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

Deliverable 6.1

Transport domain intermediate developments

Co-funded by the European Union

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

Project Details

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

Deliverable Details

Deliverable WP:	WP6
Deliverable Identifier:	D6.1
Deliverable Title:	Transport domain intermediate developments
Editor(s):	LINKS
Author(s):	Daniele Brevi (LINKS), Edoardo Bonetto (LINKS), Federica Schena (LINKS), Aruna Prem Bianzino (UC3M), Stefano Giolito (Bellantenna), Achille Montanaro (Bellantenna), Gabriele Scivoletto (NXW), Daniele Ronzani (HPE), Mauro Agus (TIM), Alessandrto Trogolo (TIM), Dionysia Triantafyllopoulou, Shahab Ehsanfar, Hefdh Allah Sakran, Timo Pech, David Kühnert, Stephan Enhuber (TUC), Orestis Tsolakis (CERTH), Robert Gdowski, Maria Safianowska (ISRD), Nina Slamnik-Kriještorac, Raul Cuervo Bello, Xhulio Limani (IMEC), Andreas Georgakopoulos (WINGS), Dionysios Deligiannopoulos (ACROMOVE), Antonis Kalis (ACROMOVE), Orestis Tsolakis (CERTH), George Barboutidis (CERTH), Sofoklis Dais (CERTH), Vasileios Giannoudis (CERTH), Aikaterini Batzou (CERTH)
Reviewer(s):	Timo Sukuvaara (FMI), Eirini Tserga (ThPA S.A.), Christos Papadopoulos (ThPA S.A.)
Submission Date:	29/05/2026
Dissemination Level:	PU

Disclaimer
The information and views set out in this deliverable are those of the author(s) and do not necessarily reflect the official opinion of the European Union. Neither the European Union institutions and bodies nor any person acting on their behalf may be held responsible for the use which may be made of the information contained therein.

Table of Contents

List of Figures

List of Tables

List of Acronyms and Abbreviations

TERM	DESCRIPTION
3GPP	3rd Generation Partnership Project
5GC	5G Core Network
ADF	Application and Data Fusion Layer
AFS	Autonomous Fitness Score
AI	Artificial Intelligence
AI/ML	Artificial Intelligence / Machine Learning
AMF	Access and Mobility Management Function
AMR	Autonomous Mobile Robot
ANW	Aruba Networking
APN	Access Point Name
API	Application Programming Interface
AUSF	Authentication Server Function
B5G	Beyond 5G
B5G/6G	Beyond 5G / Sixth Generation
CAAS	Compute-as-a-Service
CAM	Cooperative Awareness Message
CAMARA	CAMARA network API initiative
CAPIF	Common API Framework
CDT	Compute Digital Twin
CEF	Compute Exposure Function
CIR	Channel Impulse Response
C-ITS	Cooperative Intelligent Transport Systems
CLE	Cooperative Localization Engine
CN	Core Network

CNC	Core Network Controller
CNN	Convolutional Neural Network
COM	Communication Layer
COTS	Commercial Off-The-Shelf
CP	Control Plane
CPM	Collective Perception Message
CRA	Compute Resource Allocation
CRI	Context Risk Index
CRO	Compute Resource Orchestration
CSO	Centralised Service Orchestrator
C-V2X	Cellular Vehicle-to-Everything
DL	Downlink
DMRS	Demodulation Reference Signal
DP-REP	Deployment Policy Repository
DRB	Data Radio Bearer
DT	Digital Twin
DT-DC	Digital Twin Data Collection
DT-DMO	Digital Twin Decision Modelling and Optimisation
DT-F/W	Digital Twin Framework
DT-REG	Digital Twin Registry
DT-REP	Digital Twin Repository
DT-SM	Digital Twin Simulation and Modelling
DT-VI	Digital Twin Visualisation and Interaction
DTVE-TSDB	Digital Twin Virtual Entity Time Series Database
DU	Distributed Unit
EC	European Commission
ECPRI	Enhanced Common Public Radio Interface

EMBB	Enhanced Mobile Broadband
ESO	Edge Service Orchestrator
ETSI	European Telecommunications Standards Institute
EU	European Union
FR1	Frequency Range 1
FR2	Frequency Range 2
FRMCS	Future Railway Mobile Communication System
GA	Grant Agreement
GNB	gNodeB
GNSS	Global Navigation Satellite System
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HMI	Human-Machine Interface
HPE	Hewlett Packard Enterprise
IEEE	Institute of Electrical and Electronics Engineers
IMU	Inertial Measurement Unit
IOT	Internet of Things
IP	Internet Protocol
IQ	In-phase and Quadrature
ISAC	Integrated Sensing and Communication
ISRD	IS-Wireless R&D
K8S	Kubernetes
KPI	Key Performance Indicator
KPM	Key Performance Measurement
KVI	Key Value Indicator
LCX	Leaky Coaxial Cable
LIDAR	Light Detection and Ranging

LOS	Line of Sight
MAC	Medium Access Control
MCDM	Multi-Criteria Decision-Making
MCP	Mobility Context Provider
MEC	Multi-access Edge Computing
MIMO	Multiple-Input Multiple-Output
ML	Machine Learning
MLOPS	Machine Learning Operations
MNO	Mobile Network Operator
MOCN	Multi-Operator Core Network
MQF	Message Queue Factory
MRC	Minimal Risk Condition
NDT	Network Digital Twin
NEF	Network Exposure Function
NF	Network Function
NFV	Network Functions Virtualisation
NLOS	Non-Line of Sight
NR	New Radio
NRF	Network Repository Function
NSA	Non-Standalone
NSAAS	Network Slice as a Service
NSSAI	Network Slice Selection Assistance Information
NSSF	Network Slice Selection Function
NSSO	Network Slice Service Orchestrator
NWEF	Network Exposure Function
OAM	Operations, Administration and Maintenance
OBU	On-Board Unit

O-CU	O-RAN Centralised Unit
O-DU	O-RAN Distributed Unit
O-RAN	Open Radio Access Network
O-RU	O-RAN Radio Unit
ODD	Operational Design Domain
OFDM	Orthogonal Frequency-Division Multiplexing
PAE	Position Acquisition Engine
PCF	Policy Control Function
PDU	Protocol Data Unit
PEI	Permanent Equipment Identifier
PFCP	Packet Forwarding Control Protocol
PHY	Physical Layer
PRB	Physical Resource Block
PTP	Precision Time Protocol
PU	Public
QOD	Quality on Demand
QOS	Quality of Service
RAN	Radio Access Network
RCS	Remote Control Station
RDT	Resource Digital Twin
REF	Resource Exposure Function
REST	Representational State Transfer
RF	Radio Frequency
RIC	RAN Intelligent Controller
RLDT	Risk Level Digital Twin
RLC	Radio Link Control
RRC	Radio Resource Control

RSU	Road Side Unit
RT	Real Time
RTK	Real-Time Kinematic
RTT	Round-Trip Time
RU	Radio Unit
SA	Standalone
SBA	Service-Based Architecture
SIM	Subscriber Identity Module
SLA	Service Level Agreement
SMF	Session Management Function
SNS JU	Smart Networks and Services Joint Undertaking
S-NSSAI	Single Network Slice Selection Assistance Information
SPE	ServerPack Edge
SPE5G	ServerPack Edge 5G
SPR	Sensing Processing Layer
SRS	Sounding Reference Signal
SSB	Synchronisation Signal Block
STFT	Short-Time Fourier Transform
STS	Ship-to-Shore
TCP	Transmission Control Protocol
TF	Transformation Function
TFS	Tele-operation Fitness Score
TLS	Transport Layer Security
TNC	Transport Network Controller
TRL	Technology Readiness Level
TUC	Technische Universität Chemnitz
UC	Use Case

UDM	Unified Data Management
UDR	Unified Data Repository
UE	User Equipment
UIC	International Union of Railways
UL	Uplink
UP	User Plane
UPF	User Plane Function
URLLC	Ultra-Reliable Low-Latency Communications
USIM	Universal Subscriber Identity Module
USRP	Universal Software Radio Peripheral
V2X	Vehicle-to-Everything
VAM	VRU Awareness Message
VM	Virtual Machine
VNF	Virtual Network Function
VPN	Virtual Private Network
VRU	Vulnerable Road User
WP	Work Package
XAPP	Near-RT RIC Application
YAML	YAML Ain’t Markup Language
ZSM	Zero-touch network and Service Management

Executive Summary

The WP6 “Transport domain” is responsible for executing the activities related to the definition, design, and implementation of the Transport Use Cases (UCs). In Task 6.1 “Definition and Design”, the UCs' description was detailed, and the system architecture was defined. Furthermore, this task identified the KPIs and KVIs to be further evaluated in WP7. The implementation activities are being carried out in Task 6.2 “Implementation”. A key activity of this task is the integration of the B5G/6G technology enablers, which have been developed in other WPs.

In the AMAZING-6G project, five Transport UCs have been selected as representative scenarios for demonstrating the effectiveness of the developed B5G/6G enablers.

The first UC “Protection of Vulnerable Road Users” focuses on visually impaired people who are alerted about risks in the streets, caused by other road users (e.g., vehicles, bicyclists). The vertical service application of T1 exploits the data provided by infrastructure-side sensors for identifying such risks. The application of this UC relies on the information of the localization enablers providing the position of VRUs, vehicles, and other road actors. The UC further demonstrates how an automated energy-aware orchestration of services can achieve an energy-efficient use of roadside resources, which are powered by a renewable energy source, to increase their operational time. Moreover, the T1 UC exploits network reconfigurability through network slicing for having a dynamic adaptation of the Quality of Service to match the desired network performance.

In the second UC, “Enhancing Urban Security with Autonomous Mobile Robot Monitoring”, the focus is also on alerting visually impaired people about possible dangers. In this case, the warnings are about the presence of obstacles on the sidewalks. An Autonomous Mobile Robot Monitoring (AMR) patrols the sidewalks and provides sensor information to the infrastructure-side application that is in charge of identifying the danger and sending notifications to the VRUs. Like T1 UC, this UC relies on the localization enablers for the position of VRUs, and on the technological enablers that allow for an automated and dynamic reconfiguration of network and computing resources, besides services orchestration. This UC aims indeed to demonstrate how to guarantee the continuity of the monitoring service performed by the AMR, thanks to the offloading of tasks from the AMR to the edge.

The target of the third UC “Wireless Signalling on Rail Tracks” is to enhance the safety of train operations, exploiting the advanced functionalities of B5G/6G technological enablers, such as Integrated Sensing and Communication (ISAC) and network slicing. These enablers are experimented with to complement the traditional hardwired infrastructure offering real-time signalling and obstacle detection functionalities. The UC demonstrates the scenario in which two trains are travelling on the same track, one after the other. The ISAC enabler is used to guarantee the train's integrity and to retrieve the position and speed of the first train, while the communication enablers provide redundant communication links to the following train via the terrestrial B5G/6G network and via a direct connection to the preceding train.

The demonstration of tele-operated driving solutions is the focus of the fourth UC, “Tele-operation as a back-up to autonomous driving”. In this UC, a B5G/6G-enabled communication platform provides a seamless transition between autonomous driving and tele-operated driving operations. The communication platform relies on 6G communication link for achieving the ultra-reliable, low-latency connectivity that is required for exchanging data and control commands between the vehicle and the control center. Furthermore, the UC exploits the network slicing enabler to achieve the desired Quality of Service, and zero-touch management enablers for providing a flexible configuration of application placements in the compute continuum, achieving the best trade-off between latency and context awareness.

The fifth UC “Port logistics and transport operations optimisation and safety” is focused on the improvement of the end-to-end port processes concerning the loading and unloading between vessels and trucks. To this end, massive digital twinning, IoT technologies, and AI/ML enablers are jointly used for supporting real-time data collection and analysis. In addition to these enablers, this UC exploits

network slicing for supporting sensor data sharing, and compute-as-a-service enablers for seamless orchestration in the compute continuum.

The outcome of this deliverable is the intermediate status of the Transport UCs development activities. At this stage, the system design of each UC is finalised, and there is a consolidated system architecture, with the details about the main components and exploited enablers. The functional components of the network infrastructure for each UC have been identified, and the setup of the testbed in the trial site is detailed. The evaluation approach has been defined by introducing the list of KPIs and KVIs to be considered for each UC.

A further outcome of this deliverable is the planning of development and testing activities until the completion of UC demonstration and validation activities. Preliminary integration activities are reported as well in this deliverable for each UC. For each UC, testing and demo activities have started, and preliminary results are included in this deliverable.

1 Introduction

1.1 Purpose of the deliverable

This deliverable reports the outcomes of the activities carried out in Task 6.1 “Definition and Design”. The objective of this task was to define the details of the UCs that have been introduced at the proposal stage. The definition includes the system architecture design of the UCs along with the identification of the needed network infrastructure and technological enablers.

The results of task T6.1 are exploited in the consecutive Task 6.2 “Implementation”, in which activities related to the development of each UC’s services and the integration of the technical enablers are performed. The goal of T6.2 is to provide an end-to-end system that can be used for demonstrating the UCs, validating the technological enablers, and performing the evaluation of the KPIs and KVIs identified in T6.1.

The content of this deliverable reflects the intermediate status of the T6.2 activities, and it anticipates some preliminary testing results from Task 6.3 “Testing” that validate specific UC components and enablers. This document will serve as a reference for the activities, trial execution and trial analysis to be performed in WP7.

1.2 Intended audience

The dissemination level of D6.1 is ‘public’ (PU). The deliverable is available to members of the consortium, the European Commission (EC) Services and those external to the project. This document is primarily intended to provide a basis for the Transport UCs of the AMAZING-6G project for all AMAZING-6G beneficiaries, in particular the partners involved in testing activities in WP6 and in the evaluation activities of WP7. It remains also available to external people who wish to learn about the AMAZING-6G Transport UCs and related developments.

1.3 Structure of the deliverable

The structure of the deliverable is as follows:

1. Section 2 presents the definition and system architecture design for each of the five Transport UCs; each subsection is devoted to a specific UC providing a general description, the deployment and the development details, the exploited network infrastructure and technological enablers, the evaluation targets (i.e., KPIs, KVIs), and the setup of the trial site for the UC.
2. Section 3 introduces the planning of the five UCs;
3. Section 4 provides the status of integration and testing activities.
4. Section 5 reports the conclusions and the next steps.

2 Use Case Definition and Design

This section introduces five transport use cases developed under Work Package (WP) 6. For each use case, the goal, main components, and high-level architecture are highlighted. An overview of the identified functional and non-functional requirements is also included.

The implementation and practical deployment of the use cases are described. The network infrastructure in use is detailed, and its functional components are highlighted. The trial setup is presented, outlining how each use case will be tested and demonstrated.

Overall, the transport use cases address multiple aspects of the vertical, spanning different modes of transportation and functionalities. This demonstrates how innovative B5G/6G features can transform the transport vertical to new achievements in terms of safety and efficiency.

2.1 Transport 1 Use Case – Protection of Vulnerable Road Users

The Transport 1 (T1) Use Case (UC) aims to assist individuals with disabilities, such as visually impaired persons, during their urban journeys. The T1 UC investigates how B5G/6G networks can support advanced applications for this scope and which advanced B5G/6G network features can be exploited by application developers. T1 UC includes the experimentation of roadside infrastructure that is battery-powered and sourced with renewable energy, while optimisation strategies for effectively using these devices by applications will also be investigated.

2.1.1 Description of the use case

The T1 UC focuses on a road intersection where a Road Side Unit (RSU) equipped with perception sensors, such as cameras and LiDARs, is installed. The RSU monitors the intersection using its sensors, and the information from the sensors is used to detect vehicles, pedestrians, and other road users and to track their movements.

Perception information is complemented with data shared by connected vehicles and pedestrians. Some vehicles can be equipped with an On Board Unit (OBU), and they can share their position, and potentially other available data, over the B5G/6G network. Pedestrians may also share their position through their smartphones using standard C-ITS messages via an app.

The information gathered from the road intersection is provided to a Digital Twin that resides on the edge server. The Digital Twin is responsible for estimating the risk level by leveraging all the collected data. The computed risk level provides feedback to the visually impaired people so as to enhance their safety in dangerous situations, such as when crossing the road. A risk notification is sent via B5G/6G network to the smartphone, which alerts the user using vibrations or sounds.

The RSU is powered by a solar panel, and it is equipped with a battery to prolong operating time. However, the RSU will likely run out of power if intensive computation runs 24/7. For this reason, in T1 UC applications, orchestration solutions are implemented. To save energy at the RSU and keep it active, sensor data processing can be offloaded from the RSU to the edge server when needed. The orchestration automatically switches applications between the RSU and the edge server based on policies that consider network, computing, and energy metrics that are continuously monitored. The orchestration preserves the application context when moving them to avoid information gaps.

Figure 1 introduces the high-level architecture of T1 UC. At the bottom, the roadside actors involved in the UC are depicted. The RSU acts as an IoT gateway for the camera and LiDAR that represent the IoT sensors. The RSU is connected to the B5G/6G system through a mm-Wave connection link. Other UEs involved in this UC are the smartphone, which is used by visually impaired people for receiving alert notifications, and 5G OBUs on connected vehicles. These UEs will either use the public or the private network in a seamless way, depending on availability. The B5G/6G system for this UC comprises the 5G public network, the 5G private core network, the mm-wave RAN and the MEC server. The MEC server exposes UEs' location information through the MEC Location API and makes available V2X messages through a dedicated API. Another main component of the architecture of T1 UC is the orchestrator that includes all the enablers which supervise the management and orchestration of applications, computing and network resources. Enablers related to networking aspects are the Network performance monitoring and control enabler, which oversees the monitoring of the overall B5G/6G system, and the Network Digital Twin, which creates a digital replica of the B5G/6G system to be used for simulations and analysis. Further enabler is the Advanced Localization and Positioning Systems enabler, which fuses information from MEC Location API and from V2X messages for performing a cooperative and hybrid localization approach. The interactions of the application layer with the enablers are performed through a set of CAMARA APIs that include Device Location API, Edge Cloud API, Energy Footprint Notification API, QoD API, Connectivity Insights and Connectivity Insights Subscriptions APIs. In T1 UC, the main application is the Risk-level Digital Twin that is responsible for the monitoring of the intersection and the triggering of alerts.

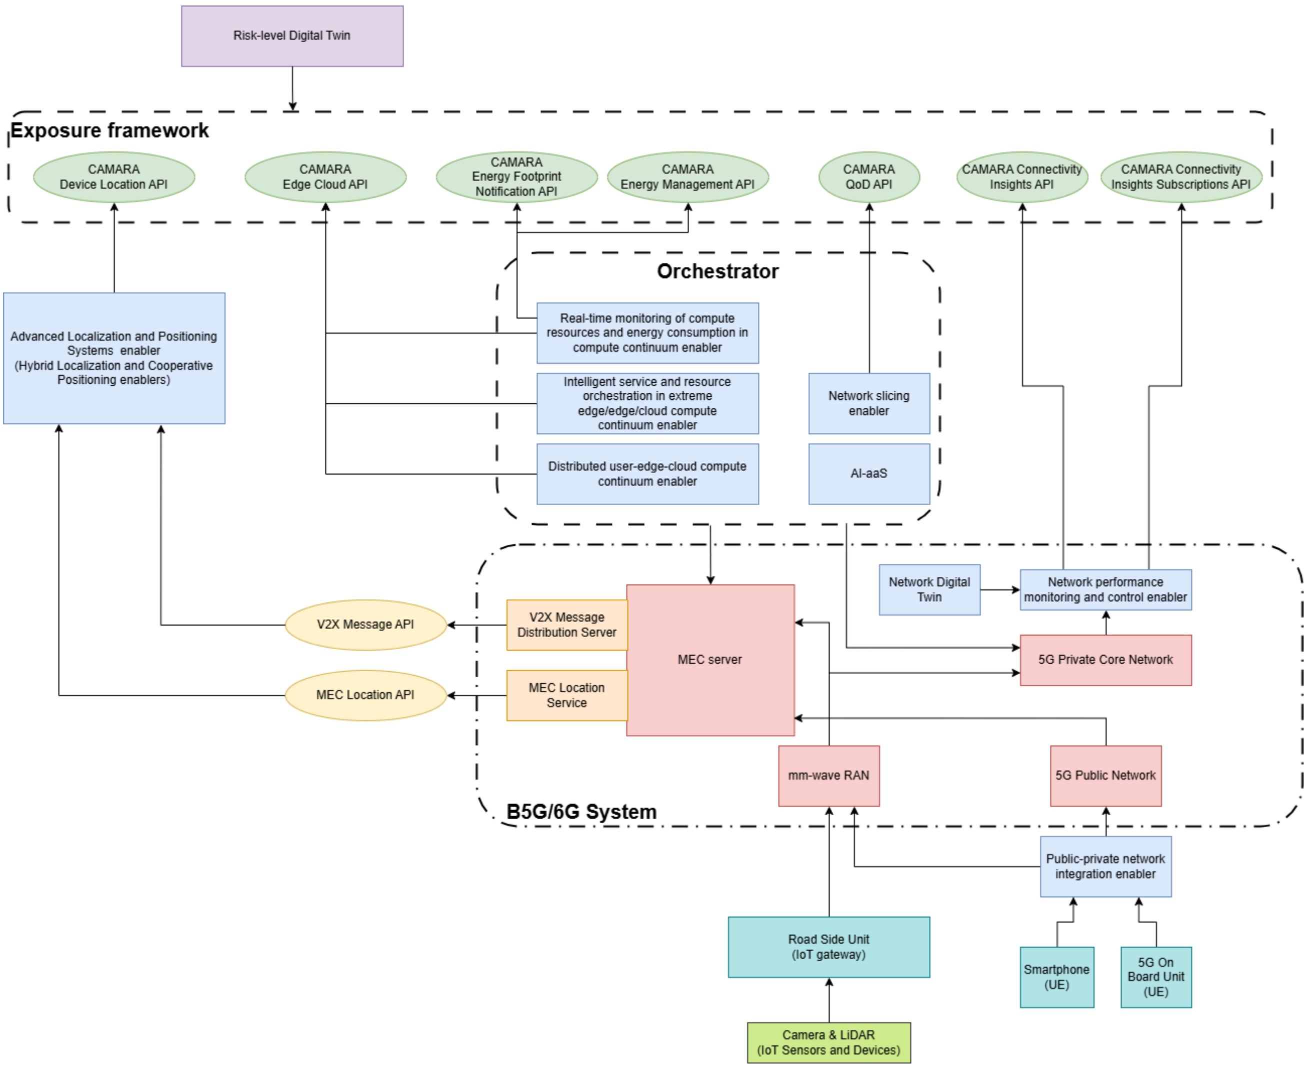

Figure 1: High-Level Architecture of T1 UC

Concerning functional requirements, the T1 UC requires the B5G/6G system to support dynamic deployment of vertical services at the edge according to policies based on metrics such as available bandwidth, RSU energy level, and latency. The applications at the edge server should receive information from the entities in the field (e.g., RSU, vehicles) with very low latency, thanks to the provisioning of specific network slices. The communication among the different applications should be secure and reliable. Network, computing, and energy should be monitored continuously.

The T1 UC identifies the following non-functional requirements to ensure system performance, reliability, and availability.

- End-to-end latency from sensor data acquisition to notification reception should remain below 200 ms.
- The system should ensure adequate QoS and resilience across the edge–server path, maintaining performance under adverse conditions.
- Scalability should be demonstrated with stable behaviour when the number of devices and/or the data volume increases by ≥100% (at least doubling the load).

The key components of T1 UC are summarised in Table 1.

Table 1: T1 Use Case High-Level Architecture Description

Domain & Enablers Reference	Component	UC Related Functionalities
Edge infrastructure	RSU	Collection and forwarding of sensor data; processing of sensor data
End-user device	Smartphone	Position sharing, risk level notification retrieval
End-user device	OBU	Position and other data sharing with the Digital Twin
Application enabler	Risk Level Digital Twin	Risk Level computation based on gathered data
Orchestration enabler	Orchestrator	Solution for automatic application deployment and management of network and computing resources
Connectivity infrastructure	B5G/6G system	System comprising an MEC server for application deployment, and a private and public network for providing the communication capabilities

B5G/6G Relevance

The proposed use case relies on B5G/6G capabilities to support the exchange of safety-critical information among vulnerable road users, roadside infrastructure, connected vehicles, and edge applications. The scenario requires ultra-reliable and low-latency communications (URLLC) to ensure that risk notifications generated by the Risk-Level Digital Twin are delivered to visually impaired pedestrians within strict timing constraints, and data exchanged by the OBU and the safety information coming from the RSU requires the same low-delay channel with high reliability.

The use case involves the transmission of heterogeneous data flows with different QoS requirements. Video flows from the RSU can be sent (in case the battery is low), requiring an Enhanced Mobile Broadband (eMBB) channel, again transporting safety-related data.

B5G/6G network slicing and QoS-aware traffic management are therefore essential to guarantee service continuity and prioritisation of safety-related communications.

Another key B5G/6G aspect addressed by the use case is the tight integration of edge computing and intelligent orchestration. Since the RSU operates using renewable energy sources and battery storage, computation tasks may be dynamically offloaded between the RSU and the MEC infrastructure. This requires advanced orchestration mechanisms exploiting B5G/6G-native support for distributed edge services, real-time monitoring, and context-aware resource management.

Large-scale deployment of the pilot

The scenario is implemented in a realistic urban deployment with an RSU with a camera and LiDAR connected to an mmWave cell. The sensed area includes a complex junction and a vehicle with an OBU, and a visually impaired person with a smartphone who moves into the area. Moreover, V2X data of real vehicles will be exchanged thanks to the RSU working as a gateway from the short-range V2X connectivity and the B5G/6G cell.

The experimentation will therefore assess the B5G/6G capability to work in a relevant complex environment, and the expected TRL of the solution is 7. Data will be collected from the private mmWave cell as well as the TIM public network, testing how they can work together.

The vertical application and the technical B5G/6G enablers will also be tested in severe conditions by stressing the interaction with the enablers, having a large number of emulated devices, and loading the network with artificially generated traffic.

2.1.2 Deployment and implementation details

2.1.2.1 Design of the application

The system architecture of T1 UC is presented in Figure 2. The Risk Level Digital Twin (RLDT) is the main application of T1 UC. This application exploits the information received from the B5G/6G system to alert road users about dangerous situations. Specifically, the T1 UC aims to demonstrate how the Risk Level Digital Twin can support visually impaired people to cross the street in a safe manner.

The RLDT application invokes a set of CAMARA-based APIs to retrieve the needed information from the B5G/6G system. The CAMARA-based APIs are contained in an exposure framework, and each API is deployed jointly with a transformation function that is responsible for the interaction with the B5G/6G system. The transformation functions are the implementation of the previously identified enablers that allow an easy interaction with the B5G/6G system to retrieve data or to perform actions.

A main component of the system architecture is the Resource and Service orchestrator. In T1 UC, the orchestrator, jointly with the Closed Loop platform, is responsible for the deployment and migration of application functions. One of the objectives of T1 UC is indeed to implement an offload strategy between the RSU and the edge server of the application functions that perform object detection and tracking on the sensor data, gathered by the RSU. The offload strategy considers the RSU battery level to decide if the detection and tracking task should be migrated to the edge server or, vice versa, from the edge server to the RSU.

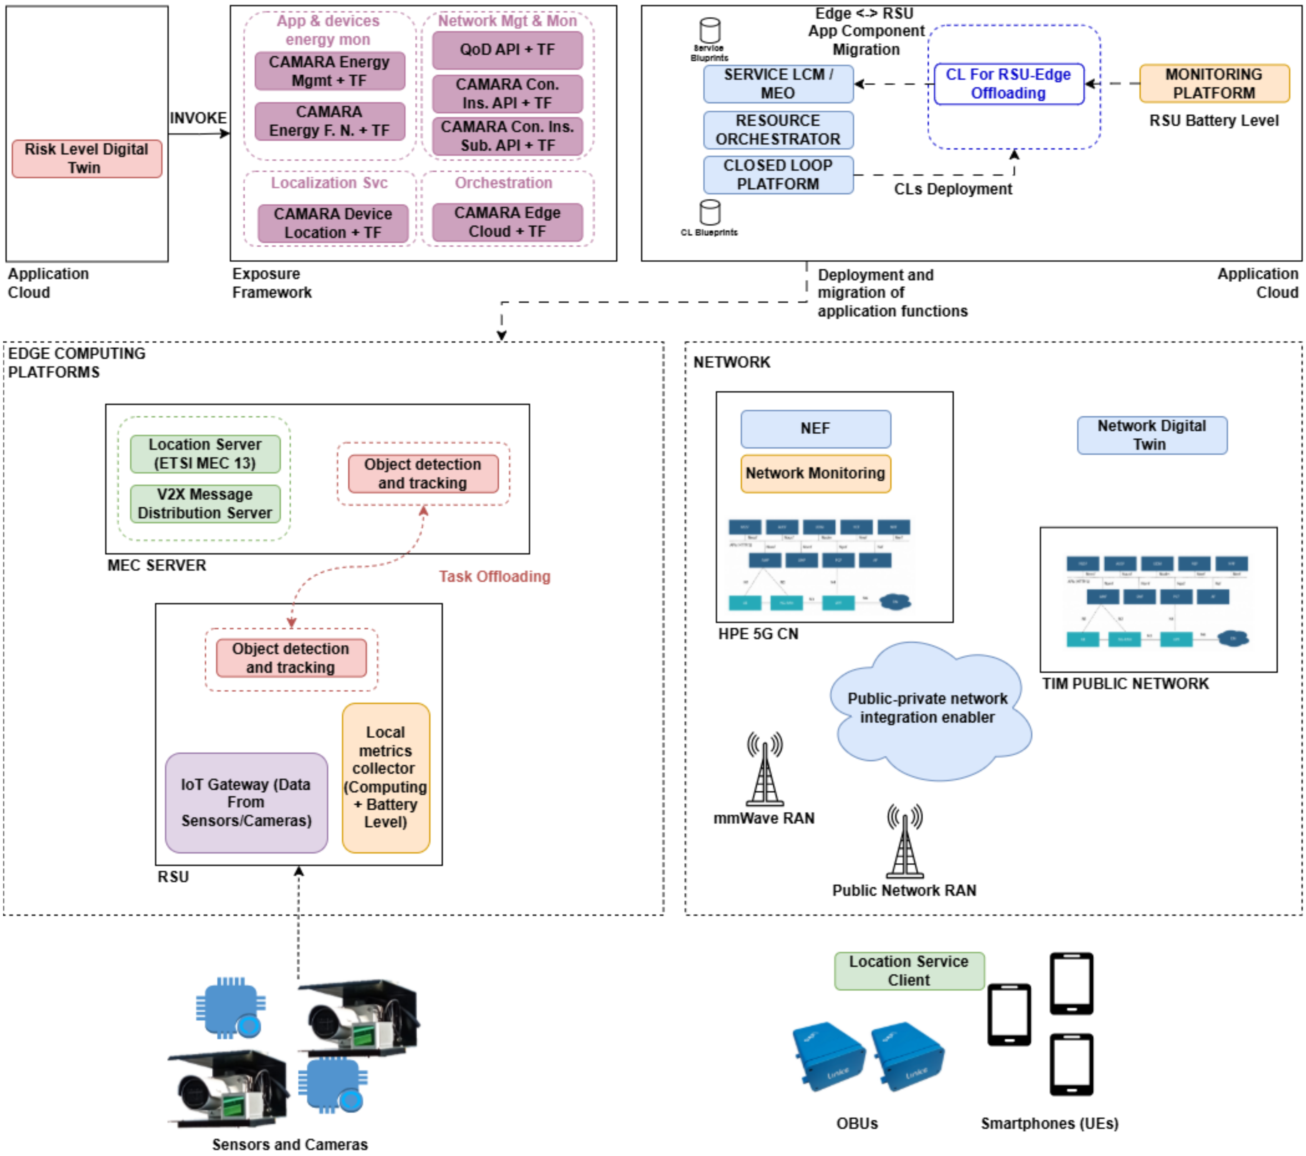

Figure 2: System Architecture of T1 UC

The offload strategy relies on the measurements performed at the RSU side by a local metrics collector that gathers the status of the battery level, the energy produced by the solar panel, and the computing resources used by the RSU. This information is made available to the Closed Loop function by a monitoring platform that centrally provides the RSU metrics, as well as network metrics coming from the Network Monitoring component. Indeed, the offload decision also needs to consider the network conditions, as sensor data must be transmitted from the RSU to the Edge server using the mm-wave connection.

The RLDT application relies on the CAMARA Device Location APIs to receive information about the road users present and their position. The transformation function of the CAMARA Device Location implements the functionalities expected by the Advanced Localization and Positioning Systems enabler. This CAMARA transformation function exploits at the southbound layer the APIs exposed by the Location server and by the V2X Message Distribution server.

The Location server exposes ETSI MEC Location APIs (i.e., ETSI MEC 13), which provide position information about the UEs. As the B5G/6G system does not support geographical location functionality, the UE itself provides its GNSS position to the Location server through the Location Service Client running at the UE side.

The V2X Message Distribution server provides V2X messages that are generated by connected vehicles, pedestrians or by the RSU. These V2X messages can provide information about the position and dynamics of the road user that generates them, e.g., Cooperative Awareness Message (CAM) for vehicles, VRU Awareness Message (VAM) for pedestrians. Another type of V2X message, namely the Collective Perception Messages (CPM), contains the processed information of sensor data gathered by local sensors on the vehicle or on the RSU. The CPM provides a list of the identified objects detailing additional information such as the type of object, its position, and dynamics. A typical implementation of the V2X Message Distribution server is based on a pub/sub approach in which a message broker is used to collect and distribute the V2X messages. The European platform C-ROADS defines the implementation details of this approach, and, specifically, it indicates using the AMQP 1.0 application protocol for the pub/sub mechanism.

Position information from UEs coming from ETSI MEC Location API and from V2X messages through the V2X Message Distribution Server is then received by the CAMARA Device Location transformation function that fuses the received data for providing device position with improved accuracy.

The RLDT application computes the risk level exploiting the received information, and it triggers an alert when a user of the RLDT service is in a dangerous situation. It is envisaged that the notification is sent to the user who receives it on a dedicated app running on the smartphone. To increase the availability of the service, the T1 UC will experiment with having the smartphone seamlessly connected to the private network using the mm-wave-based RAN or to a public mobile network. The choice of the connection depends on the effective coverage of the private network in the current user’s position.

The workflow of T1 UC is illustrated in Figure 3. The RLDT application starts by subscribing to the CAMARA Device Location APIs for receiving notifications about UEs entering the area of relevance of the RLDT. After each notification, the RLDT checks if the UE is registered for the RLDT service (steps 1-3). If the UE is registered, it means that the UE is a visually impaired person, and the RLDT starts the procedure of deploying the object detection and tracking application function at the RSU. This involves the RLDT invoking the Edge Cloud APIs, which consequently interacts with the orchestrator for the actual deployment (steps 4-5).

The workflow is then made up of two subprocesses. The first one is related to battery monitoring and task offloading, while the second one focuses on the actual risk management application.

The subprocess on battery monitoring and offloading is a periodic process in which the battery level of the RSU is verified (step 6), and based on its status, the offloading is performed or not. In case the RSU battery level is low, and the object detection and tracking application function is being executed at the RSU, the offload to the MEC is triggered (step 7). This involves the termination of the execution at the RSU, the deployment of the application function at the MEC server, and the instantiation of an eMBB slice to ensure enough communication bandwidth to the RSU for the forwarding of the raw sensor data to the MEC server (steps 7.1-7.5). Whereas the object detection and tracking application function is currently in execution at the MEC server, and the RSU has a battery level over a given threshold, the opposite migration takes place (step 8), moving the deployment of the function from the MEC server to the RSU and releasing the eMBB slice.

The subprocess on risk management begins with the subscription of the RLDT to UE Location updates for receiving information about all UEs in the relevant area (step 9). The CAMARA Device Location API sends the UE's position to the RLDT periodically (step 10), which exploits this information for computing the risk level for the visually impaired person (step 11). In the case that the computed risk is above a threshold, the RLDT push the notification to the VRU (step 12).

When the visually impaired person exits the area, the RLDT is informed of the event (step 13), and it asks the Edge Cloud API to terminate the execution of the object detection and tracking application function (step 14).

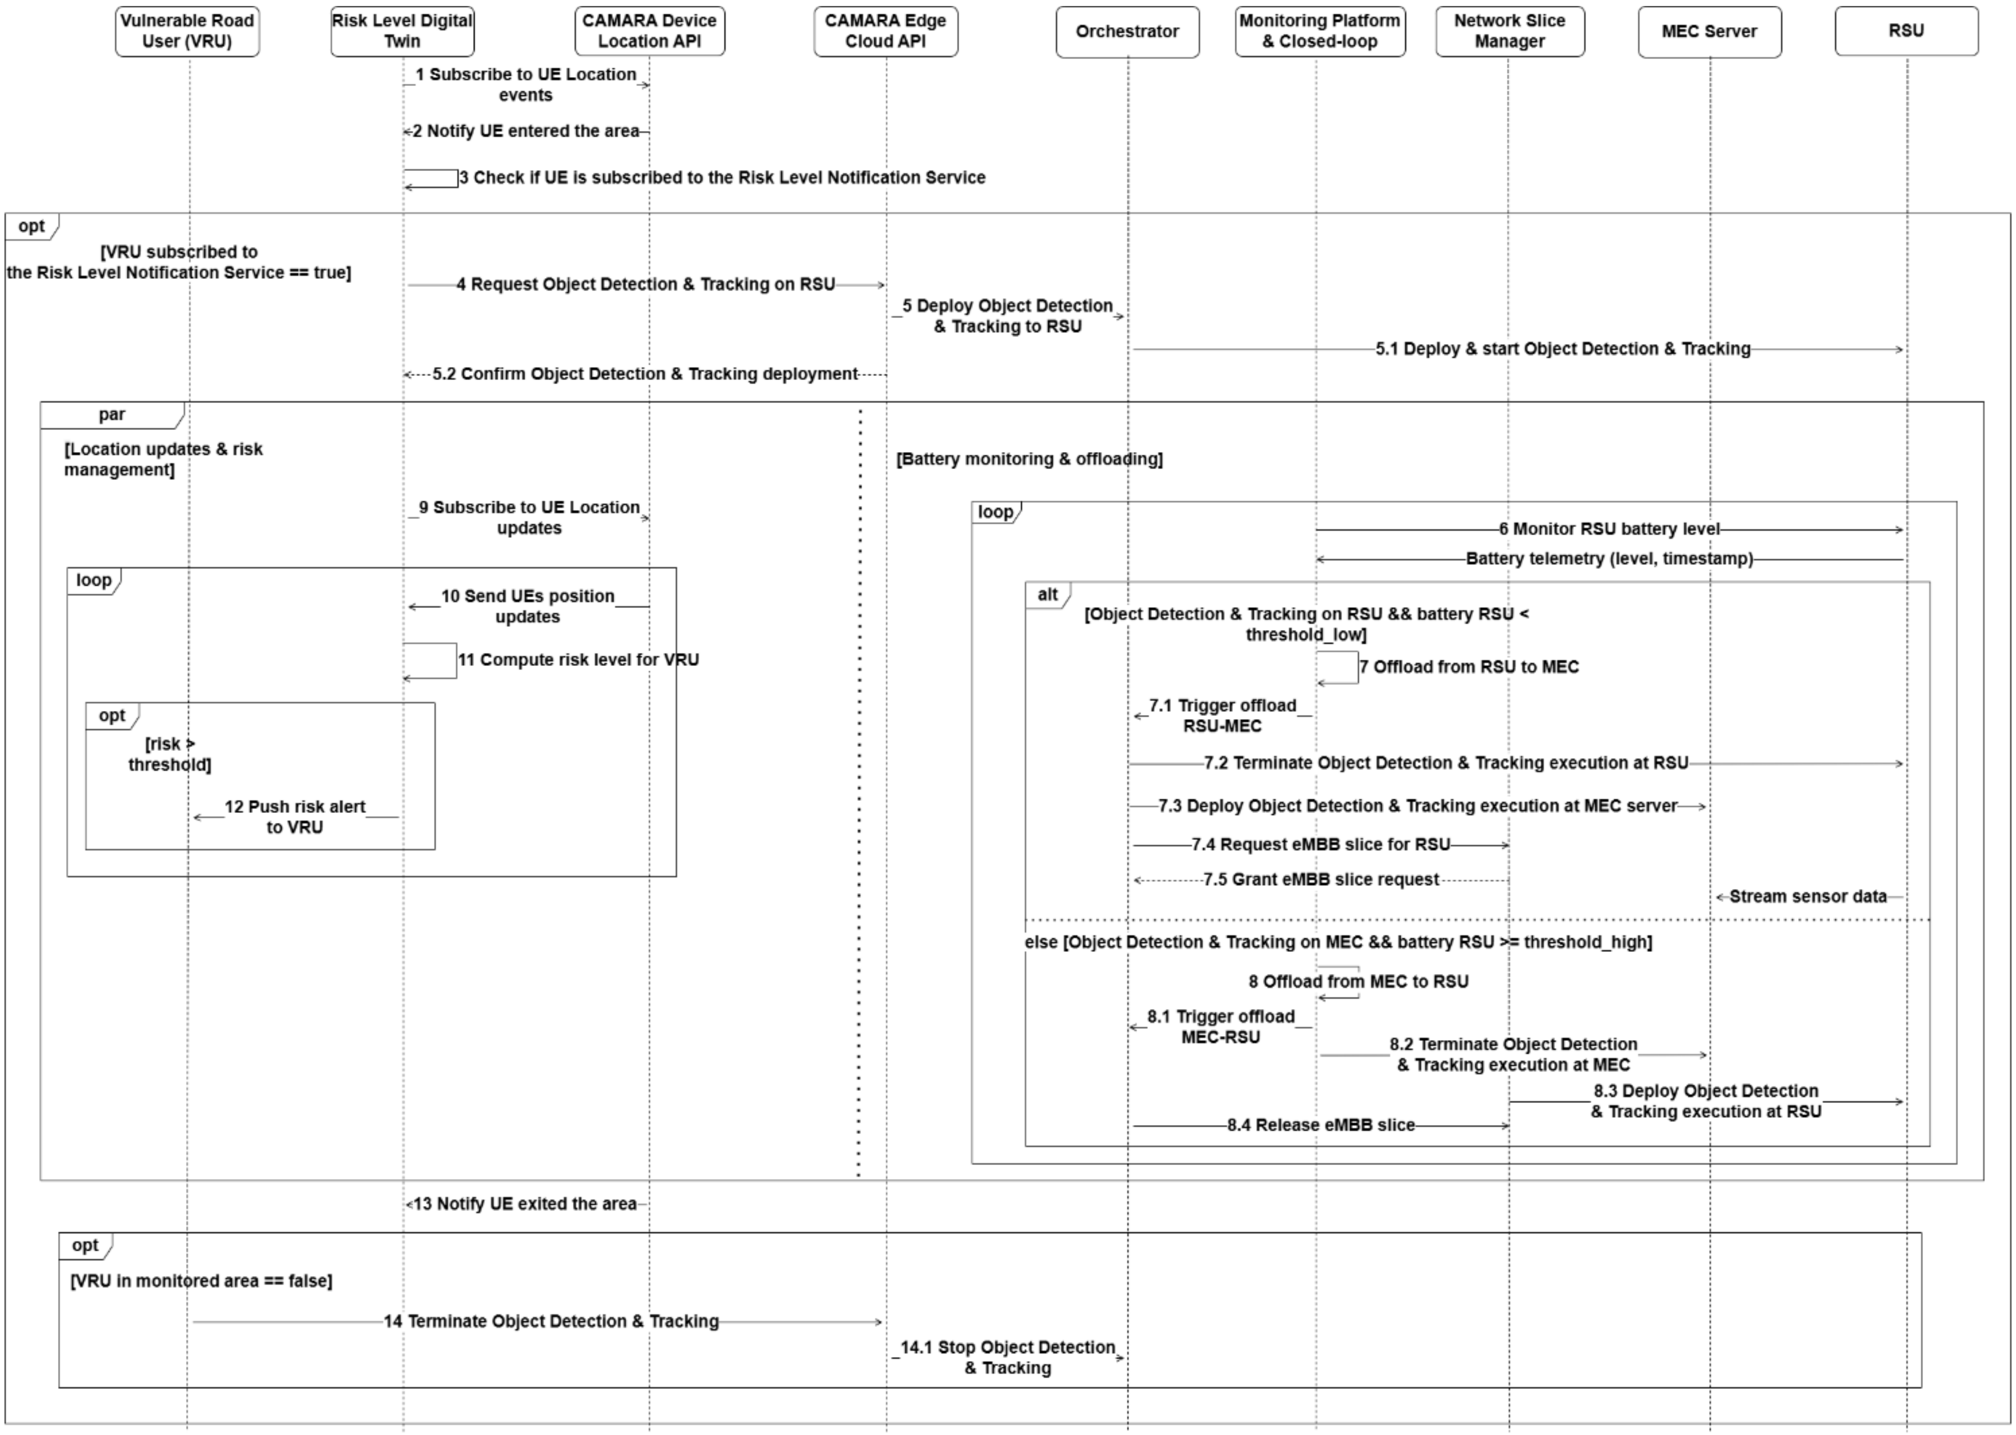

Figure 3: Workflow of T1 UC

2.1.2.2 *Devices, sensors and supporting equipment*

In Table 2, the devices and sensors employed in T1 UC are introduced.

Table 2: T1 Devices, sensors and supporting equipment

Category	Component	Technical Details	Units
5G connectivity module	5G mmWave CPE	5G SA/NSA, mmWave, Sub-6GHz, 4G LTE	1 x OBU/RSU
5G connectivity module (under investigation)	Waveshare Quectel RM530N-GL	IoT 5G Global Band Module, 5G Sub-6 & MmWave, M.2, w/ 3GPP 5G R16	1 x OBU/RSU
Positioning module	uBlox F9R 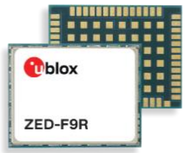	RTK-enabled multi-band GNSS receiver, Dead Reckoning	1 x OBU
End-user device	Samsung S24 	256GB AI Smartphone, Unlocked Android, 200MP, 100x Zoom Cameras, Fast Processor, Long Battery Life, Edge-to-Edge Display, S Pen, US Version, 2024, Titanium Black	1
Perception sensor	VIVOTEK IB9391-EHTV-V2 camera 	Visual intersection monitoring, H.265, H.264, MJPEG, 30 fps @ 3840x2160, Smart IR	1
Perception sensor	LIVOX AVIA LiDAR 	3D scene perception, 450 m maximum detection range, 70.4° * 77.2° FOV for Non-repetitive Scanning, ≥ 240,000 Point Rate, 2 cm	1
Vehicle equipment	OBU 	5G connectivity, GNSS positioning, C-ITS message generation	1
Roadside infrastructure	RSU	Solar-powered roadside unit with battery, local compute	1

		resources, 5G modem, energy monitoring	

2.1.3 Network infrastructure and functional components

The T1 UC comprises, as shown in Figure 2, the following network infrastructure components: mmWave RAN, public network RAN, private core network, and public core network. In the remainder of this section, these network components are introduced. Other functional components of T1 UC are the Service and Resource orchestrator, the Closed Loop platform, the monitoring platform, and the edge computing platforms. These components are introduced as well in the following.

