Richard Mackenzie (BT) Ory Eger (PW) Anna Umbert (UPC) Mir Ghoraishi (GIGASYS)
Work Package	WP5
Target Dissemination Level	Public

This work is supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101097083, BeGREEN project. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS-JU. Neither the European Union nor the granting authority can be held responsible for them.

Revision History

Revision	Date	Editor / Commentator	Description of Edits
0.1	2025-06-09	Keith Briggs (BT)	Template made ready
0.2	2025-06-19	Miguel Catalan-Cid (I2CAT)	First version of PoC#1
0.3	2025-06-27	Israel Koffman (REL), Keith Briggs (BT)	First version of PoC#5
0.4	2025-07-02	Ory Eger (PW)	First version of PoC#4
0.5	2025-07-07	Germán Castellanos (ACC)	First version of PoC#3
0.6	2025-07-08	Vladica Sark (IHP)	First version of PoC#2
0.7	2025-07-08	Miguel Catalan-Cid (I2CAT)	Revision of PoC#1
0.71	2025-07-14	Miguel Catalan-Cid (I2CAT)	Revision of PoC#1
0.8	2025-07-15	Vladica Sark (IHP)	Revision of PoC#2
0.9	2025-07-16	Germán Castellanos (ACC)	Revision of PoC#3
0.91	2025-08-01	Vladica Sark (IHP)	Revision of PoC#2
0.92	2025-08-07	Keith Briggs (BT)	Final Editing of the document
0.93	2025-08-07	Jesús Gutiérrez (IHP)	First revision of the document
0.94	2025-08-08	Vladica Sark (IHP)	Second revision of the document
0.94	2025-09-08	Jesús Gutiérrez (IHP)	Final technical revision
0.95	2025-09-10	Mir Ghoraihi (GIGASY)	Final revision of the document
1.00	2025-09-12	Simon Pryor	Submission to the Participant Portal

Table of Contents

List of Figures	5
List of Tables.....	6
List of Acronyms	7
Executive Summary.....	8
1 Introduction.....	9
1.1 Purpose of the document	9
1.2 Recap of the BeGREEN vision	9
1.3 Structure of the validation campaign	9
1.4 KPI Framework.....	9
1.5 Document overview.....	10
1.6 Intended audience	10
1.7 Disclaimer and acknowledgements.....	10
2 PoC1 BeGREEN Intelligence Plane	11
2.1 UC1: Baseline	11
2.2 UC2: ML-driven cell on/off switching	14
2.3 Validation of UC2	18
2.4 UC3: Conflict Management.....	23
2.5 UC4: Model selection.....	25
3 PoC2 Sensing-assisted communications using RIS	29
3.1 PoC2 updated description	29
3.2 ISAC final improvements and KPI estimation	30
3.3 Sensing data exposure to RIC	32
3.4 RIS-assisted ISAC	32
3.5 PoC2 conclusion	37
4 PoC3 Energy-efficient CU and O-RAN RIC	39
5 PoC4 Energy-efficient DU implementation using HW acceleration	41
6 PoC5: 6G RU PA blanking test	44
7 Summary and Conclusions	46
8 Bibliography.....	48

List of Figures

Figure 2-1: Distributed PoC1 implementation: involved locations and detailed set-up	11
Figure 2-2: PoC1-UC1: EuCNC24 demo architecture.....	12
Figure 2-3: PoC1-UC1: EuCNC24 demo results.....	12
Figure 2-4: PoC1-UC1: BeGREEN's 2 nd review demo architecture	13
Figure 2-5: PoC1-UC1: BeGREEN's 2nd review results	13
Figure 2-6: PoC1-UC2: Architecture of the UC3 demonstration	15
Figure 2-7: PoC1-UC2: Training and test data (cell ON)	19
Figure 2-8: PoC1-UC2: Training and test data (cell OFF)	19
Figure 2-9: PoC1-UC2: Model performance in the testing phase (cell ON and predicting ON)	19
Figure 2-10: PoC1-UC2: Performance during the demo (15% of QoS Threshold).....	21
Figure 2-11: PoC1-UC2: Performance of the model during unseen week (throughput when cell is ON).....	21
Figure 2-12: PoC1-UC2: Performance of the model during unseen week (throughput when cell is OFF).....	22
Figure 2-13: PoC1-UC2: Performance of the rApp under different QoS thresholds.....	22
Figure 2-14: PoC1-UC2: Energy savings under different QoS thresholds.....	23
Figure 2-15: PoC1-UC2: Outage penalty under different QoS thresholds.....	23
Figure 2-16: PoC1-UC3: Conflict manager architecture	24
Figure 2-17: PoC1-UC3: Example of Conflict manager working, and messages exchanges across xApps.	25
Figure 2-18: PoC1-UC4: Model selection function	26
Figure 2-19 PoC1-UC4: Energy consumption during model inference.....	27
Figure 2-20: PoC1-UC4: CPU usage during model inference.....	28
Figure 2-21: PoC1-UC4: CPU usage during model inference.....	28
Figure 3-1 Sub-6 ISAC testbed @ EuCNC'25 BeGREEN Exhibitor	29
Figure 3-2: Sensing-assisted communications demo - use case 2.....	30
Figure 3-3: Sensing-assisted communications demo - use case 3.....	30
Figure 3-4: Heat map of a scenario with 2 persons standing in front of the ISAC system in anechoic chamber	31
Figure 3-5: Heat map of a scenario with a single person standing in front of the ISAC system in anechoic chamber	31
Figure 3-6 Integration of Sub-6 GHz ISAC Prototype with sensing SM @ EuCNC'25.	32
Figure 3-7: Anechoic chamber configuration	33
Figure 3-8: Anechoic chamber setup, view from the ISAC system towards RIS.....	34
Figure 3-9: Anechoic chamber setup, view from RIS towards ISAC system and corner reflector.....	34
Figure 3-10: Heat maps for a) 80, b) 85, c) 90, d) 95 degrees reflection angles with respect to the RIS surface	35
Figure 3-11: Stronger reflection received back with use of PA and high gain antenna	36
Figure 3-12: Alternative setup for sensing over RIS PoC.....	36
Figure 3-13: Sensing through the RIS - obtained heat maps.....	36
Figure 3-14: Setup for sensing through the RIS.....	37
Figure 4-1: PoC3 set-up implementation	39
Figure 4-2: Average power consumption comparison for ARM and x86	40
Figure 5-1: Linear MIMO receiver	41
Figure 5-2: Sphere Decoder.....	41
Figure 5-3: Sphere Decoder Illustration	41
Figure 5-4: Implementations being tested in the PoC.....	42
Figure 5-5: PoC setup	42
Figure 5-6: Sphere Decoder wireless performance.....	43
Figure 6-1: RunEL RU installed in the lab pointing outside the Brunel University	44
Figure 6-2: Hourly traffic in an average day	45

List of Tables

Table 1-1: KPI summary table.....	10
Table 2-1: Objectives, tasks and results of PoC1-UC1	14
Table 2-2: Objectives, tasks and results of PoC1-UC2	18
Table 2-3: PoC1-UC2: Model accuracy	20
Table 2-4: Objectives, tasks and results of PoC1-UC3	25
Table 3-1: Updated Planning of PoC2.....	37
Table 4-1: Objectives, tasks and results of PoC3	40
Table 6-1: Measurement results of the outdoor testes at Brunel University	45

List of Acronyms

3GPP	3rd Generation Partnership Project
5GC	5G Core
5G NR	5G New Radio
AI	Artificial Intelligence
ASIC	Application Specific Integrated Circuit
COTS	Commercial Off-The-Shelf
CP	Control Plane
CPE	Customer-Premises Equipment
CU	Central Unit
DME	Data Management and Exposure
DoW	Description of Work
DU	Distributed Unit
EIRP	Equivalent Isotropic Radiated Power
HW	Hardware
IP	Internet Protocol
ISM	Industrial Scientific and Medical
KPI	Key Performance Indicator
KPM	Key Performance Measurement
L2	Layer 2
LoS	Line-of-Sight
MAC	Medium Access layer
MAE	Masked autoencoders
ML	Machine Learning
MNO	Mobile Network Operator
mmWave	Millimetre Wave
near-RT	near Real-Time
NLoS	Non-Line-of-Sight
non-RT	non-Real-Time
NR	New Radio
O-RAN	Open RAN
PHY	Physical
PRB	Physical Resource Block
QoS	Quality of Service
RAN	Radio Access Network
PDU	Power Distribution Unit
PoC	Proof-of-Concept
RAN	Radio Access Network
RB	Resource Block
RIC	RAN Intelligent Controller
RIS	Reconfigurable Intelligent Surface
RU	Radio Unit
SDR	Software Defined Radio
SHO	Smart Handover
SLA	Service Level Agreement
SME	Service Management and Exposure
TDD	Time Division Duplex
TGW	Telemetry Gateway
UE	User Equipment
XDP	eXpress Data Path
WP	Work Package

Executive Summary

BeGREEN D5.3 documents the consortium's final system-level demonstrations and quantitative KPI assessment after thirty months of R&D. Five Proofs-of-Concept (PoCs) were built, at BT Adastral Park (UK), Accelleran, IHP, NEC Labs, and RunEL premises, and partner universities, showcasing a holistic, AI-driven, energy-efficient 6G Radio Access Network (RAN) architecture.