2.1.3.1 mmWave RAN

The private 5G network implemented for the use cases consists of a **millimetre-wave (mmWave) Radio Unit (RU)** deployed close to the operational area, and a **COTS server** running the **baseband software stack** (gNB implemented in a virtualised environment). The mmWave RU integrates **phased-array antennas** supporting dynamic beamforming and is connected to the server through a **dedicated point-to-point optical fibre link** (fronthaul, eCPRI/IEEE 1914.x), typically operating at 10–25 GbE over single-mode fibre. Time and frequency synchronisation are ensured via **PTP (IEEE 1588v2)**, with optional GNSS or a local grandmaster clock.

The **baseband layer (DU/CU)** executes on the server, leveraging standard x86 hardware and network interface cards with low-latency acceleration. The **5G Core Network (CN)** – including AMF, SMF, UPF, UDM, and related functions – can be **co-located** on the same server (in an all-in-one edge deployment for minimal latency) or **deployed on a nearby edge server**, maintaining the UPF close to the RAN to optimise uplink throughput and end-to-end latency.

The infrastructure supports **logical network slicing and QoS profiles**, isolating critical traffic flows (e.g., telemetry, uplink video, real-time control) from best-effort services. The **user plane** allows both **local breakout** towards edge applications and secure **backhaul routing** to external networks, depending on the service configuration. End-to-end security is ensured through **encryption on N2/N3 interfaces**, hardened server configurations, L2/L3 segmentation, and role-based access control.

This architecture provides:

- **Low latency** through dedicated optical fronthaul and local UPF placement.
- **Scalability** via containerised network functions (CNFs/VNFs) and horizontal scaling of the server.
- **Reliability** with redundant PTP/GNSS synchronisation, optical point-to-point links, and service watchdog mechanisms.
- **Deployment flexibility**, allowing co-located or distributed core functions based on the use case requirements.

2.1.3.1.1 mmWave Energy Efficiency Considerations

Although mmWave operation typically requires higher peak transmit power to overcome propagation losses, it can offer **better energy efficiency at the system level** compared to Sub-6 GHz in high-capacity scenarios.
When using **200 MHz of mmWave bandwidth**, the system can achieve the same user throughput as an

80 MHz Sub-6 GHz channel with **lower total power consumption per delivered bit**, due to the **shorter transmission time**, **narrower beamforming**, and **reduced interference domain**. This efficiency arises from **higher spectral efficiency**, **spatial reuse**, and **short-range link optimisation**, which lowers the active RF duty cycle of both the RU and the baseband processor.

2.1.3.2 Public Network RAN

TIM commercial 5G network is based on the 3rd Generation Partnership Project (3GPP) release Rel-15 NSA 80 MHz at 3.7 GHz, and a multi-layer 4G network (anchor layers). It ensures wide coverage to extend the use cases area.

2.1.3.3 Private Core network

HPE contributes to the project with its HPE Aruba Networking Private 5G Core Network (HPE 5GC), a fully virtualised Standalone system that follows 3GPP (starting from Release 16) and ETSI standards. It manages authentication, control signalling, QoS rules, mobility and traffic forwarding between the base stations and external networks. As shown in Figure 4, it includes all main network functions such as AMF, SMF, UPF, UDM, AUSF, UDR, PCF, NRF and NSSF. It also has an Equipment Identity Register that improves security by checking the IMEI or PEI of each device to decide if it can join the network. The system supports all standard interfaces from N1 to N6. Through the N5 interface, trusted third-party applications can change network parameters and control end-to-end QoS for important and sensitive services.

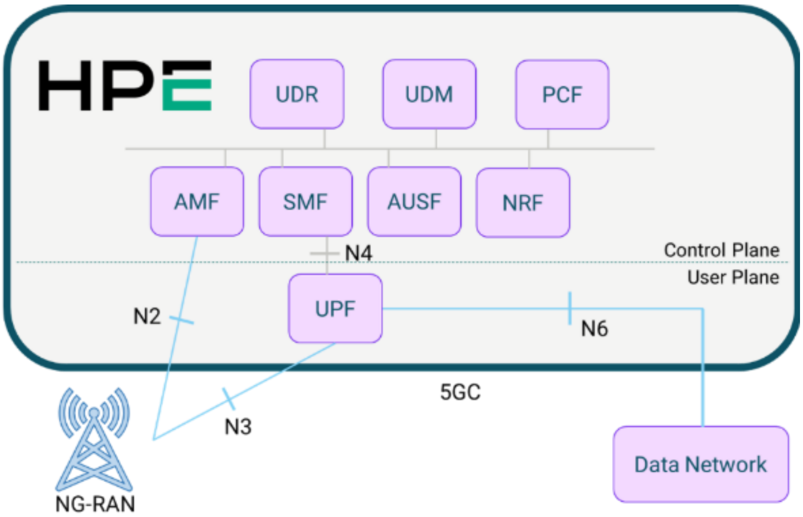

Figure 4: HPE ANW Private 5G Core Network high-level architecture.

The HPE 5G Core also provides REST APIs that follow the OpenAPI 3.0 specification. These APIs allow external OAM systems to configure network functions, deploy new settings, manage users, and adjust the system behaviour automatically. They also allow the collection of real-time and historical data, including resource usage, 3GPP signalling and traffic information, as well as metrics from virtual machines such as CPU, memory, disk and network. Monitoring data is exposed through Prometheus. All APIs are protected by TLS certificates, which ensure secure access, prevent data interception and confirm the identity of the server.

Moreover, it is possible to collect real-time and historical monitoring data or network performance metrics via API. This enables the HPE 5GC's proactive maintenance and problem resolution, as well as service control and optimisation. Network function information is collected and exposed using Prometheus. They encompass resource consumption data, 3GPP-related metrics for signaling and traffic data, and monitoring information about the state of the network. Further, the HPE 5GC exposes via API a set of parameters for monitoring its virtual machines, including metrics such as CPU utilisation, memory allocation, disk usage, and networking data.

For the T1 and T2 use cases, the HPE 5GC is deployed at LINKS premises, especially on OpenStack, which acts as the underlying virtualisation and orchestration platform. The deployment is aligned with the setup already adopted in the ENVELOPE project[1], ensuring architectural consistency across projects. In the current configuration, the deployment consists of three distinct VMs. One VM hosts the Control Plane (CP) functions, including the main signalling and service-based architecture (SBA) components (e.g., AMF, SMF, NRF, AUSF, PCF, and UDR). User Plane (UP) functionalities are deployed on two separate VMs, each running an independent UPF.

2.1.3.4 Public Core network

The TIM commercial core network, based on the Ericsson Dual Mode Core platform, provides secure and managed IP connectivity between customer access points and the target applications hosted in the data center. It ensures controlled routing, traffic segregation, and service availability, enabling reliable access to the hosted platforms accessible through commercial Access Point Names (APNs); private APNs are also supported.

2.1.3.5 Management and Orchestration Layer

The Management and Orchestration Layer provides a modular and integrated framework designed to coordinate computing and network resources alongside services deployed in virtualised environments. This system automates the deployment, scaling, and optimisation of applications to ensure that the infrastructure can satisfy the stringent requirements of transport use cases. It's composed of several components:

- **The Service Orchestrator** manages the end-to-end orchestration of services. It coordinates resource allocation through the Resource Orchestrator, sets up service components, and activates monitoring tasks. It supports runtime service management, enabling both manual and automated reconfiguration and scaling through the following key functional endpoints:
    - Service Instance Creation and Instantiation: Handles the initial setup and lifecycle of new service instances.
    - Scaling and Reconfiguration: Dynamically adjusts resources or components of active services based on demand.
    - Service Migration: Supports the movement of service components between different execution nodes to maintain performance.
    - Lifecycle Management: Provides tools for retrieving service details, listing active instances, and terminating services when no longer needed
- **The Resource Orchestrator** manages multi-cluster infrastructures spanning the entire edge-to-cloud continuum. It supports heterogeneous nodes with varying characteristics, such as mobility and compute capacity, using a modular architecture:
    - Platform Manager: Handles the registration of edge/cloud platforms and tracks the status of applications deployed on them.
    - Service Deployer: Manages the deployment lifecycle, including creating, updating, and monitoring the runtime information of services.
    - Resource Manager: Maintains a detailed inventory of clusters and nodes, allowing for precise allocation across the infrastructure.
    - Service Template Catalogue: Stores and manages reusable service blueprints to streamline deployment.
- **The Closed-Loop Platform** enables full automation of network and service management by provisioning and governing AI-driven functions. These loops consist of functional components that

1 Grant Agreement No. 101139048

continuously exchange data to analyse system behaviour, make decisions, and execute actions to maintain a desired state. The platform facilitates:
  - Dynamic Provisioning: Automated creation and termination of closed-loop instances.
  - Autonomous Operation: Analysis and execution of actions based on pre-defined descriptors and real-time data.
  - Runtime Updates: The ability to modify loop configurations without interrupting the management cycle.
- The Monitoring Platform provides the essential data foundation for orchestration and closed-loop functions by collecting a continuous stream of infrastructure metrics. The platform is composed of several open source and proprietary tools, including:
  - Telegraf: Utilised to collect data from various sources throughout the infrastructure through a comprehensive set of plugins.
  - InfluxDB: Acts as the time-series database for historical data storage and retrieval, allowing the system to analyse trends over time.
  - Kafka: Serves as the message bus for efficient and scalable data distribution between system components
  - Config-manager: developed by NXW, it dynamically cross-configure all the tools, ensuring data consistency among all the components

2.1.3.6 *Exposure framework*

The Exposure Framework (depicted in Figure 5) provides secure and standardised access to network and computational resources through open APIs. It serves as the interface between application developers and the underlying infrastructure, enabling application deployment, orchestration, and QoS management across diverse domains.

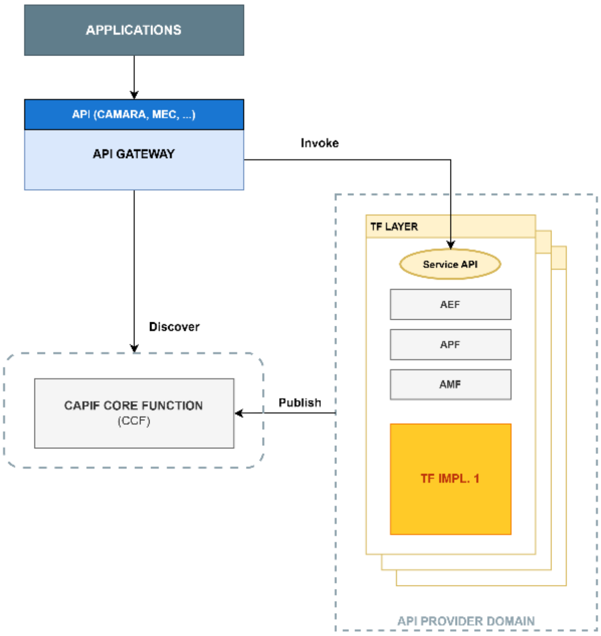

Figure 5: Exposure Framework Architecture

The framework is built upon the Common API Framework (CAPIF) specification to manage a registry of available transformation functions and exposed APIs. Its main architectural components include:

- **API Gateway:** Acts as the unified entry point for all application requests, responsible for routing them to the appropriate backend service through standardised interfaces. It performs context-aware selection based on user identity, request location, and network conditions.
- **CAPIF Core Function:** Provides core capabilities such as API discovery, authentication, authorisation, and usage control. It maintains an up-to-date registry of all available Transformation Functions.
- **Transformation Functions (TFs):** Backend services that abstract and simplify access to network and computing infrastructure by encapsulating business logic and managing southbound interactions with resources.

The Exposure Framework will specifically expose the following families of CAMARA-compliant APIs to support transport vertical requirements:

- **CAMARA Energy Footprint Notification:** Provides insights into the energy consumption and environmental impact of network and compute resources.
- **CAMARA Quality on Demand (QoD):** Enables applications to request specific network performance levels (e.g., increased bandwidth or low latency) for mission-critical tasks.
- **CAMARA Device Location:** Offers real-time geographical position information and geofencing capabilities for connected users and vehicles.
- **CAMARA Connectivity Insights:** Allows applications to register performance thresholds and query the network's confidence in meeting these requirements for specific devices.
- **CAMARA Edge Cloud (Edge Application Management):** Facilitates the discovery of edge zones and manages the onboarding, instantiation, and migration of applications across the compute continuum

2.1.3.7 Edge computing platforms

The edge computing platforms comprise of the MEC server and the RSU. The MEC server provides computing capabilities at the edge of the network for the execution of tasks that are offloaded from the RSU. Furthermore, the MEC server hosts the Location Service and the related APIs as defined in ETSI GS MEC 013, and the V2X message distribution server. The RSU is a far-edge node which offers the processing resources as close as possible to the end user for executing applicative tasks. The RSU also implements a local monitoring system to provide metrics about the resource utilisation and the energy level to the central monitoring system. The RSU also represents the IoT gateway for the IoT sensors (i.e., camera and LiDAR) that are connected to the RSU.

2.1.4 Use case enablers

The enablers exploited in T1 UC are presented in Table 3.

Table 3: List of enablers and sub-enablers for T1.

Category	Technology enablers	Sub-enablers
Communication enablers	Network slicing support	Network Slice as a Service (NSaaS)
Communication enablers	Network exposure and monitoring	Network Exposure Function (NWEF)
Communication enablers	Network Digital Twin (NDT)	Network Digital Twin
Communication enablers	Public-Private integration	
Compute-as-a-Service enablers	Intelligent service and resource orchestration	Centralised Service Orchestrator (CSO)

Compute-as-a-Service enablers	Intelligent service and resource orchestration	Edge Service Orchestrator (ESO)
Compute-as-a-Service enablers	Compute resource management	Compute Resource Orchestration (CRO)
Compute-as-a-Service enablers	Compute resource management	Compute Resource Allocation (CRA)
Compute-as-a-Service enablers	Compute continuum	Resource Exposure Function (REF)
Compute-as-a-Service enablers	Compute continuum	Compute Exposure Function (CEF)
Compute-as-a-Service enablers	Compute continuum	Compute Digital Twin (CDT)
Compute-as-a-Service enablers	Compute continuum	Resource Digital Twin (RDT)
Compute-as-a-Service enablers	Compute continuum	Deployment Policy Repository (DP-Rep)
Application enablers	Vertical enabler	Digital Twin Framework (DT-F/W)
Application enablers	Vertical enabler	Data Collection (DT-DC)
Application enablers	Vertical enabler	Simulation and Modelling (DT-SM)
Application enablers	Vertical enabler	Decision Modelling and Optimization (DT-DMO)
Application enablers	Vertical enabler	Visualization and Interaction (DT-VI)
Application enablers	Vertical enabler	Digital Twin Registry (DT-Reg)
Application enablers	Vertical enabler	Digital Twin Repository (DT-Rep)
Application enablers	Vertical enabler	Digital Twin-Virtual Entity Time Serie Database (DTVE-TSDB)
IoT and localization enablers	Localization and positioning	Position Acquisition Engine (PAE)
IoT and localization enablers	Localization and positioning	Mobility Context Provider (MCP)
IoT and localization enablers	Localization and positioning	Cooperative Localization Engine (CLE)

IoT and localization enablers	Vertical enabler	Message Queue Factory (MQF)
IoT and localization enablers	Vertical enabler	IoT-Exposure Function (IoT-EF)
IoT and localization enablers	Vertical enabler	IoT Resource Registry (IoT-RReg)
IoT and localization enablers	Vertical enabler	IoT Service Registry (IoT-SReg)

The communication enablers that are exploited in T1 UC are the Network Slice as a Service (NSaaS), which supports the network slicing capabilities, the Network exposure and monitoring, which provides insights about network status to the monitoring platform, and to the Network Digital Twin. The T1 UC exploits the Compute resource management and Compute continuum among the available Compute-as-a-Service enablers for achieving the migration of tasks among the UC computing resources.

The main vertical application of the T1 UC is the RLDT service that is based on the set of sub-enablers related to digital twin functionalities. This set includes the sub-enablers providing data collection functions, simulation and modelling as well as functionalities for listing and storing digital twin images.

In T1 UC, the Localization and Positioning enabler is exploited for retrieving the position of the actors of this UC (e.g., the VRUs, vehicles). Specifically, the Position Acquisition Engine (PAE) is used for retrieving the available position information of the UEs from the available different sources. The Cooperative Localization Engine (CLE) ingests the information from the PAE for fusing the collected data and ensuring a highly accurate positioning. The CLE is supported by the Mobility Context Provider (MCP), which provides mobility details to the CLE for enhancing its functionality.

In the Vertical enabler category, the sub-enablers Message Queue Factory (MQF), IoT-Exposure Function (IOT-EF), IoT-Resource Registry (IOT-RReg), and IoT-Service Registry (IOT-SReg) are used for collecting information from IoT devices and for making such information available to upper-layer components and enablers.

2.1.5 KPIs definition

Table 4 introduces the KPIs that have been identified for T1 UC. The KPIs will be assessed during the trial execution of T1 UC. The measurements for the KPIs computation will be obtained at the point of presence at which ad-hoc logging services will be executed. The application-level KPIs will be measured at the UE, MEC and cloud level. The network KPIs will be based on measurements collected at the UE, RAN, and core network segments. The availability of the enablers and the migration time will be measured at the exposure framework and at the Service and Resource orchestrator.

Table 4: List of KPIs for T1

KPI name	Description/KPI definition	KPI target
End-to-end application latency	End-to-end latency of the RLDT service from sensor data provisioning to VRU notification	< 200 ms
Network latency	Latency of the communication service in provisioning the feedback from the RLDT service toward VRU	< 20 ms

Average RU power consumption	Average electrical power drawn by the mmWave RU during data transfer.	<100W
Power consumption reduction vs Sub-6 GHz baseline	Relative decrease in power usage for equal throughput conditions (200 MHz mmWave vs 80 MHz Sub-6 GHz)	Target value: 20% Expected range: 15–25%, depending on beamforming efficiency, traffic burstiness, deployment density, and link conditions.
Watt per Throughput	Power-to-performance efficiency metric	20mW/Mbps depending on beamforming efficiency, traffic burstiness, deployment density, and link conditions.
UL throughput	Maximum achievable data rate in the uplink connection when the RSU is granted for improved quality	150 Mbps
Increase in operational time of RSU	RSU lifetime gain enabled by energy-aware offloading and orchestration	+30%
Application migration latency	Time needed to migrate the task between the RSU and the MEC server	< 1 s
Positioning accuracy	Positioning accuracy of UC actors	< 60 cm
Reliability	Percentage of time that the enablers supporting the RLDT service are fully operational and able to deliver expected functionalities	99.999%
Scalability	RLDT service is available in stressed scenarios (e.g., number of devices and/or data volume increases)	Stable behaviour with load increase ≥100%
API Availability	The percentage of time the Exposure Framework and its CAMARA APIs are operational and accessible to developers for resource orchestration	>90%
API Response Time	The total latency measured from the moment a request hits the API Gateway to the delivery of the processed response from the Transformation Functions	<5s

2.1.6 KVIs definition

The KVIs of T1 UC are introduced in Table 5. The defined KVIs aim to assess the environmental and societal impacts of T1 UC. Two main areas of interest have been identified during the definition of the KVIs. The first one is related to the impact of the intelligent migration of tasks between the RSU and the MEC server on the carbon footprint. This aspect is strictly connected to the reduction of operational costs obtained thanks to the energy saving achieved. The other area of interest is the effectiveness of the RLDT service. The KVIs are related to the reduction of accident risks of VRUs thanks to the service offered, and connected to this aspect, it is the reduction of the costs in healthcare assistance. Further, KVI identifies the impact of T1 UC on the potential increase in revenue of MNO for offering innovative services. The KVIs are all estimated through expert assessment, as it is impossible to achieve quantitative data for a precise computation during trial testing, as the aspects covered by the KVIs are related to benefits over long time periods.

Table 5: List of KVIs for T1

KVI name	Description/KVI definition	KVI target
Reduced carbon footprint	Reduction of carbon footprint through optimised deployment of applications and consequent effective resource usage of B5G/6G system elements	Reduction of 15%
Reduced TCO for MNOs	Reduced operator cost through lower energy expenditure and more efficient resource usage.	Reduction of 30%
Increased revenues for MNOs	Opportunity to offer value-added energy-aware and safety-related innovative services	Increase of 20%
Accident reduction	Reduction of accident risk thanks to higher awareness and support for VRUs	Reduction of 40%
Reduced healthcare expenditure	Economic benefit associated with accident prevention involving VRU	Reduction of 40%
Increased Perceived Safety and Confidence	The degree to which visually impaired users feel safer when crossing intersections thanks to timely alerts.	≥ 30% improvement in perceived safety
Improved Mobility Independence	The extent to which the system reduces reliance on human assistance for road crossing.	≥ 25% reduction in need for assistance
Trust in System Reliability	User trusts that alerts are timely, correct, and non-intrusive.	≥ 80% trust rating
Accessibility and Ease of Use	How easy and intuitive the alerting app is for visually impaired people.	≥ 85% accessibility rating
Reduced Cognitive Load for Drivers	Lower driver stress or mental load due to improved awareness of VRUs.	≥ 20% reduction in cognitive load

Increased Driver Compliance with Safe Behaviours	Observable increase in safe driving patterns when alerts are received.	≥ 15% improvement
Enhanced VRU Awareness Coverage	Percentage of events where vehicles became aware of hidden VRUs through digital twin fusion.	≥ 90% VRU awareness coverage
Improved Urban Accessibility Index	Contribution to municipal accessibility KPIs for inclusive mobility.	Qualitative improvement validated by experts
Reduced Need for Physical Infrastructure Upgrades	Value gained from digital augmentation instead of physical redesigns.	15% estimated cost avoidance
Enhanced Emergency Response Efficiency	Faster detection of risky situations enables better deployment of safety forces.	≥ 20% faster response

2.1.7 Description of the trial setup

The trial setup of T1 UC is introduced in Table 6.

Table 6: Description of T1 trial setup

<table>
<tr><th>T1 Trial</th><th colspan="2"></th></tr>
<tr><td>Trial Infrastructure / Venue</td><td colspan="2">The trial infrastructure is located in Torino (IT), Corso Castelfidardo, close to LINKS premises.</td></tr>
<tr><td>Trial Description</td><td colspan="2">Real-time intersection monitoring for risk estimation and notification to visually impaired users. Energy-aware migration of object detection and tracking tasks between RSU and MEC.</td></tr>
<tr><td rowspan="2">Technical Setup</td><td>System Components</td><td>• Solar-powered RSU with camera and LiDAR
• OBU-equipped vehicles
• Smartphone app
• mmWave RAN
• public network RAN
• private 5G core network
• public 5G core network
• MEC server
• Resource orchestrator and Closed Loop platform
• Network Digital Twin
• Risk-Level Digital Twin</td></tr>
<tr><td>System Configuration</td><td>• mmWave RAN operating at frequency 24,25-27,5 GHz
• URLLC and eMBB network slices configured
• MEC server with location and V2X services configured</td></tr>
</table>

Pre-Trial execution	**Pre-conditions**	• IoT sensors connected to the network through the IoT gateway (i.e., RSU) • Local monitoring available at the RSU • Private-public integration ready • RSU installed with camera and LiDAR • Closed Loop platform, Resource and Service orchestrator integrated • mmWave RAN operational
	Trial Steps	• Connect RSU and other UEs • Subscribe VRU to the RLDT service • Start the monitoring of the RSU
	Methodology	Check connectivity among the system components Verify the availability of the system services
	Technical Enablers	*Communication enablers*: verification of network slicing and network exposure, Public-Private integration *Compute-as-a-Service enablers*: verification service deployment on continuum *Application enablers*: RLDT service up and ready *IoT and localization enablers*: IoT exposure, MQF, PAE
	Expected Outcomes	• Connectivity verified • Availability of RLDT service verified • Monitoring data available
Trial Execution	**Pre-conditions**	• RSU with camera and LiDAR calibrated • MEC, RAN and core operational • Smartphone app and OBU connected • RLDT service is up and running • VRU registered to the RLDT service
	Trial Steps	• VRU enters the relevance area of the RLDT service • RLDT detects a dangerous situation due to the presence of other road user(s) • RLDT pushes the notification to the smartphone app of the VRU • RSU battery level is above the low threshold • Closed Loop triggers the migration of the task from the RSU to the MEC server • RSU recovers energy, and the battery level is above the high threshold • Closed Loop triggers the migration of the task from the MEC server to the RSU

	Methodology	Validation of the RLDT notification to the VRU Validation of the energy-aware migration service
	Complementary Measurements	• UL throughput • Application and network latencies • Power consumption at RSU • Positioning accuracy
	Calculation Process	• Edge-cloud-UE latencies computation • Consumption of RSU computed on timeseries data • GNSS to Fused position comparison
	Technical Enablers	*Communication enablers*: network slicing, network exposure, Public-Private integration, and network digital twin *Compute-as-a-Service enablers*: compute continuum, compute resource management *Application enablers*: RLDT operational *IoT and localization enablers*: IoT enablers, localization enablers
Expected Outcomes	• Verify the effectiveness of notification to the VRU in case of immediate danger • Validate energy-aware migration of tasks between RSU and MEC server • Verified the capability of the network to provide the desired QoS to the RSU through network slicing	

2.2 Enhancing Urban Security with Autonomous Mobile Robot Monitoring – Transport 2 Use Case

Transport 2 (T2) UC aims to support municipalities in monitoring urban walkability and accessibility for citizens with reduced mobility, as sketched in Figure 6. The solution is based on Autonomous Mobile Robot (AMR), including dual-use last-mile delivery robots, that detect and report obstacles or unsafe pavement conditions along pedestrian paths. T2 UC investigates how B5G/6G networks enable offloading perception and analysis tasks to the edge, improving monitoring accuracy while extending AMR battery life.

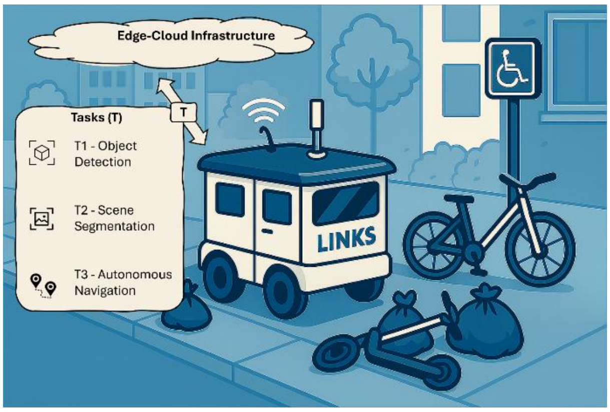

Figure 6: T2 UC overview

2.2.1 Description of the use case

T2 UC targets urban areas where continuous monitoring of walkability and accessibility is required to support citizens with reduced mobility. Sidewalks, pedestrian zones, and shared spaces often present semi-permanent obstacles, such as long-lasting construction sites, or temporary impediments including bike- or scooter-sharing vehicles, waste bins, or other objects obstructing passage. These barriers can severely affect the daily mobility of disabled or elderly individuals.

The main actor of T2 UC is an AMR, which autonomously patrols assigned pedestrian routes and collects data about obstacles, sidewalk conditions, and accessibility issues. AMRs may operate as dedicated municipal units or as dual-use platforms integrated into autonomous last-mile delivery fleets; in the latter case, the robots augment their primary mission with public-utility monitoring services. Citizens and municipal operators can be notified of critical issues through dedicated interfaces, while municipal authorities can use aggregated data to identify recurring misuse of public space or to enforce compliance among sharing-service providers.

The AMR is equipped with an advanced sensing and localization suite. Two 360° LiDAR units provide full-perimeter monitoring for safety, enabling the robot to immediately stop when people, animals, or children approach too close. Precise localization is achieved through sensor fusion of GNSS RTK, visual odometry, IMU, and wheel encoders, strengthening the position estimate even when GNSS coverage is degraded or unavailable (e.g., in urban canyons or underpasses). A solid-state LiDAR is used to enable fine-grained collision avoidance in constrained environments. An on-board unit (OBU) hosts an autonomous navigation system and provides C-ITS-enabled V2X communications. This allows the AMR to interact with connected traffic-light intersections, enabling safe, signalised road crossings and broadcasting its presence to other connected vehicles or vehicles equipped with standard C-ITS OBUs. A stereo depth camera is employed for image segmentation to delineate walkable pavement areas and to support ancillary public-utility services such as object detection (e.g., bicycles, e-scooters, waste bins) and assessment of pavement quality.

Scene understanding and surface-condition analysis are computationally intensive tasks that may significantly impact AMR energy consumption. To extend operational time, the T2 UC considers offloading these tasks from the AMR to an edge server when network conditions allow. Task migration is handled by an intelligent orchestrator that evaluates energy consumption, computational load, and network performance. When tasks such as image recognition, obstacle classification, or road-surface analysis are moved to the edge, additional uplink bandwidth is required to transmit sensor data from the AMR. The orchestrator ensures seamless migration and balanced resource utilisation across the communication and computing infrastructure.

The use case also considers scenarios in which the AMR requires teleoperation, for example, when navigating a narrow sidewalk or recovering from a blocked position. These operations necessitate a reliable, low-latency communication channel to enable real-time remote control.

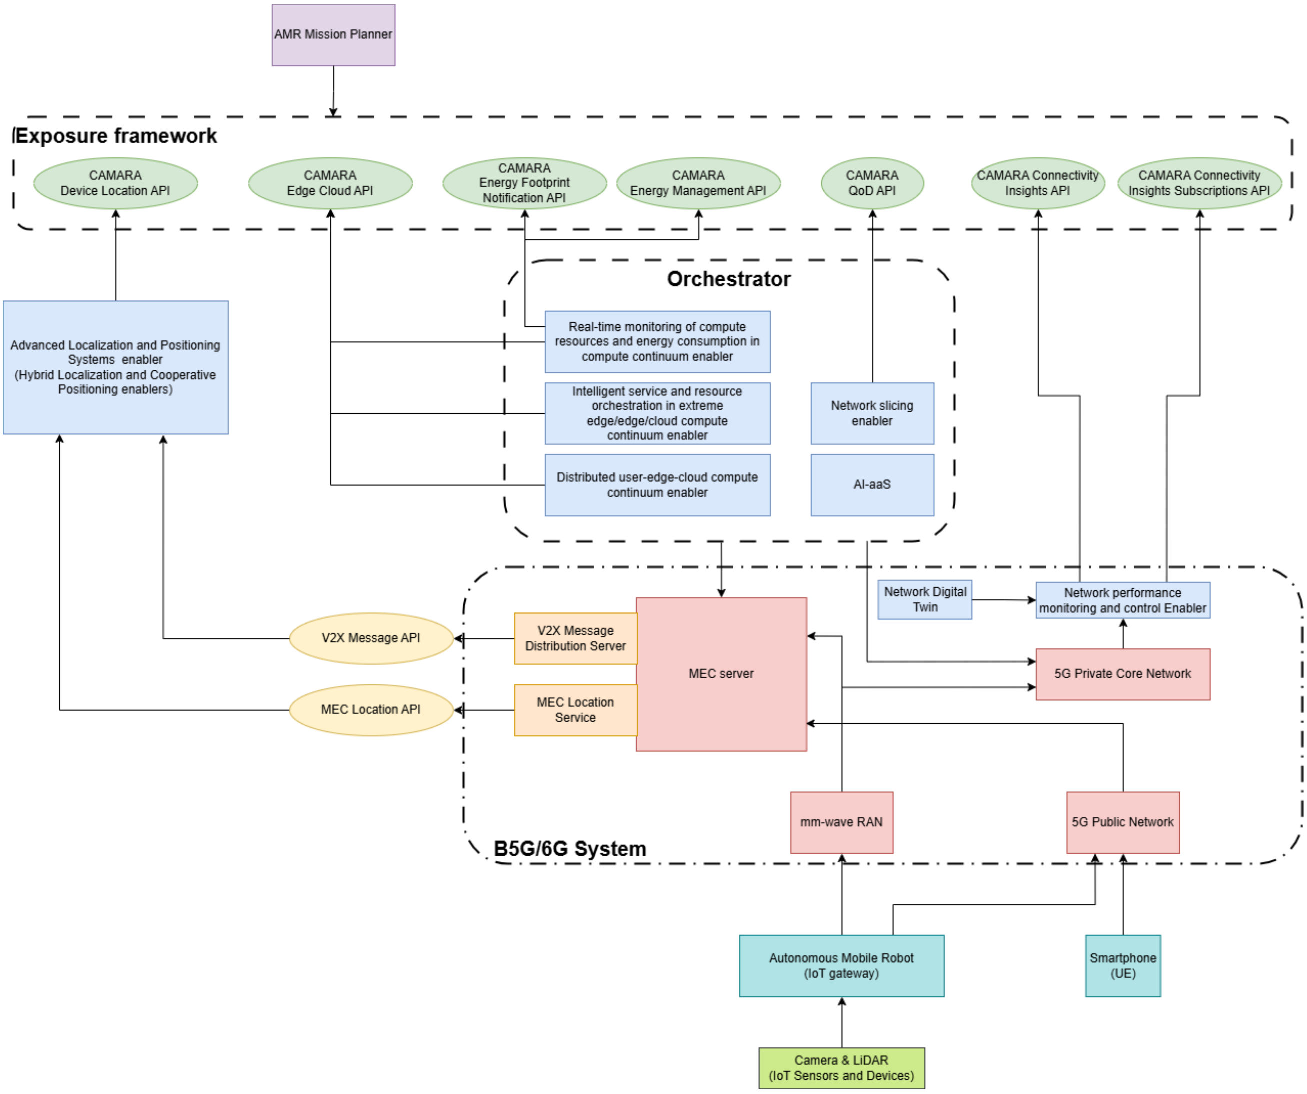

Figure 7: High-Level Conceptual Architecture of T2 UC

Figure 7 depicts the high-level architecture of the T2 UC. At the bottom, the AMR continuously monitors the urban environment and transmits sensor and telemetry data. The AMR acts as the IoT gateway for the sensors that are installed on the AMR. The AMR is connected to the B5G/6G system through an mm-Wave connection link. Another UE involved in this UC is the smartphone, which is used by visually impaired people for receiving notifications about potential obstacles. The AMR is connected to the B5G/6G system using a mm-wave connection to reach the private network. Alternatively, the AMR can be connected to a public network. Tests will be done in both configurations to assess the feasibility of T2 UC with the performance that can be offered by the two types of networks. Similar to T1 UC, the B5G/6G system for this UC includes, apart from the 5G public network and the 5G private core network, the mm-wave RAN and the MEC server. In T2 UC, the MEC server behaves as in T1 UC, exposing the MEC Location and V2X messages APIs. The orchestrator is also in T2 UC, a key element of the architecture, as it implements the enablers required for the management and orchestration functionalities. The set of network-related enablers of T2 UC is the Network performance monitoring and control enabler and the Network digital twin. An additional enabler is the Advanced Localization and Positioning Systems enabler that provides similar functionalities to T1 UC. The architecture also includes the set of CAMARA APIs for easing the interaction between the applications and the B5G/6G system. This set is made by the Device Location API, the Edge Cloud API, the Energy Footprint Notification API, the QoD API, the Connectivity Insights and Connectivity Insights Subscriptions APIs. The AMR Mission Planner is the application of T2, which is in charge of overseeing the service offered by the AMR.

The functional requirements of this UC include very low-latency communication between the AMR and edge services to guarantee the timely delivery of sensor data and the rapid detection of accessibility issues. Additionally, the B5G/6G system should support the dynamic and on-demand deployment of vertical services at the edge and offer ultra-low-latency slices with intent-based provisioning to achieve deterministic performance during operation.

The T2 UC identifies the following non-functional requirements to ensure system performance, reliability, and availability:

- End-to-end latency across the AMR–edge–server chain should target <100 ms to sustain continuous monitoring and responsive operation.
- High uplink throughput to support frequent transmission of sensor data (images, depth maps, environmental readings).
- High reliability and resilience to maintain service continuity across large urban areas.

The solution comprises several key components as summarised in Table 7.

Table 7: T2 Use Case High-Level Architecture Description

Domain & Enablers Reference	Component	UC Related Functionalities
End-user device	AMR	Monitoring of pedestrian areas, acquisition and forwarding of multi-sensor data (LiDARs, stereo camera, GNSS/IMU fusion), onboard processing for localization, perception, and collision avoidance, V2X communication for safe road crossings
End-user device	Smartphone	Position sharing, risk level notification retrieval
Application enabler	AMR Mission Planner	Supervision of AMR and obstacle warning assessment
Orchestration enabler	Orchestrator	Solution for automatic application deployment and management of network and computing resources
Connectivity infrastructure	B5G/6G system	System comprising an MEC server for application deployment, and a private and public network for providing the communication capabilities

B5G/6G Relevance

The T2 use case will rely on two main capabilities of the B5G/6G network. The first is the low-latency and reliable connectivity needed to offload critical tasks from the AMR to the edge. Different tasks require different QoS, but they all demand strict latency requirements (e.g. for AMR command data) and bandwidth requirements (e.g. for the flows coming from AMR sensors). This is even more true considering that the AMR will move on a real street with obstacles (e.g. trees) that can heavily affect the mmWave connectivity.

The second one is, as for T1, the orchestration features that should be tightly integrated with the network to provide an effective and reliable solution.

Large-scale deployment of the pilot

The scenario, as in T1, is implemented in a realistic environment with real pedestrians. The experimentation will therefore assess the B5G/6G capability to work in a relevant complex environment. Data will be collected from the private mmWave cell, and a comparison with the TIM public network performances will be performed to assess the different capabilities. Also in this UC, stress tests of the main components will be performed.

2.2.2 Deployment and implementation details

2.2.2.1 Design of the application

The focus of T2 UC is the AMR, and the main application is the AMR Mission Planner, as shown in the architecture of Figure 8. The AMR Mission Planner oversees the operations of the AMR and it interacts on the CAMARA Edge Cloud APIs for asking the deployment of specific application functions on the AMR, and, possibly, it invokes the CAMARA QoD API to ensure high-bandwidth communication link to the AMR in case that raw sensor data forwarding from the AMR to the AMR Mission Planner is required (e.g., in case human intervention is needed).

The system architecture of T2 UC is very similar to T1 UC. At the upper level, the items are the exposure framework, which is used to offer access to the CAMARA-based APIs, and the Resource and Service orchestrator, which manages the deployment and migration of application functions according to the outcomes of the instantiated closed loops. Also in this UC, an offload strategy is devised for several tasks of the AMR. A closed loop will decide, based on the AMR battery level, when and which tasks are to be offloaded to the MEC server. In the UC, different offloading approaches are experimented with depending on the tasks to be offloaded. The selected tasks are Object detection, Semantic Scene Segmentation, and Autonomous Navigation and Control. The level of complexity of migrating these tasks increases as the latency requirements become stricter, from object detection to navigation and control tasks. The other architectural components of the B5G/6G system are basically the same as T1 UC; namely, the MEC server, which hosts the Location server and the V2X Message Distribution server, besides the application functions when migrated from the AMR, and the B5G/6G network, which is made by the private and public core networks, and by the mm-Wave RAN and the public RAN.

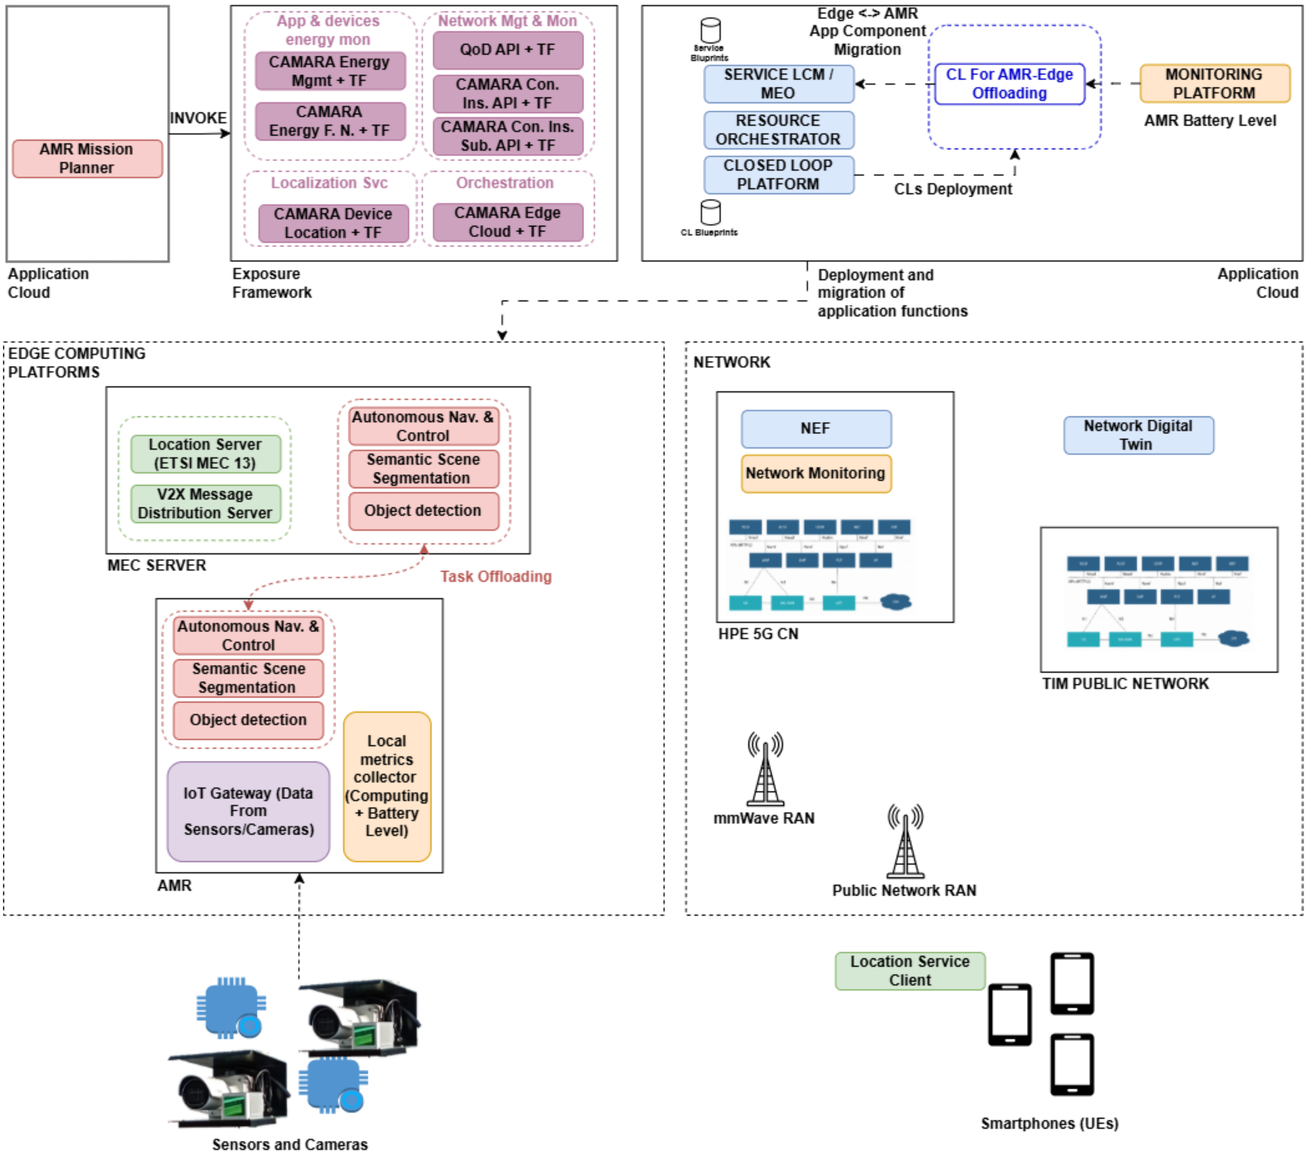

Figure 8: System Architecture of T2 UC

The workflow of T2 UC is introduced in Figure 9. The AMR Mission Planner requests the CAMARA Edge Cloud API to deploy some tasks at the AMR (step 1). The transformation function of the CAMARA Edge Cloud forwards the requests to the Orchestrator and, when tasks have been deployed, the AMR Mission Planner is acknowledged (step 2). The next step is the assignment of the mission to the AMR by the AMR Mission Planner (step 3).

At this point, two subprocesses start in the workflow:the first one is related to the main application of UC2, which is the detection of obstacles and the provision of warnings to visually impaired people, while the second one is related to the battery monitoring of the AMR.

This first subprocess begins with the AMR Mission Planner that subscribes to the CAMARA Device Location APIs for receiving updates about the UE position (step 4). The AMR then continues to perform its mission (step 5), providing information about the presence of obstacles to the AMR Mission Planner (step 6). Based on the notifications from the CAMARA Device Location API (step 7), the AMR Mission Planner verifies if a VRU is close to the identified obstacles, if any. In this case, the AMR Mission Planner notifies the VRU of the potential danger due to the obstacle's presence (step 8). If the AMR is in a critical situation (e.g., it is stuck), the AMR Mission Planner requests additional bandwidth for the AMR to the CAMARA QoD API (step 9). As the request is satisfied by the network slice manager by instantiating an eMBB slice (step 9), the AMR Mission Planner requests the sensor data from the AMR to allow human intervention (step 10-11). When the AMR is in a safe condition, the slice is released (step 12).

The subprocess on battery monitoring initiates with a periodic check of the AMR battery level (step 13). Based on the current level, the tasks are migrated from the AMR to the MEC server (step 14) or from the MEC server to the AMR (step 15).

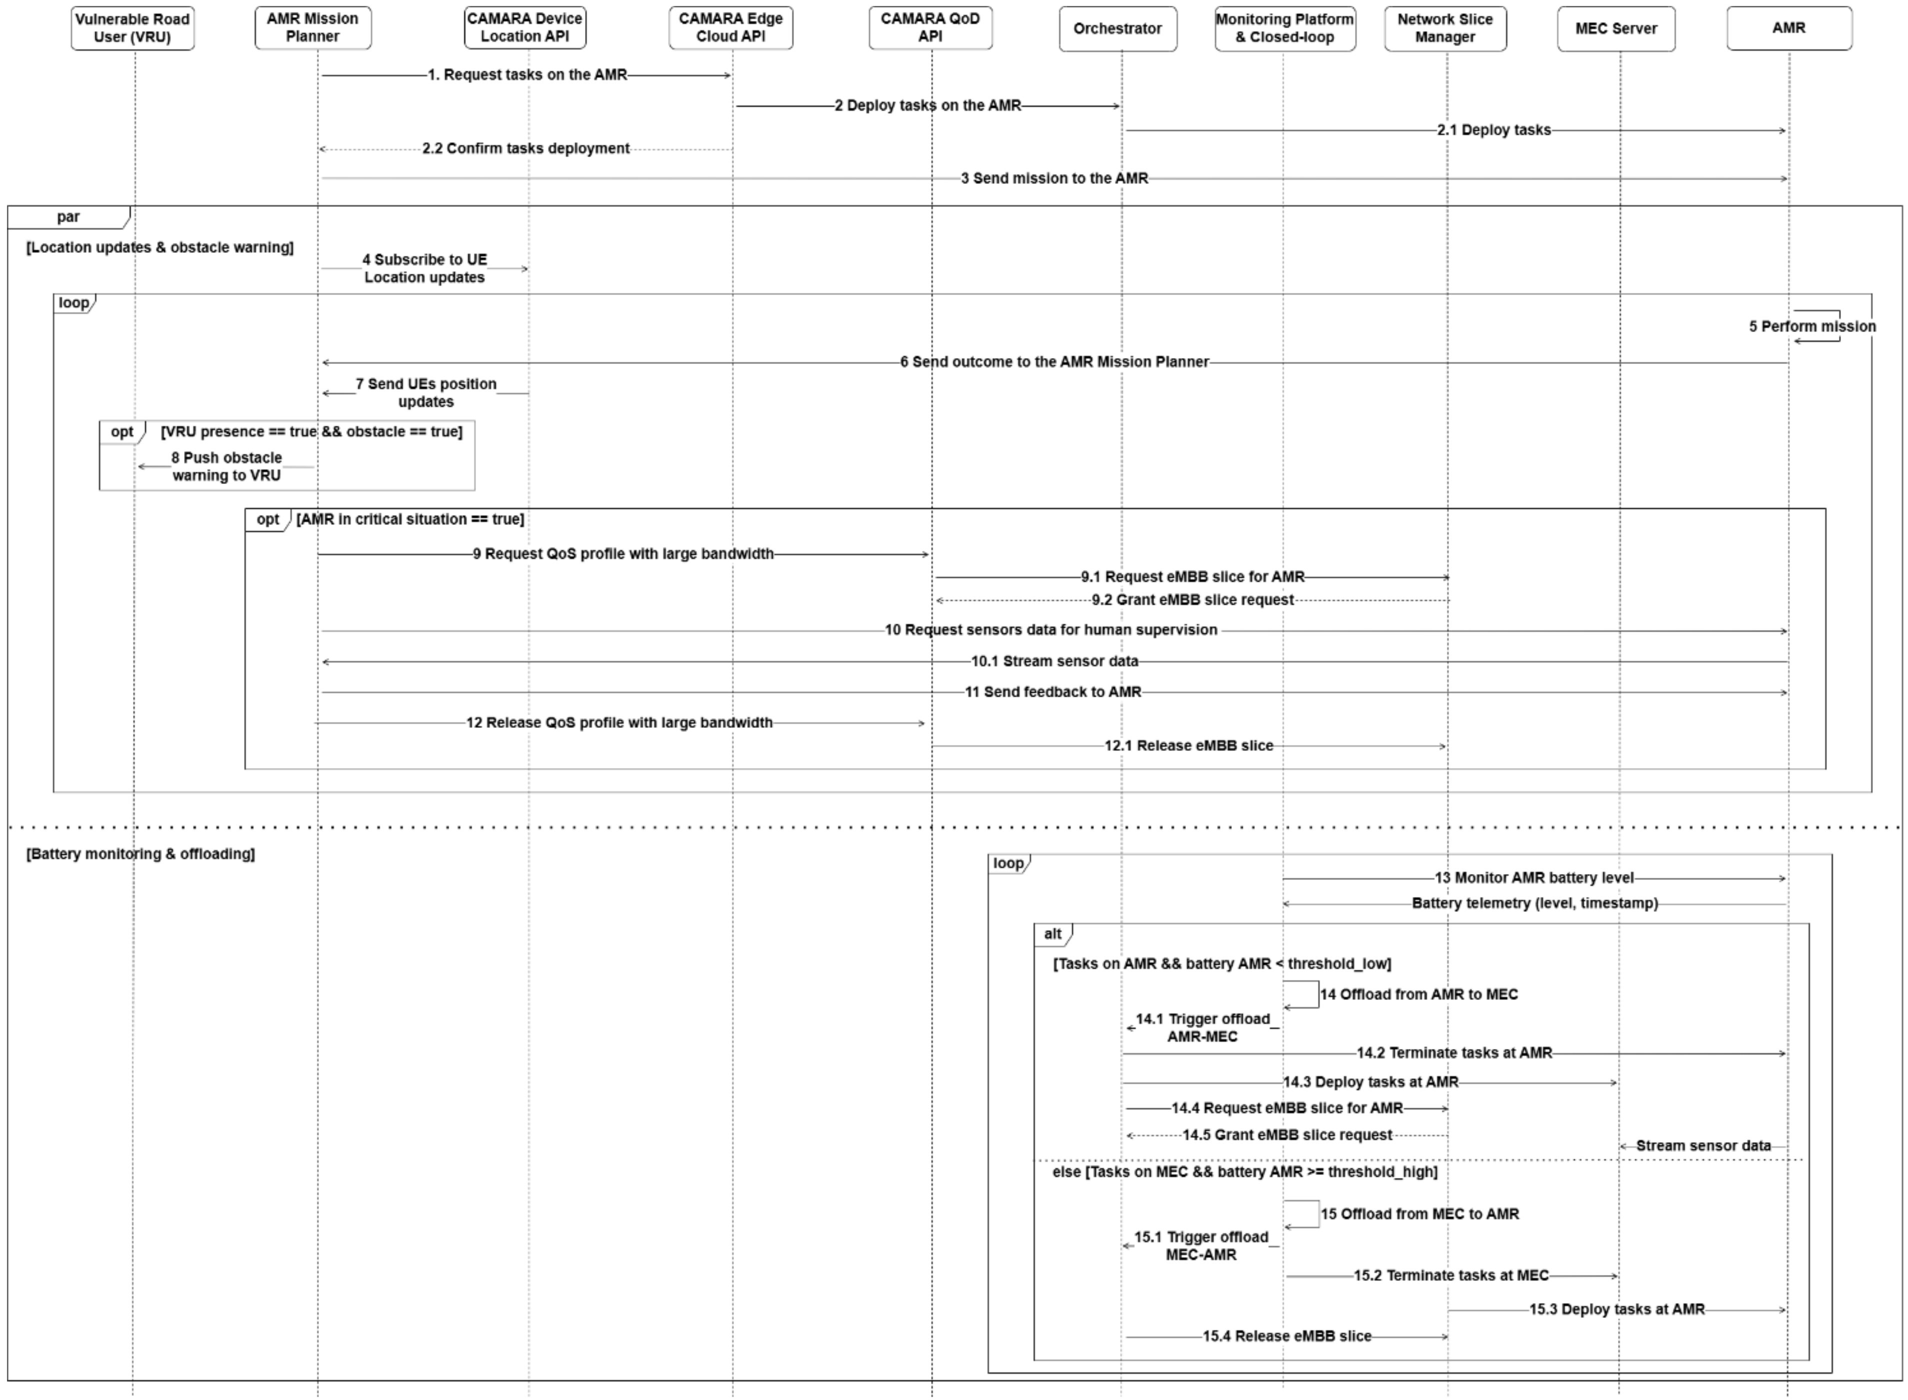

Figure 9: Workflow of T2 UC

2.2.2.2 Devices, sensors and supporting equipment

In Table 8, the devices and sensors employed in T2 UC are introduced.

Table 8: T2 Devices, sensors and supporting equipment

Category	Component	Technical Details	Units
5G connectivity module	5G mmWave CPE	5G SA/NSA, mmWave, Sub-6GHz, 4G LTE	1 x OBU
5G connectivity module (under investigation)	Waveshare Quectel RM530N-GL	IoT 5G Global Band Module, 5G Sub-6 & MmWave, M.2, w/ 3GPP 5G R16	1 x OBU
Positioning module	uBlox F9R 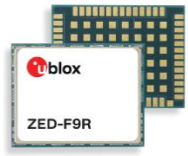	RTK-enabled multi-band GNSS receiver, Dead Reckoning	1 x OBU
End-user device	Samsung S24 	256GB AI Smartphone, Unlocked Android, 200MP, 100x Zoom Cameras, Fast Processor, Long Battery Life, Edge-to-Edge Display, S Pen, US Version, 2024, Titanium Black	1
AMR equipment	AMR Processing Unit	Local processing of perception tasks when offloading is not feasible; real-time decision-making; pre-processing of sensor data before transmission to edge services.	1
AMR equipment	On-Board Unit (OBU) with Autonomous Navigation	Autonomous path following; V2X communication for safe crossing at connected traffic lights; broadcasting AMR presence to nearby connected vehicles; reception of road-safety messages.	1
AMR equipment	AMR Communication Module	Uplink of sensor streams (LiDAR, camera, telemetry) to the edge; downlink reception of orchestration and navigation commands; support for ultra-low-latency teleoperation when required.	1

IoT sensors	360° Safety LiDARs	Perimetric monitoring for safety; human/animal proximity detection; immediate stop when obstacles approach too close to ensure safe coexistence in pedestrian areas.	1
IoT sensors	Solid-state LiDAR	High-resolution obstacle detection and collision avoidance; navigation support in constrained or cluttered urban environments.	1
IoT sensors	Stereo Depth Camera	Image segmentation for identifying walkable pavement; detection of temporary obstacles (bikes, scooters, waste bins); assessment of pavement quality (holes, cracks, unevenness).	1
IoT sensors	GNSS RTK + Visual Odometry + IMU + Wheel Encoders	High-precision localization through sensor fusion; robustness in poor GNSS coverage scenarios (urban canyons, underpasses).	1

2.2.3 Network infrastructure and functional components

The network infrastructure of T2 UC is the same as that of T1 UC. The network components are introduced in Figure 8. They consist of a mmWave RAN, a public network RAN, a private core network, and a public core network. Please refer to Section 2.1.3 for a detailed description. The Service and Resource orchestrator, the Closed Loop platform, and the monitoring platform are additional functional components. These components are introduced as well in the previous section. In T2 UC, the edge computing platforms are the MEC server and the AMR, which acts in a similar way to the RSU in T1 UC.

2.2.4 Use case enablers

As presented in deliverable D3.1, E1 enablers are presented in.

Table 9: List of enablers and sub-enablers for T2.

Category	Technology enablers	Sub-enablers
Communication enablers	Network slicing support	Network Slice as a Service (NSaaS)
Communication enablers	Network exposure and monitoring	Network Exposure Function (NWEF)
Communication enablers	Network Digital Twin (NDT)	Network Digital Twin
Compute-as-a-Service enablers	Intelligent service and resource orchestration	Centralized Service Orchestrator (CSO)

Compute-as-a-Service enablers	Intelligent service and resource orchestration	Edge Service Orchestrator (ESO)
Compute-as-a-Service enablers	Compute resource management	Compute Resource Orchestration (CRO)
Compute-as-a-Service enablers	Compute resource management	Compute Resource Allocation (CRA)
Compute-as-a-Service enablers	Compute continuum	Resource Exposure Function (REF)
Compute-as-a-Service enablers	Compute continuum	Compute Exposure Function (CEF)
Compute-as-a-Service enablers	Compute continuum	Compute Digital Twin (CDT)
Compute-as-a-Service enablers	Compute continuum	Resource Digital Twin (RDT)
Compute-as-a-Service enablers	Compute continuum	Deployment Policy Repository (DP-Rep)
Application enablers	Vertical enabler	Mission planning and control application
IoT and localization enablers	Localization and positioning	Position Acquisition Engine (PAE)
IoT and localization enablers	Localization and positioning	Mobility Context Provider (MCP)
IoT and localization enablers	Localization and positioning	Cooperative Localization Engine (CLE)
IoT and localization enablers	Vertical enabler	Message Queue Factory (MQF)
IoT and localization enablers	Vertical enabler	IoT-Exposure Function (IoT-EF)
IoT and localization enablers	Vertical enabler	IoT Resource Registry (IoT-RReg)
IoT and localization enablers	Vertical enabler	IoT Service Registry (IoT-SReg)

The set of enablers used in T2 UC is basically the same as the enablers used in T1 UC. The categories considered comprise the communication enablers (i.e., network slicing, network exposure and network digital twin), the Compute-as-a-Service enablers (i.e., intelligent service and resource orchestration, compute resource management and continuum), the localisation and positioning enabler and the IoT enabler for data collection and provisioning. In T2 UC, the vertical enabler is the Mission planning and control application, implementing the main application-level logic of the UC.

2.2.5 KPIs definition

The KPIs of T2 UC are listed in Table 10. The evaluation of these KPIs will be based on the measurements performed during the execution of T2 UC trials. The same measurement approach of T1 UC, which has been illustrated in Section 2.1.5, will be adopted.

Table 10: List of KPIs for T2

KPI name	Description/KPI definition	KPI target
End-to-end application latency	End-to-end latency of the Object warning service from sensor data provisioning to VRU notification	< 100 ms
Network latency	Latency of the communication service in provisioning the feedback from the Object warning service toward VRU	< 20 ms
Average RU power consumption	Average electrical power drawn by the mmWave RU during data transfer.	<100 Watt
Power consumption reduction vs Sub-6 GHz baseline	Relative decrease in power usage for equal throughput conditions (200 MHz mmWave vs 80 MHz Sub-6 GHz).	Target value: 20% Expected range: 15–25%, depending on beamforming efficiency, traffic burstiness, deployment density, and link conditions.
Watt per Throughput	Power-to-performance efficiency metric.	20mW/Mbps
UL throughput	Maximum achievable data rate in the uplink connection when AMR is granted for improved quality	150 Mbps
Increase in operational life of AMR	Operational lifetime gain enabled by energy-aware offloading and orchestration	+30%
Application migration latency	Time needed to migrate the task between the AMR and the MEC server	< 1 s
Positioning accuracy	Positioning accuracy of UC actors.	< 60 cm
Reliability	Percentage of time that the enablers supporting the Object Warning service are fully operational and able to deliver expected functionalities	99.999%
Scalability	Object warning service is available in stressed scenarios (e.g., number of devices and/or data volume increases)	Stable behaviour with load increase ≥100%

API Availability	The percentage of time the Exposure Framework and its CAMARA APIs are operational and accessible to developers for resource orchestration.	>90%
API Response Time	The total latency measured from the moment a request hits the API Gateway to the delivery of the processed response from the Transformation Functions.	<5s

2.2.6 KVIs definition

The KVIs of T2 UC are listed in Table 11. Similar to T1 UC, the KVIs of T2 UC address environmental and societal aspects. The environmental impact is evaluated by considering the reduced carbon footprint related to energy optimisation strategies that impact, as well as the reduced operational cost of the MNOs, thanks to a better use of the network infrastructure. Other KVIs are related to the service experimented with in T2 UC. The object warning service can reduce the accident risk of VRU, resulting in a reduction in expected healthcare expenses. Also in this UC, the KVIs are assessed using expert assessment, as it is not feasible to measure these items in a quantitative way.

Table 11: List of KVIs for T2

KVI name	Description/KVI definition	KVI target
Reduced carbon footprint	Reduction of carbon footprint through optimised deployment of applications and consequent effective resource usage of B5G/6G system elements	Reduction of 15%
Reduced TCO for MNOs	Reduced operator cost through lower energy expenditure and more efficient resource usage.	Reduction of 30%
Increased revenues for MNOs	Opportunity to offer value-added energy-aware and safety-related innovative services	Increase of 20%
Accident reduction	Reduction of accident risk thanks to higher awareness and support for VRUs	Reduction of 40%
Reduced healthcare expenditure	Economic benefit associated with accident prevention involving VRU	Reduction of 40%
Improved Walkability Perception	Users feel sidewalks and pedestrian areas are safer and more navigable thanks to AMR monitoring.	≥ 30% improvement in perceived walkability
Increased Independence in Daily Mobility	Reduction in reliance on assistance to navigate obstacles or hazardous areas.	≥ 25% reduction in assistance requests

Trust in Obstacle Notifications	Confidence that alerts are timely, correct, and relevant.	≥ 80% trust level
Improved Accessibility Awareness	Users perceive the system as contributing to a more accessible city.	≥ 75% positive perception
Improved Urban Obstacle Management	Reduction of persistent obstacles in walkways thanks to AMR detection.	≥ 20% reduction in recurring obstacles
Better Compliance of Shared Mobility Operators	Improved adherence of scooter/bike-sharing operators to proper parking rules.	≥ 25% increase in compliance
Enhanced Planning for Urban Accessibility	Use of AMR-generated datasets for long-term urban mobility planning.	Qualitative improvement confirmed by planners

2.2.7 Description of the trial setup

The trial setup of T2 UC is introduced in Table 12.

Table 12: Description of T2 trial setup

<table>
<tr><th>T2 Trial</th><th colspan="2"></th></tr>
<tr><td>Trial Infrastructure / Venue</td><td colspan="2">The trial infrastructure is located in Torino (IT), corso Castelfidardo, close to LINKS premises.</td></tr>
<tr><td>Trial Description</td><td colspan="2">Road monitoring for obstacle detection via AMR and notification to visually impaired users. Energy-aware migration of tasks between AMR and MEC.</td></tr>
<tr><td rowspan="2">Technical Setup</td><td>System Components</td><td>• AMR
• Smartphone app
• mmWave RAN
• public network RAN
• private 5G core network
• public 5G core network
• MEC server
• Resource orchestrator and Closed Loop platform
• Network Digital Twin
• AMR Mission planner</td></tr>
<tr><td>System Configuration</td><td>• mmWave RAN operating at frequency 24,25-27,5 GHz
• URLLCand eMBB network slices configured</td></tr>
</table>

<table>
<tr><td></td><td></td><td>MEC server with location and V2X services configured</td></tr>
<tr><td rowspan="5">Pre-Trial execution</td><td>Pre-conditions</td><td>AMR with integrated sensors
Local monitoring available at the AMR
Private-public integration ready
Closed Loop platform, Resource and Service orchestrator integrated
mmWave RAN operational</td></tr>
<tr><td>Trial Steps</td><td>Connect AMR and other UEs
Subscribe to the VRU for the obstacle warning service
Start the monitoring of the AMR</td></tr>
<tr><td>Methodology</td><td>Check connectivity among the system components
Verify the availability of the system services</td></tr>
<tr><td>Technical Enablers</td><td>Communication enablers: verification of network slicing and network exposure, Public-Private integration
Compute-as-a-Service enablers: verification of tasks migration and service deployment on the continuum
Application enablers: AMR Mission Planner up and ready
IoT and localization enablers: IoT exposure, MQF, PAE</td></tr>
<tr><td>Expected Outcomes</td><td>Connectivity verified
Availability of AMR Mission Planner service verified
Monitoring data available</td></tr>
<tr><td rowspan="2">Trial Execution</td><td>Pre-conditions</td><td>AMR operational
MEC, RAN and core operational
Smartphone app connected
Obstacle warning service up and running
VRU registered for the obstacle warning service</td></tr>
<tr><td>Trial Steps</td><td>AMR Mission Planner initiates the AMR mission
The AMR starts to monitor a road segment
VRU enters the area monitored by the AMR
The AMR detects an obstacle
AMR Mission Planner sends the alarm to the smartphone app of the VRU
AMR battery level is above the low threshold
Closed Loop triggers the migration of the task from the AMR to the MEC server
AMR is stuck
AMR Mission Planner requests sensor data sharing for human intervention</td></tr>
</table>

	Methodology	Validation of the Obstacle Warning Service Validation of the energy-aware migration service
	Complementary Measurements	• UL throughput • Application and network latencies • Power consumption at AMR • Positioning accuracy
	Calculation Process	• Edge-cloud-UE latencies computation • Consumption of AMR computed on timeseries data • GNSS to Fused position comparison
	Technical Enablers	*Communication enablers*: network slicing, network exposure, Public-Private integration, and network digital twin *Compute-as-a-Service enablers*: compute continuum, compute resource management *Application enablers*: AMR Mission Planner operational *IoT and localization enablers*: IoT enablers, localization enablers
Expected Outcomes	• Verify the effectiveness of obstacle warning to the VRU • Validate energy-aware migration of tasks between AMR and MEC server • Verified the capability of the network to provide the AMR with the desired QoS through network slicing	

2.3 Wireless Signalling on Rail Tracks – Transport 3 Use Case

2.3.1 Description of the use case

This use case demonstrates how B5G/6G technologies can enhance railway signalling by complementing the operation of traditional hardwired infrastructure with advanced wireless communication and sensing capabilities. The objective is to ensure safe, reliable, and resilient train operations through real-time signalling and obstacle detection. The platform integrates multiple technological enablers, including Integrated Sensing and Communication (ISAC), network slicing, and point-to-point communication between trains, deployed within the TUC testbed infrastructure.

The system comprises of two trains travelling one behind another on the same track. The front train is equipped with ISAC modules, enabling continuous monitoring of its own position, speed, and integrity. The train on the back maintains safe separation through redundant communication links, receiving signalling data via both the terrestrial B5G/6G network and a direct connection to the preceding train.

Track-side LCX cables support both communication and sensing functions by transmitting 5G signals and analysing their reflections to detect train presence and track occupancy. The O-RAN units deployed along the line act as redundant access points, ensuring reliable connectivity in case of coverage variations.

Figure 10 depicts the high-level architectural view of the T3 use case. The TUC network supports the railway communication and sensing functions, enhancing their operation with advanced capabilities.

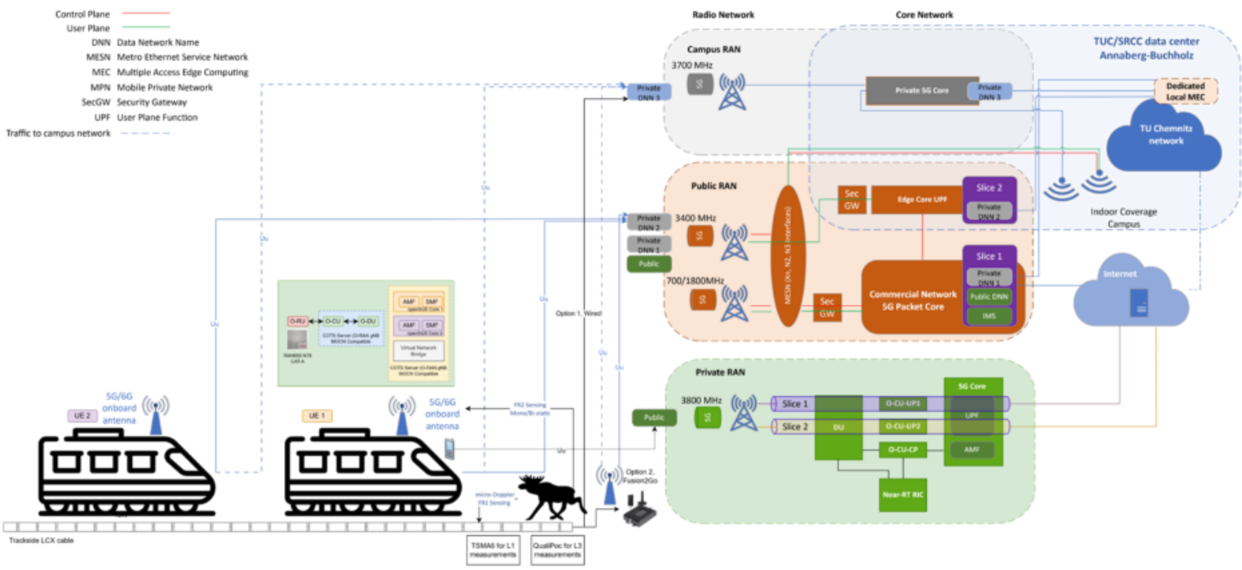

Figure 10: High-Level Architecture of T3 UC

B5G/6G Relevance

B5G/6G is particularly useful for this railway use case because it combines reliable low-latency communications with advanced sensing and real-time coordination, which are all essential for safe and resilient rail operations. In this use case, the railway system does not rely only on conventional hardwired signalling, but also on wireless train-to-ground and train-to-train communication, LCX-based distributed sensing, and ISAC-based processing to detect obstacles, monitor track occupancy, and support train integrity awareness. B5G/6G capabilities are therefore valuable because they provide the communication performance needed to exchange safety-relevant information in real time, while also enabling sensing functions to be integrated with the same radio infrastructure. In addition, features such as network slicing, edge/cloud support, and intelligent monitoring improve service continuity, scalability, and operational resilience, making the railway system more responsive to abnormal situations and better suited to future digital railway environments.

Large-scale deployment of the pilot

The large-scale character of the T3 use case is expressed through its system complexity, infrastructure distribution, and operational realism. In terms of infrastructure, T3 integrates a heterogeneous railway architecture that combines on-board train subsystems, track-side LCX sensing infrastructure, O-RAN-based radio access, 5G core functions, ISAC processing, and train-to-train / train-to-ground communications within a single operational setup. It also validates an integrated system, rather than an isolated technical component, since the use case jointly addresses and validates obstacle detection, track occupancy awareness, train integrity monitoring, and resilient safety-message delivery.

Moreover, T3 trial activities will span multiple pilot sites in more than one country, with pre-trial activities in Poland and trial activities in Germany, giving the experimentation a broader geographic and operational scope.

The T3 scenario is embedded in a realistic railway workflow involving two trains on the same track, safety-critical obstacle/event detection, and the transmission of alerts from the leading train to the following train with strict reliability and latency requirements. Moreover, the use case spans multiple deployment configurations, notably FR1 LCX-based distributed sensing and FR2/mmWave communication and sensing, which broadens the technical scope and demonstrates applicability across different railway environments. The architecture also includes O-RAN monitoring and optimisation functions, which support assessment of scalability and runtime adaptation.

In addition, the large-scale character of T3 is supported by the fact that the trials are not limited to functional validation, but also include the measurement of KPIs and KVIs, allowing the project to quantify communication performance, sensing effectiveness, reliability, and the broader operational impact of the solution.

The T3 UC identifies the following non-functional requirement:

- The system must ensure end-to-end performance according to the Vertical requirements, which include appropriate QoS for the networking aspects and scalability/resilience.

The solution comprises several key components as summarised in Table 13.

Table 13: T3 Use Case High-Level Architecture Description

Domain & Enablers Reference	Component	UC Related Functionalities
Application enabler	Train sub-system	On-board ISAC scanning, communication, and obstacle detection. Data exchange with the following train.
Edge Infrastructure	Track-side infrastructure	LCX cables for sensing. O-RAN units for redundant connectivity
Application enabler	ISAC µService	Real-time data fusion for obstacle detection and train integrity monitoring
Connectivity Infrastructure	5G/6G modem	UL/DL communication between the network and the train, supporting low latency and ensuring reliability

Connectivity Infrastructure	5G core	Provides network slicing and real-time coordination between trains, edge and cloud components.

2.3.2 Deployment and implementation details

2.3.2.1 Design of the application

2.3.2.1.1 ISAC-based obstacle detection

The T3 application is designed as a sensing-assisted railway communication and control service that reuses standard 5G NR communication signals to jointly provide reliable connectivity and environment perception capabilities. The communication stack remains fully compliant with 3GPP NR specifications, while sensing functionalities are realised through non-invasive observation and processing of communication signals without introducing dedicated sensing waveforms or requiring modifications of PHY or MAC layers. In this way, sensing capabilities are integrated on top of standard communication procedures, ensuring compatibility with existing infrastructure.

The considered deployment supports operations in both FR1 (sub-6 GHz) and FR2 (mmWave) frequency ranges, enabling complementary communication and sensing capabilities adapted to railway deployment environments. Communication remains the primary function, supporting continuous connectivity between trains and infrastructure, while sensing outputs enhance operational awareness and safety.

The operational scenario involves two trains travelling on the same track. The front train is equipped with ISAC capabilities, enabling continuous monitoring of track conditions ahead. When obstacles or abnormal events are detected, this information is transmitted to the following train via point-to-point communication in addition to network-assisted signalling, allowing sufficient reaction time for safety actions. Hence, communication and sensing jointly support safe and resilient train operations.

FR1 Deployment – LCX-Based Distributed Sensing

In FR1 deployments, communication and sensing functionalities are supported through Leaky Coaxial (LCX) cables installed along railway tracks. LCX infrastructure extends the communication coverage inside tunnels and along track segments where conventional portal-based antenna systems provide limited coverage. These distributed radiating cables simultaneously enable communication and passive sensing capabilities.

During deployment, 5G transmission and reception are activated via the LCX infrastructure to support continuous communication along the railway line. In parallel, sensing measurements are obtained by analysing variations of the received signals and generating micro-Doppler spectrograms reflecting train motion or the presence of objects along the track.

For model training, an auxiliary ground-truth generation approach is temporarily deployed to provide reference information on train presence and location. This may include automated sensing systems (e.g., camera-based detection) or manual annotation of recorded data. The resulting ground-truth labels are synchronised with LCX-derived sensing measurements to train supervised detection models.

In parallel, network KPIs are monitored using measurement equipment such as the R&S®TSMA autonomous mobile network scanner to ensure compliance with communication performance requirements.

Once training and validation are completed, the auxiliary ground-truth generation system is removed, and sensing operation relies solely on LCX-assisted measurements. Subsequent trials demonstrate autonomous train detection and track occupancy monitoring using LCX-based sensing only, while maintaining uninterrupted communication services.

FR2 Deployment – Microamp mmWave Platform

In FR2 deployments, sensing and communication experiments are conducted using a Microamp mmWave setup in which a standard 5G NR over-the-air connection is established between a Microamp CPE acting as UE, and a RAN server acting as gNB through beamformed radio links. Communication performance is first verified under standard 3GPP-compliant operation to ensure stable high-data-rate connectivity.

FR2 communication provides gigabit-class data rates using highly directional beamforming, supporting both infrastructure connectivity and data exchange between railway elements. At the same time, these directional transmissions and large bandwidth signals enable sensing capabilities with high spatial resolution. Depending on geometry and deployment configuration, the system can operate in monostatic or bistatic sensing modes, detecting objects in front of trains or within infrastructure coverage areas.

Sensing measurements are realised by passively logging uplink communication data at the RAN side while maintaining normal communication services. Frequency-domain IQ samples and PHY-layer metadata are exported through logging interfaces without interrupting communication sessions. Controlled operational scenarios are executed, including UE motion, obstruction events, LoS and NLoS conditions, and static or dynamic scenes.

Captured communication data is processed offline to extract channel observables from standard reference signals such as DMRS, SRS, and CSI-RS aligned with frame and slot timing. From these observables, sensing features such as temporal channel variations, Doppler-related signatures, and motion indicators are derived. The results are correlated with known scenario events to validate sensing performance.

Experiments are repeated using different beam configurations, UE positions, and traffic loads to evaluate robustness and reproducibility. Longer-duration runs are finally conducted to demonstrate continuous sensing feasibility during normal communication operation without requiring additional sensing hardware.

2.3.2.1.2 Point-to-Point Communication

The train-to-train communication for the purpose of obstacle notification is achieved using a private 5G Standalone (SA) network. The objective is to ensure low-latency propagation of obstacle detection events from the leading train to the following train through a 5G core-mediated communication path.

The system supports real-time obstacle detection event generation in the leading train, the transmission of structured safety messages via 5G NR (n78 band), core-network-mediated routing to the following train, and finally, the reception, interpretation, and execution of safety response (e.g., braking trigger).

The application is deployed in a distributed manner, see Figure 11, below.

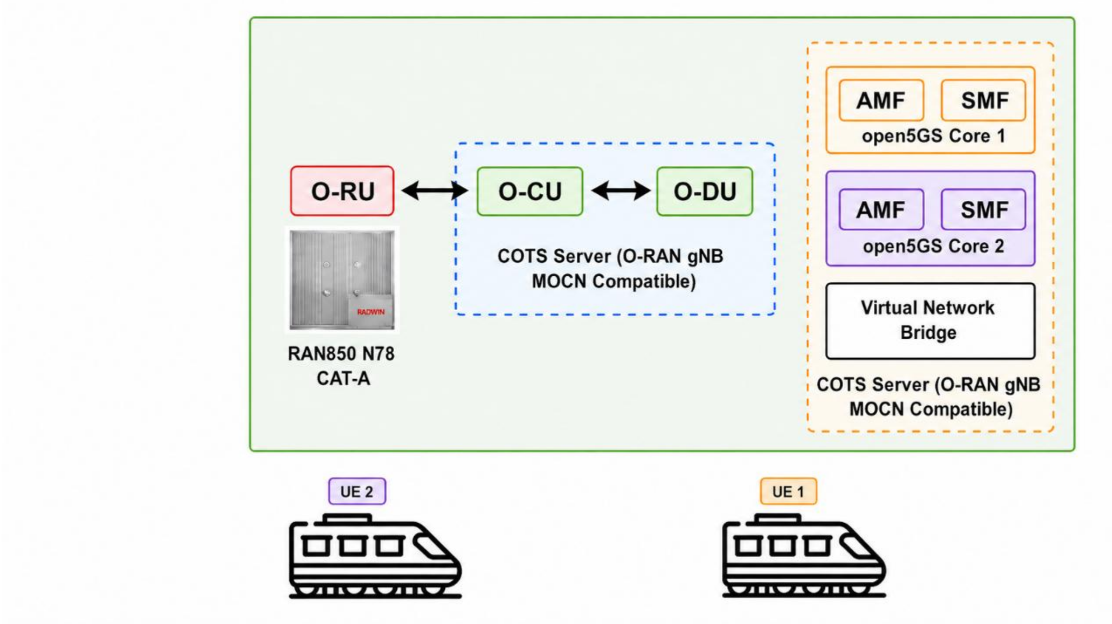

Figure 11: Point-to-Point Communication System

The leading train hosts: i) two 5G Cores (Open5GS), ii) the O-RAN Benetel RAN650 Radio Access Network (RAN) shared by the two cores (MOCN functionality), and iii) one Commercial off the shelf (COTS) User Equipment (UE) device. The following train hosts a COTS UE device.

2.3.2.1.3 Liquid RAN

This section contains the description of IS-Wireless proprietary solution: Liquid RAN. Liquid RAN consists of an O-RAN Alliance compliant RAN network and near-RT RIC with xApps, which serve as an application for the analysis and optimisation of the RAN network and can be utilised as an API to connect 3rd party enablers.

The IS-Wireless Near-RT RIC (part of the Liquid RAN product) is an O-RAN–compliant Near–Real-Time RAN Intelligent Controller designed to enhance and optimise RAN performance. As illustrated in Figure 12, it connects to E2 nodes—including O-DUs and O-CUs—as well as to xApps and the Non-RT RIC through standardised O-RAN interfaces, namely E2, xApp APIs, and A1. The Near-RT RIC itself is provided as a containerised software package based on Docker. By default, it is deployed using Docker Compose, with support for alternative orchestration platforms such as Docker Swarm and Kubernetes, enabling flexible, scalable deployment across diverse environments.

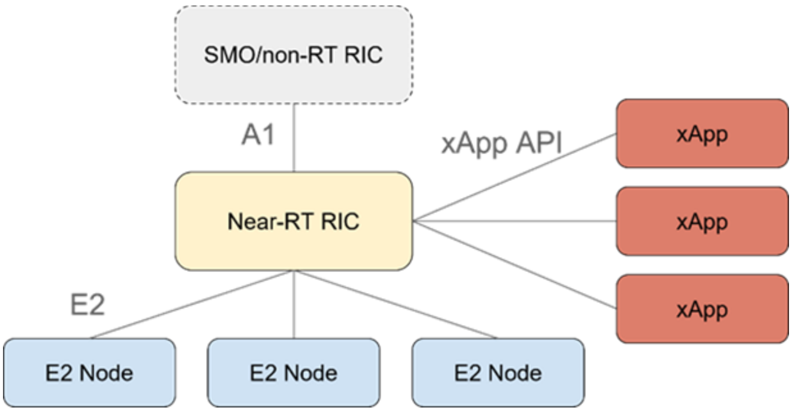

Figure 12: IS-Wireless Near-RT RIC high-level interfaces.

The basic containerised deployment of Liquid RIC utilises Docker containers orchestrated via Docker Compose. In a single-node deployment, management is handled through a single YAML file, which significantly simplifies network configuration. The deployment workflow involves pulling the Near-RT RIC Docker images from DockerHub and configuring the system using a Docker Compose YAML file. This file specifies container names, inter-container connectivity, and the overall topology of the Near-RT RIC environment.

Liquid RIC can be executed either in console mode or in detached mode, where containers run in the background. All containers, including the Liquid RIC container, generate logs that can be accessed using standard Docker commands. These logs are written both to the container's system console and to dedicated log files.

Liquid RIC also provides a Graphical User Interface (GUI) for managing the Near-RT RIC and the KPM xApp. Through the GUI, users can monitor Near-RT RIC and RAN performance and perform xApp lifecycle operations such as adding or removing xApps. The GUI offers multiple views to explore different aspects of Liquid RIC, including Home, xApps, xApp, E2 Nodes, E2 Node, Cells, Cell, Logs, and Settings. **Figure 13** presents the Home view, which serves as the main dashboard.

Figure 13: Home view of Liquid RIC GUI

xApps are modular, cloud-native applications that run on the Near-Real-Time RAN Intelligent Controller (Near-RT RIC) and are designed to perform near–real-time radio access network optimisation and control. They use standardised O-RAN interfaces, such as E2, to collect RAN data and apply control actions with latency requirements typically ranging from 10 ms to 1 s. By implementing specific use cases—such as traffic steering, interference management, mobility optimisation, and performance monitoring—xApps enable intelligent, flexible, and vendor-neutral RAN optimisation. In the T3 use case, ISRD provides three xApps: KPM xApp, Energy Efficiency xApp and PRB Quota xApp.

The KPM xApp is a native component of the Liquid Near-RT RIC. It subscribes to O-RAN standardised Key Performance Measurements (KPMs) from all cells, such as throughput or PRB utilisation, which are stored in the Valkey database. Both the KPM collection and reporting intervals are set to 1 second. The collected KPMs can be visualised through the Liquid Near-RT RIC GUI as well as the dedicated KPM xApp

GUI. The KPM xApp interface is based on Grafana and is accessible via a web browser, offering multiple dashboards tailored to different RAN vendors. An example of such a dashboard is shown in **Figure 14**.

Figure 14: Grafana dashboard with ISRD KPMs

The Energy Efficiency xApp is designed to improve RAN energy efficiency by dynamically managing server operation based on traffic conditions. Operating within the Near-RT RIC, the xApp continuously monitors commonly available RAN indicators, such as time information and the number of active UEs per gNB. Based on these inputs, it determines when network load is low and triggers power-saving actions by transitioning part of the RAN infrastructure into a reduced-power state. When traffic demand increases, the xApp restores normal operation to ensure service continuity. The xApp's primary outputs are control actions that enable or disable server operation, supporting adaptive and energy-efficient RAN management using standard O-RAN network parameters.

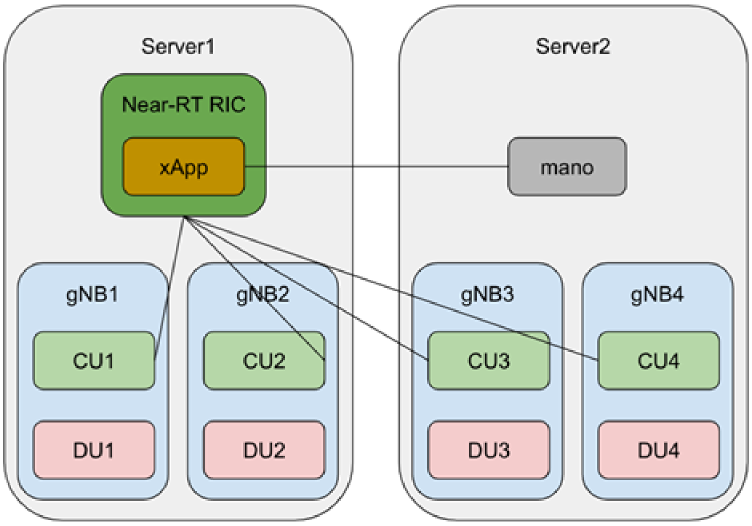

Figure 15: Energy Efficiency xApp architecture

The Slice-level PRB quota xApp adjusts the number of Physical Resource Blocks (PRBs) allocated by the MAC scheduler to UEs belonging to different network slices. It performs this function by issuing RAN

Control messages in accordance with the Control Service Model, specifically Control Service Style Type 2 and Control Action ID 6 for slice-level PRB quota control.

To support its operation, the xApp relies on performance indicators collected via the KPM service model using REPORT Service Style Type 1. These indicators include DRB.UEThpDl.SNSSAI for downlink UE throughput, DRB.UEThpUl.SNSSAI for uplink UE throughput, RRU.PrbTotDl for downlink PRB utilisation, and RRU.PrbTotUl for uplink PRB utilisation. In addition, the xApp requires information about the network slices present in the E2 node, identified by their S-NSSAI values, which is obtained from the NetInfo service within the Near-RT RIC.

2.3.2.1.4 Workflow

The workflow of T3 UC is illustrated in Figure 16, below. The interaction starts with the Leading Train, which registers the mission and requests activation of the service from the ISAC µService. In parallel, the ISAC µService requests a railway QoS profile from the Main 5GC, which in turn configures the RAN resources for prioritised railway traffic. At the same time, the Leading Train requests the setup of a P2P safety communication path from the Leading Train onboard 5GC. The onboard 5GC then configures the Leading Train onboard RAN for P2P safety communication. Once this is completed, the onboard RAN establishes the communication session with the Following Train, which acknowledges the establishment of the session. In parallel, the Main 5GC confirms to the ISAC µService that the QoS profile and connectivity are ready.

The workflow then includes an optional commissioning and validation phase. In this phase, the sensing infrastructure provides additional reference inputs to support system training and validation. In particular, the FR1 sensing chain provides ground-truth labels and KPI samples, while the FR2 sensing chain provides reference measurements. These data are used only during validation and commissioning and are not part of the normal operational loop.

During normal operation, the workflow enters a continuous monitoring loop. In this loop, the Leading Train sends its telemetry and onboard ISAC observations to the ISAC µService. In parallel, the FR1 sensing subsystem provides LCX distributed measurements, while the FR2 sensing subsystem provides NR IQ samples and PHY metadata. The ISAC µService fuses these heterogeneous inputs and extracts sensing-related features, such as channel characteristics and micro-Doppler signatures. Based on this fused information, it performs detection functions for obstacles, tracks occupancy, train presence, and train integrity issues. In parallel, the RAN exports KPI and QoS metrics to the Near-RT RIC / monitoring entity, which may respond with PRB quota adjustments and energy-efficiency optimisation actions.

The decision stage is represented as an alternative branch with two possible outcomes. In the first case, no anomaly is detected. The ISAC µService then issues a periodic status update indicating that no anomaly is present, and the workflow returns to the monitoring loop to continue normal supervision. In the second case, an obstacle or abnormal event is detected. The ISAC µService generates a structured safety event/obstacle alert and sends it to the Leading Train onboard 5GC. The onboard 5GC then activates the onboard P2P delivery mechanism through the Leading Train onboard RAN, which transmits the obstacle notification to the Following Train. The Following Train validates the received event and returns an acknowledgement/validation message. Based on the notification, the Following Train can then execute the appropriate safety response, such as braking or speed reduction, while the monitoring entities keep logging and supervising the communication conditions.

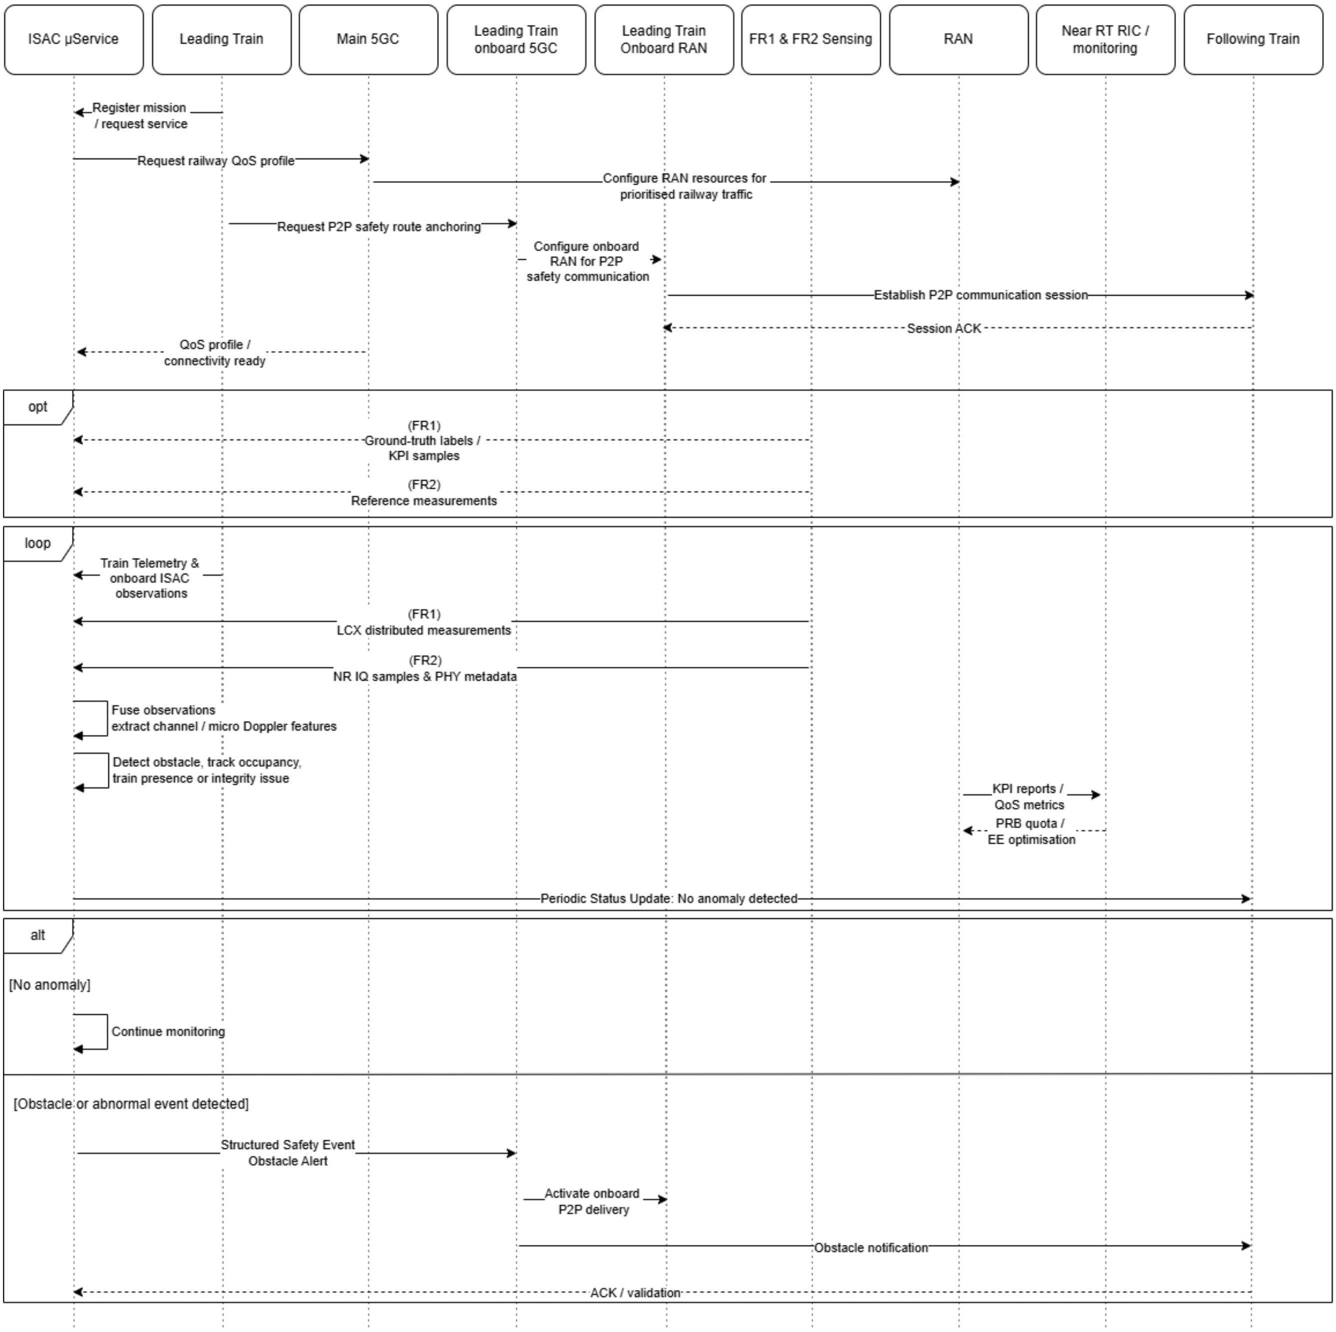

Figure 16: T3 UC workflow

2.3.2.2 Devices, sensors and supporting equipment

In Table 14, the devices and sensors employed in T3 UC are introduced.