Key results reported in this deliverable are:

- **PoC1: Intelligence-Plane maturity** - Non-RT and Near-RT RICs, AI Engine and rApp/xApp suite reached TRL-6 and executed a closed-loop cell on/off strategy that reduced radio-access energy by up to 30 % with <15 % Quality of Service (QoS) impact. A conflict manager ensured >30% energy-efficiency gain while safeguarding mobility-related xApps, proving the feasibility of multi-objective orchestration. Model-Selection logic trimmed Machine Learning (ML)-inference energy by 40–68% while maintaining accuracy through adaptive time-granularity.
- **PoC2: Reconfigurable Intelligent Surface (RIS)-assisted Integrated Sensing and Communication (ISAC)** - Sub-6 GHz sensing achieved 0.7 m range resolution and >90% person-count precision. First-of-its-kind demonstration of “sensing-through-RIS” verified user detection behind obstacles, enabling coverage-aware RU sleep modes.
- **PoC3: Energy-aware CU & O-RAN RIC** - Porting to ARM servers yielded up to 10% savings under high traffic, with 33% lower idle power versus x86, validating edge-grade, low-carbon cloud-native RAN.
- **PoC4: Hardware (HW)-accelerated DU** - Graphics Processing Unit (GPU)-based sphere decoder cut execution time by 11× and improved energy efficiency from 9 kbit/J (CPU) to 100 kbit/J—evidence that off-the-shelf accelerators are pivotal for green Layer-1 processing.
- **PoC5: Power Amplifier (PA)-blanking in 6G RU** - Field trial at Brunel University confirmed 64 % PA power savings, dropping RU PA draw from 5.41 W to 1.94 W using scheduler-driven blanking.

Collectively, the PoCs demonstrate an end-to-end 6G blueprint capable of achieving 20-35 % system-level energy reduction without degrading user experience, supporting the project's vision of a carbon-aware mobile network stack. The results are fed into O-RAN Alliance and 3GPP work items.

1 Introduction

1.1 Purpose of the document

Deliverable D5.3 “Final Demonstration and KPI Assessment” serves three complementary goals.

1. It **describes the experimental setups** used to validate the architectural innovations of BeGREEN in realistic, multi-vendor environments.
2. It **presents the measurement results and KPI attainment**, linking them back to the quantitative targets defined in BeGREEN D5.1 [1].
3. It **captures lessons learned** and provides technology-readiness evidence for transfer to standards, open-source communities and commercial exploitation.

By consolidating WP5 results, the deliverable closes the feedback loop between architectural design (WP2-WP4) and empirical validation, ensuring the BeGREEN’s contributions are both scientifically sound and practically deployable.

1.2 Recap of the BeGREEN vision

BeGREEN targets a step-change reduction in RAN energy consumption through four orthogonal levers:

- **AI-native control loops** decouple optimisation logic from vendor-specific base-band implementations, enabling rapid deployment of green policies.
- **Resource disaggregation** pushes compute-intensive tasks to the most energy-appropriate node—edge, far-edge or device—while maintaining QoS.
- **Programmable RF subsystems** (e.g., RIS, PA-blanking) allow the physical (PHY) layer itself to adapt its footprint dynamically.
- **Commodity accelerators and low-power silicon** provide a pragmatic path toward sustainable 6G infrastructure without bespoke Application Specific Integrated Circuits (ASICs).

WP5 translates these concepts into tangible PoCs, each mapped to measurable KPIs—energy savings (%), inference latency (ms), throughput loss (%), and sensing accuracy (m/deg).

1.3 Structure of the validation campaign

1. **PoC1** stresses the intelligence plane under live traffic, integrating Accelleran’s dRAX, BT’s AI Engine and Viavi’s TeraVM emulator across BT, i2CAT and Accelleran labs.
2. **PoC2** merges NEC’s 100-element RIS with IHP’s Sub-6 GHz ISAC platform inside an anechoic chamber, later ported to field trials.
3. **PoC3** ports CU & RIC code-base to ARM servers, adding eXpress Data Path (XDP) for packet-fast-path acceleration.
4. **PoC4** benchmarks Sphere-Decoder on NVIDIA RTX-A4500 versus Intel Xeon to explore compute-vs-energy trade-offs.
5. **PoC5** evaluates DU-aware Power Amplifier (PA)-blanking under outdoor 100 MHz New Radio (NR) carrier in Brunel University’s campus.

1.4 KPI Framework

Table 1-1 maps project-level KPIs to their corresponding PoC and measurement methodology. These KPIs collectively confirm that the BeGREEN architecture meets, and in several cases exceeds, Horizon-Europe SNS call benchmarks for sustainable 6G systems.

Table 1-1: BeGREEN PoCs KPI Summary Table

KPI Category	Target	Measurement Node	PoC	Achieved Value
RAN energy per bit	–25 % vs. baseline	Viavi TeraVM power probe	PoC 1	–30 % (cell sleep)
AI loop latency	<1 s non-RT; <50 ms near-RT	dRAX telemetry logs	PoC 1/3	0.8 s / 12 ms
Sensing range resolution	≤1 m	Heat-map peak separation	PoC 2	0.7 m
CU/RIC idle power	–20 %	Smart-PDU	PoC 3	–33 %
L1 decode energy	10× vs. CPU	On-board wattmeter	PoC 4	11×
PA consumption	–50 %	Keysight DAQ	PoC 5	–64 %

1.5 Document overview

The remainder of D5.3 is organised as follows:

- **Chapter 2** details **PoC1**, subdivided into four use cases—baseline setup, ML-driven cell on/off, conflict management, and model-selection efficiency.
- **Chapter 3** presents **PoC2**’s RIS-assisted sensing results, including chamber calibration data and real-world reflections.
- **Chapter 4** reviews **PoC3**’s ARM-based CU/RIC benchmarking and XDP acceleration insights.
- **Chapter 5** explains **PoC4**’s GPU off-load methodology and energy measurements.
- **Chapter 6** outlines **PoC5**’s PA blanking field trial and day-weighted energy analysis.
- **Chapter 7** synthesises cross-PoC conclusions and gives recommendations for standardisation and exploitation.

1.6 Intended audience

While the deliverable is publicly released, it targets three primary readerships:

1. **Network operators and integrators** evaluating near-term green upgrades to 5G-Advanced and early 6G trials.
2. **Standards bodies (O-RAN, 3GPP, ETSI)** looking for empirically backed contributions on AI-native energy management.
3. **Academic and industrial researchers** seeking replicable baselines for follow-up work on RIS, ISAC, or accelerator-enabled RAN.

For quick orientation, readers may consult the executive summary and KPI table, then deep-dive into the PoC chapter aligned with their domain.

1.7 Disclaimer and acknowledgements

This work was co-funded by the Smart Networks and Services Joint Undertaking (SNS-JU) under Horizon-Europe Grant No 101097083. The opinions expressed are those of the authors and do not necessarily reflect those of the European Union. The consortium thanks Viavi, NVIDIA, Keysight and Brunel University for in-kind equipment support and test-site access.

2 PoC1 BeGREEN Intelligence Plane

This PoC demonstrated one of the key elements of the BeGREEN architecture, the Intelligence Plane, which is designed to enable AI-driven decision-making for rApps and xApps towards enhanced energy efficiency. To this end, this PoC was focused on one specific mechanism, 5G cell or carrier on/off switching, which is one of the relevant energy-efficiency use cases promoted by O-RAN Alliance [2]. PoC1 was structured according to three different use cases. First, UC1 focused on demonstrating the Intelligence Plane baseline implementation comprehending the AI Engine, the RICs and the RAN emulator. Then, in UC2, the end-to-end on/off switching mechanism supported by the Intelligence Plane was validated. Finally, we considered an additional UC3 which reports the basic functionality in this use case of the Conflict Management framework developed within the project. Additionally, in this report we include a validation of the Model Selection feature introduced in deliverable D4.3 [3], also under the umbrella of the on/off cell switching mechanism.

As shown in Figure 2-1, the PoC successfully integrated testbeds of 3 different partners: i2CAT, Accelleran, and BT.

2.1 UC1: Baseline

Results of UC1 were demonstrated during the EuCNC24¹ event and during BeGREEN 2nd review. The following is a short description of the demos:

- **EuCNC24:** The demo showcased how ML models are deployed in the AI Engine and exposed to the Non-RT RIC through the serverless framework and using the AI Assist (AIA) rApps developed within the project, as illustrated in Figure 2-2. It also details the utilisation of the R1 interface Data Management and Exposure (DME) and Service Management and Exposure (SME) features to enable model serving.

The use case was elaborated according to two XGBoost predictors, cell load and cell energy consumption, plus the energy score function, which leveraged data from real operator. Results were exposed in real-time in Grafana, as shown in Figure 2-3.