Table 14: T3 Devices, sensors and supporting equipment

Category	Component	Technical Details	Units
5G communication unit	1U type rack server	Microamp CellBox Air integrated RAN server	1

5G Communication Unit	RU	Microamp CellBox Air 5G mmWave Radio Unit (n257/n258)	1
5G Communication Unit	PC server	CN server + Display	1
User Equipment	CPE	Microamp M202 5G mmWave CPE (n257/n258)	1
5G communication unit	O-RU	Benetel 550 or Benetel 650	1 (Liquid RAN), 1 (Point-to-Point Communication)
Hardware	Server	U1/U2 type rack server	1 (Liquid RAN), 2 (Point-to-Point Communication)
User Equipment	Smartphone	e.g., Huawei P40 Pro, OPPO Reno4 Z 5G	1
User Equipment	SIM card	Commercial SIM Card	1 (Liquid RAN), 2 (Point-to-Point Communication)

2.3.3 Network infrastructure and functional components

2.3.3.1 ISAC-based obstacle detection

Communication Layer (COM)

The standard 5G NR RAN (PHY/MAC/RLC/PDCP) transmits legacy NR synchronisation and reference signals, including SSB (PSS/SSS/PBCH), DMRS, and CSI-RS, as part of normal operation. No modifications to frame structure, scheduling, or waveform generation are introduced. Communication services remain fully standard-compliant and support both infrastructure connectivity and inter-train information exchange, including the transmission of sensing alerts from the front train to the following train.

Measurement and Data Extraction Layer (MDE)

A non-invasive measurement interface passively taps frequency-domain IQ samples and/or channel estimation outputs together with timing and configuration metadata such as slot and symbol indices, PRB allocation, beam identifiers, antenna ports, and TDD DL/UL state. The interface operates without affecting NR transmission behaviour or QoS guarantees, ensuring that sensing data acquisition remains transparent to communication services.

Sensing Processing Layer (SPR)

A parallel sensing chain processes reused NR signals to derive sensing products while communication continues normally. Processing adapts to the operating frequency range:

- **FR1 operation (sub-6 GHz):**
  Sensing primarily exploits temporal channel variations and Doppler signatures, enabling micro-Doppler-based detection and tracking of obstacles, moving objects, and train integrity-related signatures. DMRS and, when available, CSI-RS signals are reused to provide sufficient temporal density for Doppler analysis, particularly in LCX-based distributed sensing deployments along railway tracks and tunnels.
- **FR2 operation (mmWave):**
  Sensing leverages large bandwidth and highly directional beamforming, enabling CSI-based super-resolution in delay and angular domains. CSI-RS beam sweeps and multi-antenna channel snapshots are reused to support high-resolution range, angle, and blockage-aware sensing, including MUSIC- and ESPRIT-based processing for obstacle and motion detection ahead of trains.

Application and Data Fusion Layer (ADF)

Sensing outputs originating from train-mounted and track-side sensing units are fused in real time to support obstacle detection, track occupancy estimation, and train integrity monitoring. Detected hazards or anomalies are delivered to railway control systems and communicated to following trains via P2P or network-assisted communication, enabling timely safety reactions. Application interfaces are exposed via REST and gRPC APIs for monitoring, configuration, and KPI reporting, allowing integration with higher-layer railway management and operational systems.

2.3.3.2 Point-to-Point Communication

The system for the point-to-point train communication is based on a 5G Standalone (SA) private deployment operating in the n78 band.

It consists of:

- srsRAN 5G Enterprise (gNB implementation)
- Multi-core Open5GS deployment
- O-RAN compliant RAN devices (Benetel RAN650)
- Commercial off-the-shelf (COTS) UEs with customised USIM cards

The 5G Core is implemented using Open5GS and deployed in a multi-core configuration. Each core includes Access and Mobility Management Function (AMF) and Session Management Function (SMF). Their capabilities include independent operation, multi-core switching scenarios and MOCN support. The core is physically located in the leading train.

The RAN includes: srsRAN 5G Enterprise gNB, and Benetel RAN650 (DU/RU configuration). It supports O-RAN compliant interfaces, distributed unit testing, flexible deployment topologies and coverage ensuring second train connectivity.

Two UEs are deployed. UE 1 is located on the leading train, while UE 2 is located on the following train. They are both COTS devices, with customised USIM cards, registered to the private 5G core.

The second train remains within network coverage and communicates with the first train via the core network.

2.3.3.3 Liquid RAN

Figure 17 presents the O-RAN functional architecture within which the Liquid RAN is positioned. It consists of several core components: the 5G O-DU (Open RAN Distributed Unit), the 5G O-CU (Open RAN Centralized Unit), the 5G Near-RT RIC (Near–Real-Time RAN Intelligent Controller) hosting multiple xApps, and the 5G Core Network (5GC).

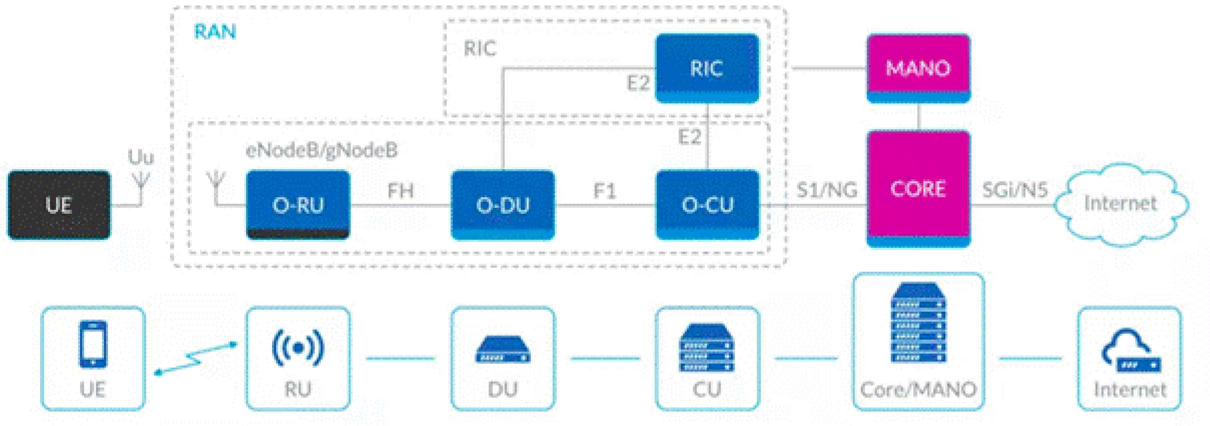

Figure 17: O-RAN functional architecture

The ISRD approach enables flexible allocation and dynamic migration of 5G NR protocol stack functions between the O-DU and O-CU. The ISRD Liquid RAN O-CU and O-DU are advanced implementations of the O-RAN Centralized and Distributed Units. In addition, the ISRD Liquid RAN Near-RT RIC is an O-RAN–compliant Near–Real-Time RAN Intelligent Controller that provides advanced logical control and optimisation of RAN components, including the O-DU and O-CU. This is achieved through comprehensive data collection and timely control actions over the E2 interface, supporting near–real-time management of RAN performance and resource utilisation, as described above.

The deployment configuration depends on the specific requirements of the end user, but typically includes multiple instances of the O-DU and O-CU. The 5G Core (5GC) and either a commercial RU or a USRP-based RU can also be included when needed. In addition, the deployment may incorporate the Near-RT RIC along with relevant xApps, such as the RAN KPM xApp. All Liquid RAN components are delivered as containerised applications and can be deployed either directly on bare-metal servers or as pods within a Kubernetes environment.

In this architecture, the RAN components interconnect with other 5G network elements as defined below. Liquid RAN supports a wide range of commercial User Equipment (UE) brands such as Oppo, Samsung, and OnePlus. UE can connect to either a commercial RU (for example, Benetel 550) or a COTS RU equipped with a USRP (b210) serving as the radio front end, with UEs typically placed inside an RF shield box for test isolation. The O-DU interfaces with commercial RUs via the Open Fronthaul (Open-FH) interface, while communication with COTS RUs occurs over a proprietary fronthaul link. The O-CU connects to the 5G Core Network (5GC) through standardised N2 and N3 interfaces. Additionally, xApps expose external APIs that provide access to RAN-related metrics, enabling enhanced monitoring and control capabilities.

Liquid RAN supports multiple configuration options, including frequency bands n77, n78, and n79, bandwidths ranging from 10 to 100 MHz, and antenna setups such as SISO and 2×2 MIMO in the downlink. Full deployment requires servers with adequate computing resources, with capacity primarily determined by the selected bandwidth.

Another feature of the Liquid RAN is the 5G network slicing, which enables multiple virtual networks to run on the same physical infrastructure, each tailored to specific service requirements such as latency, throughput, reliability, and security. This allows operators to support very different use cases simultaneously, for example, enhanced mobile broadband, massive IoT, and ultra-reliable low-latency communications. Within the AMAZING-6G project, Use Case T3 will include a single base station supporting two slices for different groups of users, e.g. train passengers and railway staff. 5G base station, through RAN Intelligent Controller's (RIC) xApp will have the capability to assign different amounts of radio resources per slice depending on the instantaneous traffic demand per slice. This setup is presented in Figure 18.

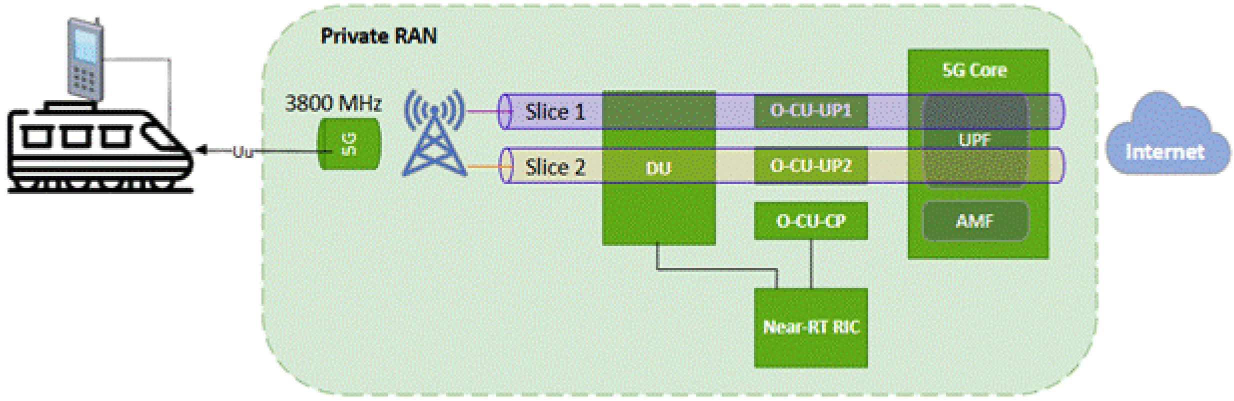

Figure 18: T3 Slicing architecture

2.3.4 Use case enablers

As presented in deliverable D3.1, T3 enablers are presented in Table 15.

Table 15: List of enablers and sub-enablers for T3.

Category	Technology enablers	Sub-enablers
Communication enablers	Network Slicing	
Communication enablers	Integrated Sensing and Communication	Reuse of NR reference signals for sensing; CSI/IQ-based channel sensing; cooperative train & trackside sensing; object/obstacle detection.
Communication enablers	Multiple-RAT connectivity	Sidelink
IoT and localization enablers	Data ingestion and telemetry	Data collection from sensors and devices
Communication enablers	RAN network programmability (ISRD)	
Communication enablers	Network Performance monitoring and control (ISRD)	
CaaS enablers	Energy-efficient compute management (ISRD)	

CaaS enablers	Cloud native service design (ISRD)	

2.3.5 KPIs definition

Table 16: List of KPI for T3 UC

KPI name	Description/KPI definition	KPI target
Service RTT	The total of "Service latency" and "AI Service RTT", where applicable, for obstacle detection	<100ms
Energy Efficiency	Energy consumption reduction due to the active management of the sleeping/idle mode of the RAN equipment.	>20%
Policy enforcement delay	The delay between policy request (Near-RT RIC) and its implementation (RAN) with regards to slice - PRB quota association	<5s
User experienced datarate	DL datarate	>100 Mbps
User experienced datarate	UL datarate	>20 Mbps

2.3.6 KVIs definition

Table 17: List of KVIs for T3 UC

KVI name	Description/KVI definition	KVI target
Minimal disruption of rail services	Assesses the reduced disruption time that rail services experience with the use of 5G/6G for signalling	12–20% disruption-time reduction
Lower cost of restoration/repairs in case of equipment fault or theft	Assesses the variation of the cost of restoration/repairs in case there are faults in the equipment or theft	8–15% restoration-cost reduction
Lower loss of revenue for the railway operator due to the minimal service disruption	Assesses the financial benefit as a result of the reduced disruption time in rail services	5–12% revenue-loss reduction

2.3.7 Description of the trial setup

2.3.7.1 ISAC-Based Obstacle Detection

Table 18: Description of T3 trial setup for ISAC-Based Obstacle Detection

T3 Trial		
Trial Infrastructure / Venue	TUC testbed (Annaberg-Buchholz) FR1 – LCX and mmWave FR2	
Trial Description	This trial demonstrates seamless, efficient, and secure data communication between trains, infrastructure, and ground systems to support railway operations beyond safety-critical functions. It enables real-time telemetry, system monitoring, and operational data exchange across the railway environment, addressing selected data collection and telemetry use cases defined by the International Union of Railways (UIC). These capabilities are aligned with the objectives of the Future Railway Mobile Communication System (FRMCS), the future worldwide railway telecommunication system specified by UIC, which is designed to provide a secure, IP-based, and broadband-capable platform for high-volume telemetry and non-critical operational data exchange. The objective is to ensure continuous reporting of status, diagnostic, and performance information from trains and infrastructure to ground systems, while also supporting the distribution and transfer of non-critical data such as monitoring and operational records. These communications emphasise automation and efficiency with minimal human intervention, and they optimise network resource usage by scheduling bulk data transfers during periods of low demand, for example, when trains are stationary or near stations. Overall, this approach supports the digital transformation of railways by improving situational awareness and operational efficiency while preserving interoperability with legacy systems such as GSM-R.	
Technical Setup	**System Components**	• LCX cables, 5G CN, RAN server, R&S®TSMA, Signal and spectrum analyser, USRP, Microamp CPE (UE), FR2 radio units, camera-based object detection system.
	System Configuration	n78 band, mmwave bands (n257, n259, n261),40MHz and 100 MHz bandwidth (n78 band), 800 MHz bandwidth (mmwave bands), uplink-based sensing, beamformed links.
Pre-Trial execution	**Pre-conditions**	Equipment (LCX, Microamp CPE, servers, RU, UE, 5G CN) and software were deployed and configured in the lab.
	Trial Steps	Lab deployment: **FR1** **Step 1:** Install the LCX along the measurement area. **Step 2:** Connect the LCX to the gNB.

		Step 3: Verify stable communication performance. **Step 4:** Deploy an auxiliary ground-truth generation system covering the LCX-monitored area (e.g., camera-based sensing or manual annotation). This is used during the training phase only. **Step 5:** Capture 5G signals through the LCX. **Step 6:** Measure 5G KPIs, channel parameters (e.g., CIR) using the R&S®TSMA, and raw Doppler-related data (i.e., complex IQ samples) using a spectrum analyser and/or USRP. **Step 7:** Capture signal reflections caused by moving objects (e.g., humans, animals, trains). **Step 8:** Generate micro-Doppler spectrograms from the collected data. **Step 9:** Use ground-truth data to label the sensing measurements and train models to identify consistent patterns corresponding to different object types. **Step 10:** Perform LCX-only sensing trials using a trained model for autonomous detection and track occupancy monitoring. **FR2** **Step 1:** Establish an FR2 5G NR over-the-air connection between the Microamp CPE (UE) and the RAN server (gNB) using the deployed radio units and beamformed links. **Step 2:** Verify stable communication performance under standard 3GPP-compliant operation. **Step 3:** For uplink-based sensing measurements, enable passive logging of frequency-domain IQ samples and PHY-related metadata at the RAN side while maintaining normal communication service. **Step 4:** Define and execute controlled operational scenarios in the test environment (e.g., pilot object, UE movement, obstruction events, LoS vs. NLoS conditions, static vs. dynamic scenes). **Step 5:** Continuously capture FR2 communication data during scenario execution via the logging/export interface for offline processing.

		Step 6: Offline, extract channel-related observables from reference signals (e.g., DMRS/SRS/CSI-RS) aligned to frame/slot timing. **Step 7:** Offline, derive sensing features (e.g., temporal channel variations, Doppler-related signatures, motion indicators) from the extracted observables. **Step 8:** Correlate sensing outputs with known scenario events and timelines (e.g., movement phases, blockage events). **Step 9:** Repeat experiments with different configurations (e.g., beam settings, UE positions, traffic load) to assess robustness and reproducibility. **Step 10:** Conduct longer-duration runs under regular communication operation to demonstrate the feasibility of continuous sensing without auxiliary sensing equipment.
	Methodology	**FR1:** LCX signals are processed to extract channel impulse responses (CIR) on a per-symbol basis. Temporal variations between consecutive channel estimates are analysed to derive Doppler components. These components are further transformed into micro-Doppler spectrograms, represented as matrices or image-like features suitable for machine learning models. In parallel, an auxiliary ground-truth generation approach is used during the training phase to monitor the same area covered by the LCX. The ground-truth data provides object presence flags and, optionally, the object position in image coordinates. **FR2:** Uplink IQ samples and metadata are logged at the gNB during normal operation. After the session, data is processed offline to extract channel responses from reference signals. Doppler and temporal features are derived and correlated with controlled object movements.
	Technical Enablers	
	Expected Outcomes	• Successful generation of micro-Doppler spectrograms from LCX signals for FR1. • Successful offline extraction of channel observables. • Identification of Doppler signatures caused by objects.
Trial Execution	**Pre-conditions**	• Equipment (LCX, Microamp CPE, servers, RU, UE, 5G CN) and software are deployed. • UE, Microamp UE and TSMA6B are installed. • Beamformed FR2 link operational.

	Trial Steps	The intended trial steps: **FR1** **Step 1**: Activate 5G transmission and reception over the LCX cable in the tunnel. **Step 2**: Training phase only, deploy and initialise the auxiliary ground-truth generation system covering the LCX-monitored area **Step 3**: Start LCX channel measurement and raw data recording. **Step 4:** Collect 5G KPIs, channel measurements (e.g., CIR) using the R&S®TSMA, and IQ samples using a spectrum analyser and/or USRP. **Step 5:** Process the recorded data offline to compute Doppler information and generate micro-Doppler spectrograms. **Step 6**: Operate the train through the real test section. **Step 7**: Apply the trained object detection model to the spectrograms. **Step 8**: Conduct longer-duration runs to classify objects around the LCX (e.g., human, animal, train) based on our model. **Step 9**: Log sensing outputs and communication performance metrics. **FR2** **Step 1:** Establish an FR2 5G NR over-the-air connection between the Microamp CPE (UE) and the RAN server (gNB) using the deployed radio units and beamformed links. **Step 2:** Verify stable communication performance under standard 3GPP-compliant operation. **Step 3:** For uplink-based sensing measurements, enable passive logging of frequency-domain IQ samples and PHY-related metadata at the RAN side while maintaining normal communication service. **Step 4:** Define and execute controlled operational scenarios in the test environment (e.g., pilot object, UE movement, obstruction events, LoS vs. NLoS conditions, static vs. dynamic scenes). **Step 5:** Continuously capture FR2 communication data during scenario execution via the logging/export interface for offline processing. **Step 6:** Offline, extract channel-related observables from reference signals (e.g., DMRS/SRS/CSI-RS) aligned to frame/slot timing.

		Step 7: Offline, derive sensing features (e.g., temporal channel variations, Doppler-related signatures, motion indicators) from the extracted observables. **Step 8:** Correlate sensing outputs with known scenario events and timelines (e.g., movement phases, blockage events). **Step 9:** Repeat experiments with different configurations (e.g., beam settings, UE positions, traffic load) to assess robustness and reproducibility. **Step 10:** Conduct longer-duration runs under regular communication operation to demonstrate the feasibility of continuous sensing without auxiliary sensing equipment.
	Methodology	• LCX channel measurements collected during train movement are processed to estimate channel impulse responses and derive Doppler components. Micro-Doppler spectrograms are generated and used as input to the trained detection model. • Uplink IQ samples and metadata are logged at the gNB during normal operation. After the session, data is processed offline to extract channel responses from reference signals. Doppler and temporal features are derived and correlated with controlled object movements.
	Complementary Measurements	-
	Calculation Process	-
Technical Enablers	**Communication Enablers**	Integrated Sensing and Communication
	Compute-as-a-Service (CaaS) Enablers	
	IoT and localization enablers	Data ingestion and telemetry
Expected Outcomes	• Demonstration of object detection using LCX- sensing. • Demonstration of FR2-based object sensing feasibility. • Validation of trained models in a real environment. • Characterisation of micro-Doppler signatures of different object types. • Assessment of detection accuracy and robustness.	

2.3.7.2 Liquid RAN

Table 19: Description of T3 trial setup for Liquid RAN

Trial ID & Name	Data Collection and Telemetry	
Trial Infrastructure / Venue	SRCC/Annaberg-Buchholz ISRD lab for pre-trials	
Trial Description	This trial demonstrates seamless, efficient, and secure data communication between trains, infrastructure, and ground systems to support railway operations beyond safety-critical functions. It enables real-time telemetry, system monitoring, and operational data exchange across the railway environment, addressing selected data collection and telemetry use cases defined by the International Union of Railways (UIC). The objective is to ensure continuous reporting of status, diagnostic, and performance information from trains and infrastructure to ground systems, while also supporting the distribution and transfer of non-critical data such as monitoring and operational records. These communications emphasise automation and efficiency with minimal human intervention, and they optimise network resource usage by scheduling bulk data transfers during periods of low demand, for example, when trains are stationary or near stations. Overall, this approach supports the digital transformation of railways by improving situational awareness and operational efficiency while preserving interoperability with legacy systems such as GSM-R.	
Technical Setup	**System Components**	UE (2x), O-RU, O-DU, O-CU, Near RT RIC, 5GC
	System Configuration	n77/78 band, MIMO 2x1, 40MHz bandwidth
Pre-Trial execution	**Pre-conditions**	Equipment (servers, RU, UE) and software (O-DU, O-CU, Near RT RIC, 5GC) were deployed and configured in the IS-Wireless lab.
	Trial Steps	Lab deployment: **Step 1**: RAN server moves back from the sleeping mode, and RAN becomes operational. **Step 2**: User Equipment 1 (UE1) attaches to the network and establishes PDU sessions at slice 1, **Step 3**: UE1's application initiates data transfer, **Step 4**: User Equipment 2 (UE1) attaches to the network and establishes PDU sessions at slice 2, **Step 5**: UE2 application initiates data transfer,

		Step 6: UE2 attachment and/or application data transfer triggers a change in DRB allocation per slice (via Near RT RIC xApp) **Step7**: UE1 and UE2 stop transmission **Step 8**: After transmission is over, the RAN server moves into a sleeping mode
	Methodology	**Operational approach**: 1. UE1 and UE2 will be static (no mobility), 2. UE1 and UE2 will be mobile - the trial operation (except server sleeping mode) will happen when UEs are in range of the base station. **Success criteria**: The success of the trial will be considered in the case of: • Established data transfer between UE1, UE2 and the data network, each UE shall be assigned to a separate slice, • Successful DRB allocation change per slice that corresponds to the trial logic, • Successful power management transfer of RAN server to sleeping mode and back
	Technical Enablers	• Network Slicing • RAN network programmability • Network Performance monitoring and control • Energy-efficient compute management • Cloud native service design
	Expected Outcomes	The trial proves the capability of K8s-deployed RAN to: • Assign different users to the appropriate slice and provide data transfer for those users • Change PRB quota of each slice at run time based on RAN metric analysis, • Provide energy-efficient gains derived from the guided server's EE mode transfer
Trial Execution	**Pre-conditions**	• O-RAN, Near RT RIC and 5GC are deployed in the train track vicinity, • Internet connectivity to the RAN server is available • UE1 and UE2 are on the train
	Trial Steps	The intended trial steps: **Step 1**: Train initiates the movement and starts collecting telemetry data - UEs gather payload to be sent to the application server **Step 2**: Once the train reaches the offloading area (e.g. station or depot), it transmits the telemetry data to the on-ground data centre via a terrestrial O-RAN link

		Step 2.1: User Equipment 1 (UE1) attaches to the network and establishes PDU sessions at slice 1, **Step 2.2**: UE1's application initiates data transfer, **Step 2.3**: User Equipment 2 (UE1) attaches to the network and establishes PDU sessions at slice 2, **Step 2.4**: UE2 application initiates data transfer, **Step 2.5**: UE2 attachment and/or application data transfer triggers a change in DRB allocation per slice (via Near RT RIC xApp) **Step2.6**: UE1 and UE2 stop transmission **Step 3**: Once the train passes the offloading area, the offloading of data is stopped, and the RAN server moves into sleeping mode.
	Methodology	**Operational approach**: UE1 and UE2 will be located in the train and move along the train tracks. The data offloading will happen when UEs are in the range of the base station. **Success criteria**: The success of the trial will be considered in the case of: • Established data transfer between UE1, UE2 and the data network, each UE shall be assigned to a separate slice, • Successful DRB allocation change per slice that corresponds to the trial logic, • Successful power management transfer of RAN server to sleeping mode and back
	Complementary Measurements	-
	Calculation Process	-
Technical Enablers	**Communication Enablers**	• Network Slicing • RAN network programmability • Network Performance monitoring and control
	Compute-as-a-Service (CaaS) Enablers	• Energy-efficient compute management • Cloud native service design
Expected Outcomes	The trial proves the capability of K8s-deployed RAN to: • Assign different users to the appropriate slice and provide data transfer for those users • Change PRB quota of each slice at run time based on RAN metric analysis,	

	• Provide energy-efficient gains derived from the guided server's EE mode transfer

2.4 Tele-operation as a back-up to autonomous driving – Transport 4 Use Case

2.4.1 Description of the use case

The T4 use case demonstrates how tele-operation can serve as a reliable backup when an automated vehicle encounters conditions where self-driving is no longer safe or effective (e.g. adverse weather, construction zones, or complex traffic). The concept relies on a B5G/6G-enabled communication platform that ensures seamless switching between autonomous and tele-operated modes, supported by real-time video streaming, V2X situational awareness, and distributed AI-based decision logic.

The tele-operation system integrates vehicle sensors, edge/cloud computing resources, and a remote control centre where a human operator can safely take control when needed. Communication between the vehicle and the operator is maintained via a low latency 6G link supported by the TUC testbed, while additional context awareness is provided through V2X data fusion and event detection.

Figure 19 depicts the architectural view of T4.

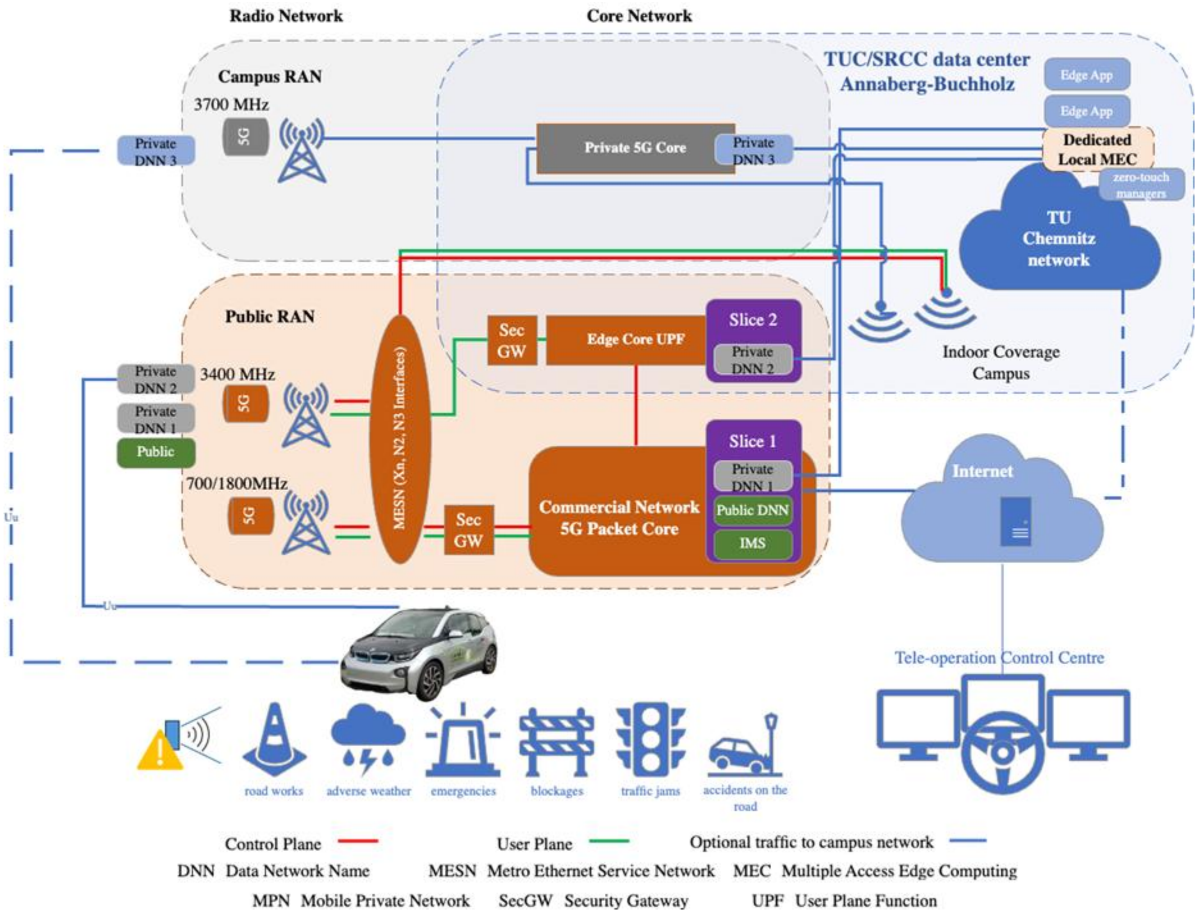

Figure 19: High-Level Architecture of T4 UC

B5G/6G Relevance

B5G/6G is particularly useful for the T4 use case as it provides the combination of ultra-reliable low-latency communication, dynamic QoS support, distributed intelligence, and context-aware connectivity that is needed to enable safe switching between automated driving and remote operation. In this use case, the vehicle must continuously exchange telemetry, situational information, and V2X-enhanced context data, while also being able to activate remote-operation support whenever the onboard driving stack can no longer guarantee safe operation. B5G /6G is therefore valuable because it can support low-latency video, control, and feedback loops between the vehicle and the remote operation centre, while also enabling service-specific prioritisation, flexible placement of functions across the vehicle–edge–cloud continuum, and real-time monitoring of connectivity conditions. In addition, advanced V2X support improves situational awareness by complementing onboard perception with cooperative road-context information, allowing the decision logic to react earlier and more safely to hazardous or uncertain conditions. Overall, B5G /6G is useful in T4 because it makes tele-operation a credible and scalable safety backup for automated vehicles, increasing continuity, resilience, and operational safety.

Large-scale deployment of the pilot

Large-scale deployment will be addressed in the pilot by the combination of several dimensions of scale that go beyond a small lab proof-of-concept. First, they span multiple pilot sites in more than one country, with pre-trial activities in Belgium and trial activities in Germany, which already gives the experimentation a broader geographic and operational scope. Second, they rely on a complex and realistic infrastructure stack, integrating B5G/6G networking, edge/cloud compute, decision logic, dynamic slice orchestration, and zero-touch service management, rather than validating a single isolated component. Third, the trials are planned around realistic operational workflows for remote operation, including dynamic service deployment across the compute continuum, real-time video and sensor traffic, runtime slice changes, and service relocation under network stress and energy-related constraints. This means the system is being exercised under conditions that resemble real deployment complexity rather than simplified demonstrations.

The large-scale character of T4 is also supported by its emphasis on integrated system validation, scalability, and measurable impact. Stress testing with emulated ToVs (Teleoperated Vehicles) will be performed in order to assess the robustness of the slice orchestrator and ZSM framework, demonstrating the intention to test how the solution behaves when scaled technically. In addition, the trials include extensive KPI and KVI measurements, together with the evaluation of different network-slice configurations and different service-placement and relocation decisions, so that the impact of the enablers on QoS and service continuity can be quantified. This is further reinforced by the employment of end-to-end slice management across RAN, transport, and core domains, dynamic orchestration under load, and realistic teleoperation traffic from cameras and sensors, which are all indicators of a trial that aims at deployment-relevant validation at the system level.

The T4 UC identifies the following functional requirements:

- The system must provide AI/ML support to the vertical for the training and inference tasks. In addition, those models must be able to be deployed dynamically (see CaaS requirements) according to the Vertical preferences, and the nature of the tasks involved (training, re-training, merging, inference, ...)
- The system must be able to dynamically deploy Vertical services at the most appropriate location (UE/edge/cloud/core) depending on their QoS requirements and deployment policies and ensure optimal execution of those services at all times, considering a wide range of criteria (network bandwidth, energy levels, latency, ...)
- The system should provide intent-based network slice provisioning.

The solution comprises several key components as summarised in Table 20.

Table 20: T4 Use Case High-Level Architecture Description

Domain & Enablers Reference	Component	UC Related Functionalities
Application enabler	Automated Vehicle Subsystem	Automated driving is enabled by on-board sensing and processing units.
Application enabler	Tele-operation Control Centre	Remote human operation.
Connectivity Infrastructure	6G Communication Link (TUC Testbed)	Ultra-reliable, low-latency connectivity for real-time video, control, and telemetry data between the vehicle and the control centre.
Application enabler	V2X Communication Module	V2X data exchange, enhancing situational awareness for both the automated and tele-operated modes.
Edge Infrastructure	Edge and Cloud Infrastructure	Decision making for dynamic switching between driving modes: teleoperation and autonomous driving.
Edge Infrastructure	Zero-touch network service management	Flexible placement of applications across the compute continuum to balance latency and context awareness.
Connectivity Infrastructure	Network slice orchestrator	Dynamically (re)configuring network slices based on the experienced QoS and SLA violations, ensuring required levels of QoS for safe remote operation over 6G

2.4.2 Deployment and implementation details

2.4.2.1 Design of the application

The system architecture of T4 is centred on a tele-operation service that complements the automated driving stack of the vehicle. The objective of the application is to ensure a safe and seamless transition from automated operation to remote human control whenever the vehicle detects conditions in which the on-board driving functions are no longer able to guarantee the required safety or service continuity. Typical situations include adverse weather conditions, temporary roadworks, degraded lane markings, complex traffic scenarios, or obstacles that cannot be handled with sufficient confidence by the automated driving functions.

The application is organised around three main logical domains: the **Automated Vehicle Subsystem**, the **Tele-operation Control Centre**, and the **B5G/6G-enabled edge/cloud platform**. In normal conditions, the vehicle operates in automated mode and continuously processes data from its on-board sensing and perception modules. In parallel, a tele-operation client running in the vehicle exchanges telemetry, vehicle status, and selected video feeds with the edge/cloud platform through the TUC communication infrastructure. This enables continuous monitoring of the vehicle state and provides the necessary situational awareness to prepare a possible intervention by the remote operator.

A central element of the application design is the decision-making logic. This function evaluates the operational context and determines the appropriate driving mode based on three key indicators: the *Autonomous Fitness Score (AFS)*, the *Tele-operation Fitness Score (TFS)*, and the *Context Risk Index (CRI)*. The AFS reflects the capability of the vehicle to safely continue automated driving, taking into account perception quality, Operational Design Domain (ODD) compliance, planner stability, and system health. The TFS represents the feasibility of safe remote operation, considering factors such as network performance, video and sensor data quality, situational awareness, and operator availability. The CRI captures the level of environmental risk, based on elements such as hazard saliency, proximity of vulnerable road users, closure rates, occlusions, and map confidence.

The decision logic continuously monitors these indicators and applies a threshold-based policy to manage transitions between operating modes. When the AFS remains above a predefined threshold, the vehicle continues automated operation. If the AFS degrades while the TFS remains above its threshold, the system initiates a controlled transition to tele-operation. In cases where both AFS and TFS fall below acceptable levels, particularly under high CRI conditions, the system triggers a minimal risk condition to ensure safety. The computation of these indicators may rely on AI/ML models deployed across the vehicle–edge–cloud continuum, enabling context-aware and adaptive decision making. This structured approach ensures that mode transitions are both timely and safe, while maintaining service continuity under varying environmental and network conditions.

The tele-operation procedure is supported by the **Tele-operation Control Centre**, where the remote operator is provided with a real-time view of the vehicle's surroundings and status. This includes live video streams, vehicle telemetry, localisation data, and contextual information received through the V2X Communication Module. The control centre sends driving commands back to the vehicle through the same low-latency communication path. In this phase, the application must ensure strict coordination between uplink media streams and downlink control commands, while preserving service continuity and avoiding unstable transitions between operation modes.

The T4 application relies on the underlying B5G/6G platform through a set of service exposure interfaces. In particular, APIs are used to request connectivity and compute resources according to the needs of the running function. These interfaces may include REST-based service APIs for telemetry exchange and service orchestration, and low-overhead interfaces such as gRPC or message-broker mechanisms for time-sensitive internal communication between application modules. Through the exposure framework, the application can interact with enablers related to network slicing, connectivity monitoring, service orchestration, and compute resource allocation. This allows the dynamic placement of selected functions at the most appropriate point of the compute continuum, i.e., on-board the vehicle, at the edge, or in the cloud.

Another important element of the design is the **distributed processing strategy**. Safety-critical and hard real-time vehicle control functions remain on-board, while functions such as video optimisation, event detection, operator support, logging, and decision support can be placed at the edge or in the cloud, depending on latency and resource constraints. This approach reduces the computational burden on the vehicle while preserving the responsiveness required for tele-operation. Edge deployment is particularly relevant for video pre-processing, V2X data fusion, and AI inference functions that support the operator with enhanced situational awareness.

The application also includes a **monitoring and feedback loop**. Network quality indicators, application latency, packet loss, video quality, and compute resource usage are continuously monitored and fed to the orchestration layer. Based on these observations, the system can adapt the deployed functions, modify the resource allocation, or request a network slice with more suitable QoS characteristics. In this way, the application design is tightly integrated with the project enablers for intelligent service orchestration and intent-based network provisioning.

Overall, the T4 application demonstrates how autonomous driving and tele-operation can be integrated into a single service framework supported by B5G/6G capabilities. The design enables a progressive transition between autonomous and remote-assisted operation, while leveraging edge/cloud resources, V2X context information, and AI-based decision support to increase safety, resilience, and service continuity.

The workflow of T4 UC is illustrated in Figure 20. The interaction starts with the initialisation of the tele-operation support chain and the deployment of the required vertical service components. In this initial phase, the Decision Logic requests the deployment of the T4 vertical service from the Zero-touch Service Management (ZSM) Framework, including the decision-logic function itself, video-streaming support, and remote-control functions. The ZSM Framework performs compute and network monitoring and applies decision-making for service placement, after which it confirms that the vertical service is available. In parallel, the Decision Logic requests an initial QoS profile and service connectivity from the 6G End-to-End Slice Orchestrator, while the Remote Operation Control Centre provides information on remote-operation readiness and operator availability. The orchestrator then confirms that connectivity is ready. At the same time, the Automated Vehicle Sub-System provides the vehicle-state, localisation, and perception summary to the Decision Logic, which forms the initial basis for subsequent mode-selection decisions.

The workflow then enters a continuous monitoring loop. During this loop, the V2X Communication Module provides the Decision Logic with V2X context, cooperative alerts, and road information, while the 6G End-to-End Slice Orchestrator reports KPI status to the ZSM Framework. In parallel, the ZSM Framework continues to compute and perform network monitoring and may perform decision-making for service relocation if the observed resource conditions require adaptation of the deployed service components. Based on the continuously updated vehicle-state, localisation, perception, and V2X context information, the Decision Logic performs two internal functions: first, it assesses the fitness of automated driving, remote operation, and contextual risk, and second, it evaluates the corresponding mode transition. In this way, the decision function continuously determines whether the vehicle can safely remain in automated mode, should prepare for remote operation, or must trigger a safe fallback state.

The first decision branch corresponds to the case in which the automation fitness remains above the relevant threshold. In this situation, the system keeps the vehicle in Automated Driving, and the workflow returns to the monitoring loop for continued supervision. The second branch is activated when the automation fitness falls below its threshold while the remote-operation fitness remains above its threshold. In this case, the Decision Logic instructs the Automated Vehicle Sub-System to prepare for remote operation and requests a remote-operation slice from the 6G End-to-End Slice Orchestrator. Once the orchestrator confirms that the low-latency remote-operation session is ready, the Remote Operation Control Centre receives a takeover request. If the remote operator confirms takeover, the Decision Logic commands the Automated Vehicle Sub-System to switch to remote operation.

After this transition, the workflow enters a second alternative branch corresponding to the remote-operated driving state. If the remote-operation fitness remains above the required threshold, the control centre sends remote control commands to the vehicle, while the vehicle sends video feeds and telemetry back to the control centre to provide the operator with the required situational awareness. In parallel, the ZSM Framework may determine that service relocation is needed and execute the corresponding relocation actions to preserve service continuity and the required performance. Conversely, if the remote-operation fitness falls below the threshold, the Decision Logic triggers the Minimal Risk Condition. The Automated Vehicle Sub-System then executes the corresponding safe-stop or safe-fallback manoeuvre, ensuring that the vehicle transitions to a safe state when remote control can no longer be considered reliable enough.

Overall, the workflow shows how T4 UC combines vertical-service deployment and placement, continuous context monitoring, fitness-based decision making, 6G-supported remote-operation setup, service relocation support, and safe fallback logic to manage transitions between automated driving, remote-operated driving, and minimal-risk operation. This reflects the role of the T4 use case as a B5G/6G-enabled tele-operation backup solution for automated vehicles.

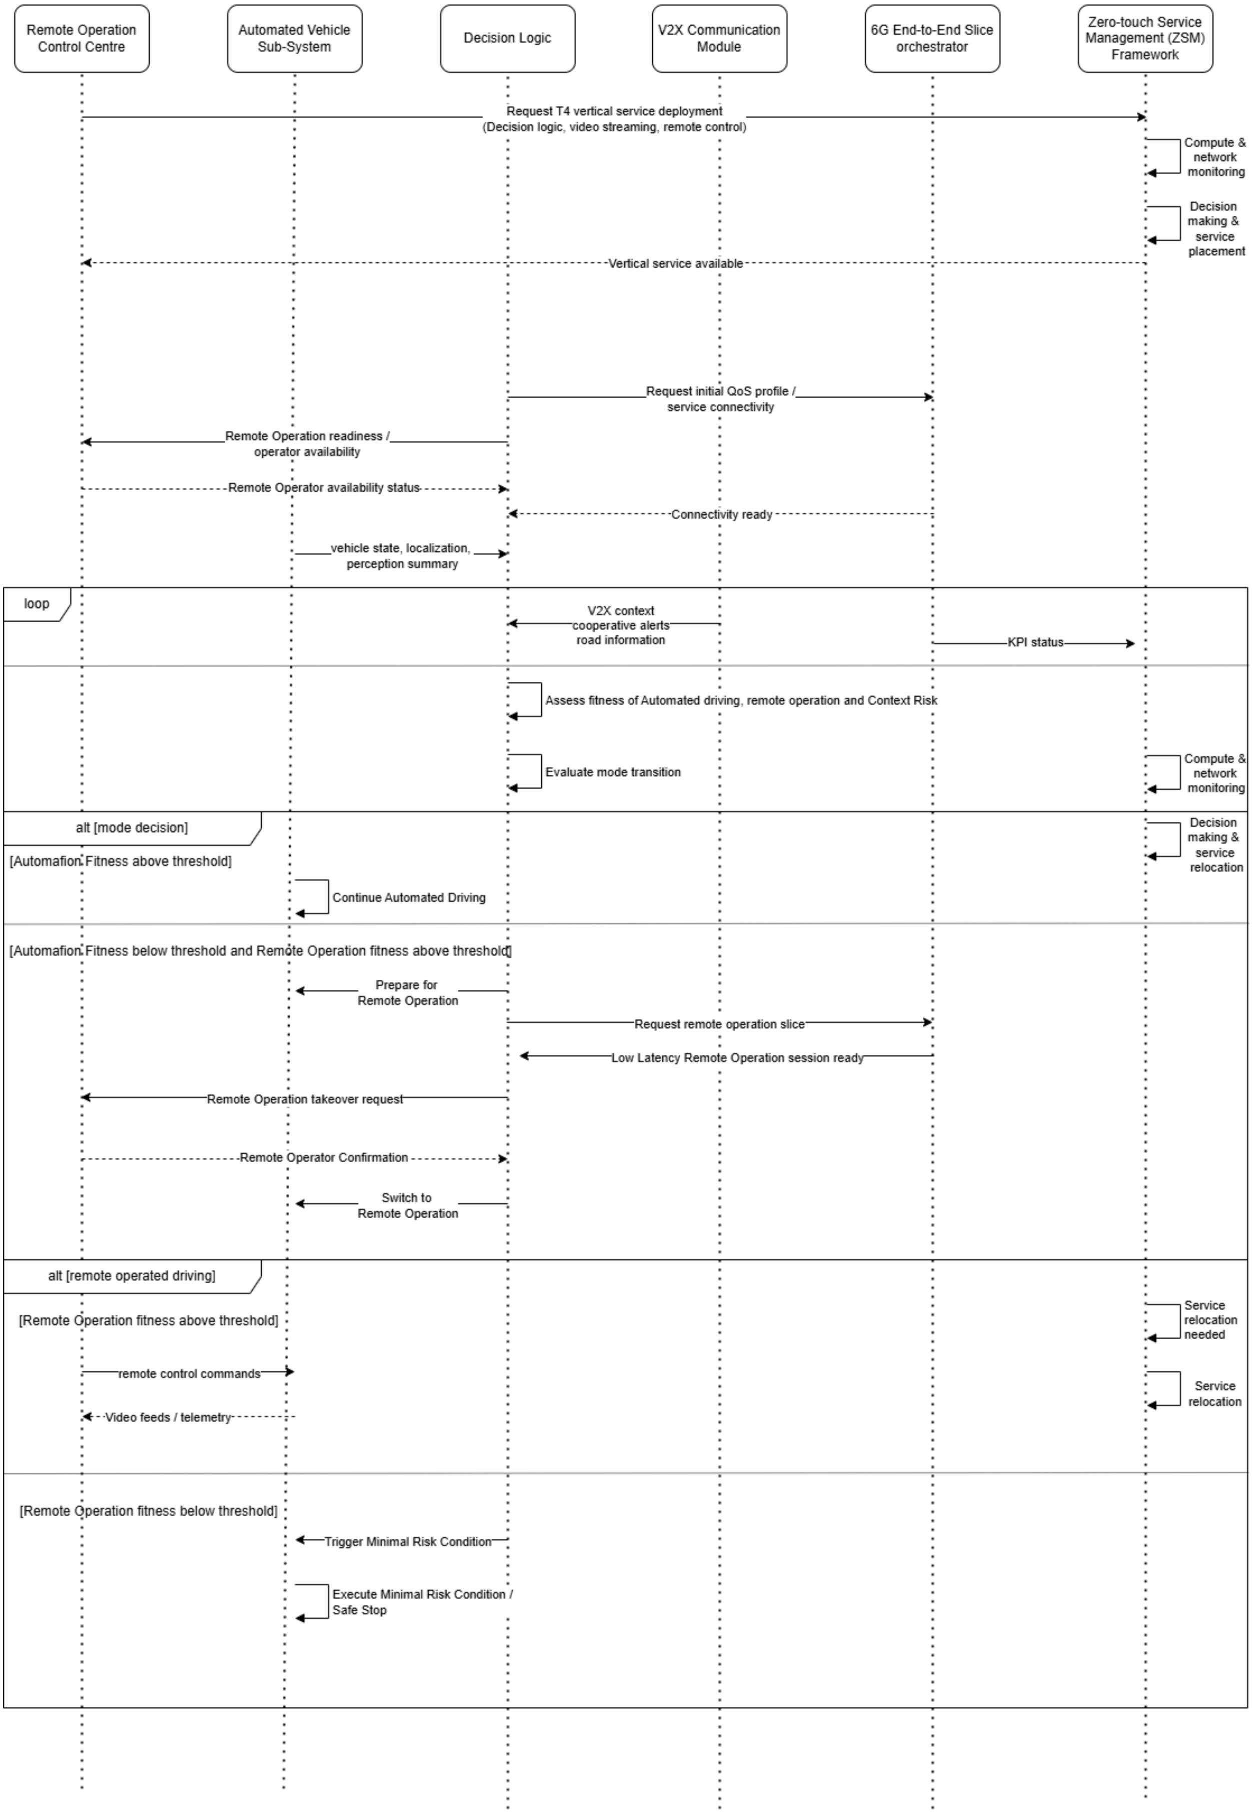

Figure 20: T4 UC Workflow

2.4.2.2 Devices, sensors and supporting equipment

In Table 21, the devices and sensors employed in T4 UC are introduced.

Table 21: T4 Devices, sensors and supporting equipment

Category	Component	Technical Details	Units
Vision System	Smart camera	Smart camera delivers detections of different objects' positions and speed (cars, trucks, pedestrians, cyclists), traffic sign information and lane marking information in front of the car	1
Vision System	360 camera view	Camera view to the front (fov 120 degrees), rear (120 degrees) and blind spot area left and right (fov 85 degrees)	4
Ranging Sensors	Automotive Radar	Delivers object information	1
Ranging Sensors	LiDAR	Delivers object and environment information	1
CAN	CAN Gateway	Obtaining ego motion and vehicle system data	1
Odometer	IMU	Ego motion data	1
Localization	GNSS	Inertial navigation system	1
Processing Units	Real-time prototyping system, car PC	Prototyping, testing, execution of control functions; sensor data acquisition, processing and data collection	1
HMI	Sensors on the steering wheel and pedals;	Measuring driver behaviour, state and control input; Information system for safety driver and observer/researcher	1

2.4.3 Network infrastructure and functional components

2.4.3.1 Test Vehicle

The test vehicle Carai3 (BMW i3) is equipped with sensors and actuators for automated driving functions. The vehicle is equipped with a generic, fully accessible and customizable automation system.

The system provides longitudinal and lateral control through emulation of the electric accelerator pedal, a mechatronic brake system, and a servo mechatronic steering system. Its modular design allows testing each control function independently. The mechatronic brake system uses a servo drive with a gearbox to provide sufficient momentum for emergency braking. Lateral control is achieved via a servo-driven steering module capable of turning the wheel at up to 1920°/s, allowing steering while stationary for parking or emergency manoeuvres. The system mounts directly to the stationary part of the steering column and connects via a gear to the cardan joint, without restricting driver adjustments.

The system supports predefined automated driving manoeuvres by executing specified longitudinal and lateral control inputs. Hybrid operation modes combining automated control with Wizard-of-Oz interventions are also possible. For instance, a compact auxiliary steering wheel mounted to the right of the passenger seat enables a concealed Wizard-of-Oz operator to perform complex driving tasks while longitudinal control remains automated. Furthermore, the test vehicle can be fully teleoperated from an operator workstation via keyboard or gamepad using Wi-Fi or 5G mobile internet connectivity. The modular design supports bidirectional communication, including video feedback, sensor data transmission, and control commands, allowing components to be added or removed as needed.

Carai3 is also equipped with sensors detecting driver actions, such as hands on the wheel or feet on the pedals. Triggering these sensors immediately disables automated driving, returning full control to the driver. Switching between automated and manual modes is possible via control elements at the front (driver) or rear (supervisor) of the vehicle.

For environment perception, the test vehicle is equipped with range sensors (Radar, LiDAR) and various vision sensors. Precise localisation is achieved using an inertial-coupled Global Navigation Satellite System (GNSS), while low-cost GNSS modules are employed with different filtering methods. To enable decentralised sensing and support research on cooperative mobility, wireless communication devices for standards such as C-V2X (G5) and 5G are integrated.

Figure 21: Test Vehicle Carai3.

2.4.3.2 NextG-Lab

During the pre-trial phase, the development and testing of 6G technological enablers used in the T4 use case are being mainly done in the lab environment, i.e., a real-life testbed environment maintained by IMEC. For the sake of completeness, we describe the testbeds briefly as follows.

The **NextG-Lab testbed** is located in Antwerp, offering indoor and temporary outdoor 5G connectivity based on two gNodeBs and multiple datacenters able to host the Network Functions (NFs) of the 5G Core Stand Alone deployment. In the context of the T4 use case, this testbed is primarily used for deploying and testing Communication (Communication resource management and Network slicing), Compute-as-a-Service (Compute resource management, Compute continuum, and Cloud-native service design), and Application and AI enablers (CAMARA APIs for Quality on Demand).

The NextG-Lab testbed supports different combinations of Open RAN and open-source solutions. IMEC has obtained a BIPT (Belgisch Instituut voor postdiensten en telecommunicatie) test license, which is valid for the 4035-4085 MHz frequency range (N77 upper), which also applies to the stretch on the road covered by the Smart Highway testbed. Currently, NextG-Lab testbed is deployed as a testbed in a box, containing both radio and core capabilities (Figure 22 and Figure 23), offering tools to pursue 5G/6G technology validation, 5G/6G Network Slicing (RAN, Transport, Core), network monitoring, integration and testing, in both office and open/outdoor environments.

On the Core side, we have Open5GS[2] and OpenAirInterface[3] (OAI). Nevertheless, the NextG-Lab testbed can support other open-source solutions, such as Free5GC[4]. The 5G New Radio (NR) is based on OAI, with the following specifications: RTX 3060 GPUs for AI/ML deployment, Intel i7-11700K (8 Cores, 16 Threads), 62GB RAM. The 5G gNBs are equipped with different USRPs (B210[5], N310[6] and X310[7]) and are controlled using FlexRIC[8], an open-source Radio Intelligent Controller (RIC). For User Equipment (UE), we use COTS 5G modems (Quectel RM500Q[9]) connected to COTS mini-PCs (Intel NUCs with 32GB RAM, and Intel i7-10710U processor). Furthermore, we use programmable SIM cards to enable these UEs to connect to our 5G network. Finally, to enable computing capabilities of the testbed, we enabled Proxmox on Supermicro servers suitable for provisioning virtualised resources.

Due to its flexibility, this testbed is actively used for configuring dynamic end-to-end network slicing, in combination with Camara APIs for dynamic collection of requirements and QoS needs from vertical services, and it has already been actively used for testing teleoperation services over 5G. IMEC and TUC are now actively working on the extensions of the TUC testbed to enable the Open5GS Core and the OpenRAN components, such as RIC on the TUC radio, which is essential for transferring technological enablers and preparation of their integration with the use case in the trial site environment.

Figure 22: NextG-Lab testbed.

2 https://open5gs.org/

3 https://openairinterface.org/

4 https://free5gc.org/

5 https://www.ettus.com/all-products/ub210-kit/

6 https://www.ettus.com/all-products/usrp-n310/

7 https://www.ettus.com/all-products/x310-kit/

8 https://gitlab.eurecom.fr/mosaic5g/flexric

9 https://www.quectel.com/product/5g-rm50xq-series/

Another testbed we are actively using, mainly in the context of Zero-touch network and Service Management (ZSM), is the Smart Highway testbed, which is illustrated in Figure 7. The **Smart Highway** test site is built on top of the E313 highway, and it consists of a highway strip of 4 km equipped with seven Road Side Units (RSUs) in total. This testbed is designed for Vehicle-to-Everything (V2X) communication & distributed/edge computing research, thereby providing means for creating a scalable and reliable V2X system, with opportunities to validate automotive requirements for automated driving with industry in real-life trials. Besides the RSUs that are deployed along the highway, there are three Onboard Units (OBUs), which are either in-vehicle or rooftop units. The important feature of such OBUs is their flexibility and adjustability, as they can be easily mounted on any regular car. In addition to that, the testbed is equipped with two testing vehicles that can be used in experimentation activities.

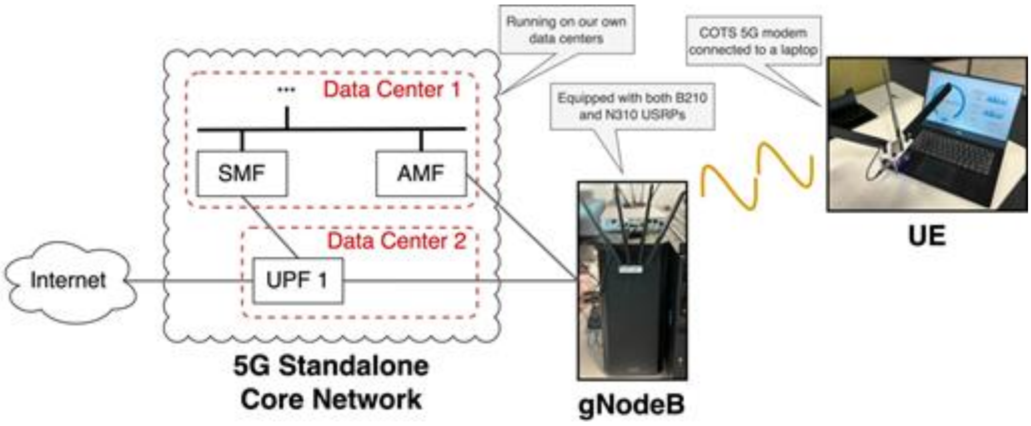

Figure 23: Example of a fully end-to-end 5G Standalone network deployed on IMEC infrastructure.

As these RSUs are mounted on top of the real highway, they are connected to the fiber network of the road operator, which makes them able to fetch information about electronic traffic signs on the highway, and to give a stable backhaul connection for edge processing on the RSU. This is useful for understanding and emulating the background traffic in case of teleoperated vehicles, as we can retrieve real-time information about the vehicles at different highway segments. This information can be used for modelling realistic scenarios in which multiple vehicles are connected to the network and consuming network services. Finally, the testbed is equipped with a testing car, a BMW X5 with an automatic gearbox from 2014 that is provided by the University of Antwerp. The car is equipped with an OBU, power system, and communication hardware. This car can be driven and used as a mobile node in a 6G V2X context, and as such, it is being used as a mobile UE mimicking the teleoperated vehicle in the context of T4 and AMAZING-6G.

Figure 24: Smart Highway testbed.

The distributed deployment of RSUs and the compute infrastructure of OBU is used as a compute continuum for hosting the T4 ML-based applications and other types of vertical services, which are automatically and intelligently orchestrated by the ZSM solution. IMEC and TUC are currently replicating this distributed compute continuum on the virtualised network infrastructure within the TUC testbed, preparing the compute environment for dynamic service deployment and intelligent orchestration.

2.4.4 Use case enablers

As presented in deliverable D3.1, T4 enablers are presented in Table 22.

Table 22: List of enablers and sub-enablers for T4.

Category	Technology enablers	Sub-enablers
Communication enablers	Communication resource management	Zero-touch network and service management
Communication enablers	Communication resource management	Network programmability
Communication enablers	Communication resource management	Automated gNB configuration
Communication enablers	Communication resource management	Network performance monitoring and control
Communication enablers	Network slicing	
Compute-as-a-Service enablers	Compute resource management	Intelligent service and resource orchestration in extreme

		edge/edge/cloud compute continuum.
Compute-as-a-Service enablers	Compute resource management	Intent-based service and resource management
Compute-as-a-Service enablers	Compute resource management	Real-time monitoring of compute resources and energy consumption in the compute continuum
Compute-as-a-Service enablers	Compute continuum	Distributed user-edge-cloud compute continuum
Compute-as-a-Service enablers	Cloud-native service design	
Application enablers and AI	OpenAPIs for Network Exposure	CAMARA APIs for Quality on Demand
Application enablers and AI	OpenAPIs for Network Exposure	CAMARA APIs for Edge/Cloud
IoT and localization enablers	Data Ingestion and Telemetry	Data Collection from sensors and devices
IoT and localization enablers	IoT Contextual Awareness Systems	

Communication resource management: The T4 use case targets scenarios where autonomous driving must hand over control to a remote operator, requiring ultra-reliable, low-latency, and resilient connectivity. It leverages key enablers such as Zero-touch Service Management (ZSM), network programmability, automated gNB configuration, and network performance monitoring to guarantee consistent QoS and sub-30 ms end-to-end latency for mission-critical teleoperation. Network programmability enables real-time path optimisation, dynamic traffic steering, edge-deployed UPFs, and isolated teleoperation sessions with guaranteed resources, while ZSM provides intent-based, autonomous orchestration of network and compute resources across domains to maintain service continuity under challenging conditions.

The underlying 5G SA architecture treats the RAN, transport network, and 5G Core as fully programmable, software-driven systems. In the RAN, O-RAN disaggregation (RU/DU/CU) and standardised interfaces, together with near-RT RICs and xApps, enable fine-grained, slice-aware radio control. In the 5G Core, the service-based architecture allows independently deployable and scalable network functions, coordinated through well-defined interfaces and APIs, with edge-located UPFs supporting low-latency services. SDN-enabled transport networks dynamically enforce deterministic paths and bandwidth guarantees. Experimental results from prior projects demonstrate that the ZSM framework can autonomously mitigate performance degradation, such as latency spikes under high load, by orchestrating service migration and resource reallocation, restoring QoS within thresholds. Building on these results, the T4 implementation will further enhance ZSM with ML-based decision logic, progressing from initial testbed validation to full end-to-end trials over the project lifetime.

Network slicing: The T4 use case requirements need to be met, especially when switching from autonomous to teleoperated driving, when the network is being actively used. This will be achieved through dynamic end-to-end network slicing, which allows teleoperation traffic to be isolated in a dedicated slice with guaranteed radio, transport, and core resources. A centralised Network Slice Orchestrator (NSSO) orchestrates this slice across the RAN, transport network, and 5G Core by coordinating domain-specific controllers via harmonised interfaces, enabling deterministic routing, prioritised radio resource allocation, and dynamic instantiation or scaling of UPF functions at the edge. Through closed-loop monitoring and automated resource reallocation, the NSSO ensures that stringent teleoperation SLAs are maintained under varying load conditions, with end-to-end validation confirming stable low latency, sufficient throughput, and high reliability in realistic deployment scenarios.

Intent-based networking: For the T4 use case, intent-based networking enables teleoperation services to specify high-level requirements, such as ultra-low latency, high reliability, and service continuity, which are automatically translated into coordinated actions across the RAN, edge, and core domains. Within a ZSM framework, real-time and historical monitoring data are combined with anomaly detection and predictive analytics to drive zero-touch orchestration, ensuring optimal placement and lifecycle management of T4 services across the edge-cloud continuum. This allows proactive, energy-aware, and latency-sensitive resource optimisation that adapts to mobility and network dynamics, guaranteeing resilient support for teleoperation and autonomous driving.

Compute continuum: Through AI-driven orchestration, cross-domain management, and intent-based automation, the framework dynamically allocates resources and maintains QoS under changing conditions. Implementation occurs in three phases: first, integrating the ZSM framework with 5G/B5G APIs (like CAMARA QoD) to gather network metrics and define QoS intents; second, deploying and validating orchestration components on Kubernetes-based infrastructure through real-world testing of handover scenarios; and third, optimising the framework with machine learning for anomaly detection, traffic forecasting, and energy-aware orchestration. The project's first two phases span the first and second years, while the optimisation phase concludes by mid-third year.

Cloud-native service design: For the T4 use case, an event-driven, cloud-native EdgeApp framework will be used to dynamically manage per-vehicle QoS across the compute continuum, leveraging CAMARA and other network exposure APIs to trigger on-demand QoS adaptations based on ML-detected anomalies and real-time network telemetry. A dedicated QoS Control EdgeApp ensures high-throughput, low-latency connectivity during tele-operation and reverts to baseline QoS in autonomous mode to avoid over-provisioning, following defined QoS profiles mapped to CAMARA QoD parameters and deployed via NFV orchestration at the network edge. Implementation steps include integrating context sources, defining base and enhanced QoS profiles, deploying the EdgeApp on edge infrastructure, validating behaviour through transition simulations, and refining it with predictive triggers and fallback logic.

QoD CAMARA API: In the T4 use case, the CAMARA QoD API enables dynamic, standard-based QoS upgrades to meet the strict latency and throughput requirements of teleoperation by letting an ML-based anomaly detection service trigger just-in-time network slice reconfiguration when autonomous driving becomes unreliable, and scale resources back down once autonomy resumes. The solution, built with Open5GS for the core and OpenAirInterface for the RAN, exposes a RESTful QoD interface that carries JSON requests from the application server through the NEF and PCF to the SMF and UPF, where updated QoS rules and PDU session parameters are installed in both core and RAN. The entire process is governed by the NSSO. End-to-end, this POST-based message flow configures enhanced radio and user-plane QoS (via PFCP, NGAP, and RRC procedures) for specific service flows such as real-time video and control traffic, ensuring safe, reliable tele-operation while optimising network and compute resource usage during normal autonomous driving.

2.4.5 KPIs definition

Table 23: List of KPIs for T4

KPI name	Description/KPI definition	KPI target
End-to-end network service latency (RTT)	Time between the source application/node sending a data packet and the destination application/node receiving a data packet in ms.	< 30ms
Service reliability	Service reliability is a period of time for which the service satisfies the required performance constraints (DL/UL throughput, end-to-end service latency, etc..)	> 99.999%
UL Throughput	Minimum amount of data that needs to be transferred from a single device (sensor or camera installed on the teleoperated vehicle) to the network over a certain period of time.	> 5Mb/s
Energy efficiency	Energy consumption reduction in the edge-cloud compute continuum, where vertical applications are running.	>20%

2.4.6 KVIs definition

Table 24: List of KVIs for T4

KVI name	Description/KVI definition	KVI target
Improved Energy Efficiency	Energy efficiency through minimisation of the time the autonomous cars are idling until they are allowed to resume their operation (e.g., in the case of roadblocks).	Enable sub-minute recovery time (≤60 s) in ≥90% of teleoperation-assisted interventions, resulting in ≥50% reduction in idle-induced energy waste.
Reduced Energy Consumption	Reduced energy consumption through energy-aware service orchestration (Zero-touch service management framework)	≥ 25% reduction in total service-related energy consumption (compared to static / non-energy-aware orchestration baseline)
Lower Data Overhead	Lower data overhead through leveraging edge computing for data processing and service runtime.	≥ 50% reduction in data transmitted to the core/cloud
Lower cost for drivers	Lower cost for drivers through the increased efficiency and safety of autonomous driving	20–30% reduction in total cost per km (compared to baseline driving without advanced autonomous + teleoperation support)

Minimised time of disruption	Minimised time of disruption in case of situations where autonomous driving is deemed unsafe or ineffective	≤1s handover time to teleoperation (from trigger to effective remote control)
Minimised risk of accidents	Assesses the improvement in the trust and acceptance of autonomous driving as a result of the minimised risk of accidents.	≥ 25% increase in willingness to use autonomous driving services e.g., % of users opting in or preferring autonomous mode) ≥ 80% of users reporting "high" or "very high" perceived safety

2.4.7 Description of the trial setup

T4 Trial		
Trial Infrastructure / Venue	TUC testbed in Chemnitz, Germany (trial activities) NextG-Lab testbed in Antwerp, Belgium (pre-trial activities, a subset of Minimum Viable System activities)	
Trial Description	The T4 trials will consist of the following segments: 1) optimal deployment of vertical services within the B5G/6G edge cloud continuum, leveraging ZSM decision-making capabilities to select deployment locations balancing the QoS requirements, compute load and energy consumption, 2) configuration of end-to-end network slices as per vertical service requirements, 3) real-time operation of vertical services over B5G/6G network, such as decision making algorithm used for switching the mode between teleoperation and autonomous driving and real-time video streaming during the teleoperation phase, and 4) dynamic network slice changes based on the QoS and real-time network conditions.	

Technical Setup	**System**
	Components
		• B5G/6G network infrastructure
		  • Advanced B5G Radio and Core capabilities with RAN Intelligent Controller (RIC), Transport Network Controller (TNC) and Core Network Controller (CNC), for dynamic network slice management
		  • Edge cloud compute infrastructure
		  • decision-making algorithm for efficient switching between two operational modes (teleoperation and autonomous driving)
		  • Test vehicle with automated driving capability
		  • Remote operation software

	System Configuration	• Reliable B5G/6G network connectivity for high-quality sensor and camera data acquisition (uplink) and real-time remote control (downlink) • Distributed compute infrastructure for hosting vertical services, AI-based applications, and other B5G services ZSM system for dynamic configuration of end-to-end network slices and services/applications
Pre-Trial execution	**Pre-conditions**	• A testing vehicle with/without teleoperating capabilities is properly configured and connected to the networks (both to the test one, i.e., NextG-Lab and Smart Highway, in Belgium and the TUC one in Germany) • ZSM framework components properly installed in virtual infrastructure environment (VMs dedicated to the TUC MEC, and Smart Highway testbed) • RIC, TNC, CNC, and NSSO (Network Slice Orchestrator) deployed in the TUC network and NextG-Lab • ZSM monitoring modules receiving real-time metrics (network performance, compute load, energy consumption) • Simple vertical application prepared and packaged for testing
	Trial Steps	• Vertical application package (simple application for testing purposes) is onboarded on the NFVI of the TUC MEC system • ZSM service orchestration system deploys the application selecting the optimal location, i.e., specific node within the compute continuum (UE, MEC, cloud) • NSSO selects the optimal network slice for the vertical application • Emulated video and sensor traffic is being transferred from a testing vehicle to the vertical application • Network and service performance are being monitored • Network stress testing is applied to trigger changes in service and slice deployments (stress tests designed to mimic real-life network overload events) • Slice changes during runtime, and performance is monitored to detect the impact of dynamic network slice orchestration on service performance

		• Energy spikes are modelled to trigger changes in service placement within the compute continuum • Performance is monitored to detect the impact of dynamic service orchestration on service performance
	Methodology	• Using simplistic video streaming and iperf3 traffic generators to mimic the behaviour of real-life teleoperation services • ZSM service orchestration employs a simple multi-criteria decision-making process to select optimal placement strategies ZSM slice orchestration uses a simple rule-based approach to decide when to trigger the slice reconfiguration process, based on the service performance requirements derived from the T4 teleoperation applications (video/sensor data streaming, remote control)
	Complementary Measurements	• All KPIs measured
	Calculation Process	• Analysis of KPIs and their interdependency and correlation to understand the impact of enablers on the T4 performance
Trial Execution	**Pre-conditions**	• Remote control center and a vehicle with teleoperating capabilities are properly configured • ZSM framework components properly installed in virtual infrastructure environment (VMs dedicated on the TUC MEC) • RIC, TNC, CNC, and NSSO (Network Slice Orchestrator) deployed in the TUC network • ZSM monitoring modules receiving real-time metrics (network performance, compute load, energy consumption) • Vertical application for AI-based processing and switching between teleoperation and autonomous driving, prepared and packaged for testing • External users are available for testing the setup and its evaluation
	Trial Steps	• AI-based application for switching between teleoperation and autonomous driving is onboarded on the NFVI of the TUC MEC system • ZSM service orchestration system deploys the application selecting the optimal location, i.e., specific node within the compute continuum (UE, MEC, cloud) • NSSO selects the optimal network slice for the vertical application

		• Video and sensor traffic is being transferred from the on-board system (cameras and sensors) on the teleoperated vehicle to the vertical application • Network and service performance are being monitored • Network stress testing is applied to trigger changes in service and slice deployments (stress tests designed to mimic real-life network overload events) • Slice changes during runtime, and performance is monitored to detect the impact of dynamic network slice orchestration on service performance; At the same time, video quality is being visually observed on the remote control side to identify the changes • Energy spikes are modelled to trigger changes in service placement within the compute continuum • Performance and video quality are monitored to detect the impact of dynamic service orchestration on service performance • Vertical application needs a higher quality of video than initially anticipated, which cannot be accommodated by ZSM, as it is not aware of new requirements. Vertical application triggers QoD via CAMARA API. Both ZSM service and slice orchestrators translate this request to resource reconfiguration processes. Video quality is observed to identify changes, and network and service performance are being monitored.
	Methodology	• Transferring real data from the cameras and sensors installed on board • ZSM service orchestration employs advanced ML-based decision-making solutions to select optimal placement strategies • ZSM slice orchestration uses a simple rule-based approach to decide when to trigger the slice reconfiguration process, based on the service performance requirements derived from the T4 teleoperation applications (video/sensor data streaming, remote control)
	Complementary Measurements	• All KPIs measured
	Calculation Process	Analysis of KPIs and their interdependency and correlation to understand the impact of enablers on the T4 performance

Technical Enablers	**Communication Enablers**	Communication resource management (ZSM, network and service performance monitoring, network programmability, automatic gNB configuration), and Network Slicing
	Compute-as-a-Service (CaaS) Enablers	Compute resource management (Intelligent service and resource orchestration in extreme edge/edge/cloud compute continuum, intent-based service and resource management, real-time monitoring of computing resources and energy consumption in compute continuum), Compute continuum, Cloud-native service design.
	Application and AI	CAMARA APIs for Quality on Demand and Edge/Cloud
Expected Outcomes	• Safe and reliable network connectivity is proactively ensured for the teleoperation process, once the switch is triggered Smooth transitions between teleoperation and autonomous driving regardless of the network conditions	

2.5 Port logistics and transport operations optimisation and safety - Transport 5 Use Case

2.5.1 Description of the use case

The T5 Use Case aims to enhance operational efficiency and safety by reducing the need for physical human–machine interaction through increased automation. At the same time, it supports the automatic monitoring and enforcement of safety measures, while contributing to the reduction of the environmental footprint of port logistics and transport operations.

It is focused on operations at the Port of Thessaloniki Container Terminal (ship-to-shore and shore-to-truck operations). Containers in the port of Thessaloniki (ThPA S.A.) are being handled 24 hours a day, 361 days a year at the designated area of the Container Terminal in the western part of Pier 6 (quay 26). It is part of the port's Free Zone and extends in a 550m long and 340m wide area (352.000 m^2) with a storage capacity of 5,200 TEU in ground locations, while it can serve vessels with a draught of up to 12 meters. The Container Terminal is fully equipped with modern container handling equipment, and it features technical facilities and an experienced technical support team. Loading and unloading of containers from and to ships is conducted by 4 STS cranes, 2 of which are specially designed New Panamax STS cranes that were recently delivered. The handling, storage and stowage of containers is conducted by straddle carriers and a series of modern container handling equipment such as reach stackers, tractors, front lifts, trailers and forklifts. The pilot implementation is broken down into two Subcases: **Subcase [1] - From Ship to shore** and **Subcase [2] - From Shore to Truck**.

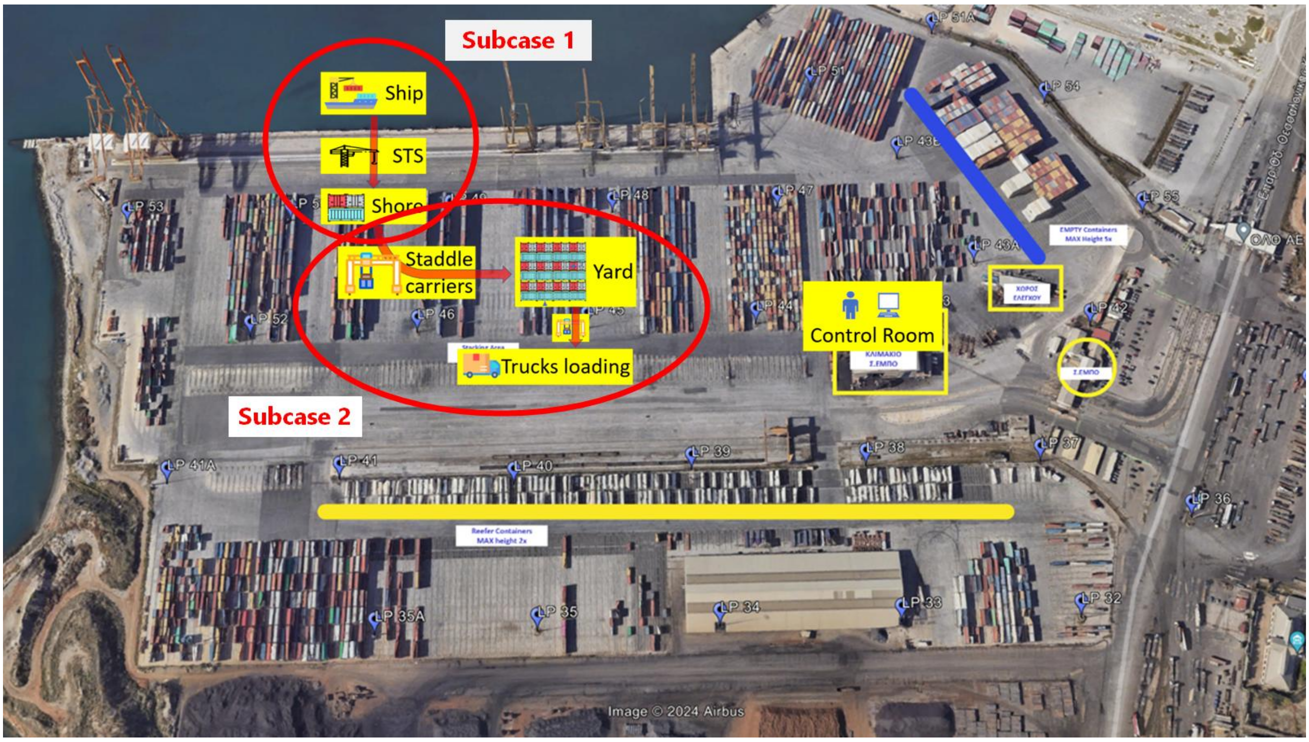

Figure 25: T5 Subcases location in the Container Terminal area

T5 Use Case will enable real-time, large-scale digital twinning by extending an existing Digital Twin (DT) platform for process monitoring, while it will also enable the teleoperation of key port equipment, specifically Ship-to-Shore (STS) cranes. The platform will be validated through a pilot covering an end-to-end port process, from vessel unloading to truck loading. In addition, the Use Case will implement algorithms and digital systems connected to IoT devices installed across port areas and equipment, i.e., cameras, to support enhanced monitoring and automated enforcement of safety measures.

The objectives of the use case in relation to the project's technical objectives are:

- Optimisation of operational efficiency through process optimisation (including cargo handling, optimised routes, and equipment operation). This optimisation leads to increased overall logistics efficiency, reduced operational downtime, and better resource utilisation.
- Real-time visibility through instant and continuous data transmission (high-speed connectivity, low latency), allowing operators to have visibility into cargo movements, levels, and conditions, ensuring effective supply chain management.
- Enhanced safety and security through real-time surveillance, monitoring, and response mechanisms based on high-definition video streaming, sensor data analysis, and quicker response times to potential security threats or emergencies.
- Improved supply chain resilience by providing robust connectivity even in remote or challenging environments to ensure uninterrupted communication and data flow, reducing the risk of disruptions.

The main innovations include the real-time collection, monitoring, and analysis of data through the combined use of Digital Twin, B5G/6G, AI/ML, simulation, and IoT technologies. They also encompass the orchestration of seamless tele-operations, as well as the integration and exposure of sensing capabilities from the network to the end user. In addition, sensors, cameras, and other devices are interconnected via B5G/6G infrastructure, with network slicing enabling dedicated slices for critical data transmission.

The pilot implementations shall address the following functional and non-functional requirements and challenges:

- **Enhanced connectivity, reliability, scalability, and high data rates:** The pilot shall support ultra-reliable, high-capacity, and scalable communications to enable real-time massive twinning, high-throughput data streams, analytics, teleoperation functionalities, and both real-time and batch data processing. The system shall support horizontal scalability to accommodate increasing data volumes, connected devices, and user demand. Low end-to-end latency, high bitrate, and high network reliability are essential to ensure the synchronisation of port operations, where delays of even a few milliseconds may affect logistics efficiency and supply chain performance.
- **Massive twinning, seamless IoT integration, and interoperability:** The pilot shall support the integration and interconnection of a large number of port assets and IoT devices, including cranes, vehicles, containers, and camera-based monitoring systems. This capability shall enable near real-time data exchange, automation, situational awareness, and comprehensive Digital Twin operation across the port environment. The system shall further provide standardised interfaces and interoperable mechanisms to ensure seamless integration and data exchange with external systems, platforms, and third-party services.
- **Security, data protection, and trustworthy operation:** The system shall ensure end-to-end data protection and trustworthy execution of all primary operations, including secure communication, authentication and authorisation mechanisms, encryption of data both at rest and in transit, and compliance with relevant cybersecurity and data protection standards. The pilot shall ensure secure handling of operational, sensor, and teleoperation data across all interconnected systems and devices.
- **Availability, resilience, and operational continuity:** The pilot shall ensure continuous operation with minimal downtime by incorporating fault tolerance, redundancy, failover, backup, and recovery mechanisms. The system shall maintain service continuity and operational reliability under varying workloads, network conditions, and potential system failures to support uninterrupted port operations and teleoperation services.
- **Support for sustainability objectives:** The pilot implementations shall contribute to environmental, societal, and economic sustainability goals by enabling more efficient, automated, safer, and resource-optimised port operations. The implemented digital services shall support the reduction of operational inefficiencies, congestion, emissions, and safety risks, while facilitating the transition towards climate-neutral and sustainable port operations.
- **Operational and technological constraints:** The pilot shall consider and address limitations related to:

    (a) port equipment and crane operation capabilities,

    (b) communication network performance requirements, including ultra-low end-to-end latency, high bitrate, reliability, and scalability, and

(c) human-factor aspects related to operator interaction, usability, safety, trust, and adoption of automated and teleoperated systems.

B5G/6G relevance

Based on the above, B5G/6G technologies are key enablers for this Use Case, as they provide the communication and computing capabilities required to support real-time, highly automated, and safety-critical port operations. The proposed pilot relies on the continuous exchange and processing of large volumes of data generated by Digital Twin platforms, teleoperated Ship-to-Shore (STS) cranes, and camera-based monitoring systems deployed across the port environment. Such operations require ultra-low end-to-end latency, high reliability, enhanced network availability, and high data throughput in order to ensure accurate synchronisation between physical and digital assets, uninterrupted teleoperation, and near real-time situational awareness. In particular, teleoperation of STS cranes demands highly reliable and responsive communication links to guarantee safe and efficient remote control, while the Digital Twin platform requires scalable connectivity to support real-time massive twinning and continuous process monitoring.

In addition, B5G/6G networks provide the scalability and device density needed to support the large-scale integration of IoT-enabled camera systems and other connected port assets within a unified digital ecosystem. Their advanced capabilities in edge computing, AI-native networking, and intelligent orchestration can further enhance automated monitoring, data analytics, and the enforcement of safety measures. These technologies also contribute to sustainability objectives by enabling more efficient coordination of logistics operations, reducing unnecessary movements and delays, improving resource utilisation, and supporting safer working conditions with reduced physical human–machine interaction. Overall, B5G/6G capabilities are essential to achieve the performance, scalability, automation, and reliability targets required for the successful deployment of next-generation smart port operations.

Large-scale deployment of the pilot

Large-scale deployment will be addressed in the pilot through both the communication infrastructure design and the scope of the implemented digital services. In particular, the deployment of the private B5G/6G network across the full port area will ensure wide-area coverage and seamless connectivity for port operations. The installed network equipment and infrastructure will be specified and dimensioned to support connectivity over large operational areas, enabling all connected assets, IoT devices, cameras, and port equipment to reliably access the B5G/6G network. This approach will allow the pilot to validate network performance, scalability, reliability, and interoperability under realistic operational conditions representative of full-scale smart port deployments.

In parallel, large-scale deployment aspects will also be addressed through the implementation of the Digital Twin platform, which will cover almost all major port processes, from vessel unloading to truck loading. Although the teleoperation pilot may initially involve only one of the four Ship-to-Shore (STS) cranes, as well as selected port equipment such as straddle carriers, the integration of these assets within the broader Digital Twin ecosystem will enable a comprehensive assessment of the operational impact of large-scale teleoperation and automation deployment. The Digital Twin environment will provide the capability to monitor, analyse, and evaluate system behaviour, operational efficiency, safety improvements, communication requirements, and potential bottlenecks under different deployment scenarios.

Furthermore, the pilot will enable the simulation and evaluation of multiple large-scale operational scenarios, including the gradual expansion of teleoperated and automated equipment across the port infrastructure. This capability will provide valuable insights into the expected gains, limitations, scalability challenges, and resource requirements associated with full-scale deployment. As a result, the pilot will not only validate the feasibility of the proposed technologies in realistic conditions but will also provide a clear understanding of the technical, operational, and sustainability impacts of future large-scale adoption of B5G/6G-enabled automation and Digital Twin technologies in port environments.

Pilot Architecture

Figure 26 below depicts the high-level architecture of T5 UC. The T5 high-level architecture for port logistics in Thessaloniki is organised in layered domains spanning port field assets, private B5G/6G connectivity, orchestration for edge/core execution infrastructure, intelligence and data services, and port operational applications for data visualisation. At the field level, cranes, cameras, localisation units, and environmental sensors generate real-time operational data and receive control commands. A private 5G SA/B5G platform with edge computing provides deterministic low-latency connectivity and local breakout for mission-critical teleoperation and analytics. The Digital Twin, AI/ML, monitoring, zero-touch and service and resource orchestration services support real-time situational awareness, optimisation, service assurance, and remote control. Finally, teleoperation HMIs and web-based dashboards for port operations expose these capabilities to operators and port stakeholders.

Figure 26: High-Level Architecture of T5 UC

For the trial implementation, several enablers will be employed after their development to ensure the successful completion of the use case implementation. These key components are summarised in Table 25 below.

Table 25: T5 Use Case High-Level Enablers Utilisation

Domain & Enablers Reference	Component	UC Related Functionalities
Communication - Communication resource management	Zero-touch network and service management	Provides entire system and network automated provisioning, monitoring, optimisation, self-healing, SLA compliance, rapid fault recovery and continuous adaptation to dynamic port conditions such as vessel arrivals and variable crane workloads, including crane teleoperation
Communication - Communication resource management	Automated gNB configuration	Assists in keeping the dedicated B5G network reliable, with ultra-low latency and continuous service, while having automated setup of gNodeB stations or even programmatically control a gNodeB's parameters for setting up frequencies, bandwidth, and interfaces
Communication - Communication resource management	Network performance, monitoring and control	STS crane remote operation requires safe and real-time operation by maintaining strict latency and throughput requirements with high-quality video and haptic feedback while enabling rapid detection of anomalies before they impact crane movements
Communication - Communication resource management	Identification and selection of backhauling	The identification and selection of backhaul links among the available ones is critical to deliver the low-latency, high-capacity, redundant, and secure connectivity that keeps port operations and STS crane teleoperation running smoothly end to end.
Communication - Public-private network integration	Public-private network integration	Enables seamless swap of backhaul connectivity of the private network between local connections and the use of public cellular networks operated by telecom providers

Communication - Multiple-RAT connectivity	Wi-Fi integration	Complements B5G by handling non-critical traffic, allowing operators to balance cost with performance by assigning the right network for the right task, achieving seamless connectivity across the port ecosystem. This allows the dedicated B5G network to be utilised exclusively for the safety-critical teleoperation of the STS crane.
Compute-as-a-Service - Compute resource management	Intelligent service and resource orchestration in extreme edge/edge/cloud compute continuum.	The Digital Twin platform that models the entire port operation from vessel unloading to truck loading relies on continuous data ingestion from IoT sensors, real-time analytics, and simulation updates. These requirements can only be met with reliable orchestration and automated service management through containerization, cloud-native orchestration and CI/CD solutions.
Application enablers and AI - AI-as-a-Service framework	ML models catalogue	AI and Machine Learning predictive modelling techniques will be developed to optimise multiple aspects of operations at the Port of Thessaloniki, hosted on a dedicated repository for better discoverability and ease of ML operations. These models will be embedded within the port's Digital Twin environment, enabling decision makers to operate at a higher level of detail and accuracy.
Application enablers and AI - AI-as-a-Service framework	MLOps	The MLOps platform will enable all lifecycle operations of the use case's ML models, such as training, testing and deployment through AI/ML pipelines.
Application enablers and AI - Digital Twin framework	Digital Twin framework	The Digital Twin will be developed as a decision-support environment focused on optimising straddle carrier operations. By replicating yard processes and container flows, the framework enables the testing of routing algorithms and operational strategies under realistic conditions.

IoT and localization enablers - Advanced Localization and Positioning Systems	Hybrid Localization (5G + GNSS)	A vehicle will be used that will be equipped with a LiDAR sensor to capture a detailed 360-degree view of its surrounding environment, enabling the system to continuously perceive and map everything within the area of operation. A Real-Time Kinematic (RTK) positioning system can be integrated into this vehicle to deliver precise geolocation data, achieving sub-centimeter accuracy in outdoor settings when needed.
IoT and localization enablers – IoT Connectivity and Infrastructure	IoT Gateways	IoT gateways are important for the connectivity of available devices/ sensors.
IoT and localization enablers – IoT Connectivity and Infrastructure	IoT Sensors and Devices	IoT sensors and devices are relevant for monitoring environmental parameters.
IoT and localization enablers - IoT Service Platforms and Management	IoT Exposure Function	IoT platforms' APIs are exposed for functionality and message handling. They are essential to provide the necessary information to the users who will participate in the trial in the Port area and provide further analytics and results of the project, while also enabling configuration capabilities, while keeping their function uninterrupted.
IoT and localization enablers - Remote Control and Operation (Actuation)	Real-Time Actuation/Teleoperation	Remote-control operation of a Ship-to-Shore crane in the Port of Thessaloniki premises will take place, utilising a dedicated B5G/6G network.

2.5.2 Deployment and implementation details

2.5.2.1 Design of the application

The T5 use case involves the remote-control operation of a Ship-to-Shore crane in its premises, utilising a dedicated B5G/6G network. Unlike semi-automated processes, teleoperated crane movements involve human-in-the-loop control, where delays, jitter, or loss of precision directly affect container handling accuracy, port efficiency, and operational safety. To achieve real-time actuation and teleoperation of an STS crane, instantaneous responsiveness, continuous situational awareness, and uncompromised safety must be reassured.

The pilot implementation is broken down into two Subcases:

[Subcase 1] From ship to shore: Ship-to-shore (STS) cranes are used to move containers between ships and the quay, handled by operators on board. In this subcase, remotely controlled and monitored operations of STS cranes are introduced (equipped with sensors and cameras), utilising the B5G/6G network; with automated cranes being able to perform additional functions autonomously (e.g., video streams may not be continuously transmitted), so that one crane operator can supervise multiple cranes simultaneously.

[Subcase 2] From shore to truck: Straddle Carriers are used to move containers from the quay to the areas for loading the trucks. Similar to STS, Straddle Carriers are also equipped with sensors and cameras to optimise the routing between the areas of (un)loading. Moreover, cameras are used for monitoring and human intervention, especially for navigating areas where people are present. This requires reliable exchange of position (and other) data collected from 3D sensors using B5G/6G connectivity. The Digital Twin platform will simulate the complex process of moving containers within the port to support the transition of these processes towards full automation. This solution is intended to replace the current equipment with Automated Guided Vehicles (AGVs). AGVs could be used in real-life implementations according to the Port's current expansion/investment plan. In that case, the addition of augmented reality (AR) will allow the representation of asset information on the ground with geolocation in real-time. This will enable improved response and efficiency in event resolution.