Figure 2-1: Distributed PoC1 implementation: involved locations and detailed set-up

¹ <https://www.youtube.com/watch?v= NOJY0Sepgc>



- BeGREEN** [SNS-JU-101097083]


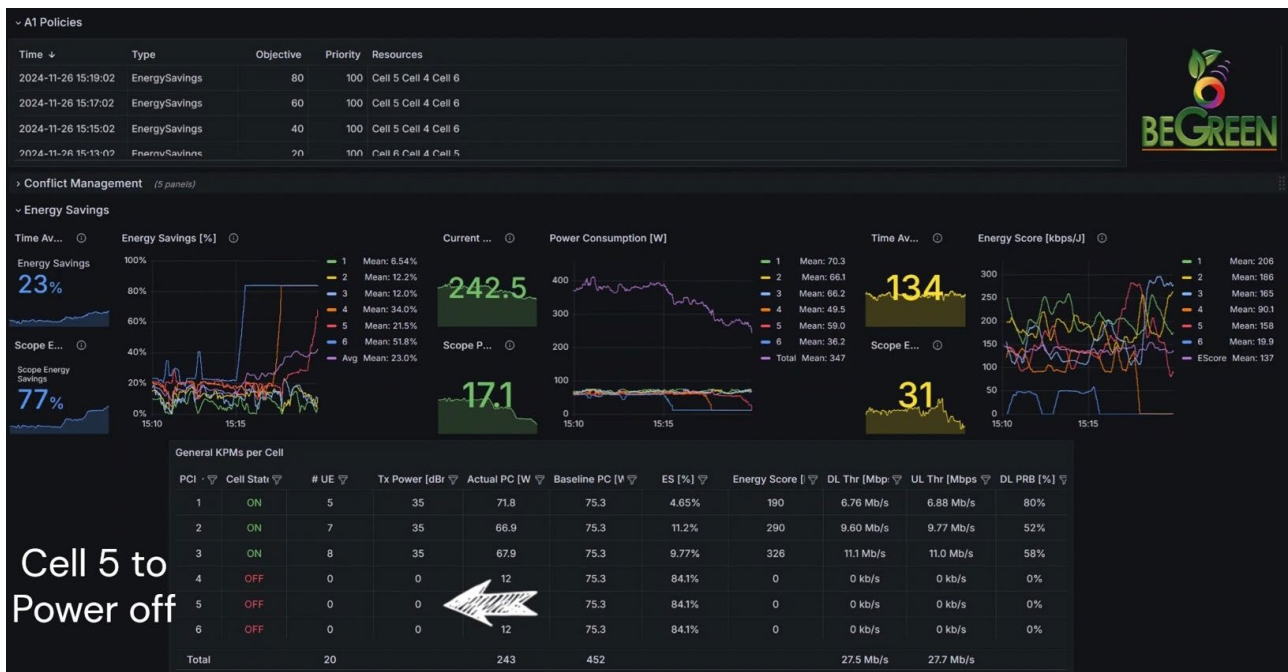
Figure 2-4: PoC1-UC1: BeGREEN's 2nd review demo architecture

Figure 2-5: PoC1-UC1: BeGREEN's 2nd review results

Table 2-1 shows the achieved objectives and results of this UC, which established the baseline for the next use cases.

Table 2-1: Objectives, tasks and results of PoC1-UC1

Objective	Tasks	Description	Month	Results
Intelligence Plane development	AI Engine development	Implement the baseline version of AI Engine, enabling serverless inference of ML models.	DONE M18	Validation in D4.2 [4] and D4.3 [3], by providing details on AI Engine framework.
	Non-RT RIC developments	Implement AIA rApps for hosted models and consumer rApps to demonstrate this feature. Integrate with near-RT RIC through A1/O1 interfaces.	DONE M22	Validation in D4.2 [4] and D4.3 [3], by providing details on how models and EE functions are exposed to rApps and benchmarking Non-RT control-loop latency (≥ 1 s).
	Near-RT RIC developments	Integrate with non-RT RIC through A1/O1 interfaces. Exposure of RAN telemetry including Energy-related KPIs.	DONE M22	Implemented A1-EE policy as described in D4.2 [4]. Demonstrated the exposure of RAN telemetry to xApps and rApps through TGW and KPM producer rApps.
Integration with TeraVM	TeraVM deployment and preparation	Deploy TeraVM and make it accessible to RICs. Enable targeted UCs and scenarios.	DONE M20	Creation of scenario based on BT's Adastral Park topology.
	Intelligence Plane integration with TeraVM	Integrate near-RT RIC with TeraVM, allowing collection of KPIs and RAN control.	DONE M22	Exposure of metrics through Telemetry Gateway and management of Energy Saving status and Handovers through E2 interface and devoted xApps.
Demonstration	AI Engine Demonstration	Demonstrate integration of AI Engine and RIC.	DONE M18	Achieved and demonstrated at EuCNC'24.
	RAN Control Demonstration	Demonstrate integration of Intelligence Plane and TeraVM.	DONE M25	Achieved and demonstrated at BeGREEN's 2 nd Review.

2.2 UC2: ML-driven cell on/off switching

The demonstration of PoC1-UC2 was performed as part of BeGREEN's final demos, validating the implementation of the Intelligence Plane and highlighting its features towards enhancing energy efficiency of the RAN segment. The validated scenario, illustrated in Figure 2-6, comprehends the following main steps:

1. The scenario emulates realistic period of two day-night cycles traffic pattern from a European Mobile Network Operator (MNO), as further described in D2.3 [5] and D4.4 [6]. Six 5G SA cells are deployed, with one acting as a capacity cell that can be switched off during low-demand periods. Since TeraVM emulates in real-time, to speed up the demo, we reduced 1 hour of real time to 5 minutes in TeraVM. Thus, every two day-night cycles lasted in simulated time for 4 hours.
2. Data exposure for model training is obtained via the dRAX TGW, the KPM Producer rApp, which exposes this data to other rApps through R1 interface, and the Dataset rApp, which exposes a dataset service to the Non-RT RIC operator. More details can be found in D4.3 [3]. We used TeraVM with different seeds, impacting the mobility and Physical Resource Block (PRB) pattern of the UEs, to obtain statistical variation.
3. An XGBoost regressor is trained to predict throughput during both ON and OFF periods, using Viavi's TeraVM to simulate user mobility diversity according to different seeds. The model uses actual cells PRBs as inputs (average of the last 30 seconds), trying to minimize the number of features, and,

accordingly, the energetic footprint of the model. The output is the predicted ON and OFF throughput for the next 30 seconds.

4. ES Control rApp determines the target cell ON/OFF status according to the operator's Service Level Agreement (SLA) (e.g., maximum 15% throughput degradation), using predictions provided by the AI Assist rApp and real-time data via KPM producer rApp. The rApp has an internal logic to avoid ping-pong effects (i.e., constant on-off loops during high variability).
5. The generated A1 ES policy is sent to the ES xApp, which activates or deactivates the cell according to it.
6. The Smart Handover (SHO) xApp manages UE handovers to avoid disconnections.
7. The automated control loop is demonstrated over multiple full day-night traffic cycles.





Figure 2-6: PoC1-UC2: Architecture of the UC3 demonstration

The resulting demo² was showcased in the following events:

² <https://www.youtube.com/watch?v=CriID1oHveA>

1. **IEEE INFOCOM'25:** Presented in the demo session, after demo paper³ acceptance (real-time demo).



2. **EuCNC'25 & 6G Summit⁴:** Presented (real-time demo) at the joint booth organized during the whole EuCNC week by the projects BeGREEN, 6G-TWIN, EXIGENCE/6Green and 6G-SENSES, with the motto "Sustainability for 6G: from Infrastructure to Services".



3. **BeGREEN final event in Adastral Park⁵:** Presented together with the rest of final BeGREEN demos (real-time demo).

³ <https://zenodo.org/records/15040326>


⁴ <https://www.sns-begreen.com/news/eucnc-2025-joint-booth-on-sustainability-for-6g-from-infrastructure-to-services>

⁵ <https://www.sns-begreen.com/news/begreen-project-final-event>



4. BeGREEN final webinar⁶: Presented within WP4 session (recorded video).

Demo of BeGREEN Intelligence Plane for AI-driven Energy Efficient O-RAN management | i2CAT



Deployment of ES rApp:
Target cell, QoS threshold (15%), loop period (30s)

Table 2-2 summarises the achieved objectives and results of this UC.

⁶ <https://www.sns-begreen.com/news/begreen-webinar>

Table 2-2: Objectives, tasks and results of PoC1-UC2

Objective	Tasks	Description	Month	Results
Energy-Saving RAN Control	Control rApp development	Develop a control rApp managing energy-saving policies according to UC logic.	DONE M25	Developed rApp which allows to set a QoS threshold to control the maximum throughput degradation is allows to maintain the target cell off. Logic to avoid ping-pong effect. Integration with R1 interface (consumer rApp).
	Control xApp development	Develop control xApps devoted to energy-saving use cases and implement control policies to manage them.	DONE M25	Developed Energy Saving xApp capable for managing the Tx power of the cells. Allows gradual modification until cells are completely off or at maximum power.
	Integration with TeraVM	Integrate rApps and xApps with TeraVM to test developed control-loop.	DONE M25	Validated after finalising UC1.
AI-driven RAN Control	ML-driven Control-loop development	Integrate ML model outputs with the control-loop, for instance load predictors.	DONE M28	Developed AI Assist rApp to generate and expose network throughput predictions from the model stored in the AI Engine. Integration with R1 interface (producer of predictions and consumer of KPMs to generate them)
	ML models development	Develop the required ML models.	DONE M28	Trained a throughput predictor based on XGBoost which considers the actual PRBs of all the cells and predicts the throughput when a specific cell is OFF and ON. Trained with data from TeraVM, obtained through a dataset service rApp and exposed as dataset in Zenodo [D4.4]
Demonstration	AI-driven ES control-loop validation	Validate developed ML-based energy-saving strategy.	DONE M29/M30	Validated in the different demos of UC3.

2.3 Validation of UC2

In this subsection we provide a validation of the evaluated scenario under different rApp QoS thresholds. Also, we present details about the developed model.

A. Model training and performance:

We performed multiple runs with different random seeds to generate sufficient data for training and testing the model. Specifically, we used 12 days of data for training and 4 days for testing, considering both the ON and OFF states of the target cell. The dataset included the PRB usage of all cells in the network and the corresponding aggregated throughput measurements. This allowed the XGBoost regressor to learn how to predict the next aggregated throughput in case the target cell is ON or OFF, based on the PRB status of all cells. Figure 2-7 and Figure 2-8 show the data used for training and testing.




Figure 2-7: PoC1-UC2: Training and test data (cell ON)




Figure 2-8: PoC1-UC2: Training and test data (cell OFF)

Figure 2-9 illustrates the model performance during the testing phase, in this case predicting the next ON throughput when the target cell status is also ON. Table 2-3 reports the accuracy of the model, which is satisfactory in all the four ON-OFF combinations, especially when the actual status of the cell is ON. According to these results, we applied this model to the use case, as will be described in the next section.




Figure 2-9: PoC1-UC2: Model performance in the testing phase (cell ON and predicting ON)

Table 2-3: PoC1-UC2: Model accuracy

Actual Cell Status	Target Cell Status	R2	MAE (%)
ON	ON	0.93	3.84
ON	OFF	0.89	4.04
OFF	ON	0.79	6.47
OFF	OFF	0.84	4.83

B. UC results:

Once the model was deployed, we exposed it through the AI Engine and the AIA rApp to the control Energy Saving rApp, which performed ON/OFF decisions according to the actual status the network and the predictions. The AIA rApp subscribed to the KPM producer rApp, obtaining this way the PRB status of the cells, which was sent as input to the model at the AI Engine to retrieve the ON and OFF throughput prediction. The prediction was generated in approximately 0.1s, enabling a Non-RT control-loop. Once received the throughput estimation, and according to the configured QoS threshold, the Energy Saving rApp decided to switch the cell ON or OFF.

We evaluated the control rApp with an unseen period of two day-night cycles. Figure 2-10 shows selected Grafana figures on how the rApp performed using a QoS threshold of 15%. As can be seen in the upper figure, the capacity cell being targeted (yellow solid line), became deactivated during off peak periods, when the predictor forecasted a throughput degradation lower than 15%. The middle figure depicts the QoS score, i.e. the throughput demanded by the UEs vs the served one. Note that during deactivation periods, there is no noticeable degradation of QoS as most of the cell throughput can be served by the remaining active cells. Indeed, only near cell activation QoS degradation seems to approach the targeted 15%, causing the control rApp to reactivate the cell. Finally, bottom figure depicts the average energy savings during the whole period. Note that even during cell deactivation, we have a baseline consumption of 15%, therefore maximum energy savings during this period reach between 30% and 20%.




Figure 2-10: PoC1-UC2: Performance during the demo (15% of QoS Threshold)

In addition to the demo, we evaluated the model and rApp performance during the unseen week using different QoS thresholds. Figure 2-11 and Figure 2-12 depict model accuracy to predict ON throughput and OFF throughput, respectively, depending on the actual status of the cell. As expected, the Masked autoencoders (MAE) is higher than in the training evaluation but remains around 10%. This causes some errors in the rApp decisions, which lead to energy saving missed opportunities or QoS outages/degradations, as is shown in the next figures. The most challenging case was OFF to ON, since we noticed that in some cases the capacity cell could not obtain back the throughput level that it had before the OFF (compared to the case where it was always ON). Nevertheless, results were satisfactory enough considering the lightweight model.




Figure 2-11: PoC1-UC2: Performance of the model during unseen week (throughput when cell is ON)




Figure 2-12: PoC1-UC2: Performance of the model during unseen week (throughput when cell is OFF)

Figure 2-13 shows the success percentage of the rApp under different QoS thresholds, compared to the optimal solution. Note that the number of energy-saving opportunities (OFF status according to the oracle) increases with the QoS threshold. Thus, at low thresholds, the missed opportunities (decided ON but should be OFF) are less in proportion. In any case, error percentage is in general aligned with the obtained MAE.




Figure 2-13: POC1-UC2: Performance of the rApp under different QoS thresholds

Figure 2-14 depicts the achieved energy savings by using the rApp under different thresholds, compared to the optimal solution. In general, excepting the 20% threshold case which underperforms, the other cases shown a similar performance of rApp compared to the optimal solution. Also, note that using a QoS threshold of 25% lead to almost 90% of OFF decisions during the analysed period.




Figure 2-14: PoC1-UC2: Energy savings under different QoS thresholds

Finally, we computed the penalty on total throughput when incurring in a policy outage (deciding OFF but should be ON). As show in Figure 2-15, the average penalty increases with the QoS threshold, since errors lead to higher throughput differences. However, if we consider the QoS threshold itself, the absolute penalty error (i.e., the deviation from the threshold) systematically decreases. For instance, in the case of a 25% threshold, erroneous decisions result, on average, in a throughput that is 3% lower than the specified by the policy.




Figure 2-15: PoC1-UC2: Outage penalty under different QoS thresholds

According to these results, we can conclude that the validation of the Intelligence Plane in this UC was successful, being able to implement and validate an AI-driven cell on/off switching strategy using the features of the AI Engine, the RICs and the RAN emulator.

2.4 UC3: Conflict Management

In PoC1, Use Case 3 (UC3) focuses on the implementation and validation of a conflict mitigation mechanism designed to ensure robust coordination between energy-saving and mobility xApps operating concurrently within the Near-RT RIC. This use case addresses the common scenario where aggressive energy-saving policies may conflict with mobility-driven decisions, such as user handovers, leading to performance degradation or service interruptions. The Conflict Manager, as presented in Figure 2-16, was integrated into the Intelligence Plane to detect and mitigate such conflicts in real-time, ensuring that policy-driven restrictions (e.g., cell deactivation for energy savings) do not disrupt mobility actions and vice versa.




Figure 2-16: PoC1-UC3: Conflict manager architecture

During the validation phase, the system was tested with and without the Conflict Manager active, as shown in Figure 2-17 where conflicts were detected and messages to restrict and avoid conflicts are presented. When disabled, the activation of a 100% energy-saving policy led to restricted cells being selected for user handover, resulting in lower network performance and an average energy score of just 90.2 Kbps/J. Once the Conflict Manager was activated, it intercepted conflicting handover commands, prevented unnecessary user transitions to restricted cells, and guided the SHO xApp to select active alternatives. This mitigation significantly improved the handover success rate and reduced user disconnection times. As a result, the network's energy efficiency improved, with the energy score rising to nearly 125 Kbps/J—an increase of more than 30%. These findings confirm that proactive conflict detection and resolution at the RIC level is a key enabler for reliable coordination between multiple optimisation goals and highlight the importance of a zero-touch intelligence layer in future multi-objective O-RAN environments.




Figure 2-17: PoC1-UC3: Example of Conflict manager working, and messages exchanges across xApps.

Table 2-4 shows the achieved objectives and results of this UC:

Table 2-4: Objectives, tasks and results of PoC1-UC3

Objective	Tasks	Description	Month	Results
Definition of CM	Definition of the Conflict management process	Define the mechanism to implement the conflict management guidance based on A1 policies and the entities needed inside the RIC.	DONE M20	The conflict manager implementation and the definitions of the conflicts was defined and presented in D4.2 Section 2.1.3.3 [3]. The development work and initial testing was presented in D4.3 section 2.3.4 [3] where more than 50% of the real conflicts between ES and SHO xApp were reduced.
Development work	Update of xApps to support Conflict Management	Implementation of functionalities inside the xApps and the Conflict Manager to account and avoid conflicts.	DONE M22	
Initial testing	Definition of testing scenarios and integration of functionalities	Integrate initial conflicts and functionalities in the system.	DONE M23	
Demonstration	Conflict management validation	Validate developed Conflict management strategy.	DONE M26	As described in D5.2 [9] and D4.3 [3], the conflict manager was able to support the more than 50% conflicts between xApps. Additionally, the CM was presented and active in the whole MWC25 demo.