- **QoS:** is essential to ensure real-time responsiveness, reliability, and safety assurance required for remote human-in-the-loop operation of STS cranes, enabling precise container handling with the same effectiveness as on-cab control. This means that high-performance connectivity is required, delivering data (advanced sensors and video systems) with low latency, high reliability and availability
- **AIaaS:** can coordinate actuation across all devices and sensor data streams, positioned in the edge, where all real-time control functions, sensor fusion, video encoding, and safety applications take place.

In Figure 27, the system architecture of the T5 Use Case is organised across the same five functional domains shown in the high-level architecture in the previous section. At the Port Field Domain, STS cranes, cameras, mobile port assets, localization units, and environmental/safety sensors generate real-time telemetry, video streams, and contextual information. These data are transported through the Private 5G SA/B5G Domain, composed of the RAN, private 5G core, and MEC/edge compute platform, which provides deterministic low-latency communication, local breakout, and support for mission-critical teleoperation and analytics.

Above the connectivity layer, the Orchestration Domain implements zero-touch network and service management, auto-gNB configuration, network performance monitoring and control, and service/resource orchestration across the extreme edge–edge–cloud continuum. The Intelligence & Data Domain integrates the Digital Twin, AI/ML analytics, MLOps/model lifecycle management, and APIs/IoT exposure services to support real-time situational awareness, simulation, optimisation, and predictive capabilities. The Applications & Visualization Domain exposes these capabilities to operators and stakeholders through teleoperation HMIs and web-based operational/KPI dashboards.

The interactions between domains are bidirectional. In particular, telemetry, video, positioning, and monitoring data flow upward from the Port Field Domain toward the Intelligence & Data and Applications & Visualization domains, while teleoperation commands, orchestration actions, configuration updates, and QoS/slice control policies flow downward toward the field assets and connectivity infrastructure. This enables a closed-loop system in which port operations are continuously monitored, optimised, and controlled in real time.

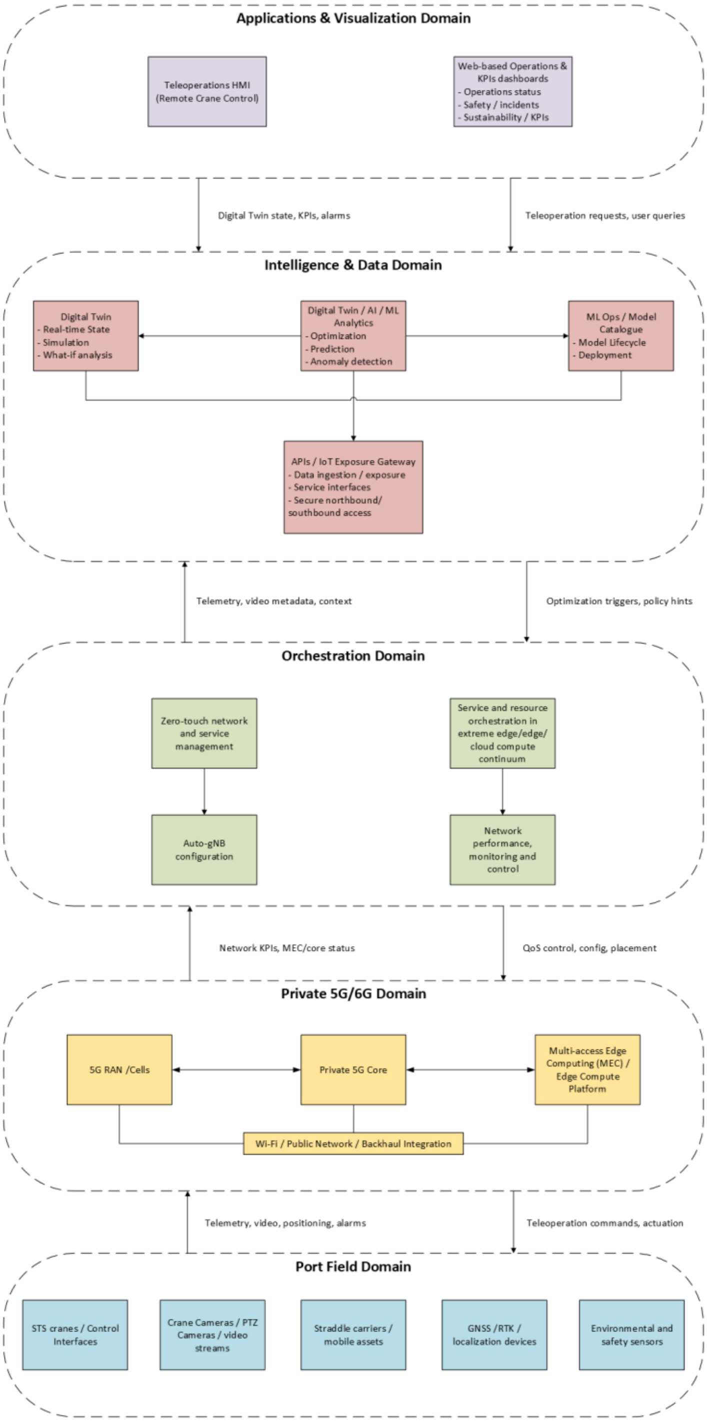

Figure 27: System Architecture of T5 UC

Digital Twin of the Port of Thessaloniki

The T5 UC will use an Anylogic Simulation Engine as the main microService, as well as supporting planning and optimisation models that will support the functionalities of the management challenges of Thessaloniki Port. The services planned at this stage are:

- **AnyLogic Simulator microService:** the Simulation Digital Twin microService, which serves as the backbone of the Port's digital model, enabling the testing of strategic process changes and operational/planning algorithm configurations. This service provides a robust environment for scenario analysis, enabling stakeholders to assess various strategic and operational decisions without disrupting real-world operations.
- **AI-based Port Operation microService:** it aims to enhance container loading and unloading efficiency by optimising the movement paths and scheduling of straddle carriers within the yard. This optimisation will lead to significant improvements in operational flow, reduced congestion, better utilisation of resources, and overall increased productivity of port activities.

The Digital Twin platform, which models the entire port operation from vessel unloading to truck loading, relies on continuous data ingestion from IoT sensors, real-time analytics, and simulation updates. These requirements can only be met with reliable orchestration and automated service management.

To achieve this, a containerised architecture is introduced to host models for predictive analytics and optimisation. These models are critical for enabling the Digital Twin platform to simulate and optimise port operations dynamically, including crane scheduling, straddle carrier routing, and AGV coordination. Containerised message brokers and subscribers facilitate real-time IoT data exchange from sensors and cameras installed on cranes, straddle carriers, and other port assets, ensuring that the Digital Twin receives accurate and timely information. Simulation services, deployed in Docker containers, allow for scalable and reproducible modelling of port processes, supporting scenario testing and optimisation.

To manage these distributed components, a cloud-native-based orchestration will be introduced, which automates deployment, scaling, and lifecycle management across heterogeneous compute resources. This orchestration is critical for handling fluctuating workloads, such as peak traffic periods in the port, while maintaining energy efficiency through dynamic resource allocation and DVFS techniques on edge devices. Secure interfaces and authenticated message queue topics ensure data integrity and privacy. Additionally, CI/CD pipelines (e.g., GitLab) enable continuous integration and automated deployment of ML/AI models, simulation services, and IoT components, reducing downtime and accelerating innovation cycles.

These capabilities are indispensable for meeting the stringent requirements of the port logistics use case, such as ultra-low latency for tele-operations, high reliability for safety-critical processes, and energy efficiency on edge devices. By enabling the coordination between IoT devices, edge nodes, and cloud resources, the above CaaS enablers ensure that the Digital Twin platform can operate in real time, supporting operational efficiency, safety, and sustainability objectives.

The workflow of Subcase A starts with the initialisation of a remote operation session for a selected STS crane or a supervised operation session for one or more MHC cranes. The crane operator, through the teleoperation human-machine interface, selects the target equipment and requests access to the corresponding telemetry, video, and control interfaces. These requests are resolved through the APIs and IoT exposure gateway, while the orchestrator instantiates the required service chain at the edge. In parallel, the network performance monitoring and control functions activate the appropriate QoS and connectivity policies, including low-latency service support and local breakout through the private 5G/MEC platform.

Once the communication and service environment is ready, the crane systems start transmitting real-time telemetry, video streams, alarms, and position-related information to the edge and data services. These data are forwarded to the Digital Twin, which maintains a synchronised representation of crane operations and feeds AI/ML analytics and dashboarding functions. The teleoperation interface then provides the operator with a real-time operational view, enriched with the relevant state, video, and performance indicators required for safe and efficient remote handling.

During the normal operation phase, the operator sends teleoperation commands through the HMI, which are transmitted through the private B5G/6G platform toward the crane control interfaces. In response, the crane sends back execution feedback, updated telemetry, and live video. This creates a continuous closed control loop between the operator and the crane system. At the same time, the Digital Twin and analytics components continuously process the incoming data to support situational awareness, detect anomalies, estimate performance, and provide operational recommendations.

For MHC cranes, an additional supervised automation mode can be supported. In this mode, selected crane functions, such as stacking-related tasks, may be executed semi-autonomously, while the operator remains in charge of overall supervision and intervention. This allows a single operator to supervise multiple cranes simultaneously, reducing the need for uninterrupted continuous video transmission from every crane, while still preserving visibility and control over critical events.

In parallel with the operational loop, the monitoring and orchestration functions continuously supervise latency, jitter, packet loss, edge load, and service performance indicators. When degradation or anomalous behaviour is detected, the orchestrator can trigger optimisation actions such as QoS reconfiguration, radio tuning, or adaptation of edge service placement. These actions ensure that mission-critical teleoperation remains protected and that the private 5G path continues to support the required operational reliability.

The workflow also includes a dedicated safety handling branch. If the Digital Twin or analytics layer identifies an unsafe state, anomaly, or collision-related condition, an alert is raised to the teleoperation interface. The operator can then pause the crane, force a safe state, or resume operation once the condition is resolved. At the end of the session, the service chain is terminated, QoS and monitoring policies are released, and the final logs, KPIs, and Digital Twin state are stored for reporting, replay, and evaluation purposes.

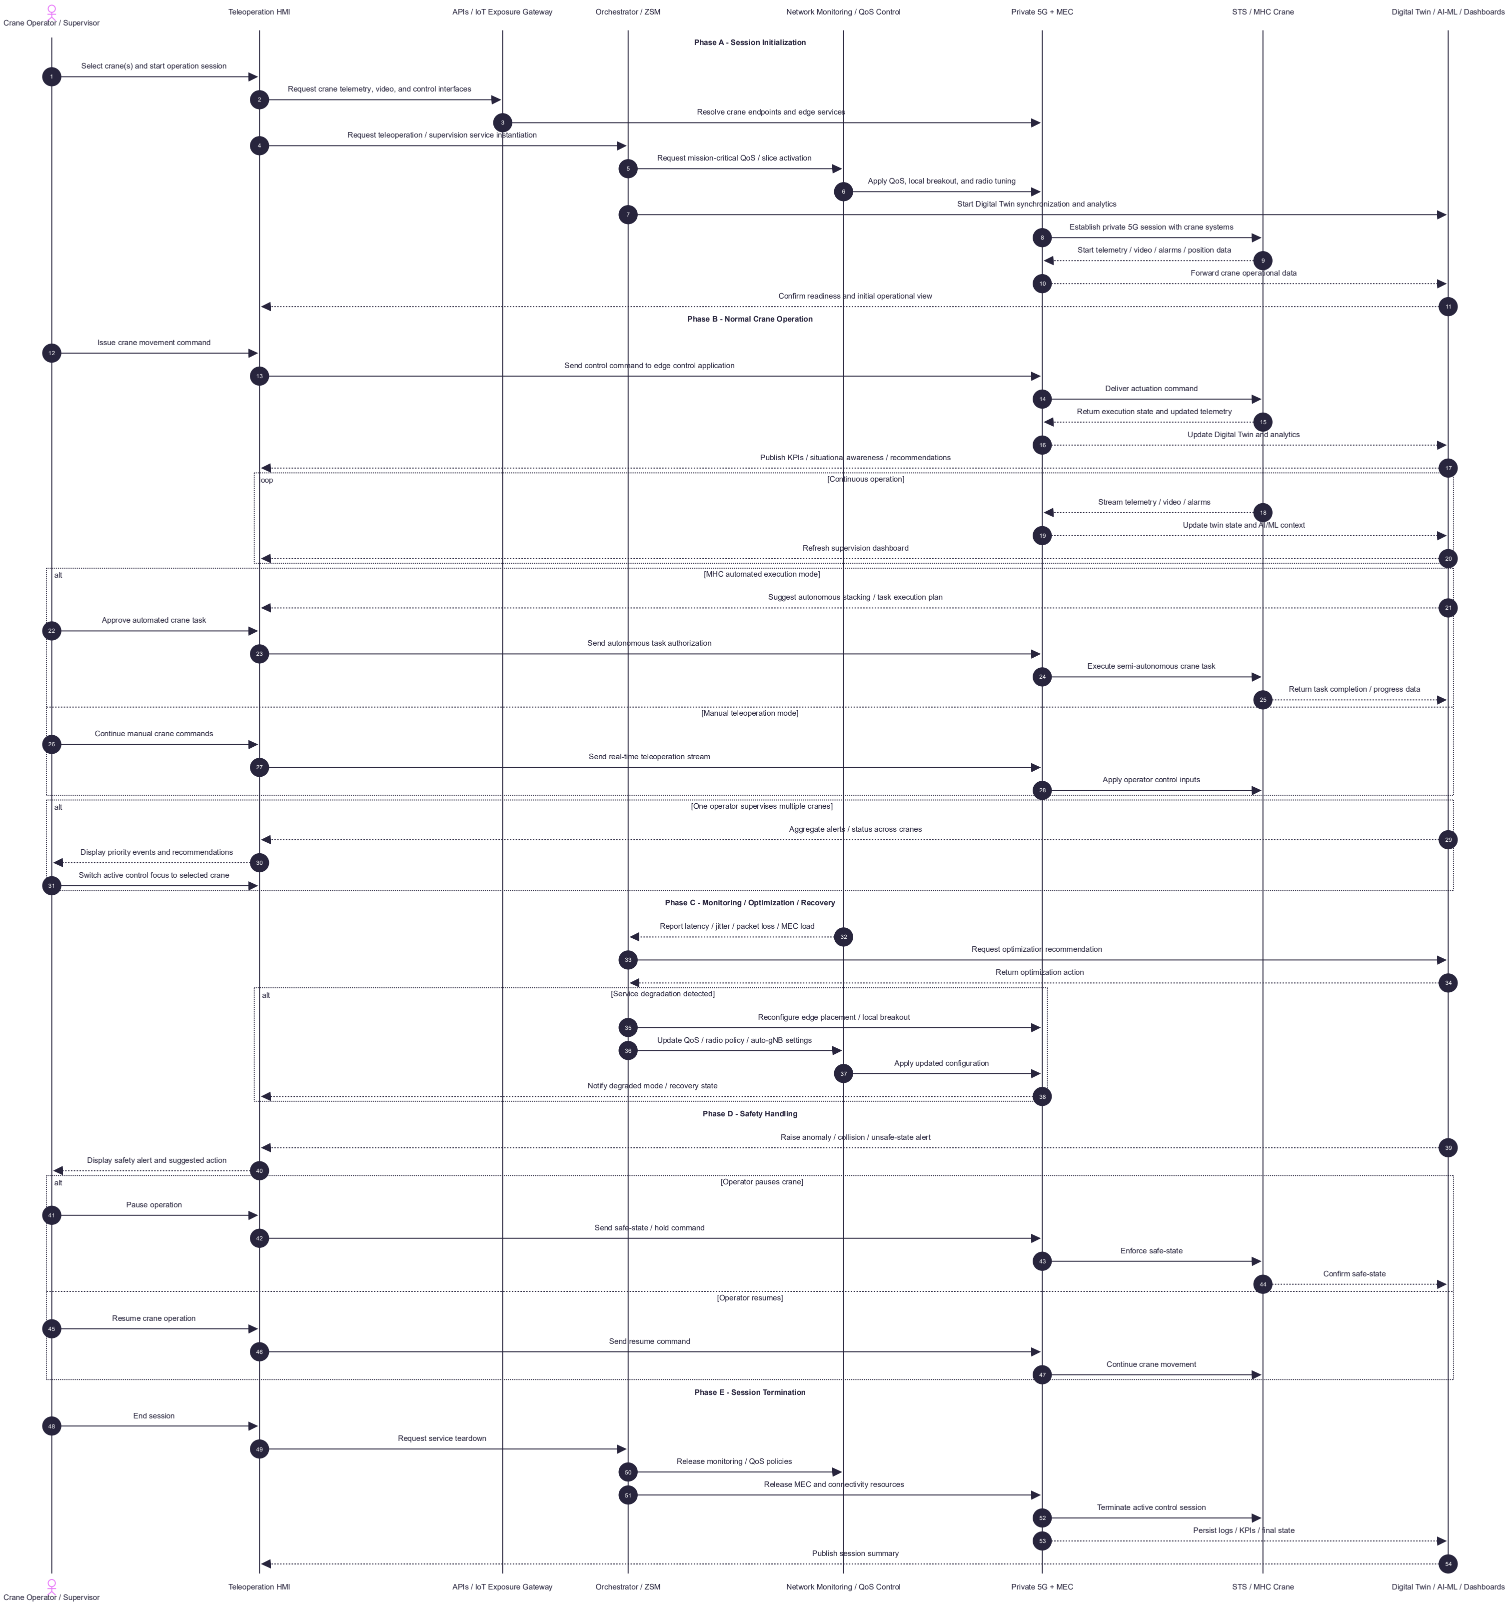

Figure 28: Work of T5 UC – Ship-to-Shore Crane Operations

The workflow of Subcase B starts with the initialisation of a container transfer mission between the quay and the truck loading area. Through the operations dashboard, the yard operator or supervisor selects the transfer task and requests access to the associated straddle carrier, route information, and available sensor data. These requests are resolved through the APIs and IoT exposure functions, while the orchestrator activates the required services for mission execution, monitoring, and route optimisation. In parallel, the network monitoring and control functions apply appropriate communication policies to support the reliable exchange of telemetry, positioning, sensor, and video data through the private B5G/6G platform.

Once the mission is activated, the straddle carrier starts transmitting position, speed, container handling state, 3D sensing data, video streams, and other operational telemetry toward the edge platform. These data are forwarded to the Digital Twin, which maintains a dynamic representation of the port transport process and continuously evaluates the state of the yard operation. Based on these inputs, the Digital Twin and AI/ML analytics components generate route optimisation suggestions, conflict detection indications, ETA updates, and operational recommendations, which are then visualised through the dashboard.

During normal operation, the selected route or mission plan is sent to the straddle carrier through the private B5G/6G platform. The carrier executes the movement and continuously updates the system with its location, load state, and environmental context. This creates a closed operational loop in which real-time positioning and sensing are combined with Digital Twin-based optimisation to improve route efficiency and reduce delays in container transfer.

A key part of this workflow is the support of human intervention in complex or mixed-traffic environments, especially when the straddle carrier must operate in areas where people may be present. In such cases, the Digital Twin or analytics functions may detect conditions requiring intervention and raise an alert to the operator dashboard. The operator can then validate a slowdown, stop, or manual override action, which is transmitted to the carrier through the B5G/6G communication layer. This ensures that optimisation and automation are always complemented by safe human supervision when needed.

In parallel, the network and service orchestration functions continuously monitor communication and infrastructure performance, including mobility-related connectivity conditions, edge resource load, and service quality indicators. If degradation is detected, the orchestrator may adapt the edge placement, update QoS policies, or modify connectivity settings in order to preserve service continuity and maintain the required performance for mission execution.

An additional objective of this workflow is to support the transition toward future AGV-based operations. To this end, the Digital Twin simulates alternative yard configurations and automation scenarios, enabling the evaluation of automated transfer processes before real deployment. In future implementations, these workflows may be extended to AGVs and enhanced with AR-based visualisation, through which geolocated information about assets, routes, and events can be displayed in real time to support faster and more efficient intervention and decision-making.

The mission ends once the container has been successfully transferred to the truck loading area. The final operational state, route efficiency metrics, safety indicators, and KPI/KVI-related measurements are then stored by the Digital Twin and supporting services. The operator closes the mission, the service chain is deactivated, and the compute and network resources are released.

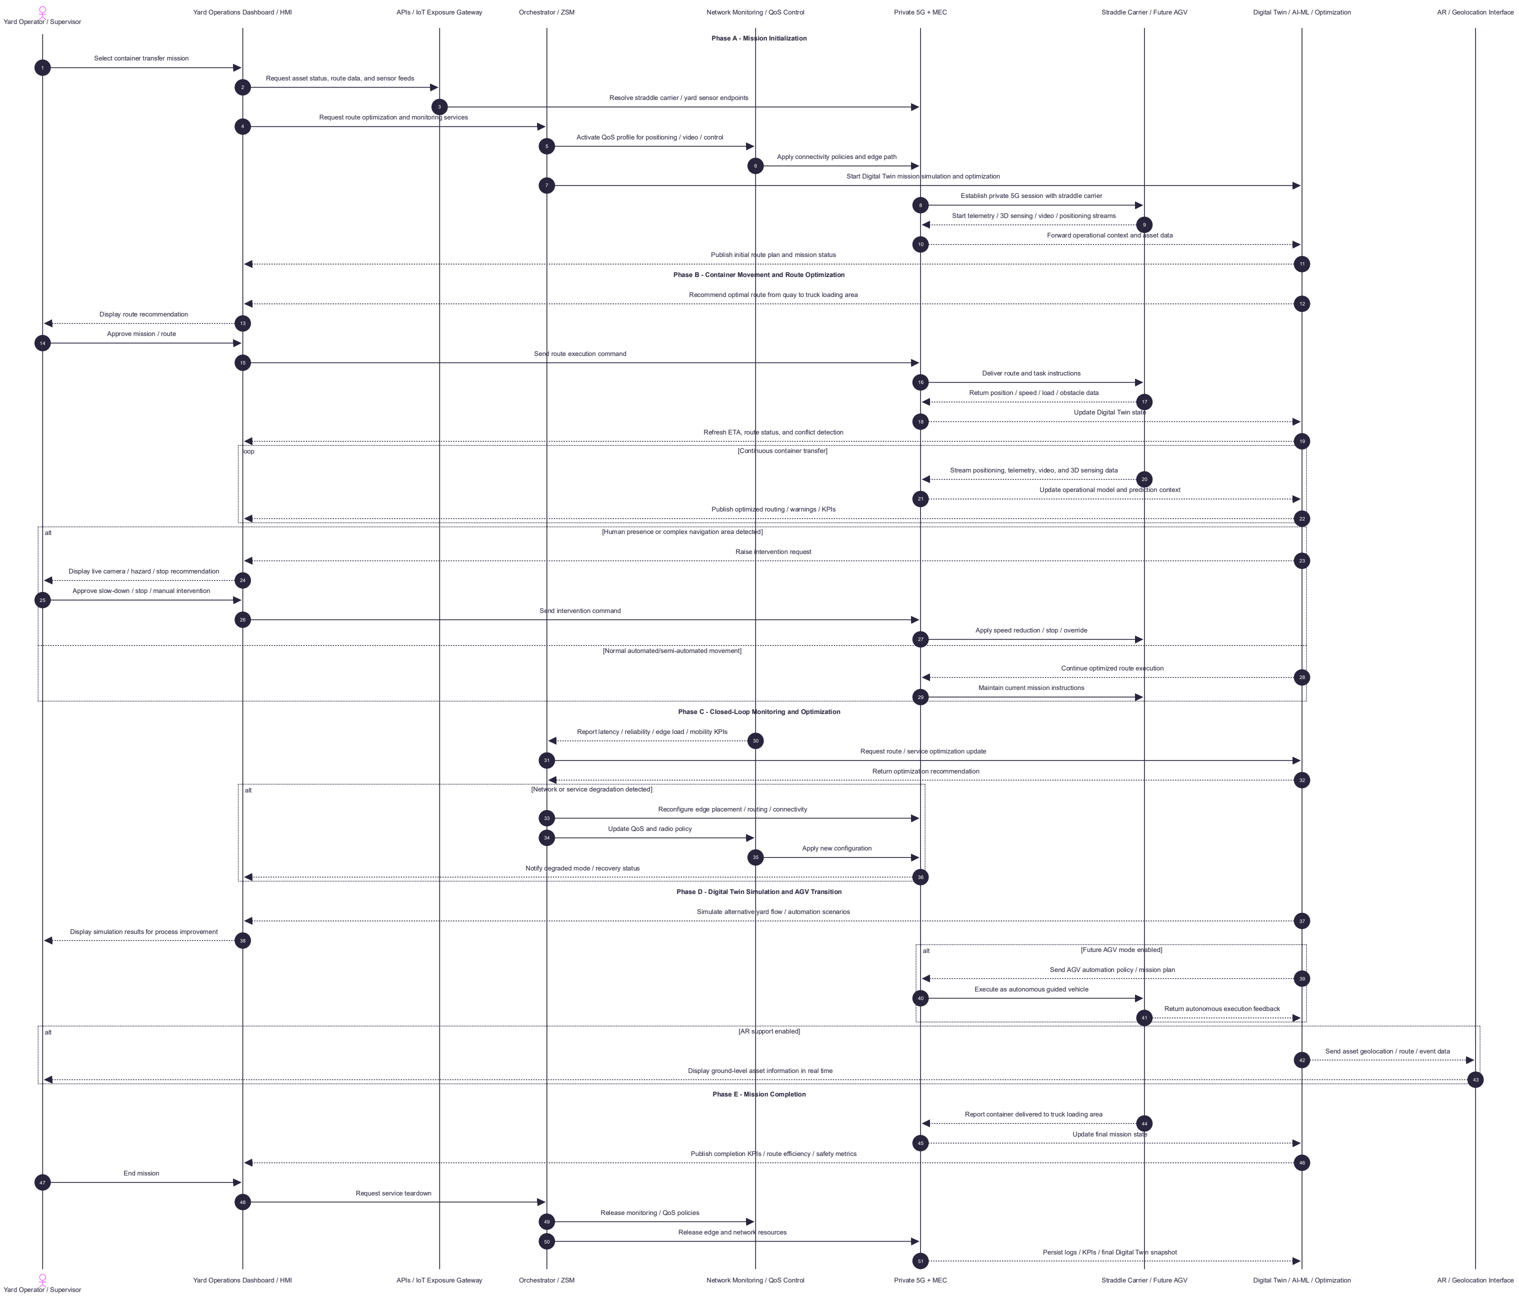

Figure 29: Workflow of T5 UC - Shore-to-Truck Container Transfer Operations

AI platform "wi.move"

The AI functionality for incident detection in moving vehicles and also for the timely identification of safety and security issues within the port area will be supported by wi.move. wi.move is a product for the management of transport infrastructure and the provision of intelligent mobility services.

Area	Configuration	Monitoring ("Accounting")	Faults	Performance	Security
1. Infrastructure (roads, bridges, …)	Visualization.	*1a. Structural health.* *1b. Traffic levels.*	Impairments prediction.	*1c. Infrastructure planning* (tuning).	
2. Stations: 2a. Vehicles 2b. Passengers 2c. Ships	Visualization. Component control.	*2a1. Where can I find a free park / charging spot? Parking spot allocation / scheduling.* *2b1. Crowd concentration / flow.*	Component failure. Passenger health hazards.	*2a2. Planning: Parking / Charging spot deployment.* *2b2. Routing: crowd evacuation.*	People / vehicle safety.
3. Vehicle management.	Location on map.	***Vehicle as a City Sensor.*** Conditions of vehicle / driver	Predictive maintenance.	*3a. Routing: vehicles, wrt various objectives.*	

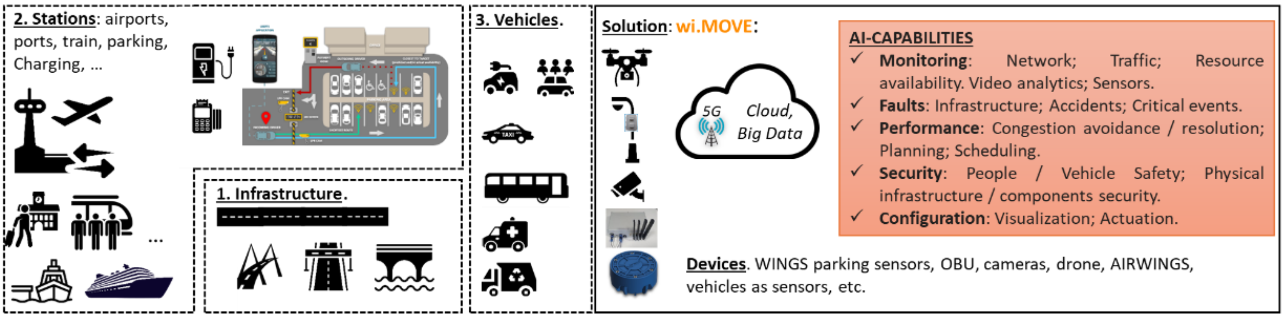

Figure 30: wi.move Platform Overview

Wi.move is a fully integrated management system for transport infrastructures that provides solutions for various stakeholders, such as public and private transport providers and infrastructure operators. Wi.move provides advanced monitoring, fault detection, performance optimisation, and configuration capabilities in the following areas:

1. **Infrastructure:** The platform provides tools for structural health monitoring, traffic level predictions, and predictive maintenance for infrastructure like roads, pavements, and bridges. By monitoring the structural health of infrastructure and analysing traffic levels, the platform can help stakeholders make decisions about maintenance and long-term planning. Additionally, predictive maintenance tools allow for more efficient infrastructure management, reduced downtime, and people's safety.

2. **Stations:**

    a. **Vehicles (Parking/Charging):** In the parking area, the platform offers advanced visualisation tools to help drivers find available parking and charging spots. It also monitors parking occupancy and can detect the failure of parking lot elements. This information can help with scheduling and booking and improve physical infrastructure and vehicle/people safety.

    b. **Passengers (Airport, Rail, Port)**: For transportation hubs like airports, ports, and train stations, wi.move provides mapping tools that show crowd concentrations and related statistics, such as flows, congestion areas, etc. Critical events and predictive maintenance tasks can also be detected, and itineraries can be provided to passengers. Additionally, insights can be provided to stakeholders, such as operators and retailers, to improve overall efficiency and people's safety, including evacuation plans.

    c. **Yachts**: Berth Management, Monitoring of passing vessels at entry/exit gates. Booking a spot, Integration with payments.

3. **Vehicle Management**: Vehicle… "as a city sensor", Infrastructure Inspection for predictive maintenance scenarios, Vehicles and In-Vehicle aspects, e.g., passenger transfer conditions,

City aspects, e.g., air quality and environmental parameters. Fleet Management for transport infrastructures. This includes real-time tracking of vehicles, traffic status monitoring, and ETA calculations to support efficient and timely operations. Also, routing optimisation is achieved through algorithms that consider vehicle capacity, maximum travelling distance, etc. This optimisation ensures that routes are planned efficiently, reducing travel time and costs. Predictive Maintenance scenarios can also be supported.

Wi.move utilises a range of **devices**, including the WINGS parking sensors (as well as 3rd party ones), the WINGS OBU (on-board unit), cameras, AMRs, drones, and more, to collect and analyse data about transportation and mobility infrastructure, parking, and transportation hubs. These devices provide real-time information about parking and charging spots, traffic, congestion, and resource availability, helping stakeholders make informed decisions about infrastructure management and maintenance.
Wi.move also incorporates advanced **IoT** and **AI** mechanisms, including video analytics, predictive analytics, and optimisation mechanisms, to improve transportation and mobility. With these tools, wi.move can predict demand and availability, conduct slot allocations (bookings, enforcement), and derive policies for infrastructure and parking management. This allows for more efficient use of resources and improved overall transportation and mobility experiences.

2.5.2.2 Devices, sensors and supporting equipment

In the following Table 26, the equipment that will be employed for the T5 UC testing is briefly presented.

Table 26: T5 Devices, sensors and supporting equipment

Category	Component	Technical Details	Units
5G Connectivity Module	5G-In-box (SPE-5G) 	A portable ruggedised edge DataCenter in-a-Box that provides a 5G bubble while running edge AI software simultaneously. Integrated within an Intel-powered Edge DataCenter system, the Private 5G Network In-a-Box combines a multicore CPU, GPU, ultra-fast NVMe storage, and remote management. This all-in-one solution delivers high-speed, low-latency communications alongside robust, on-site data processing and storage. Enhancing live events and critical operations with secure, efficient, on-premises Edge Cloud connectivity. Smartphone App for 5G easy management.	1

IoT sensors	CT cameras	AXIS PTZ cameras with 21-32x optical zoom, 360 ° visibility	5-6
Security camera	HIKVISION DS-2XS2T47G1-LDH/4G	4 MP ColorVu Solar-powered Security Camera	2
TETRA Portable Terminal	MOTOROLA MTP3550 MOTOROLA MTP5400	TEDS Ready hardware - enables 20x faster data connectivity for accessing back-office systems and databases	2
Crane control system	AC500 controller (process PLC, safety PLC, communication interfaces and I/O units)	Modbus RTU protocol with a TCP interface that runs on Ethernet	1
Remote Control Stations (RCS)	An operator workstation used to monitor and manage automatic cranes (quay, stacking, and intermodal)	Tabletop station/console with joysticks and controls that support the operator's natural workflow, featuring optimised control layouts, personalised settings, integrated displays, and support for 24/7 operations	1

2.5.3 Network infrastructure and functional components

The T5 UC network infrastructure is entirely provided by ACROMOVE through its portable Private 5G Network-in-a-Box solutions, deployed on-site at the Container Terminal of Thessaloniki Port Authority (ThPA). The infrastructure comprises two complementary platforms: the ServerPack™Edge 5G (SPE5G) and the ServerPack™Edge Ultra (SPE Ultra), which offer two alternatives for a self-contained, rapidly deployable B5G/6G edge computing and communication environment. No external mobile network operator or third-party 5G Core provider is involved; the entire radio access network, core network, and edge compute infrastructure are co-located within the ACROMOVE equipment at the port premises.

2.5.3.1 Private 5G RAN and Core Network (ServerPack™Edge 5G)

The ACROMOVE ServerPack™Edge 5G (SPE5G) is a portable, ruggedised Data-Center-in-a-Box that integrates a complete 5G Standalone (SA) network within a single shippable unit. The SPE5G houses an AMD-based server running the 5G SA Core Network functions (AMF, SMF, UPF, UDM, AUSF, PCF, NRF, NSSF), an SDN management router (MikroTik RouterOS) providing advanced routing, VPN, firewall, and DHCP/DNS capabilities, a Wi-Fi 5 dual-band access point (802.11b/g/n/ac Wave 2, MU-MIMO, up to 867 Mbit/s), and an LTE/5G cellular gateway with dual SIM auto-switch for WAN backhaul.

The radio access is provided by up to four Askey 5G NR Small Cells with integrated baseband and radio, supporting 5G NR bands N48/N77/N78 with 2×2 MIMO and wide bandwidth. Each Small Cell offers approximately 280–300 metres radius coverage (line of sight). The Small Cells connect to the SPE5G via fibre/power cables, which carry both 10 Gbit single-mode fibre data.

Time and frequency synchronisation across the Small Cells and the 5G Core is ensured via a Silicom GNSS PTP Grandmaster integrated within the SPE5G server, using Precision Time Protocol (IEEE 1588v2) with a GPS antenna for clock accuracy of 100 nanoseconds or below. PTP synchronisation enables cell-to-cell handovers (N2 and XN) between Small Cells and minimises interference from nearby base stations operating in the same frequency band.

The SPE5G is equipped with an integrated UPS comprising up to four hot-swappable Li-Ion batteries (98 Wh each), enabling battery-powered operation for up to two hours and ensuring uninterrupted service during power transitions. The system supports AC mains input (100–240V, 50/60 Hz), DC input (10V/40A to 28V/15A, compatible with solar panels and vehicle batteries), and battery operation, providing deployment flexibility across the port environment. The entire private 5G network can be deployed in approximately 15 minutes.

The infrastructure supports logical network slicing and QoS profiles, isolating safety-critical crane teleoperation traffic from IoT telemetry and best-effort services. The user plane allows local breakout towards edge-hosted applications (Digital Twin, AI/ML inference) with minimal latency, as the UPF is co-located with the RAN and edge compute on the same platform.

2.5.3.2 Edge Computing Platform (ServerPack™Edge Ultra)

The ACROMOVE ServerPack™Edge Ultra (SPE Ultra) serves as the high-performance edge computing platform for the T5 UC. The SPE Ultra is designed for low-latency, compute-intensive workloads at the edge, including AI/ML inference, Digital Twin simulation, and real-time video analytics. It is built around up to a 6th Generation Intel Xeon processor (up to 42 cores, 84 threads, 168 MB cache) with Intel AMX, AVX-512, vRAN Boost, and Data Streaming Accelerator (DSA) capabilities. The platform includes an NVIDIA L4 Tensor Core GPU providing up to 485 TOPS of compute performance (7,424 CUDA cores, 240 fourth-generation Tensor Cores), which accelerates ML model inference for port operations optimisation, predictive maintenance, and environmental monitoring.

The SPE Ultra supports up to 128 GB DDR5-6400 ECC RAM, up to 2 NVMe SSD storage (PCIe 5.0), and high-speed networking via 2×QSFP28 ports (up to 100 Gbit each) and 2×SFP28 ports (10 Gbit single-mode fibre) for Small Cell connectivity. The platform features a crash-proof and waterproof aircraft-aluminium chassis with proprietary anti-shock design, TPM 2.0 security, and operates within a wide temperature range (–40°C to 55°C ambient).

2.5.3.3 Network Management and Connectivity

The integrated management router provides the internal network backbone (preconfigured on 192.168.88.0/24), offering DHCP, DNS, advanced firewall, and VPN support (OpenVPN, IPsec, WireGuard, among others). Device management is accessible via the Wi-Fi access point, direct wired connection through 10GbE switch ports, or remotely over WAN through port forwarding. The integrated server is managed via IPMI v2.0, providing out-of-band access including KVM-over-IP, remote power control, hardware sensor monitoring, and BIOS/firmware updates independently of the operating system state.

Public-private network integration is achieved through the WAN interface (1GbE) and the LTE/5G cellular gateway, which enables seamless fallback to public cellular networks operated by telecom providers for backhaul redundancy. Wi-Fi integration complements the 5G connectivity for non-critical traffic, allowing operators to balance cost with performance by routing different traffic classes over the most appropriate access technology. This ensures that the dedicated private 5G network is reserved for safety-critical crane teleoperation, while IoT monitoring, environmental data, and management traffic can utilise Wi-Fi or public network paths.

2.5.4 Use case enablers

As presented in deliverable D3.1, T5 enablers are presented in Table 27.

Table 27: List of enablers and sub-enablers for T5

Category	Technology enablers	Sub-enablers
Communication enablers	Communication resource management	Zero-touch network and service management
		Automated gNB configuration
		Network performance, monitoring and control
		Identification and selection of backhauling
	Public-private network integration	Public-private network integration
	Multiple-RAT connectivity	Wi-Fi integration
Compute-as-a-Service enablers	Compute resource management	Intelligent service and resource orchestration in extreme edge/edge/cloud compute continuum.
Application enablers and AI	AI-as-a-Service framework	ML models catalogue
		MLOps
	Digital Twin framework	Digital Twin framework
IoT and localization enablers	Advanced Localization and Positioning Systems	Hybrid Localization (5G + GNSS)
	IoT Connectivity and Infrastructure	IoT Gateways
	IoT Connectivity and Infrastructure	IoT Sensors and Devices
	IoT Service Platforms and Management	IoT Exposure Function
	Remote Control and Operation (Actuation)	Real-Time Actuation/Teleoperation

These enablers cover all activities within the context of T5, including the deployment of a robust, scalable, and resilient communication network capable of supporting the advanced digital functionalities of the port ecosystem. This includes enabling reliable connectivity for the teleoperation of Ship-to-Shore (STS) cranes and the automation of straddle carriers, ensuring safe, efficient, and synchronised port operations.

In addition, the deployed infrastructure shall support high-capacity and continuous data streaming requirements associated with safety, security, and operational monitoring through large-scale IoT device integration and the wi.move platform. This encompasses the collection, transmission, processing, and analysis of real-time operational and environmental data generated by sensors, cameras, connected vehicles, and other smart port assets.

Finally, the enablers support the efficiency, optimisation, and sustainability enhancement of port operations through the implementation of an end-to-end Digital Twin of port processes. The Digital Twin framework enables near real-time situational awareness, predictive analytics, operational optimisation, and data-driven decision-making across the port environment, contributing to improved resource utilisation, reduced operational inefficiencies, lower emissions, enhanced safety, and the transition toward more autonomous, sustainable, and climate-neutral port operations.

2.5.5 KPIs definition

The selected KPIs have been defined to provide a comprehensive assessment of both the technical performance of the deployed enablers and their overall impact on port operations within the context of the T5 pilot.

On the one hand, the KPI framework includes technical and system-level indicators that evaluate the operational effectiveness, reliability, scalability, responsiveness, and overall performance of the deployed communication, IoT, automation, and Digital Twin enablers. These indicators measure the extent to which the implemented technologies achieve the targeted levels of functionality and performance, including aspects such as connectivity quality, latency, data throughput, service availability, resilience, and real-time operational capability.

On the other hand, considering the particular nature of the T5 pilot, which is implemented within the highly complex and dynamic environment of an operational port, the KPI framework also incorporates operational impact indicators that assess the tangible effects of the pilot implementations on port processes and activities. These KPIs evaluate improvements related to operational efficiency, safety, sustainability, resource utilisation, automation effectiveness, and environmental performance. In this way, the assessment framework extends beyond purely technical validation and captures the broader operational and strategic value generated by the pilot deployment.

Furthermore, these operational KPIs are strongly aligned with and contribute to the evaluation of the pilot's Key Value Indicators (KVIs), ensuring that the performance assessment reflects not only the technical success of the implemented enablers but also their contribution to the overall business, operational, societal, and sustainability objectives of the pilot.

Table 28: List of KPIs for T5

KPI name	Description/KPI definition	KPI target
Deployment time	Minimum amount of time for non-experts to deploy the system.	< 15 min
Service Reliability	Ability to monitor network configuration remotely.	99.99%

End-to-End network service latency (RTT)	RTT (ms) measurement	<15ms
Network availability	Maintain connectivity to the public network as a backup and/or use of public resources	99.99%
UL/DL Throughput	Uplink and downlink data rates achieved over the public-private integrated B5G network, measured between UE and edge/cloud services during port operations. Validates seamless data plane continuity across private and public network domains.	>100 Mbps DL, >150 Mbps UL
Multi-RAT handover latency	Time elapsed during handover between 5G NR and Wi-Fi connectivity within the port operational area, ensuring service continuity for teleoperation and IoT data flows. Measured as the interruption time perceived at the application layer during RAT switching.	<50 ms
AI model accuracy	Accuracy of detection of transport safety incidents (e.g. usage of seatbelts, etc.)	>80%
Operational Continuity Index (OCI)	Measures the level of gradual and linear response of the DT model to input changes with only minor deviations from linearity, under normal circumstances.	≥ 0.85
Boundary Logic Compliance Rate (BLCR)	Evaluates the fraction of degenerate or extreme scenarios for which the model produces outputs that match analytically expected results.	100%
PM10 Emissions	Measured concentration of PM10 or equivalent operational indicator associated with port transport and cargo-handling operations.	-50%
Dust emissions	Measured level of dust emissions generated by cargo-handling and transport movements in the port.	-50%
Safety levels	Number of safety-related incidents or accidents occurring in the operational environment. It captures the effect of automation, teleoperation support, improved situational awareness, and faster event detection/response.	-70%
Response time to emergencies and incidents	Elapsed time between incident detection/alarm generation and operational response initiation by relevant personnel or automated support function.	-50%
Average truck turnaround time	Average time required for a truck to complete its targeted operational cycle in the port area, from arrival/process initiation to departure/process completion. It measures the effect of routing	-30%

	optimisation, real-time visibility, Digital Twin-assisted coordination and better synchronisation between unloading, yard and loading operations.	
Average straddle carrier productivity	Average number of effective moves per hour performed by straddle carriers.	10%
Average NCR (net crane rate)	Average number of effective crane moves per hour achieved during cargo-handling operations.	15%
Average container loading/unloading time	Average time required to load or unload a container during the port processes.	-10%
Average truck loading time	Average time needed to complete the truck loading process.	-25%
Average vessel waiting time at anchor	Average time a vessel remains at anchor before berth-side process execution can begin.	-10%
Average time of crane maintenance	Average duration required to complete a maintenance intervention on the crane.	-25%

2.5.6 KVIs definition

The KVIs of the T5 UC are presented in

Table 29. They are defined to capture the societal, environmental, and economic impact of the use case. During their definition, several key areas of interest were identified.

The first area relates to operational safety, focusing on the reduction of human–machine interaction through automation and teleoperation. This is expected to limit operator exposure to hazardous tasks and environments, leading to a measurable decrease in accidents and faster response times in critical situations. A second area concerns workforce skills and digital readiness. The deployment of B5G/6G-enabled Digital Twin and AI-based solutions requires upskilling of personnel. The corresponding KVIs assess the effectiveness of training activities, both in terms of participation and user acceptance of the introduced technologies. Another area of interest is the environmental footprint of port transport operations. The KVIs in this category evaluate the reduction of emissions achieved through improved coordination of operations, optimised routing, and reduced idle and waiting times. These improvements are enabled by the availability of real-time data and enhanced situational awareness. From an economic perspective, the KVIs also address operational cost efficiency, with a focus on the optimisation of key port processes such as vessel and truck turnaround times and cargo handling operations. In addition, asset efficiency is considered through KVIs related to maintenance, where predictive approaches are expected to reduce downtime and associated costs.

The KVIs will be evaluated using a combination of KPI analysis, simulation, comparative analysis, expert assessment, and questionnaires. This approach allows both quantitative and qualitative aspects to be captured, especially for impacts that cannot be fully measured within the duration of the pilot.