2.5 UC4: Model selection

As was reported in deliverable D4.3 [3], a Model Selection function was introduced in the AI Engine to select the most efficient models devoted to the same function, e.g. traffic forecasting. Figure 2-18 shows an

example of workflow for this function, which works together with AIA rApps to assist on model selection according to the requirements of the control rApp in terms of accuracy and/or periodicity.

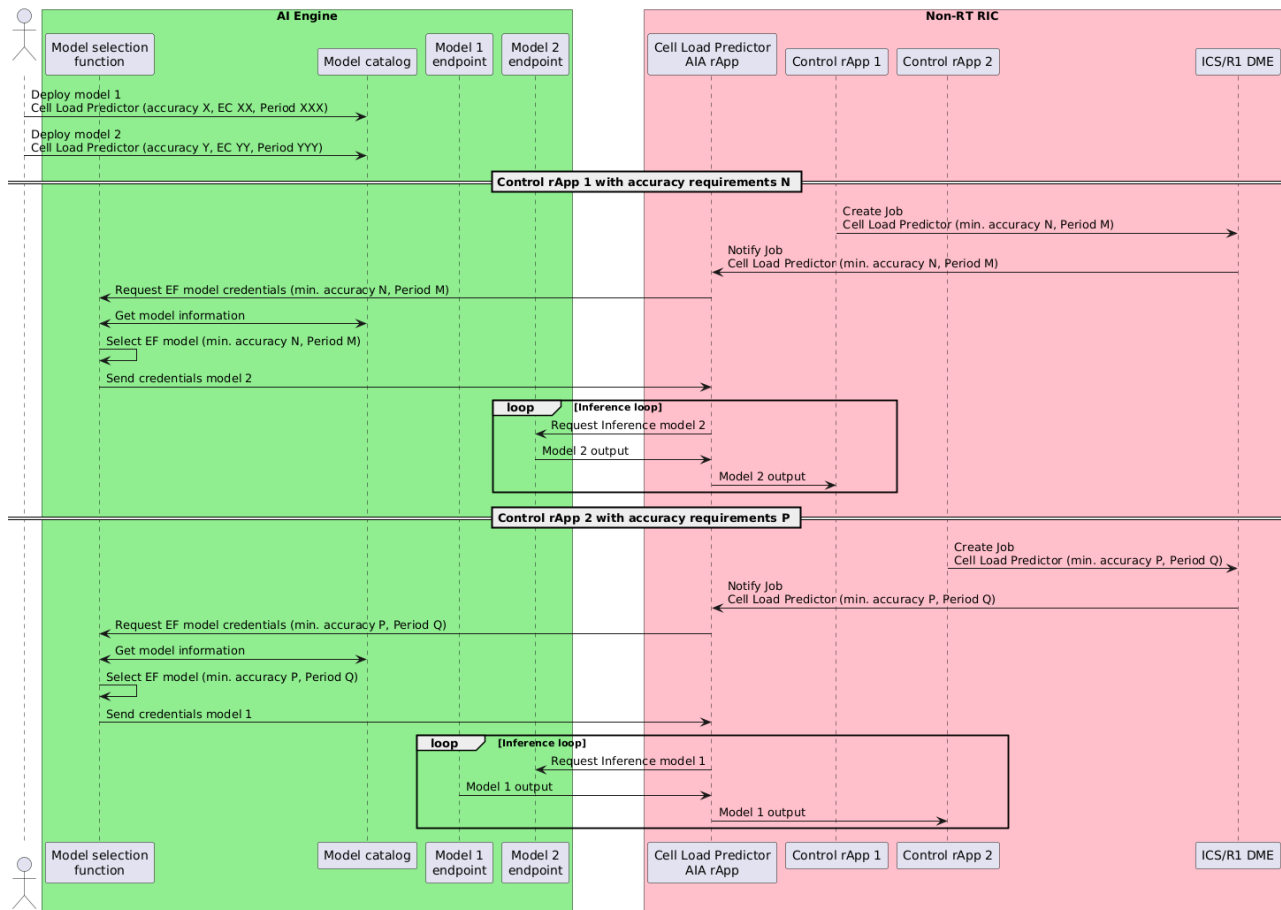


Figure 2-18: PoC1-UC4: Model selection function

While work reported in deliverable D4.3 [3] was mainly focused on reducing the number of model features, in this use case we considered the impact of changing the time granularity of the model, which determines the size of the dataset. In particular, we compared accuracy and energy consumption of a model trained using the maximum data frequency (i.e., 1 second in our case) with models using higher periods. Also, we analysed the trade-off when using the model with average values compared to the averaging the outputs of the model with the maximum data frequency. Table 2-3 shows the results for the model used in UC2 focused on throughput prediction for determining the status of the capacity cell.

Period (s)	Model with period granularity				Average of 1s Model
	Dataset reduction (%)	Training time (s / % reduction)	R2	MAE (%)	MAE (%)
1	0	59 / 0%	0.83	4.79	4.79
13	92.3	35 / 40%	0.84	5.3	6.01
30	96.5	31 / 47%	0.86	4.79	4.92
88	98.9	30 / 49%	0.88	4.68	4.26
263	99.6	19 / 68%	0.91	4.52	4.48

First, the dataset reduction achieved by increasing the input granularity significantly shortened the training time (by 40% to 68) on the server used for training. Since the CPU frequency was fixed at 2 GHz across all cases, this reduction in training time can be directly associated with a corresponding reduction in energy consumption. Secondly, increasing the input period helps to smooth out PRB fluctuations, which are typically harder to predict and can introduce noise. This leads to simpler, more stable traffic trends that result in improved prediction accuracy, as observed in 88 and 253 cases. Notably, these are also the two scenarios in which averaging the outputs of the 1s-resolution model yielded better MAE results, since the output averaging process similarly filters out short-term fluctuations.

Next figures show the impact of inferencing these models on energy consumption, resource consumption and latency, according to the number of simultaneous model consumers. Note that the CPU frequency of the server hosting the AI Engine was fixed to 2GHz, to obtain results which do not depend on the CPU governor of the Operating System. As shown in the figures, only the model using the one-second period leads to a significant increase in power consumption and CPU usage with the number of consumers. This impacts also the average inference delay per consumer, which with 25 consumers already overcomes the requested period of 1 second. The model with a period of 13s leads to some notable increase of resource consumption, but it doesn't impact the serving delay for the analysed of consumers.

These results, together with the ones related to model training benchmarking, highlight that using the 1s model when the cell on/off switching periodicity is higher than 1s, does not bring any benefit.




Figure 2-19 PoC1-UC4: Energy consumption during model inference




Figure 2-20: PoC1-UC4: CPU usage during model inference




Figure 2-21: PoC1-UC4: CPU usage during model inference

3 PoC2 Sensing-assisted communications using RIS

This PoC aims to demonstrate how integrated sensing and communication (ISAC) can be used to obtain user position, which later can be used for energy savings in future mobile networks.

There are two main entities that play a role in **PoC2**: the Sub-6 sensing system, and the RIS board.

3.1 PoC2 updated description

The scenarios considered for **PoC2** are tested in an anechoic chamber in order to obtain representative results.

For **PoC2 – use case 1** – as included in deliverable D5.1 [1], we test standalone sensing by setting up a sensing demo including only the Sub-6 sensing system (without the RIS). This use case has been initially showcased at EuCNC'24 and later an improved version on EuCNC'25 conference (see Figure 3-1), and the results associated to it have been already captured in deliverable D3.2.

In D5.2 [9] the basic functionality was demonstrated and validated. In D5.3, the quantitative evaluation of the system was performed to estimate the KPIs regarding range and angular precision.



Figure 3-1 Sub-6 ISAC testbed @ EuCNC'25 BeGREEN Exhibitor

In **PoC2** (Figure 3-2) – **use case 2** – NEC's RIS board serves as the reflective surface in a scenario where an obstacle separates two adjacent RUs that cover different areas. We consider both Sub-6 RUs having ISAC support that can provide sensing metrics related to the position and density of the users in the scenario. The RIS board is placed aside the RUs and permits control over its reflection properties.

Nevertheless, due to the nature of the RIS which was available, slight changes to this scenario needed to be made, to validate the sensing through the RIS approach.




Figure 3-2: Sensing-assisted communications demo - use case 2

Finally, **PoC2** (Figure 3-3) – **use case 3 - beam steering/tracking**: In use case 3, we extend use case 2, adding mobility to the UE that we offloaded from RU 2 to RU 1. We show how users' position/spatial density sensing data can be used to reconfigure the reflection properties of the RIS. Moving users can still be covered, and the RU can maintain its sleeping state, leading to substantial energy savings. Figure 3-3 shows use case 3 where we extend use case 2 and add mobility to the users of RU 2 that are offloaded to RU 1.

The use cases previously described will be evaluated by means of different KPIs. Use case 3 will be evaluated only theoretically, from the point of view of the sensing accuracy. The main limiting factor is the complexity of the setup and limited availability of equipment, i.e., RIS panels and software defined radios (SDRs).