Table 29: List of KVIs for T5

KVI name	Description/KVI definition	KVI target
Operational safety improvement - create a safe working environment through reduced human-machine interaction (automation and teleoperation)	Captures the positive impact of automation and reducing direct human involvement in high-risk operational tasks, minimising exposure to hazardous environments and decreasing the risk of accidents and human injuries	-70% accidents, -50% response time
Workforce skills and digital Readiness Improvement	Raise workforce skills through competence building in the use of new technologies.	≥80% of targeted staff successfully trained, average training satisfaction ≥4.0/5
Reduced environmental footprint of port transport operations - Improve the environmental footprint (reduced emissions) of port transport operations through optimal use of resources.	The extent to which the environmental footprint of port logistics and transport operations is reduced through optimised routing, reduced idle and waiting times, improved synchronisation of cargo-handling activities, and data-driven operational decisions enabled by B5G/6G, IoT and Digital Twin capabilities.	-50% PM_{10} emissions, -50% dust emissions
Operational cost efficiency through port process optimisation	Reduced cost through port operations (time) optimisation, i.e., turnaround times (ship, truck), unloading times (ship, container, truck).	-30% Average truck turnaround time, +15% Average NCR, +10% Average straddle carrier productivity, -10% Average container loading/unloading time, -25% Average truck loading time, -10% Average Vessel waiting time at anchor
Reduced maintenance downtime and cost	Reduced downtime and cost of maintenance through optimal use of resources and predictive maintenance.	-25% Average time of maintenance

2.5.7 Description of the trial setup

The T5 trial setup is described in the following Table 30.

Table 30: Description of T5 trial setup

T5 Trial		
Trial Infrastructure/Venue	Port of Thessaloniki (GR) Container Terminal	
Trial Description	The T5 trial will demonstrate a next-generation, B5G/6G-enabled Digital Twin platform for port operations, covering the full workflow from vessel unloading to truck loading. During the testing, real-time data from cranes, vehicles, and sensors will be integrated into a dynamic virtual environment, enabling remote operation, monitoring, and optimisation of key processes. The trial will validate advanced capabilities such as tele-operated cranes, intelligent routing of straddle carriers, identification of incidents and coordinated cargo handling, while assessing performance in terms of efficiency, safety, and environmental impact. Through this end-to-end pilot, the transition towards more automated, resilient, and sustainable port logistics will be explored under realistic operational conditions.	
Technical Setup	**System Components**	• **Private B5G/6G Network Infrastructure** Private 5G SA network, MEC, small cells, Wi-Fi integration, and backhaul connectivity. • **Digital Twin and AI Platform** Digital Twin framework, simulation engine, AI/ML analytics, and MLOps services. • **IoT and Sensing Systems** Cameras, IoT sensors, LiDAR, telemetry, and hybrid 5G/GNSS localization. • **Teleoperation and Automation Systems** Remote crane control, teleoperation HMIs, actuation, and automation services. • **Monitoring and Orchestration Platforms** Zero-touch orchestration, network monitoring, KPI dashboards, and wi.move platform.
	System Configuration	• **Layered Distributed Architecture** Field, connectivity, orchestration, intelligence, and application domains.

		• **Edge-Based Private Deployment** Co-located 5G Core, RAN, MEC, and edge computing at the port premises. • **Real-Time Closed-Loop Operation** Continuous exchange of telemetry, video, analytics, and control data. • **Containerized Cloud-Native Services** Microservice-based deployment of AI, IoT, and Digital Twin services. • **QoS-Aware Resilient Connectivity** Network slicing, redundancy, failover, and QoS-controlled communications.
Pre-Trial execution	**Pre-conditions**	• Availability of wireless network connectivity and power supply for the installation of equipment
	Trial Steps	• **Step 1:** Installation of equipment (e.g. camera, etc.) • **Step 2:** Installation of the platform and other relevant software • **Step 3:** Running of pre-trials and collection of statistics
	Methodology	• Trained YOLOv11 model on seatbelt detection data • Run offline on fixed camera feed • Achieve detections in vehicles of all classes: cars, buses, trucks, etc.
	Technical Enablers	• Application enablers and AI
	Expected Outcomes	• Identification of incidents
Trial Execution	**Pre-conditions**	Availability of wireless network connectivity and power supply for the installation of equipment
	Trial Steps	• **Step 1:** Installation of equipment (e.g. camera, etc.) • **Step 2:** Installation of the platform and other relevant software

<table>
<tr><td rowspan="4"></td><td></td><td>Step 3: Running of trial and collection of data/results</td></tr>
<tr><td>Methodology</td><td>Trained YOLOv11 model on seatbelt detection data
Run offline on fixed camera feed
Achieve detections in vehicles of all classes: cars, buses, trucks, etc.</td></tr>
<tr><td>Complementary Measurements</td><td>To be defined in D6.2</td></tr>
<tr><td>Calculation Process</td><td>To be elaborated for D6.2</td></tr>
<tr><td rowspan="2">Technical Enablers</td><td>Communication Enablers</td><td>According to D3.1 (T5 enablers)</td></tr>
<tr><td>Compute-as-a-Service (CaaS) Enablers</td><td>According to D3.1 (T5 enablers)</td></tr>
<tr><td>Expected Outcomes</td><td colspan="2">Validation of seamless tele-operation of port equipment
Verification of seamless handover to private B5G/6G communication system without service interruption
Validation of operational and sustainability gains from the DT implementation for port operations optimisation
Identification of incidents</td></tr>
</table>

3 Use case planning

3.1 T1 planning

Table 31: T1 planning

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q1 2026 - done
Devices	Acquisition of hardware end devices	Q2 2026
Infrastructure & Application Design	Design of infrastructure and applications for the Turin site	Q4 2025 – Q1 2026 - done
Preliminary Integration, including enablers	Initial integration of all subsystems to verify interoperability	Q1/Q2 2026 – in progress
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q3 2026
Final Application Development	Complete application implementation, optimise performance and reliability	Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyse trial results	Q2–Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2–Q4 2027

The deployment of the infrastructure is experiencing some delays due to device procurement and the need to obtain permits for the installation of the mmWave gNodeB and the Roadside Unit.

3.2 T2 planning

Table 32: T2 planning

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q1 2026 - done
Devices	Acquisition of hardware end devices	Q2 2026
Infrastructure & Application Design	Design of infrastructure and applications for the Turin site	Q4 2025 –Q1 2026 - done
Preliminary Integration, including enablers	Initial integration of all subsystems to verify interoperability	Q1/Q2 2026 – in progress
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q3 2026
Final Application Development	Complete application implementation, optimise performance and reliability	Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyse trial results	Q2–Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2–Q4 2027

The deployment of the infrastructure is experiencing some delays due to device procurement and the need to obtain permits for the installation of the mmWave gNodeB and the Roadside Unit.

3.3 T3 planning

Table 33: T3 planning

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q2-Q3 2026 - done
Devices	Acquisition of hardware end devices	Q3 2025 - done
Infrastructure & Application Design	Design of infrastructure and applications for the TUC site	Q2-Q4 2025 - done
Preliminary Integration, including enablers	Initial integration of all subsystems to verify interoperability	Q2 2026 – in progress
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q3 2026
Final Application Development	Complete application implementation, optimise performance and reliability	Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyse trial results	Q2–Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2–Q4 2027

3.4 T4 planning

Table 34: T4 planning

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q2-Q3 2026 - done
Devices	Acquisition of hardware end devices	Q3 2025 - done
Infrastructure & Application Design	Design of infrastructure and applications for the TUC site	Q2-Q4 2025 - done

Pretrial testing of technological enablers	Testing enablers in the T4 scenarios, but using mockup applications associated with teleoperation	Q3 2025 – Q1 2026 - done
Preliminary Integration, including enablers	Initial integration of all subsystems to verify interoperability	Q2-Q3 2026 – in progress
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q3 2026
Final Application Development	Complete application implementation, optimise performance and reliability	Q4 2026
Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyse trial results	Q2–Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2–Q4 2027

3.5 T5 planning

Table 35: T5 planning

Phase	Details	Schedule
UC Definition	A pre-trial phase to lay the foundations of the activities and define the Use Case	Q3 2025-done
Devices	Acquisition of hardware end devices (sensors, redcap devices, etc.)	Q2 2025-done
Infrastructure & Application Design	Design of B5G infrastructure and applications for the Bucharest site	Q2–Q4 2025-done
Preliminary Integration, including enablers	Initial integration of all subsystems to verify interoperability	Q1 2026-done
First Tests & Initial Demonstration	Execute first tests and demonstrate preliminary UC scenarios	Q2 2026-done
Final Application Development	Complete application implementation, optimise performance and reliability	Q2–Q4 2026

Full Trial Integration	Integrate all UCs, devices, and applications into the final testbed	Q1 2027
Trial Execution	Conduct full UC trials, including operational scenarios and real-time monitoring	Q2 2027
KPIs & KVIs Collection & Analysis	Collect performance indicators and analyse trial results	Q2–Q3 2027
Final Demonstrations & Dissemination	Present final UC results, publish findings, and perform end-of-project demonstrations	Q2–Q4 2027

The planning of the pilot implementation follows a phased approach, progressively advancing from the definition and design stages towards full-scale integration, trial execution, and final demonstrations. Following the completion of the Use Case definition, device acquisition, and the design of the B5G infrastructure and applications during 2025, the project moved into the preliminary integration phase in Q1 2026, where the interoperability of the different subsystems and enablers was validated.

As part of the "First Tests & Initial Demonstration" phase scheduled for Q2 2026, T5 partners successfully executed the first tests and demonstrated preliminary Use Case scenarios during the 4th plenary meeting (held in Thessaloniki), at the Port of Thessaloniki. These initial demonstrations provided an important opportunity to validate the connectivity infrastructure for supporting teleoperation functionalities and IoT-based monitoring systems under realistic port conditions. The activities also allowed partners to identify technical challenges, assess interoperability and communication performance, and collect initial feedback from operational stakeholders.

The outcomes of these first demonstrations will support the next implementation stages, including the completion and optimisation of the final applications during Q2–Q4 2026, followed by the full integration of all Use Cases, devices, and applications into the final testbed in Q1 2027. Subsequently, the project will proceed with the execution of large-scale operational trials, KPI and KVI collection and analysis, and final demonstrations and dissemination activities throughout 2027. This phased planning ensures a gradual validation process, enabling continuous refinement of the technologies and providing a clear pathway towards the final large-scale deployment scenarios.

4 Preliminary integration and results

The integration done in the different trial sites for the Transport UCs is introduced in this Section. For each UC, the status of the integration activities of the different network and system components is reported. Preliminary results achieved from testing the initial system are presented as well.

4.1 Transport 1 and Transport 2 Use Cases

The T1 and T2 UCs are demonstrated in the same Trial site using the same network infrastructure (even if in a slightly different way). The integration activities performed so far in the Italian trial site are related to network and system components exploited in both UCs. To this end, the integration status and the preliminary results are jointly presented in the following of this deliverable.

4.1.1 T1 and T2 Integration status

4.1.1.1 mmWave RAN

The mmWave RAN architecture was first validated in a laboratory environment. The setup includes an HPE 5G Standalone Core, connected via VPN to the gNodeB, together with mmWave radio and CPE devices provided by Bellantenna.

The picture below represents the 5G Standalone architecture implemented in the LINKS Lab, where the transport of the N2 and N3 interfaces between the gNodeB and the 5G Core is carried over a VPN tunnel.

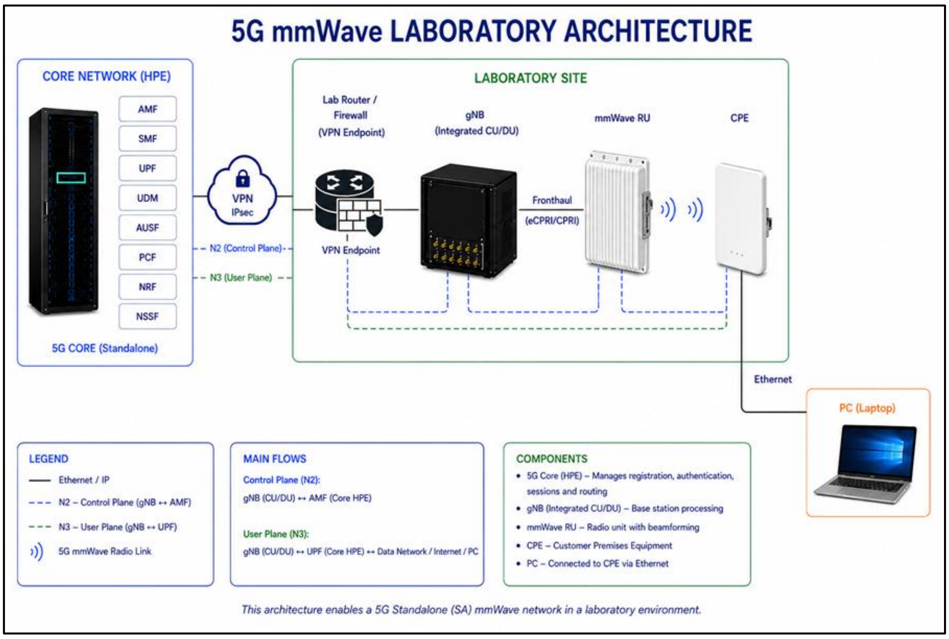

Figure 31: mmWave test Architecture

The configuration parameters are reported in the tables below

Table 36: RU parameters

Parameter	Value
Carrier frequency	26237.28 MHz
Subcarrier spacing: 120 kHz	120 kHz

Number of subcarriers: 792	792
Number of FFT: 1024	1024
Length of OFDM Symbol: 8.93 us	8.93 us
Number of OFDM Symbols:	14
Bandwidth	100 Mhz

Table 37: Test configuration

Parameter	**Value**
Core Network	HPE 5G Core (SA)
Transport	VPN tunnel (N2/N3)
gNB	Sourced by Bellantenna
Radio Unit	mmWave RU sourced by Bellantenna
CPE	mmWave CPE sourced by Bellantenna
Test Tool	iperf3
Test Duration	15 seconds
TCP Streams	1 and 2 parallel streams
UDP Target Rate	50 Mbit/s

4.1.1.2 Private Core network

The HPE Aruba Networking Private 5G Core has been integrated to enable end-to-end system operation, with support for network slicing to differentiate traffic handling. In this architecture, common control plane functions are shared across slices, while separate User Plane Functions are configured to manage different traffic profiles.

The deployed HPE Private Core Network solution is currently up and running on a physical server hosted at the LINKS data centre. Under normal operating conditions, general-purpose traffic is handled by a default UPF with best-effort characteristics, while an additional UPF can be dynamically activated to serve high-priority communications. The activation of dedicated user-plane resources and the traffic steering between slices are managed through orchestration mechanisms leveraging 5GC management APIs, allowing selected users or services to be migrated to dedicated resources when enhanced QoS is required.

4.1.1.3 Management and Orchestration Layer

The Management and Orchestration Layer has been successfully integrated into the T1 infrastructure within a dedicated Management Virtual Machine (VM) to enable automated control and coordination of system resources. This deployment includes the Service Orchestrator, which manages end-to-end service lifecycles, and the Resource Orchestrator (depicted in Figure 32), responsible for managing multi-cluster infrastructure across the edge-to-cloud continuum. Furthermore, the Closed-Loop Platform has been provisioned to govern AI-driven automation functions, working in conjunction with a centralised monitoring platform.

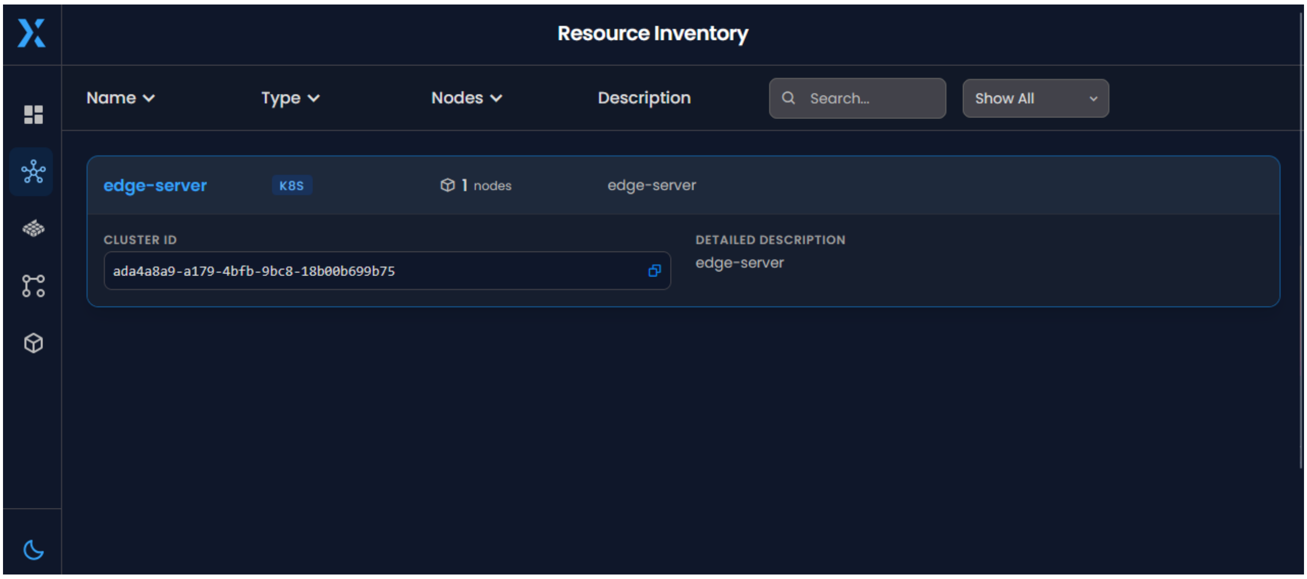

Figure 32: The Resource Orchestrator GUI

Within this monitoring platform, a new plugin has been developed and integrated to specifically track the energy consumption of Kubernetes-based services orchestrated through the existing management layer. A screenshot of the Graphical User Interface (GUI) of the InfluxDB data storage, which is collecting energy consumption data from already existing k8s-based services, is shown in Figure 33.

Figure 33: The InfluxDB data storage, part of the Monitoring Platform, which is collecting energy consumption data from k8s-based services

This addition allows the system to gather granular power metrics from containerised workloads, providing the necessary data for energy-aware offloading policies and proactive resource optimisation.

4.1.1.4 Exposure framework

The Exposure Framework has been successfully deployed within the same dedicated Management Virtual Machine as the orchestration components to provide a unified and secure interface for network and computational resources. This integration leverages the Common API Framework (CAPIF) to manage the registry of transformation functions and exposed interfaces. As part of the current integration status, several CAMARA-compliant APIs have already been onboarded to support the specific requirements of the T1 use case, as depicted in the following figure. These onboarded services enable the application layer to interact with the underlying B5G system for functions such as on-demand quality of service and edge application management. The framework utilises its internal API Gateway and core functions to ensure that all incoming requests are authenticated and routed to the appropriate backend transformation functions based on real-time network context.

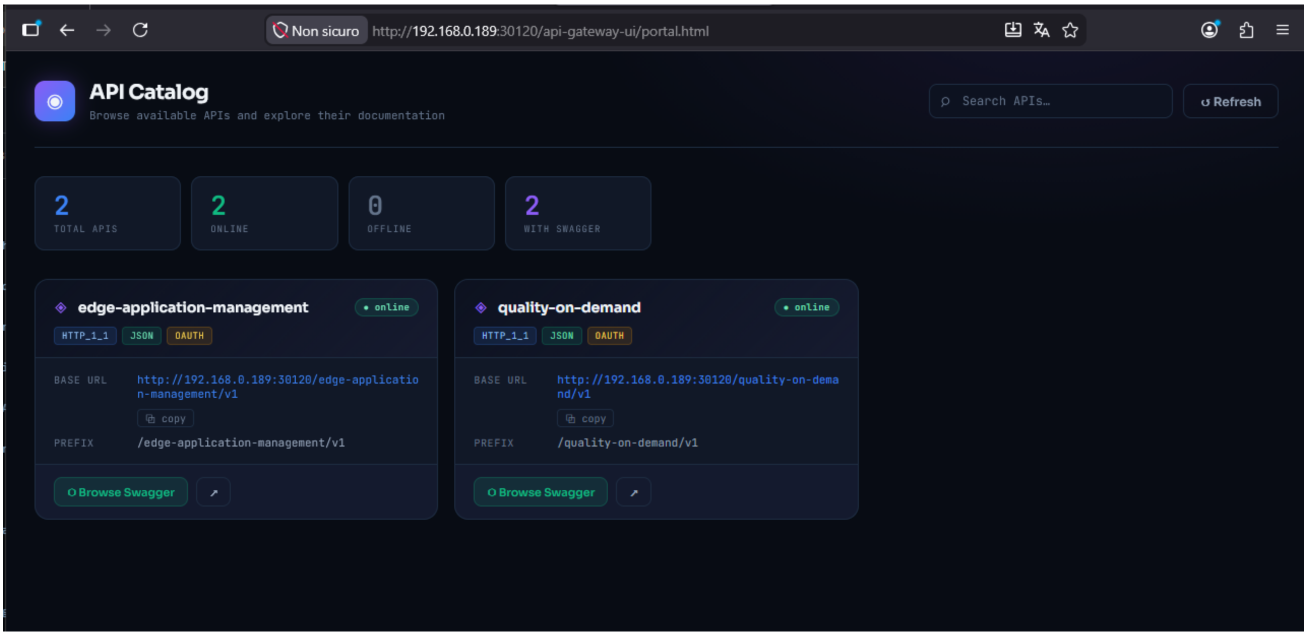

Figure 34: API Gateway GUI of the Exposure Framework

4.1.1.5 Edge computing platforms

A monitoring platform has been developed to collect system metrics from the RSU and from the AMR, which are both based on a NVIDIA Jetson Xavier NX board. The collected metrics include current and average power consumption, temperature, fan speed, GPU load, CPU usage per core, disk availability, and RAM usage. Data collection is performed using jetson-stats, an open-source package for monitoring and controlling NVIDIA Jetson devices, which exposes system metrics through a Python library called jtop.

The collected data is periodically sent (e.g., every 10 minutes) to a local PostgreSQL database via a RESTful API built with FastAPI. A cron job running on the Jetson device triggers the data collection script at each interval, which packages the metrics into a JSON payload and sends it to the API endpoint via an HTTP POST request. The API validates the payload, stores it in the database, and exposes the data through dedicated GET endpoints, allowing users to query metrics by category (RAM, disk, CPU, GPU, temperature, fan, power).

4.1.2 T1 and T2 Preliminary test results

Demonstrate the capabilities of the technology enablers from WP3 and the site capabilities. This task will allow obtaining evaluation data, which will help with the better planning and execution of large-scale trials and pilots that will take place in the context of WP7 and will support ongoing standardisation activities

4.1.2.1 Lab testing of mmWave RAN

Before moving to real deployment, the RAN mmWave architecture has been validated in a laboratory environment.

The objective of the laboratory validation was to verify:

- interworking between components
- correct network configuration
- overall traffic and performance behaviour
- preliminary system stability under controlled test conditions

Initial tests were carried out on the 5G mmWave Standalone (SA) laboratory network, focusing on:

- End-to-end latency
- TCP throughput (uplink and downlink)
- UDP performance (bandwidth, jitter, packet loss)
- Overall system stability

Traffic tests were performed using iperf3, with:

- 15 seconds duration
- both single and dual TCP streams
- and a UDP target rate of 50 Mbit per second

This setup allowed us to evaluate both throughput and stability under controlled conditions, as shown in the table below.

Table 38: Detailed first lab results

Test Category	Direction	Mode	Throughput	Latency / Jitter	Packet Loss / Retransmissions	Status
Ping	CPE → Server	ICMP	—	**14.6 ms (avg RTT)**	**0% loss**	Stable
TCP	Upload	1 Stream	**55.9 Mbit/s**	—	89 retransmissions	OK
TCP	Download	1 Stream	**76.3 Mbit/s**	—	178 retransmissions	OK
TCP	Upload	2 Streams	**56.5 Mbit/s**	—	85 retransmissions	OK
TCP	Download	2 Streams	**89.8 Mbit/s**	—	630 retransmissions	OK
UDP	Upload	50 Mbit/s	**50.0 Mbit/s**	**0.200 ms jitter**	**0% loss**	Excellent
UDP	Download	50 Mbit/s	**50.0 Mbit/s**	**0.258 ms jitter**	**0.007% loss**	Excellent

The network showed stable and responsive behaviour across all preliminary tests. The average latency was approximately **14.6 ms**, confirming proper end-to-end integration between the RAN and the Core Network over the VPN-based transport.

TCP throughput reached approximately:

- **89.8 Mbit/s in downlink**, with two parallel streams
- **56.5 Mbit/s in uplink**, with two parallel streams

UDP tests achieved the configured target rate of **50 Mbit/s**, with very low jitter and negligible packet loss.

Key highlights:

- Low and stable latency
- negligible packet loss
- stable performance over time
- expected asymmetry between download and upload
- improved downlink throughput when using parallel TCP streams

- presence of TCP retransmissions, likely related to radio conditions and/or transport tuning

Overall, the performance is consistent and reliable, as shown in Table 39.

Table 39: First lab test results

KPI	Value	Remarks
Average Latency	**14.6 ms**	Low and stable
Max TCP Uplink	**56.5 Mbit/s**	Consistent across streams
Max TCP Downlink	**89.8 Mbit/s**	Improved with parallel streams
UDP Throughput	**50 Mbit/s**	Fully achieved target rate
Packet Loss	**~0%**	Negligible in all tests
Jitter	**< 0.3 ms**	Very stable radio performance

The latency trend shown in Figure 35 confirms low and stable RTT behaviour. The average RTT is 14.6 ms, with most values converging within a relatively narrow deviation band after an initial transient phase.

The first measured samples show higher RTT values, around 18.4 ms and 17.2 ms, followed by a progressive stabilisation. From the third sample onwards, the RTT remains mostly around the average value or below it, reaching a minimum value of approximately 10.9 ms at the end of the test sequence.

This behaviour indicates that, after the initial transient phase, the system reaches a stable operating condition. The measured latency is compatible with near-real-time applications under controlled laboratory conditions. However, these results should be considered preliminary and will require further confirmation during larger-scale trials and field validation.

The TCP throughput results shown in Figure 36 highlight a clear asymmetry between uplink and downlink performance.

In uplink, the throughput remains stable when moving from one TCP stream to two parallel TCP streams, increasing only slightly from 55.9 Mbit/s to 56.5 Mbit/s. This indicates good uplink stability, although no significant gain is achieved by adding a second stream.

In downlink, the use of two parallel TCP streams improves throughput from 76.3 Mbit/s to 89.8 Mbit/s. This suggests that parallel streams help to better exploit the available radio and transport resources, likely due to improved scheduling, buffering and TCP behaviour.

Nevertheless, the measured TCP throughput remains below the theoretical expectations for a mmWave system. Further tuning and optimisation activities are still required. The next steps will therefore focus on tuning and optimisation activities, including radio alignment, link adaptation, scheduler configuration, TCP parameters, VPN transport performance and end-to-end network configuration.

UDP tests demonstrate:

- full achievement of the configured 50 Mbit/s target rate
- negligible packet loss
- very low jitter, below 0.3 ms
- stable behaviour suitable for preliminary real-time traffic validation

The presence of TCP retransmissions is not unexpected in mmWave environments, due to the high sensitivity of the radio link to propagation conditions, alignment, blockage and link adaptation mechanisms.

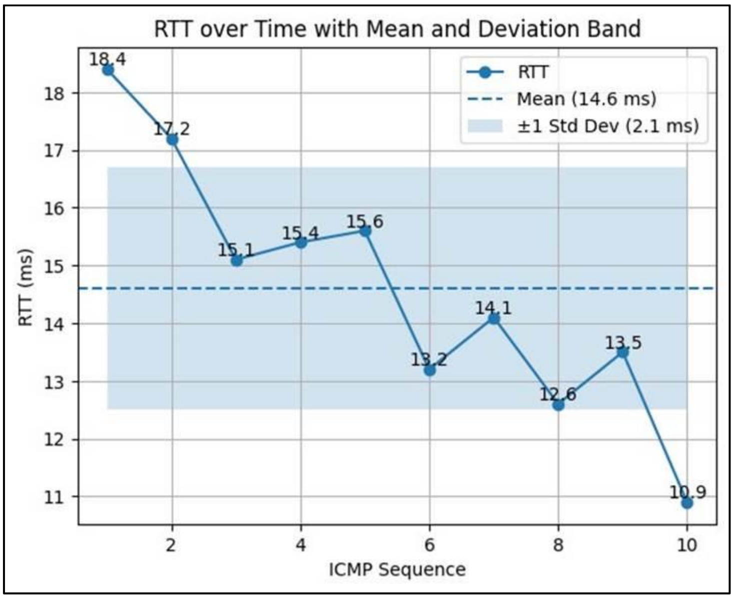

Figure 35: Lab Test Latency

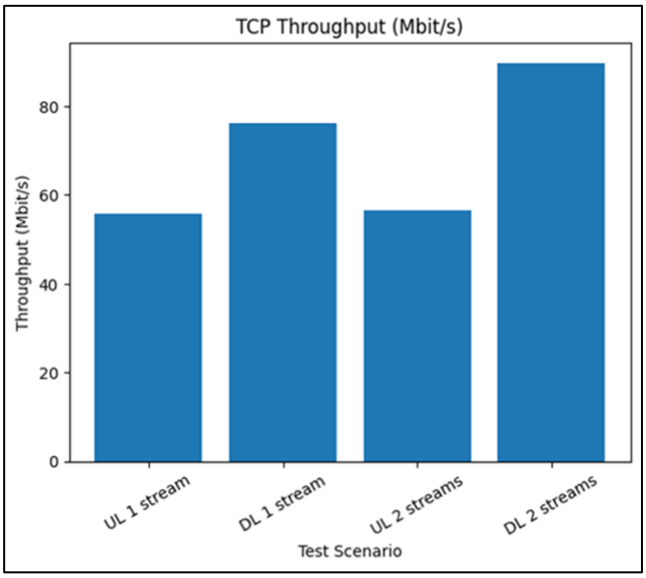

Figure 36: Lab Test Throughput

After the initial transient phase, the system stabilises quickly. This provides a positive preliminary indication for the target use cases, including:

- T1, focused on assisting individuals with disabilities during urban mobility
- T2, aimed at helping municipalities monitor accessibility and walkability

Final suitability for T1 and T2 should be confirmed through further testing in representative field conditions.

4.1.2.2 Energy-aware policy simulations

A system simulation has been realised to evaluate the energy sustainability of the target system, and later on, to feed the Digital Twin. In particular, a high-granularity simulator has been developed in the Python language. It has been calibrated using measurements from a real prototype, and the system has been evaluated under realistic seasonal solar patterns and pedestrian mobility profiles.

Figure 37 reports the high-level architecture of the simulated system, including the RSU, with battery, solar panel, LiDAR, camera, processing power, network connectivity, the edge µDataCenter, the object detection and tracking task, migrating between the RSU and the edge, the user with a smartphone, and the app backend generating and transmitting alerts.

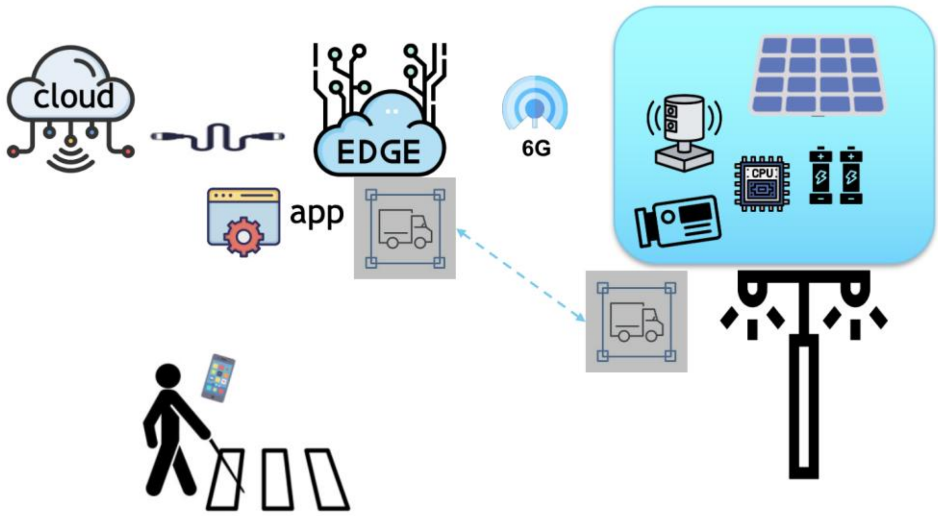

Figure 37: High-level system architecture of the simulated system.

The core logic is governed by a finite state machine (FSM) that manages the computational location of the detection task (YOLO). The system monitors the RSU's battery charge (B_{curr}) and triggers migrations based on two critical thresholds: B_{low} (when $B_{curr} < B_{low}$, the detection task is offloaded to the Edge server to preserve the remaining battery for essential sensing), and B_{high} (when $B_{curr} > B_{high}$, the task migrates back to the RSU to minimise grid energy consumption).

Each migration event incorporates a realistic transition period ($T_{mig} = 60s$), during which both the RSU and Edge run the detection task simultaneously to ensure service continuity, albeit at the cost of higher temporary power consumption.

The simulator implements a high-granularity energy model where the RSU's battery drain is calculated at each time step ($d_t = 0.1s$) as the sum of several components, namely System Overhead and Idle Power, Computational Load, Sensing Load, and 6G Transmission.

Solar recharge is modelled using time-based schedules for summer and winter seasons, with power generation sampled from a normal distribution to account for weather-related variability.

User arrivals and departures are modelled as independent stochastic processes following hourly rates defined for the Italian cluster. The simulator tracks the end-to-end delay for every detection cycle.

A prototype has been developed to test the solution functionalities. It is based on the Nvidia Xavier AGX module, equipped with a Livox Horizon LiDAR, a Vivotek ib9387-ht-a camera, a SolidRun HummingBoard module, a short-range Mobilmark antenna, a LiFePO4 12V 100Ah battery and a 1.65 × 0.95 m solar panel, providing a 300W peak.

On the developed prototype, measurements have been performed to gather the values and models for the different parameters to be used as input for the simulator. In particular, Table 40 reports the different values used as input for the simulation processes, where “load” refers to the state in which the object detection and tracking tasks run in the RSU, while “idle” refers to the state in which they run in the edge.

Table 40: Values measured on the system prototype and used as input for the simulation process.

Parameter	Value	Model
Xavier AGX Power consumption (load)	20W	Constant
Xavier AGX Power consumption (idle)	5W	Constant
Camera	4W	Constant
LiDAR	12W	Constant
RSU overhead	10W	Constant
Solar panel efficiency	80%	Constant
Antenna power consumption	11.5W	Normal distribution σ = 3.5
Object detection delay (RSU) (T_{det})	40ms	Normal distribution σ = 5
Object detection delay (Edge) (T_{det})	30ms	Normal distribution σ = 5
Delay for Video upload to Edge (T_{upld})	60ms	Normal distribution σ = 10
Backend processing ($T_{backend}$)	30ms	Normal distr. σ = 5
Alarm notification to user (T_{tx_alarm})	20ms	Normal distribution σ = 3

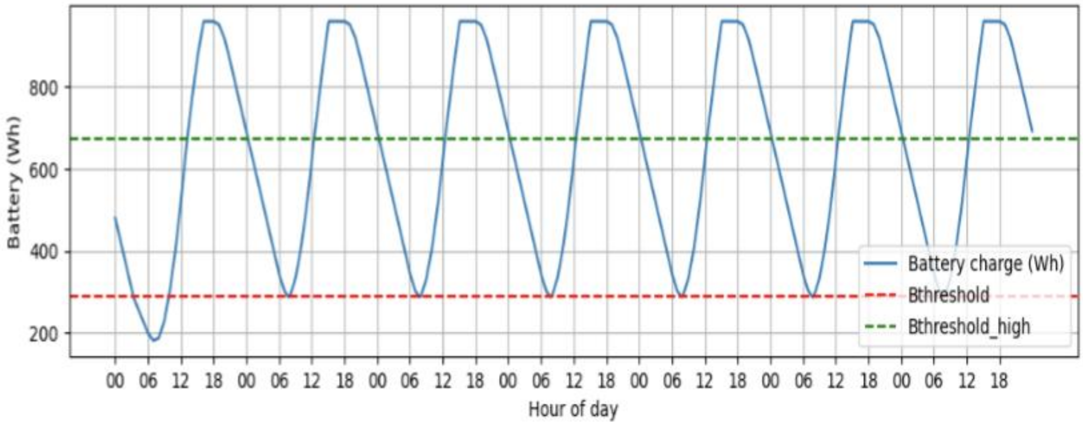

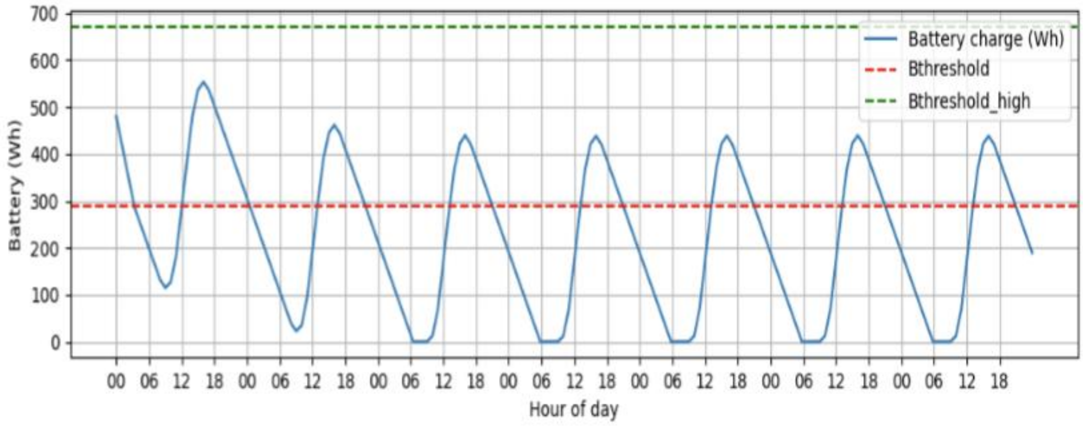

Figure 38: RSU battery charge over time, simulating the summer behaviour (Top), and the winter behaviour (Bottom).

When simulating the system over a realistic environment (i.e., user arrival/departure and Sun irradiation), the simulator allows for evaluating the battery charge. Its evolution over time is reported in Figure 38 for Summer (Top) and Winter (Bottom). As can be seen, now the system shows a daily periodicity as a consequence of the sun (and charge) periodicity and of the user arrival/departure periodicity. During the summer, the system is self-sufficient: during the last hours of the night, the battery reaches the lower threshold, and the object detection task is migrated to the edge, but the battery recharges, eventually reaching the higher threshold around noon and triggering the migration of the object detection task back to the RSU. On the other hand, during the winter, the system is not self-sufficient: after the convergence is reached, i.e., during the third day starting with a battery charged at 50%, there is a period in which the battery is completely drained to 0% charge, even after migrating the object detection task to the edge (at the very beginning, as the high threshold is never reached), lasting a few hours every day (from around 6 AM to around 9 AM, 9.2% over the whole simulated time). During this interval, the RSU cannot operate either the basic tasks or the system cannot work.

To make the system also viable in the harshest winter conditions, different energy-saving strategies have been evaluated and simulated. The best one, which was also viable accounting for the HW considered to implement the system, was to switch off the LiDAR when no users are detected in the area during a given amount of time (i.e., 30s).

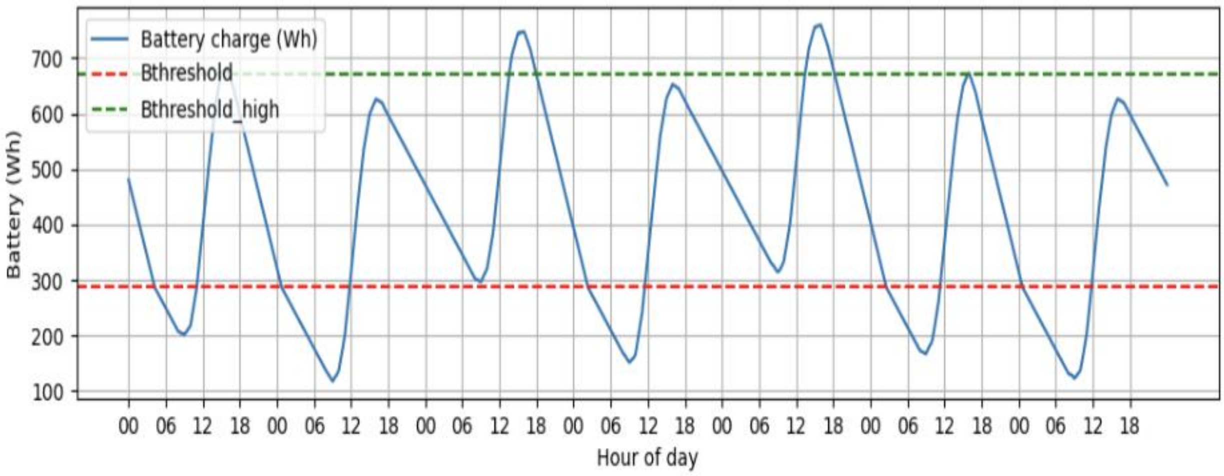

Figure 39: Applying the energy-saving strategy in the winter setting. RSU battery charge over time.

Figure 39 reports the system behaviour when applying the designed energy-saving strategy during the Winter. In particular, the plot reports the battery charge over time, and it is possible to see how, with this strategy, the system periodically migrates the object detection task between the RSU and the Edge, but never falls short of power. In detail, this strategy allows to save 11.1% of battery energy, 39.1% of grid energy, and 15.0% energy overall, making the difference between a non-viable and a viable solution in the worst scenario (i.e., Winter). Even in this case, the amount of time the solution results in a screen-to-screen delay higher than the constraint (i.e., > 200ms) is 1.1×10^{-6}, which is acceptable for the system implementation.

4.1.2.3 Energy monitoring

The energy monitoring system was tested in several scenarios to verify that the energy consumption is effectively measured. Multiple tests were conducted to analyse the performance of the NVIDIA Jetson Xavier NX board when the NVIDIA Power mode changes. The energy monitoring platform has been used to collect hundreds of samples in 1/2 hours with a sampling rate of 10 seconds. Moreover, a detection algorithm has been executed every 10 minutes to stress the system.

4.1.2.3.1 20W 2 cores

The figures present four time-series plots recorded over approximately one hour (16:30–17:30), capturing the dynamic behaviour of the NVIDIA Jetson Xavier NX board under both idle and stressed operating conditions.

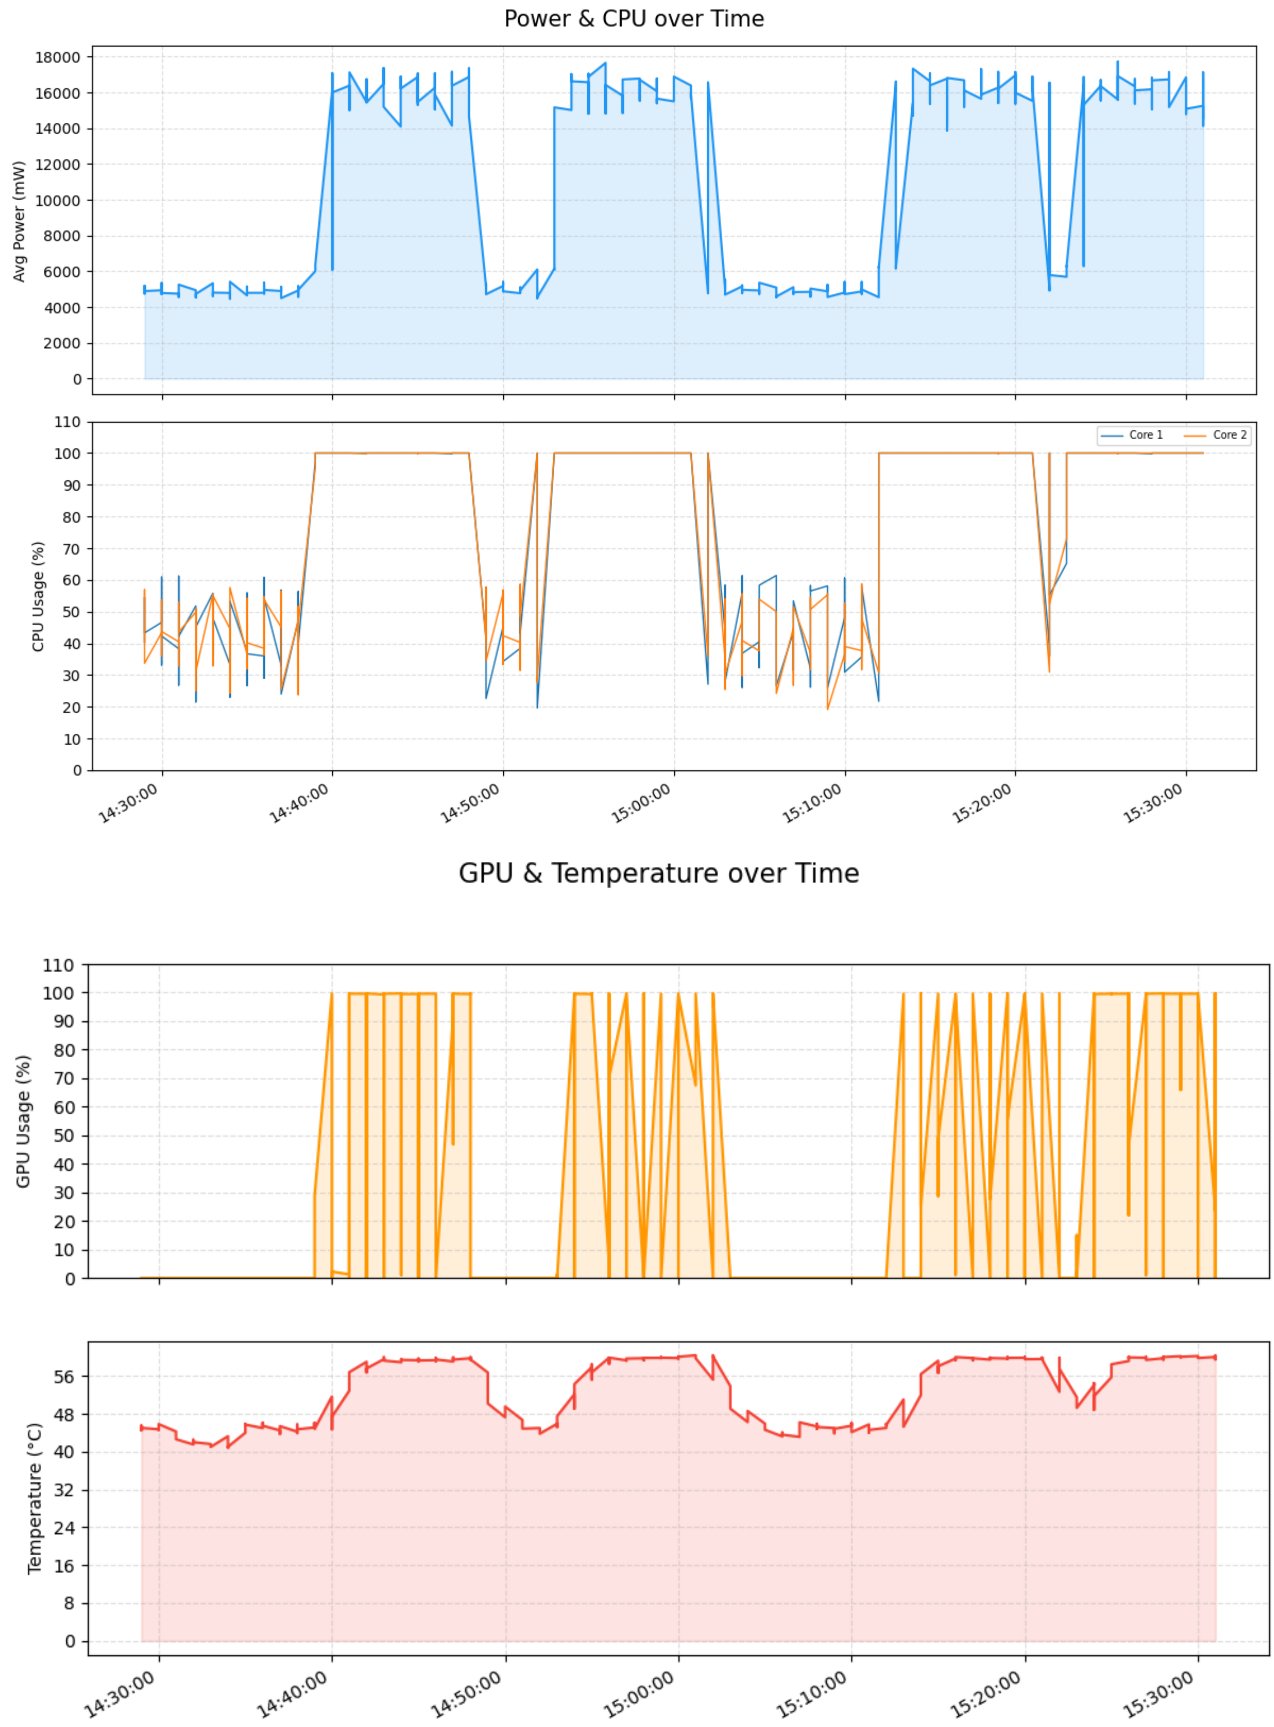

Figure 40: Monitoring results for 20 W and 2 cores mode

Top plot – Average Power Consumption (mW): The power consumption exhibits a clear bimodal pattern, alternating between two distinct regimes. During idle or standard operation, the board draws a baseline power of approximately 5,000 mW, with moderate short-term fluctuations visible as spikes in the 4,000–7,000 mW range. Under stress conditions, power consumption rises sharply to approximately 14,000–17,000 mW, approaching the 20W power limit, with noticeably high-frequency oscillations at the peak suggesting active power management or throttling behaviour. Three distinct stress episodes are visible across the recording window, each followed by a rapid return to baseline.