Figure 3-3: Sensing-assisted communications demo - use case 3

3.2 ISAC final improvements and KPI estimation

The initial working version of the ISAC system was shown in deliverable D5.2 [9]. Additional improvements were required to upgrade the ISAC system to the final one reported in this document.

One of the few issues was the need of manual background modelling. The background modelling is required to remove the static clutter and background, thus being able to distinguish between static and moving objects. Initially, the background was modelled and used during the operation of the ISAC system, assuming that it will stay the same for an extended period of time. Nevertheless, as noticed on different conferences and events where the demo was shown, this is usually not the case. Namely, the background can change over time and then remain the same for some period. This is due to relocation of some objects, furniture,

etc. Therefore, an adaptive algorithm for background modelling was developed and implemented. Using this algorithm, performing of background modelling periodically (manually) is not needed anymore.

Further improvements were made by adding an additional PA to the transmitter. Namely, the output power of the used SDRs in the 5 GHz band is less than 15 dBm. Additionally, there is a cable to the transmit antenna which can add up to 3 dB of loss, making the output power not more than 12 dBm. This is far below the legal limits and greatly limits the sensing range. Therefore, an external PA was added to the system, pushing the output power to the legal limit and enabling sensing in a larger area.

To estimate the resolution for detection of a person, i.e., a potential user, experiments in an anechoic chamber were performed. These experiments were performed with one and two persons standing in front of the sensing system. The system generated a heat map as shown in Figure 3-4 and Figure 3-5.




Figure 3-4: Heat map of a scenario with 2 persons standing in front of the ISAC system in anechoic chamber




Figure 3-5: Heat map of a scenario with a single person standing in front of the ISAC system in anechoic chamber

From Figure 3-4 and Figure 3-5 the **range resolution** of the system was estimated to be approximately 0.7 meters. While this largely depends on the specific scenario, the collected data and the theoretically obtained values suggest that it is a realistic estimate. Based on average human body dimensions [10], the range resolution is sufficient to distinguish multiple individuals in close proximity. Only in specific scenarios this will not be possible, like for example in crowded trains, busses or, eventually, concerts. In these cases, a heuristic approach can be used where a given density can be assigned to areas which are densely populated with individuals. This range resolution enables highly accurate detection of the number of individuals, with precision better than 90%.

The **angular resolution**, on the other hand, is relatively limited due to the small number of antennas used, in this case only 8. The beam width of the phased antenna array is approximately 14 degrees. This means that a single individual can be seen under angle of 14 degrees or larger, depending on the distance to the sensing system. Therefore, in the angular dimension, the system will have a limited resolution. This can be an issue if two individuals are at the same distance from the ISAC system, close to each other. In this case additional analysis can be used to estimate the return power and, based on that, to estimate the number of individuals. Nevertheless, this is out of the scope of the project and was not investigated.

3.3 Sensing data exposure to RIC

As shown in BeGREEN deliverable D4.3 [3] in section 2.1.5, the BeGREEN ISAC system has been provided with E2 Node capabilities and this implementation has been deployed in a real environment as a PoC in a radar-based localization scenario at the Exhibition area of EuCNC'25, as it is show in Figure 3-6.




Figure 3-6 Integration of Sub-6 GHz ISAC Prototype with sensing SM @ EuCNC'25.

The setup in Figure 3-6 presents the sensing system comprising four USRP N321 devices, in addition to a PC that is responsible for the transmission, reception, and processing of the data. Once processed, this data is sent to the E2 agent, which is responsible for communicating with the RIC itself. The data is transmitted through the E2 interface using a novel sensing Service Model (SM) dedicated to sensing data. Next, the near-RT RIC is responsible for orchestrating the collection of this information in real-time via an xApp, which is responsible of storing this data in a database for later analysis and visualization. Finally, using Grafana and querying the database, the heatmap generated by the xApp itself can be visualized.

3.4 RIS-assisted ISAC

To showcase PoC2 use cases, the technologies developed in WP3 were integrated in the context of WP5 and evaluated. In this WP, concerning PoC2, both the developed RIS (from NEC) and the ISAC (from IHP) systems have being integrated together in a single system. With this approach, a single application can control both the RIS and the ISAC systems.

To be able to perform sensing behind an obstacle, the sensing system is configured to perform sensing over the RIS. The main idea is to transmit a signal through the RIS and to receive a reflection from an object, again through the RIS. To allow performing sensing a wider area, the RIS is controlled to adjust the angle between the transmitter and the receiver.

The system was integrated and tested initially in a controlled environment, such as an anechoic chamber. This was necessary, since the developed system works in the 5 GHz Industrial Scientific and Medical (ISM) band, and strong multipath is expected in an indoor environment. Having a strong multipath would make hard drawing any meaningful conclusions and to characterize the behaviour of the system.

The setup for testing the integrated RIS/ISAC system is shown in Figure 3-7. The blue triangles represent the absorbers in the anechoic chamber. The “SDR sensing node” marked in the figure, is the developed ISAC system using SDRs. The RIS is placed in front of the ISAC sensing system at a distance of 3.4 meters. Additionally, a corner reflector is placed a bit off side, hidden behind a wall made of absorbing material and completely undetectable by the sensing system. A photo of the setup is shown in Figure 3-8 and Figure 3-9.

This setup was initially tested and the corner reflectors, behind the absorber wall were not detected by the sensing system.

To verify the functionality of the RIS, it was configured to reflect the signal back to the ISAC. Therefore, the signal which was transmitted from the ISAC system was reflected back. This reflection produced a strong signal in the heat map indicating that the RIS is reflecting the signal in the right direction. This strong reflection can be seen in Figure 3-10.

With the initial tests in D5.2, detection of the reflectors behind an obstacle, in this case a wall made of absorbers, using the RIS was not possible. A few different issues were suspected as the leading cause. It was suspected that the transmit power of the radio was not sufficient, due to the large path loss, as well as the loss due to the RIS and the reflectors. In order to solve this, a power amplifier and a phased array antenna was used in order to achieve high equivalent isotropic radiated power (EIRP). This was immediately to be noticed as a relatively strong reflection in the heat map used to represent the position of the incoming reflections. One such heat map is shown in Figure 3-11.




Figure 3-7: Anechoic chamber configuration




Figure 3-8: Anechoic chamber setup, view from the ISAC system towards RIS




Figure 3-9: Anechoic chamber setup, view from RIS towards ISAC system and corner reflector




Figure 3-10: Heat maps for a) 80, b) 85, c) 90, d) 95 degrees reflection angles with respect to the RIS surface

The use of additional PAs did increase the detection range and the returned signal strength, but it did not solve the issues with the detection of the corner reflectors behind the obstacle. With a thorough investigation it was concluded that the issue lies in the RIS itself. Namely, the RIS type used was only capable of reflecting the signal in one direction but not back to the ISAC system. The RIS is capable to reflect the signal in both directions, but in time division duplex (TDD) mode. Nevertheless, using TDD mode for sensing will not be sufficient to perform sensing. To solve this issue, an additional RIS panel was needed for simulations reflection in the return direction. Since no additional RIS was available, the tests were performed in a slightly different way.

To showcase sensing through the RIS, a setup as in Figure 3-12 was proposed and built. Namely, the transmit antenna is installed on a pole and placed approx. 2 meters from the receive antenna. The transmit antenna is pointed towards the RIS, and can be assumed that it represents the corner reflector, reflecting back the signal towards the RIS. The angle between the transmit antenna and RIS and the receive antenna and RIS is approximately 30 degrees.




Figure 3-11: Stronger reflection received back with use of PA and high gain antenna




Figure 3-12: Alternative setup for sensing over RIS PoC

To test the possibility of sensing through the RIS, the sensing is started and the RIS is controlled in such a way that it scans the area by changing the angle between the incident and the transmitted wave. The setup for this experiment is given in Figure 3-14 and in Figure 3-13 are given the results. It can be seen that the reflection for the angle of 25 degrees is the strongest. A small error in the estimated angle is to be noticed (the angle should be 30 degrees), and this is probably due to small misalignments in the setup. Existence of reflections for the other angles is due to the reflections coming from the stand on which the RIS is mounted as well as from the RIS, since it is a non-ideal device.




Figure 3-13: Sensing through the RIS - obtained heat maps




Figure 3-14: Setup for sensing through the RIS

3.5 PoC2 conclusion

In this PoC we have shown that the sensing functionality can be used to detect potential users for the sake of improving the network coverage and optimal assignment of network resources, leading to improvement of energy efficiency of the network. Non-3GPP wireless access technologies, such as the BeGREEN Sub-6 ISAC system, can provide useful sensing data that can be conveyed to the RIC element.

PoC2 has proven as well that performing sensing through the RIS is possible. The obtained results are probably not ideal, due to the slightly different carrier frequencies of the RIS and the ISAC system, as well as the slightly different channel bandwidths (see Unforeseen Risks 7 and 10, Table 5-1 [9]). Nevertheless, it was observed that the RIS can be used to sense potential users located behind an obstacle.

Table 3-1: Updated Planning of PoC2

PoC 2 UC	Objective	Tasks	Description	Month
UC1 (v1)	Sub-6 ISAC system development	Development of Sub-6 HW	Sensing demo with the sub-6 GHz ISAC system. To be used as a prerequisite for the sensing assisted communications demo with RIS.	DONE M16
		Development of Sub-6 HW		DONE M17
		Initial test of Sub-6 system in the lab		DONE M18
UC2 (v2)	Sensing-assisted demo using the RIS board for coverage extension	Setting up the necessary demo equipment for RAN functionality and sensing (Sub-6 ISAC system and RIS)	The sensing metrics from RU 1 and 2 are pushed to an accessible database. It is possible to retrieve the position of potential users from the sensing metrics	DONE M25 - M27
		Ensure that the sensing metrics can be retrieved		
		Calibrate the RIS and reconfigure it to work in the area of the scenario	An rApp reconfigures the reflection properties of a RIS to offload a UE in the coverage area of RU 2 and puts it to sleep mode to save energy.	DONE M27 - M30
		Implementation of the rApp		

UC3 (v3)	Beam tracking by reconfiguring the reflection properties of the RIS	The available RIS does not support angle reciprocity. Therefore, implementing this UC and showing a demo is not possible. A minimum of 2 RISs are needed, and they were not available. Additionally, the carrier frequency of the RIS is slightly offset to the carrier frequency of the ISAC system. The available bandwidth is also much smaller compared to the bandwidth of the ISAC. This was also a limiting factor that led to implementing only parts of this use case and not integrating all of them.	Periodically reconfigure the RIS reflection properties.	Partially DONE M28 - M30
-----------------	---	---	---	-----------------------------------

KPI ID	KPI Technical Description
KPI-04.1	Precision of the developed sensing algorithm for detecting potential users
KPI-04.2	Sensing assisted beam search – 20% performance improvement with respect to extensive search and hierarchical search.
KPI-04.3	Detection of users in order to estimate the presumed network load – at least 50% accuracy of estimation of potential mobile users

Measurement	Measurement description	BeGREEN KPI	Unit
Sensing accuracy	Sensing accuracy refers to the precision and reliability with which a sensor or measurement system can detect and report the actual value of the parameter it is designed to measure. Sensing accuracy is crucial in various applications such as scientific research, industrial monitoring, and healthcare. It ensures that the data collected is trustworthy and can be used for decision-making and control.	KPI-04.2, KPI-04.3	m

4 PoC3 Energy-efficient CU and O-RAN RIC

PoC3 focused on implementing and validating energy-efficient enhancements to the Central Unit (CU) and the RIC by leveraging low-power ARM-based platforms and advanced packet processing techniques. The primary objective was to demonstrate that carrier-grade RAN software, such as Accelleran's dRAX components, can operate efficiently on ARM servers without compromising functionality or performance. To this end, the CU and RIC components were ported from x86 to ARM, enabling more sustainable RAN deployments tailored for edge environments or private networks. Figure 4-1 presents the PoC3 set up disclosing the servers used.




Figure 4-1: PoC3 set-up implementation

Beyond porting, PoC3 introduced performance optimisations through the integration of XDP in the CU, offering accelerated packet handling and reduced CPU usage. These enhancements were thoroughly tested using traffic scenarios that reflect real-world operating conditions, ensuring that the ARM-based CU could meet the demands of various User Equipment (UE) loads. Additionally, testing frameworks and visualisation tools, such as Grafana dashboards, were developed to monitor and validate the RIC's behaviour and responsiveness on the ARM platform.

The PoC culminated in a full-system integration and demonstration, combining the optimised CU and RIC components into a complete RAN setup with single and multiple UE traffic. The outcomes of the demonstration confirmed the feasibility of ARM-based O-RAN deployments, with clear indications of energy-saving potential and scalable performance. The work concluded with a comprehensive result analysis, paving the way for further standardisation and industrial adoption of energy-aware RAN implementations.

Initial measurements confirmed that the CU could operate effectively on both x86 and ARM servers under realistic RAN conditions. In single UE scenarios, x86 showed approximately 22% better energy efficiency, making it suitable for light-load environments as discussed on D3.2 section 2.3 [7]. However, when scaling to higher UE densities, the ARM platform outperformed x86, achieving up to 10% energy savings and demonstrating better scalability under sustained traffic as discussed in D3.2 Section 2.3 [7]. This validated ARM's suitability for dense deployments where cumulative energy consumption becomes critical.

Performance testing of the Near-RT RIC on both architectures revealed distinct behavioural differences. ARM servers exhibited significantly lower average power consumption (up to 33% less) during computational tasks such as matrix multiplication and hashing, with consistent scaling across cores, as discussed in D3.3 section 2.3 [8], and as show in Figure 4-2. Conversely, x86 systems delivered lower latency and better performance

for small, time-sensitive operations. These trade-offs highlight a design space where ARM can support energy-sensitive functions, while x86 remains advantageous for ultra-low-latency workloads.




Figure 4-2: Average power consumption comparison for ARM and x86

Integrated evaluation of energy-saving algorithms (PoC1) on the PoC3 testbed confirmed that ARM-based RIC setups react more dynamically to A1 policies, with power consumption varying in direct correlation to policy changes as discussed in D3.3 section 3.1 [8]. In contrast, x86 power usage remained relatively static, even under heavy reconfiguration by xApps such as SHO or TGW. Under full energy saving policies, ARM servers showed up to 5% additional savings, reinforcing their potential in AI-driven, policy-controlled RAN scenarios. This was also presented in the MWC'25 and EuCNC'25 demos.

Table 4-1: Objectives, tasks and results of PoC3

Objective	Tasks	Description	Month	Results
CU platform & testing	- Extend CU to ARM platform - - Develop XDP function for CU - CU testing with traffic scenarios	Implementation and performance enhancement of the dRAX CU on ARM servers, including XDP-based acceleration features. A complete set of traffic tests was executed to validate its behaviour under different load conditions.	DONE M20	CU successfully ported to ARM, demonstrating functional parity with x86. XDP integrated for accelerated packet processing. ARM outperformed x86 in high-traffic scenarios, achieving up to 10% energy savings. In low-traffic (1 UE), x86 remained more efficient by ~22%, and showcased D3.2 Section [7].
RIC platform & testing	- Extend RIC to ARM platform - - Develop RIC testing and visualisation tools	Porting of the dRAX RIC to ARM platforms followed by development of a tailored testing framework, including algorithms and a Grafana dashboard for data visualisation.	DONE M23	RIC ported to ARM with testing xApp for performance and latency benchmarking. ARM showed ~33% lower average power consumption, while x86 showed lower latency and better small-task performance and showcased on D3.3 Section 2.3 [8].
Full demo integration	Integration and multi-UE support	Integration of all enhanced capabilities including HW acceleration and validation with both single and multiple UE scenarios to ensure end-to-end performance.	DONE M25	Full system validated with integrated CU, RIC, and energy-saving xApps. Tests confirmed scalability under multi-UE loads and system stability across architectures. Demonstrated at MWC25 and EuCNC'25 and presented in D3.3 [8].

5 PoC4 Energy-efficient DU implementation using HW acceleration

In the final DU acceleration demo, we compare the power consumption of a sphere decoder running on CPU and on GPU. In general, a MIMO link utilizes the spatial diversity in the channels for increasing throughput by multiplexing several spatial streams (also called layers) over the same time/frequency resources. This works well in case the spatial correlation between the layers is low and then a simple linear receiver shown in Figure 5-1 may be used. However, when this is not the case, using a linear receiver may result in poor performance and reduced throughput. A sphere decoder shown in Figure 5-2 attempts to find the symbols which were most likely transmitted by the UE, given the received signal. It can achieve good performance also when spatial correlation is higher. This is done by looking at the multidimensional lattice created by the possible transmitted symbols for each layer. If we go over all combinations that would amount to a very large number. For example, in this PoC the constellation is 64QAM, the number of layers is 4 and the number of RBs is 10. For this case we will need to go through $120 \cdot 64^4 = 200$ million different combinations, which would be unfeasible. With the sphere decoder we only look at the combinations contained within a sphere with a configurable radius. For the case of the PoC this reduces the number of combinations we check to one million. An illustration for this concept for one layer is shown in Figure 5-3.

Note that one million combinations are still a very large number which requires significant resources and power. To improve it, we can perform parallelization over subcarriers and also within the subcarriers. CPU has some parallelization capability, but it is limited. On the other hand, GPU has a much larger parallelization capability and the purpose of this PoC is to check how this capability can be used for reducing power consumption. We have implemented the sphere decoder in both CPU and GPU and then in this demo we compare the power consumption for these two implementations as shown in Figure 5-4.




Figure 5-1: Linear MIMO receiver




Figure 5-2: Sphere Decoder




Figure 5-3: Sphere Decoder Illustration




Figure 5-4: Implementations being tested in the PoC

We use a power distribution unit (PDU) to measure the power of each implementation as shown in Figure 5-5.




Figure 5-5: PoC setup

The sphere decoder advantage in wireless performance for the scenario being tested in the PoC is shown in Figure 5-6, where we can see a 2-dB improvement over the basic linear receiver for the case where we have correlation of 0.03 for the Tx antennas and the Rx antennas. This applies to both the CPU and the GPU implementations. As for power consumption, when running the 120 Resource Blocks (RBs) sphere decoder scenario on CPU the power consumption is 26W, and when running on GPU the power consumption is 24W, which is similar. However, the running times are dramatically different. It takes the CPU 5.5 ms to complete the task where the GPU completes it in 0.5 ms. Hence, the GPU energy score is 100 kbit/joule and the CPU energy score is 9 kbit/joule, resulting in a 10x improvement. To summarize, this demo shows that for complex L1 tasks, such as the sphere decoder, GPUs have the potential to significantly reduce power consumption.




Figure 5-6: Sphere Decoder wireless performance

6 PoC5: 6G RU PA blanking test

The purpose of this PoC was to test the power saving that can be achieved in the 6G RU using the PA blanking algorithm. This algorithm has been explained in detail in BeGREEN D3.3 paragraph 3.2 [8], and it involves interrelated processes as follows:

- a. Checking the downlink stream arriving from the DU to the RU to verify if there are symbols that do not contain data and if the answer is positive the RU PA is turned off.
- b. Configuring the DU Scheduler in such a way that it will maximise the number of symbols that do not contain data.

The PA blanking module has been tested in the following scenarios:

1. Test 1: **Indoor test over cables** (no RF transmission) at EuCNC 2024 in Antwerp, Belgium.
2. Test 2: **Indoor test over the air** at RunEL Lab – September 2024, Rishon Lezion, Israel.
3. Test 3: **Final Outdoor test over the air** at Brunel University – November 2024- London, UK.

The results of Test 1 and Test 2 have been reported as well in deliverable D3.3 paragraph 3.2 [8]. This chapter refers to the results obtained from PoC5 Test-3: Final Outdoor Test over the air.

In this test the RU was installed in the lab pointing outside the Brunel University, as shown in Figure 6-1. The RU was operating at 3.5GHz with 100MHz bandwidth, and it was connected to a server that contains the DU + CU + 5G Core (5GC). The RU power consumption was tested with and without the PA blanking algorithm in four different use cases, simulating: a) no traffic; b) low traffic; c) average traffic and d) peak traffic. The results of the measurements from the four use cases are shown in Table 6-1.



Figure 6-1: RunEL RU installed in the lab pointing outside the Brunel University

Table 6-1: Measurement results of the outdoor testes at Brunel University

Use Case	DL Throughput (Mbps)	PA Current with PA Blanking (mAmps)	PA Current with PA Blanking (mAmps)	RU PAs Power with PA Blanking (Watts)	RU PAs power without PA Blanking (Watts)
No Traffic	No Data	21	240	0.42	4.8
Low Traffic	50	25	241	0.5	4.82
Average Traffic	100	60	240	1.2	4.8
Peak Traffic	1000	243	361	4.86	7.22

The results of the measured data from Table 6-1 have been analysed assuming that, in an average day, the cellular traffic contains approximately 6 hours of low traffic, 6 hours peak traffic and 12 hours of average traffic, as can be shown in Figure 6-2, representing the cellular traffic behavior in a dense urban area in Western Europe.

**Figure 6-2: Hourly traffic in an average day**

The analysis of the results shows that without the PA Blanking module, the power consumptions of the PAs for the small cell is 5.41 Watts. With the PA Blanking module, the power consumption of the PAs for small cell is 1.94 Watts. The power saving using PA blanking module is 3.47 Watts. We conclude that the PA Blanking Module saves 64% of the power amplifier power consumption. A video of the Final PA Blanking Test that was displayed at BeGREEN stand at EuCNC25 in Poznan shows some additional details of the tests⁷.

⁷ PoC5 PA Blanking video @ EuCNC'25

7 Summary and Conclusions

The **BeGREEN** project has successfully demonstrated a comprehensive suite of innovations aimed at enhancing energy efficiency across multiple layers of the Radio Access Network (RAN) architecture. Through five distinct Proof-of-Concepts (PoCs), the project has validated the feasibility, performance, and impact of AI-driven control mechanisms, HW optimisations, and integrated sensing technologies in real-world and emulated environments.

PoC1 showcased the Intelligence Plane as a cornerstone of **BeGREEN**'s architecture, enabling dynamic and intelligent control of RAN components. The use cases demonstrated the deployment of machine learning models for cell on/off switching, conflict management between xApps, and model selection strategies. The AI Engine, integrated with Non-RT and Near-RT RICs, proved capable of executing real-time decisions based on predictive analytics, achieving significant energy savings while maintaining service quality. The validation of the ML-driven control loop, particularly under varying QoS thresholds, highlighted the robustness and adaptability of the system. Moreover, the Conflict Manager ensured harmonious operation between energy-saving and mobility functions, reinforcing the importance of coordinated multi-objective optimization in future O-RAN deployments.

PoC2 explored the potential of Integrated Sensing and Communication (ISAC) technologies, particularly in conjunction with Reconfigurable Intelligent Surfaces (RIS). The project demonstrated that Sub-6 GHz sensing systems can accurately detect user presence and density, with range resolution below one meter and angular resolution sufficient for most deployment scenarios. The RIS-assisted sensing experiments, although constrained by HW limitations, confirmed the feasibility of detecting users behind obstacles, paving the way for energy-aware coverage strategies. These findings underscore the value of ISAC in enabling context-aware network operations and resource allocation.

PoC3 validated the deployment of Central Units (CUs) and RICs on low-power ARM platforms, demonstrating that carrier-grade RAN software can operate efficiently without compromising performance. The integration of XDP for packet acceleration and the use of Grafana-based monitoring tools enabled detailed performance analysis. Results showed that ARM platforms offer superior energy efficiency under high traffic loads, while x86 remains optimal for latency-sensitive tasks. These insights provide a clear roadmap for hybrid deployments, where energy and performance trade-offs can be balanced according to network demands.

PoC4 focused on Layer 1 processing enhancements, specifically the implementation of a sphere decoder on GPU versus CPU. The GPU-based solution achieved a tenfold improvement in energy efficiency, demonstrating the critical role of parallelization in reducing power consumption for complex signal processing tasks. This PoC highlights the importance of HW acceleration in future 6G systems, where computational demands will continue to rise.

PoC5 addressed energy savings at the PHY layer through Power Amplifier (PA) blanking. The final outdoor tests confirmed that the PA blanking algorithm can reduce power consumption by up to 64%, with minimal impact on performance. This technique, when combined with intelligent scheduling at the DU, offers a practical and impactful method for reducing energy usage in dense urban deployments.

The **BeGREEN** project has delivered a robust and multi-faceted framework for energy-efficient RAN design, integrating AI, sensing, and HW innovations. Each PoC has contributed unique insights and validated technologies that collectively advance the state-of-the-art in sustainable mobile networks. The results presented in D5.3 not only demonstrate technical feasibility but also provide a foundation for future standardization and industrial adoption.

As mobile networks evolve toward 6G, the principles and solutions developed in **BeGREEN** will be instrumental in shaping architectures that are not only high-performing but also environmentally responsible. The project's emphasis on modularity, interoperability, and real-world validation ensures that its outcomes

are both scalable and adaptable to diverse deployment scenarios.

BeGREEN has proven that energy efficiency can be achieved without compromising user experience, and that intelligent, context-aware control is key to unlocking the full potential of next-generation networks.

8 Bibliography

All BeGREEN deliverables are online at <https://www.sns-begreen.com/deliverables>.

- [1] BeGREEN D5.1, “Use Case Identification and Demonstration Plan”, December 2023.
- [2] O-RAN Alliance Working Group 1 (Use Cases and Overall Architecture), “O-RAN Use Cases Analysis Report”, O-RAN.WG1.TR.Use-Cases-Analysis-Report-R004-v17.00, June 2025.
- [3] BeGREEN D4.3, “Final architecture and evaluation of BeGREEN O-RAN Intelligence Plane, and AI/ML algorithms for NFV user-plane and edge service control energy efficiency optimization”, May 2025.
- [4] BeGREEN D4.2, “Initial evaluation of BeGREEN O-RAN Intelligence Plane, and AI/ML algorithms for NFV user-plane and edge service control energy efficiency optimization”, October 2024.
- [5] BeGREEN D2.3, “Energy efficient RAN architecture and strategies”, March 2025.
- [6] BeGREEN D4.4, “Datasets Used to Train AI/ML Models”, June 2025.
- [7] BeGREEN D3.2, “Initial development and evaluation of the proposed enhancements and optimization strategies”, September 2024.
- [8] BeGREEN D3.3, “Final evaluation and benchmarking of the implemented solutions”, March 2025.
- [9] BeGREEN D5.2, “Solution integration and validation”, March 2025.
- [10] H. W. Jürgens, I. Matzdorff, J. Windberg, “International Anthropometric Data. International Anthropometric Data for Work-Place and Machinery Design”, 1. Auflage. Bremerhaven: Wirtschaftsverlag NW Verlag für neue Wissenschaft GmbH, 1998.