Second plot – CPU Usage (%): The CPU utilisation closely mirrors the power consumption pattern. Core 1 (blue) and Core 2 (orange) both saturate at 100% during stress episodes, confirming full CPU load. During idle periods, both cores fluctuate in the 25–60% range, with frequent short bursts indicating intermittent background activity. Notably, Core 2 consistently shows slightly higher utilisation than Core 1 during transitional and idle phases, suggesting a minor load imbalance in the task scheduling.

Third plot – GPU Usage (%): The GPU usage displays a highly dynamic, pulsed pattern throughout the entire recording, in contrast to the smoother CPU behaviour observed previously. During stress episodes, the GPU rapidly alternates between 0% and 100% utilization in rapid succession, producing the dense vertical oscillation pattern visible across the plot. This suggests the GPU is being driven in a burst or duty-cycle fashion rather than sustaining a steady load. During idle periods, GPU usage drops to near 0%, with only occasional brief spikes, indicating the GPU is largely inactive outside of stress events.

Bottom plot – Temperature (°C): The thermal response is notably smoother and more gradual than the GPU usage, reflecting the thermal inertia of the hardware. At baseline, the board maintains a temperature of approximately 48–50°C. During stress episodes, temperature rises to a plateau of around 55–57°C, with the increase lagging slightly behind the onset of GPU load — consistent with expected thermal dynamics. Between stress episodes, the temperature partially recovers but does not fully return to the initial baseline within the observation window, suggesting cumulative heat buildup over the session. No thermal throttling threshold appears to be reached, as temperature remains stable below 60°C throughout.

4.1.2.3.2 20W 6 cores

The figures present four time-series plots recorded over approximately one hour (17:40–18:40), characterising the NVIDIA Jetson Xavier NX board's behaviour under the same 20W power cap but now with 6 cores enabled.

Figure 41: Monitoring results for 20 W and 6 cores mode

Top plot – Average Power Consumption (mW): The power consumption pattern remains consistent with the previous configurations, alternating between a stable idle baseline of approximately 5,000–6,000 mW and stress peaks of 15,000–17,500 mW, again approaching the 20W cap. Four stress episodes are visible, each with the characteristic high-frequency oscillations at peak load. The transitions between idle and stressed states remain sharp, indicating rapid load response regardless of core count. No meaningful increase in baseline power is observed compared to the 4-core case, suggesting diminishing overhead per additional core at idle.

Second plot – CPU Usage (%): With 6 cores active (Core 1–6), only Core 1 (blue) and Core 2 (orange) saturate at 100% during stress episodes, while Cores 3, 4, 5, and 6 remain in their typical idle range of approximately 10–35% throughout the entire recording. This indicates that the stress workload is not parallelised across all available cores but instead confined to two threads — consistent with what was observed in the 2-core configuration. The remaining four cores appear to handle only background system tasks. This suggests that the application under test is either single-threaded or limited to two execution threads, and that enabling additional cores beyond 2 provides no direct benefit for this specific workload in terms of CPU-level parallelism.

Third plot – GPU Usage (%): GPU behaviour remains consistent with previous configurations — rapid burst-like alternation between 0% and 100% during stress episodes, and near-zero activity during idle periods. Notably, the GPU stress episodes appear slightly more confined in time compared to the 4-core case, though this is likely attributable to differences in the test schedule rather than any effect of core count. The GPU utilisation pattern continues to appear independent of CPU topology.

Bottom plot – Temperature (°C): The thermal profile closely resembles the 4-core case. Baseline temperature sits at approximately 47–50°C, rising to peaks of 55–62°C during stress. The temperature tracks the stress episodes with the expected thermal lag, and partial recovery is visible between events. The maximum observed temperature remains consistent with the 4-core configuration, suggesting that enabling additional cores does not significantly worsen the thermal envelope under a fixed 20W power limit — likely because the power budget constraint limits total heat dissipation regardless of core count.

4.1.2.3.3 10W 4 cores

The figures present four time-series plots recorded over approximately one hour and ten minutes (09:40–10:50), characterising the NVIDIA Jetson Xavier NX board's behaviour under a further reduced 10W power limit with 4 cores enabled.

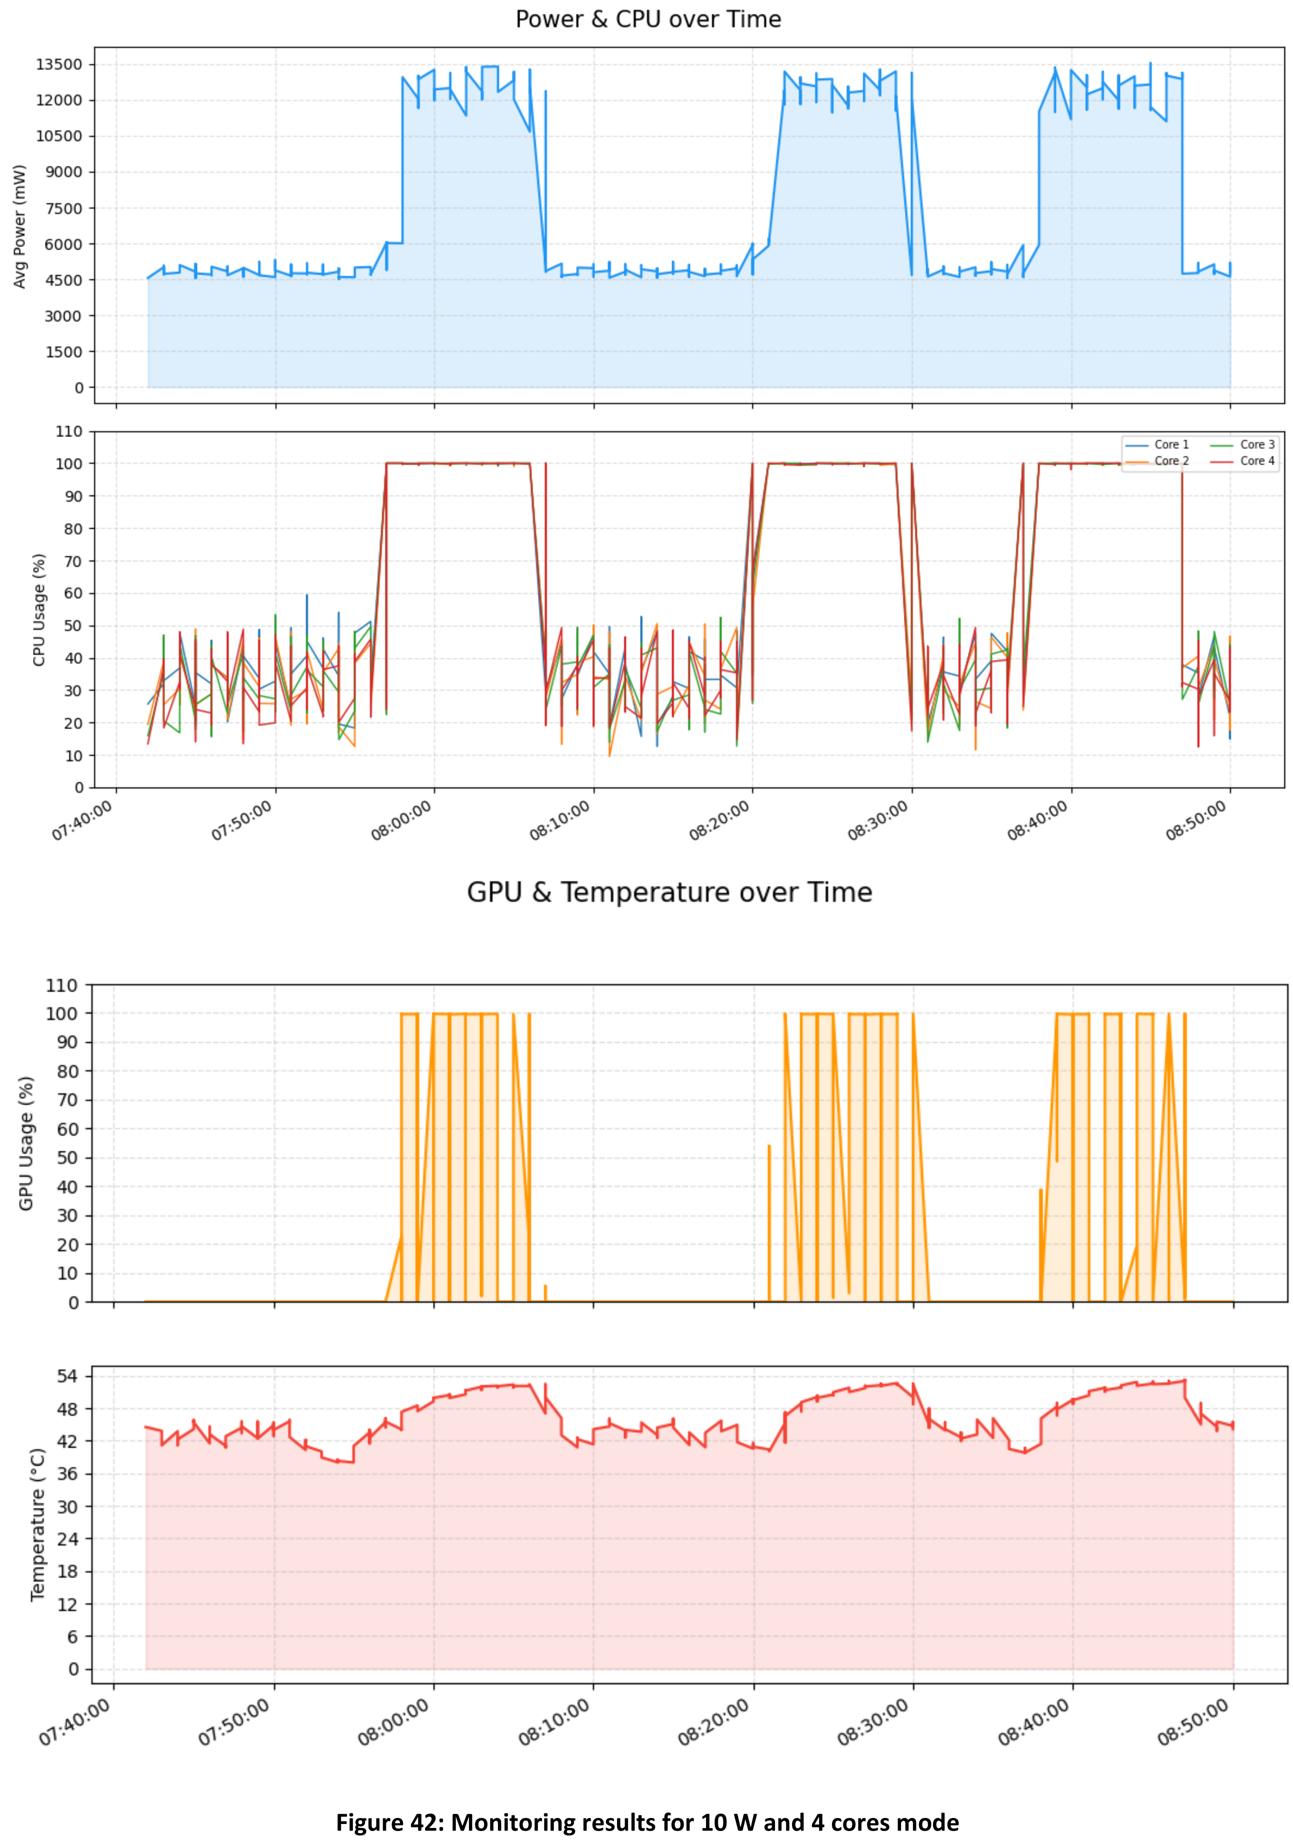

Figure 42: Monitoring results for 10 W and 4 cores mode

Top plot – Average Power Consumption (mW): The bimodal power pattern persists, but with notably reduced absolute values compared to the 15W and 20W configurations. At idle, the RSU draws approximately 4,500–5,500 mW, consistent with previous baselines. Under stress, peak power is limited at approximately 11,500–12,500 mW — a significant reduction from the ~15,000 mW and ~17,500 mW peaks observed at 15W and 20W, respectively — confirming effective enforcement of the 10W limit. The characteristic high-frequency oscillations at peak load remain visible, indicating active power throttling near the budget ceiling. Three stress episodes are observed across the recording window.

Second plot – CPU Usage (%): All four cores (Core 1–4) saturate uniformly at 100% during stress episodes, consistent with the 15W 4-core configuration and confirming that the workload fully utilises all available threads regardless of the power limit. During idle periods, all cores fluctuate in the 15–50% range, with occasional short bursts reaching higher values, and no significant imbalance between cores. The idle per-core utilisation appears slightly more variable than in previous configurations, possibly reflecting increased scheduling activity under tighter power constraints.

Third plot – GPU Usage (%): The GPU maintains its characteristic burst-like alternation between 0% and 100% during stress episodes, fully consistent with all previously tested configurations. Idle GPU usage remains near 0%. The robustness of this pattern across 10W, 15W, and 20W limits confirms that GPU utilisation behaviour is entirely independent of the power budget.

Bottom plot – Temperature (°C): The thermal profile shows a notably lower baseline compared to higher power limit configurations, sitting at approximately 42–44°C at idle. During stress episodes, temperature rises gradually to peaks of approximately 48–52°C — the lowest stress-peak temperatures observed across all configurations tested — directly reflecting the reduced power dissipation enforced by the 10W limit. Temperature recovery between stress episodes is smooth and complete, with the system returning close to its idle baseline before each subsequent stress event, suggesting more effective thermal management at this lower power budget.

4.1.2.3.4 Conclusions

In conclusion, idle power remains consistent at around 4,500–6,000 mW using 20W or 10W modes, confirming that the power limit only affects the stressed regime. Peak power scales proportionally with the limit, from 15,000–17,500 mW at 20W down to 11,500–12,500 mW at 10W. CPU utilisation is unaffected: all four cores saturate at 100% under stress regardless of the power limit, confirming that workload parallelism is fully preserved across all configurations. GPU burst behaviour is likewise invariant. Temperature is the metric that benefits most visibly from a lower power limit, with peak values dropping from 58–62°C at 20W to 48–52°C at 10W, a reduction of approximately 10°C, at the cost of a lower available power budget.

Power consumption is remarkably stable using 2, 4 or 6 cores, suggesting that enabling additional cores contributes negligible overhead at the system level. CPU utilisation shows the most notable difference: the 2-core and 4-core configurations fully utilise all available cores under stress, while the 6-core configuration stresses only Cores 1 and 2, indicating that the workload does not scale beyond 2 active threads in that setting, possibly due to scheduler or affinity behaviour. Idle per-core utilisation decreases progressively with core count, from 25–60% with 2 cores to 10–35% with 6 cores, as background tasks are distributed more broadly. GPU behaviour and temperature remain consistent across all core counts, confirming their independence from CPU topology.

Overall, the 4-core configuration offers the best balance: it fully exploits available parallelism without the scheduling anomaly observed at 6 cores, and its thermal profile remains acceptable across all tested power limits.

4.2 Transport 3 Use Case

4.2.1 T3 Integration status

The current integration status of the T3 work on LCX-based sensing can be summarised as follows. An FR1 deployment has been set up in the TUC wireless laboratory using a 5G NR signal with a carrier frequency of 3.75 GHz. At this stage, the sensing chain relies on standard antenna-based transmission and reception, enabling over-the-air communication under controlled laboratory conditions. The received signal is captured using a Rohde & Schwarz FSW instrument, which allows the recording of complex IQ samples for subsequent signal analysis. On top of these measurements, a Short-Time Fourier Transform (STFT) processing stage has been integrated in order to generate micro-Doppler spectrograms. The resulting framework already supports the detection and differentiation of moving objects, such as humans and animals, demonstrating the feasibility of the sensing approach in the current FR1 laboratory setup.

In the context of FR2 sensing, end-to-end mmWave measurements were carried out inside the TUC wireless laboratory using a 5G NR processing chain with OFDM demodulation and DMRS-based channel estimation. The setup enables performance evaluation and channel analysis in time, frequency, and delay domains under controlled indoor conditions. The experiments cover downlink and uplink transmission as well as channel behaviour under LoS, metallic obstruction, and human blockage scenarios, providing a reproducible baseline for further ISAC-related evaluations. The mmWave platform combines RF front-end, IQ data acquisition, and MATLAB-based baseband processing in a unified workflow. Resource grid reconstruction (1584 subcarriers, slot-based structure) and DMRS-based channel estimation are fully operational. Data processing is currently performed offline, with further steps focusing on improving automation and execution efficiency.

The equipment for the train-to-train communication, as described in section 3.3.3.2, is fully integrated, tested and operational in the TUC lab.

The current integration status of the T3 work on Liquid RAN can be summarised as follows. The setup has already been deployed in the ISRD laboratory for pre-trial testing and is being used to demonstrate the LiquidRAN™ solution together with network monitoring, energy-efficiency functions, and slicing support with dynamic PRB allocation. The current system includes the LiquidRAN™ software stack, a Benetel radio unit, a GNSS antenna, an Open5GS core network, and two smartphone-based UEs. In the final trial configuration, the first UE will be placed on the train to support train-critical communications, such as telemetry offloading, while the second UE will be located on the ground to handle non-critical traffic related to staff and passenger communications.

4.2.2 T3 Preliminary test results

4.2.2.1 Preliminary channel measurements for FR1 and FR2 ISAC

In this sub-section, the results obtained from processing measured uplink IQ data using a 5G NR signal processing pipeline are presented. In addition to the signal processing analysis, a performance evaluation of the measurement setup has been carried out.

Measurement Setup

The measurement setup consists of three main components: a RAN server, a CellBox Air radio unit (RU), and an M202 CPE device acting as user equipment (UE). These components are interconnected to enable end-to-end uplink transmission and channel measurements.

The radio unit is connected to the RAN server, which performs baseband processing and controls the communication. The CPE generates uplink traffic toward the radio unit, enabling realistic measurement conditions.

The Figures below illustrate the physical setup and the main hardware components used in the experiment.

a) Overall setup view (CellBox Air radio unit and RAN server)

b) CPE (user equipment)

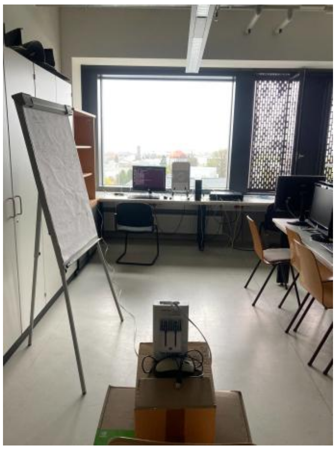

c) LoS scenario (no obstruction)

d) Metallic plane obstruction using whiteboard

Figure 43: T3 physical setup

Preliminary Results

The downlink measurement shows an aggregated throughput of approximately 1.04 Gbit/s (two parallel streams of ~520 Mbit/s each), with stable transfer rates over time and no visible packet loss.

In the uplink, a sustained data rate of about 250 Mbit/s is achieved over more than 30 seconds, with jitter close to 0 ms and 0% packet loss, indicating reliable transmission under laboratory conditions.

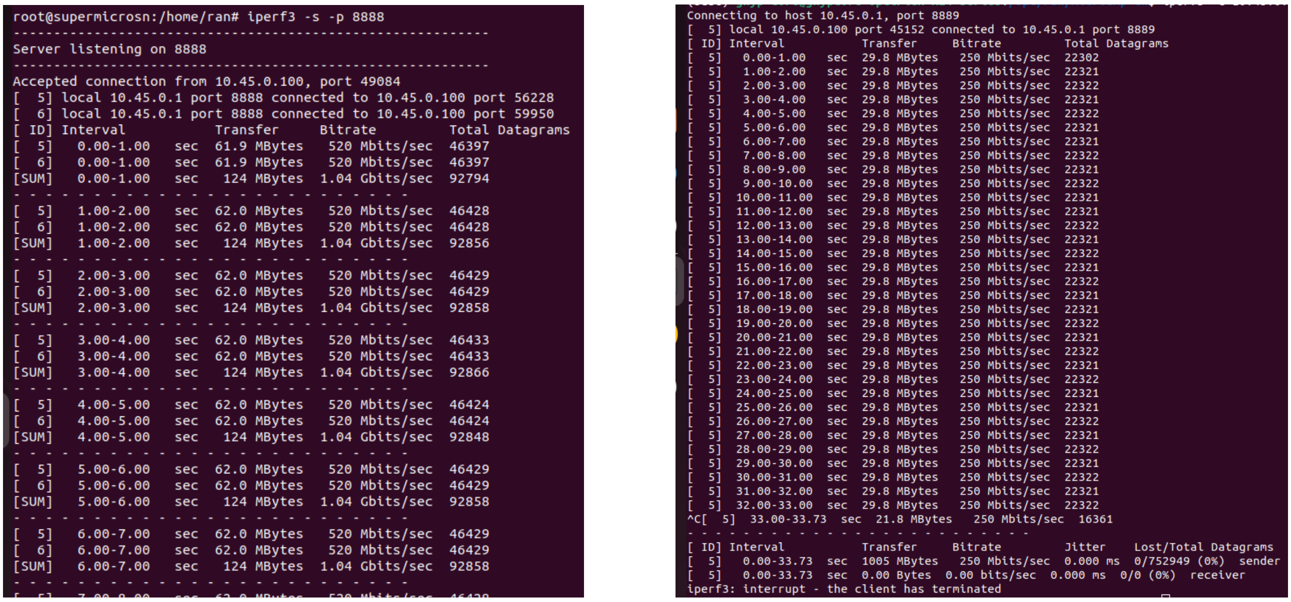

a) Data-rate performance of downlink **b) Data-rate performance of uplink**

Figure 44: Downlink and uplink data-rate performance

The 2D channel magnitude plots (1584 subcarriers × 14 OFDM symbols) reveal frequency-selective fading across the full bandwidth. Power variations along the frequency axis are on the order of 10–20 dB, while temporal variations within a slot remain moderate.

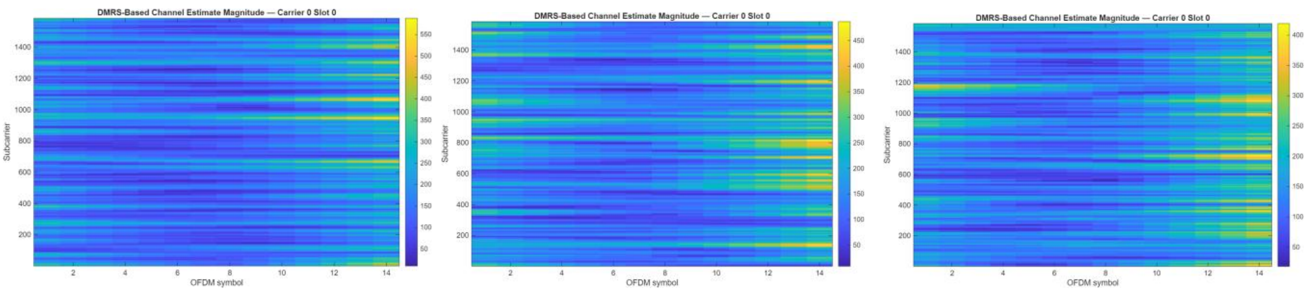

a) Line of Sight **b) Thin metallic plane** **c) Human**

Figure 45: Channel magnitude

A noticeable increase in magnitude towards later OFDM symbols (symbols ~10–14) appears in all measurements, suggesting a systematic effect (e.g., front-end gain drift or estimation bias) rather than propagation changes.

Comparing scenarios, metallic obstruction leads to stronger local peaks and sharper fluctuations, while human blockage results in a reduction of overall magnitude by several dB. Despite these effects, the dominant propagation component remains clearly visible in all cases.

Future work will move toward object-detection-oriented ISAC by incorporating beam index–aware and angle-dependent channel analysis, enabling spatial discrimination of targets across azimuth and elevation. This includes mapping channel responses per beam to extract angular signatures and extending the current grid-based analysis with high-resolution parameter estimation. In particular, super-resolution techniques based on CIR eigen-decomposition (e.g., subspace methods such as MUSIC/ESPRIT) will be applied to resolve closely spaced paths in delay–angle domains, improving target separability beyond FFT limits. These steps will allow transitioning from bulk channel characterisation to explicit extraction of object-related features such as direction, reflection strength, and potentially motion.

In addition to the mmWave platform, we also focus on FR1-based sensing using 5G NR signals. An experiment was carried out in the TUC wireless laboratory environment using a 5G NR signal transmitted at a carrier frequency of 3.75 GHz using an antenna which lies in the FR1 band. The received signal was captured using a Rohde & Schwarz FSW, enabling the recording of complex IQ samples. These samples were then processed in MATLAB, where a Short-Time Fourier Transform (STFT) was applied to generate the micro-Doppler spectrogram. This processing chain establishes a complete pipeline from signal transmission to motion signature extraction.

The experimental results demonstrate that a micro-Doppler spectrogram can be successfully obtained using the described setup. When a person moves in front of the antenna, the spectrogram exhibits periodic Doppler components and a symmetric frequency spread, confirming the detection of human motion as presented in the figure below.

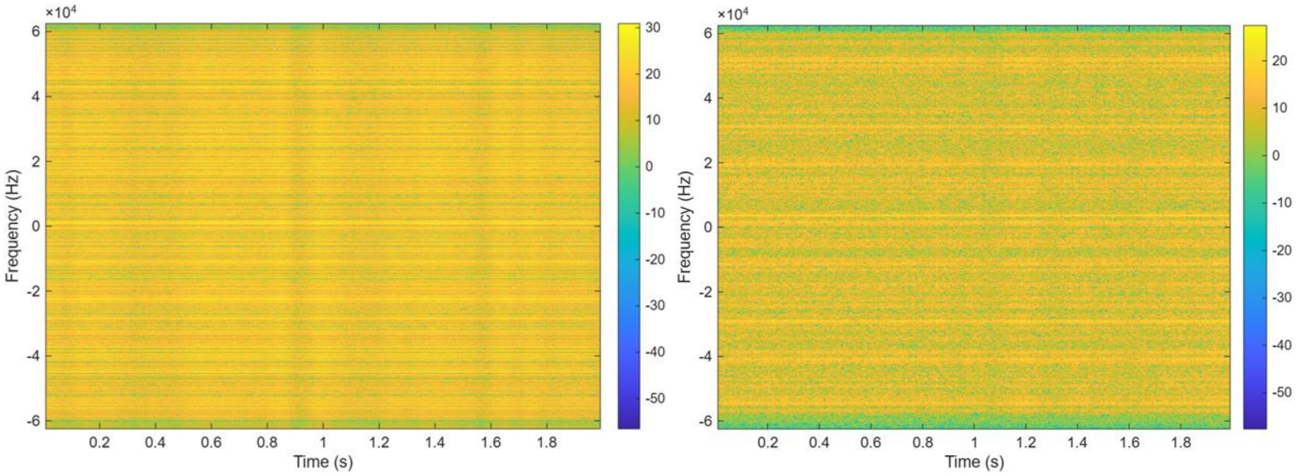

a) Human motion **b) No motion**

Figure 46: Spectrogram plots

Several limitations in the experimental setup were noted: the sensitivity to small or weakly reflecting objects is limited, and the quality of the micro-Doppler signature strongly depends on parameter selection, such as sampling rate, measurement duration, and frequency resolution.

As part of the future work, the use of software-defined radio platforms such as USRP will enable real-time IQ data acquisition. Also, channel parameters using the R&S TSMA6 will allow the extraction of channel state information, including delay, power, and angular characteristics, which are helpful for object detection based on propagation effects. Another important direction involves the use of LCX as a distributed sensing medium, providing continuous spatial coverage in environments such as tunnels and corridors. By combining Doppler-based motion features with spatial channel information, it becomes possible to detect the objects and their movement.

4.2.2.2 Liquid RAN

The following activities have been completed by ISRD. The new enablers have been developed in the form of xApps. Specifically, the Energy Efficiency for Train Scenario has been developed (Energy Monitoring enabler) and Slicing with PRB Quota for Train Scenario has been developed (Slicing and Network Programmability enablers). Figure 47 shows the LiquidRAN™ RIC environment with the train timetable, which is used by the xApps for scheduling the respective functionality, i.e. changing the PRB quota of the slices or decreasing the energy usage of the servers.

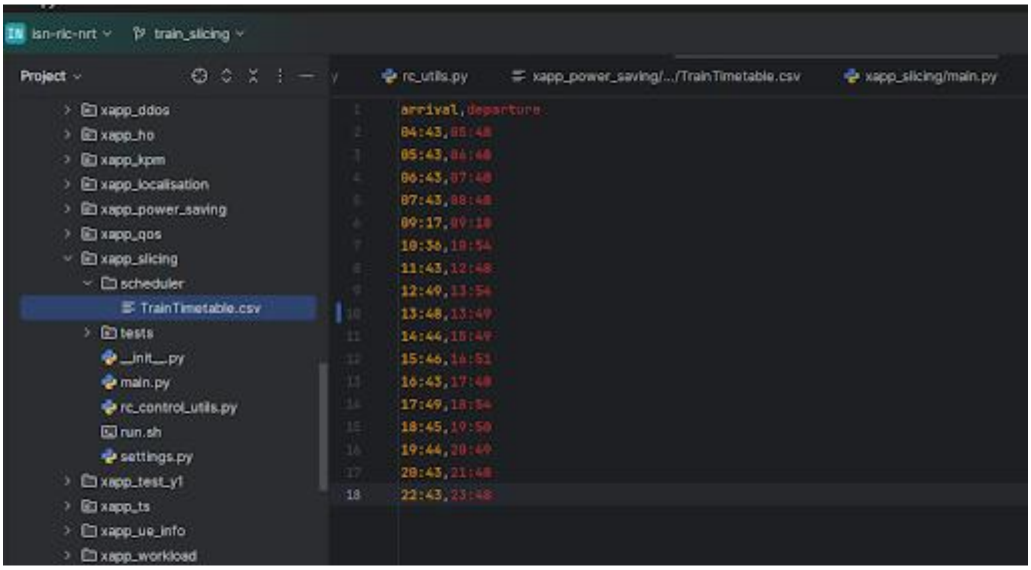

Figure 47: LiquidRAN™ RIC environment showing the train timetable used in Energy Efficiency and Slicing with PRB Quota xApps

The setup of the final demo deployment at TUC has been agreed between both partners, and the High Level Design of the setup has been provided as shown in Figure 48. The system will demonstrate all ISRD enablers from WP3, specifically the cloud native LiquidRAN™ with Network Monitoring, Energy Efficiency and Slicing with Dynamic PRB Quota features.

The system includes the following components:

- Software: LiquidRAN™, including RIC and Network Monitoring, Energy Efficiency for Train Scenario and Slicing with Dynamic PRB Quota for Train Scenario xApps
- Hardware: Benetel RU, GNSS antenna, Core Network: Open5GS, UEs: 2 smartphones

In the final setup, the first UE will be located on the train to transmit train-critical communication (e.g. Telemetry offloading), and the second UE will be on the ground for non-critical staff and passenger communication.

Figure 48. High-Level Design of the final demo deployment at TUC premises in Annaberg Buchholz

The dedicated equipment has been procured and installed in the ISRD lab, where the system will undergo initial testing before the deployment at TUC. The next steps before the mid-term review comprise recording a demo of the ISRD setup showcasing the enablers and KPI measurements of the above setup, as well as reporting the results in the relevant deliverables.

4.3 Transport 4 Use Case

4.3.1 T4 Integration status

Automated Driving and Remote Operation integration: As described above in section 4.2.1, the TUC testbed comprises the test vehicle Carai3 (BMW i3), which is equipped with sensors and actuators for automated driving functions. The vehicle is equipped with a generic, fully accessible and customizable automation system. Moreover, the test vehicle can be fully teleoperated from an operator workstation via keyboard or gamepad using Wi-Fi or 5G mobile internet connectivity. The modular design supports bidirectional communication, including video feedback, sensor data transmission, and control commands, allowing components to be added or removed as needed. Finally, multiple experiments have taken place in order to establish the remote operation capability for parking scenarios and lane change/obstacle avoidance.

Network integration: The T4 UC trials will take place in the TUC campus, using an open5G core and an O-RAN Benetel RAN650. This setup allows for full implementation and experimentation flexibility, and access to IMEC to integrate the end-to-end network slicing and ZSM framework. The 5G core and RAN are already deployed, tested and fully functional in the TUC lab (see Figure 49, below). The 5G connectivity of the test vehicle has been evaluated in a wide range of scenarios.

Figure 49: The 5GC and RU to be used in the T4 UC trials.

Network Slicing integration: Due to internal policies and regulatory constraints across the participating partners, access to the required infrastructure at the German trial site has not yet been granted and is still in progress. As a result, full integration activities have not been completed at this stage. Nevertheless, a clear integration strategy has been defined. The approach foresees an initial deployment within a controlled testbed environment (IMEC testbeds in Belgium, used as a complementary trialling setup), enabling validation of the end-to-end T4 system, focusing mainly on the technological enablers. In this context, the network components already deployed within the IMEC infrastructure (NextG-Lab testbed) can be progressively migrated to the TUC infrastructure within the German trial site, once access becomes available. This phased approach ensures integration continuity and readiness for full integration while mitigating current access limitations.

Service orchestration integration: As previously reported for the Network Slicing integration progress, access to the required infrastructure is currently being processed, pending approval in line with partner policies and regulatory requirements. In the meantime, development and validation of service orchestration tasks have been performed in IMEC testbeds. Once access is granted, the orchestration components tested in IMEC facilities will be migrated to the target infrastructure to continue with development, aiming for full integration. The same step-wise approach is adopted in the case of orchestration as well, but for this case, the Smart Highway testbed is employed, due to the availability of distributed edge computing nodes that can be used for intelligent service placement.

4.3.2 T4 Preliminary test results

In this section, we briefly show the progress made with the two technological enablers: Network slicing and Zero-touch network and service management. The preliminary results are part of pre-trial activities, executed within the IMEC testbed environment, which serve as preparation for the full trials that will be performed on the German trial site. Moreover. We introduce the decision-making algorithm that is responsible for determining the vehicle's mode of operation.

Network deployment and network slicing aspects

A decentralised 5G SA architecture has been designed and deployed, featuring a centralised control plane and multiple UPFs, with one UPF instantiated per network slice. This design ensures strict traffic isolation between slices and supports service-specific performance guarantees. For the teleoperation use case, a dedicated UPF has been deployed at the network edge, enabling proximity to the end user and supporting URLLC requirements. This configuration enables achieving end-to-end latency targets aligned with teleoperation constraints (i.e., below 30 ms under optimal conditions), while maintaining high reliability and throughput.

As design and development progress, an orchestration component, referred to as the Cross-Domain Controller (CDC), has been implemented to enable end-to-end slice management across the RAN, TN, and 5GC. The CDC establishes a complete end-to-end network slice, spanning from the RAN to the edge-hosted application server where the teleoperation service is executed. The CDC performs coordinated resource allocation across all domains, including Radio resources (e.g., Physical Resource Blocks in the RAN), Transport capacity (e.g., backhaul bandwidth and routing paths), and Core network resources (e.g., compute and connectivity for UPFs). Furthermore, the CDC implements a closed-loop control mechanism based on continuous monitoring of slice-specific KPIs. When SLAs are not met (e.g., latency, throughput, or packet loss thresholds), the orchestrator dynamically triggers reconfiguration actions across the network to dynamically change the slice resources in runtime. Actions may include increasing radio resource allocation, rerouting traffic in the transport network, or scaling compute resources in the core network.

The experimental validation was conducted using the NextG-Lab testbed setup, mimicking a teleoperation scenario, where a vehicle generates uplink video traffic and receives control commands via a 5G SA network. The evaluation focused on assessing the effectiveness of end-to-end network slicing and dynamic resource allocation under realistic traffic conditions. The considered KPIs, aligned with teleoperation requirements, are end-to-end latency, throughput, and packet loss ratio.

Two network slices were instantiated:

- **teleoperation slice**, supporting latency-sensitive and high-throughput traffic
- **best-effort slice**, generating background traffic to emulate network congestion

The experiment starts from an intentionally under-provisioned configuration for the teleoperation slice (~8 Mbps), representing a baseline where SLA requirements are not met. This setup enables assessing the ability of the CDC to detect performance degradation and progressively restore service quality. The control loop operates at fixed intervals (30 seconds), allowing sufficient time to observe the system response after each reconfiguration.

As shown in Figure 50, the system evolves across three adaptation points ("1 Change", "2 Change", and "3 Change"), corresponding to successive CDC interventions. Following 1 Change, an initial increase in radio resources leads to partial performance improvement. At 2 Change, further adjustments, such as link switching and radio resource configurations, enhance throughput and reduce congestion effects. After "3 Change", the teleoperation slice reaches its target performance, achieving ~25 Mbps throughput, near-zero packet loss, and a latency of around 4 ms.

These results demonstrate the effectiveness of the closed-loop, cross-domain orchestration, where resource allocation is progressively refined based on observed KPIs. The stepwise adaptation confirms that periodic control actions enable controlled convergence toward SLA compliance, while maintaining strict isolation between the teleoperation and best-effort slices.

After the integration of the network slicing framework within the German trial site is complete, the same set of slices will be replicated in the TUC network to support the teleoperation service and corresponding T4 applications.

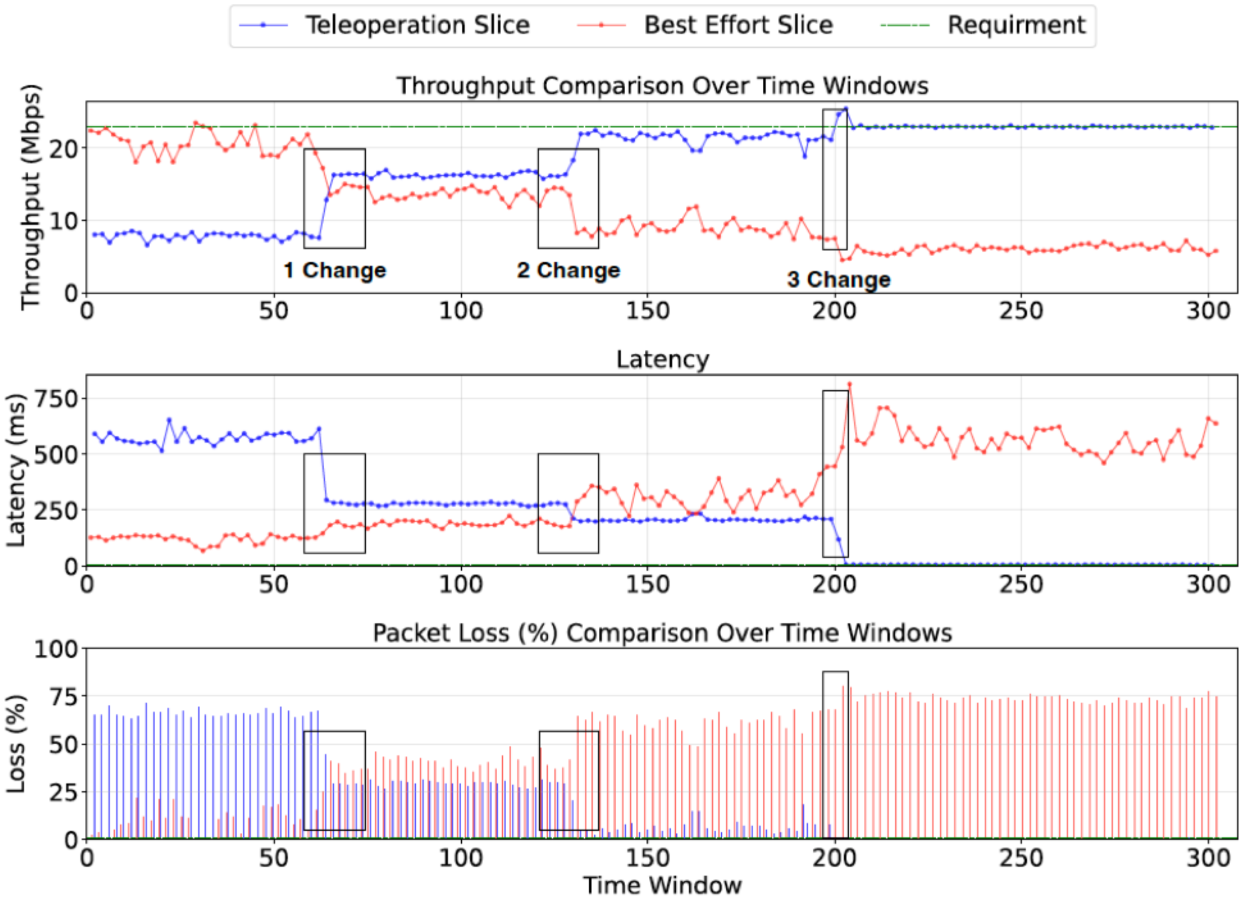

Figure 50: Cross-Domain Controller performance in the context of T4 network slices.

Zero-touch network and Service Orchestration

An Adaptive Hybrid Decision-Making mechanism has been designed and validated within a Zero-touch Network and Service Management (ZSM) framework, enabling dynamic arbitration between deterministic rule-based methods and stochastic Machine Learning (ML) models for autonomous orchestration of 6G compute continuum resources. The architecture integrates a Distributed Decision Engine (DDE) with Adaptive Logic Arbitration, capable of dynamically switching between decision methods based on real-time performance metrics, service requirements, and infrastructure constraints. The overall closed-loop structure aligns with ETSI ZSM principles, where Analytics and Intelligence blocks collaborate to observe, orient, decide, and act upon managed resources.

The framework operates through a closed-loop sequence spanning telemetry ingestion (Figure 2), data preparation, hybrid decision-making, and decision enforcement. A Domain Control Wrapper (DCW) serves as the abstraction layer between the physical infrastructure and the orchestration framework, normalising heterogeneous telemetry streams from distributed compute nodes. A ZSM Analytics module aggregates and dispatches this telemetry according to a dynamic Mapping Manifest, which associates each Managed Domain with a designated Decision Engine and decision method. This configuration supports hot-reload capabilities, allowing decision models to be updated without interrupting the control loop. A Performance Supervisor continuously monitors the health of active Decision Engines and autonomously updates the routing manifest upon detecting degradation or failure.

Each DDE instance comprises four submodules (Figure 3): an Interface Module, a Method Selector, a Method and Algorithm subsystem hosting both deterministic policies and ML agents, and a local Model Repository for low-latency inference. The framework supports flexible M:N mappings between Managed Domains and Decision Engines, enabling deployment topologies such as domain isolation, resource consolidation, and simultaneous validation of competing decision methods.

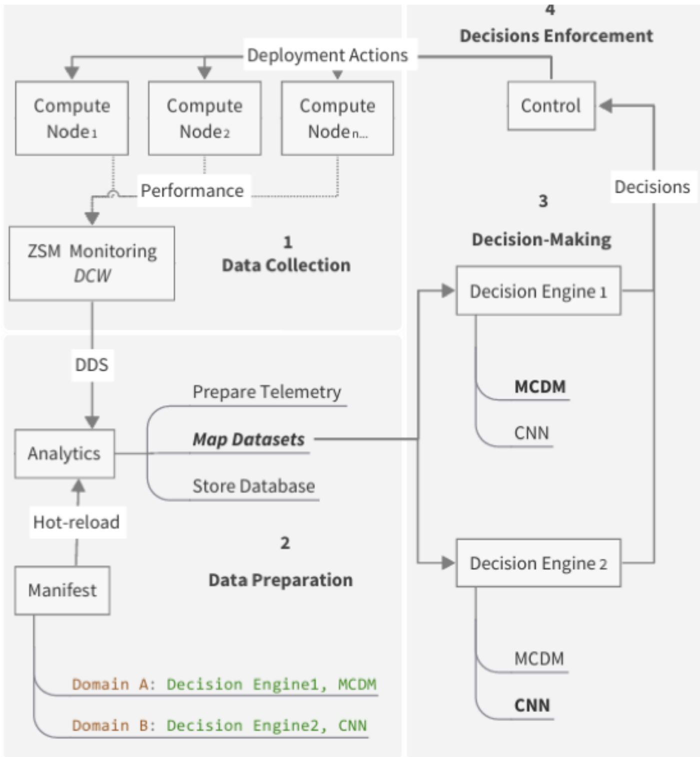

Figure 51: Operational sequence of the Adaptive Hybrid Decision-Making mechanism, illustrating data flow from telemetry ingestion through autonomous decision execution.

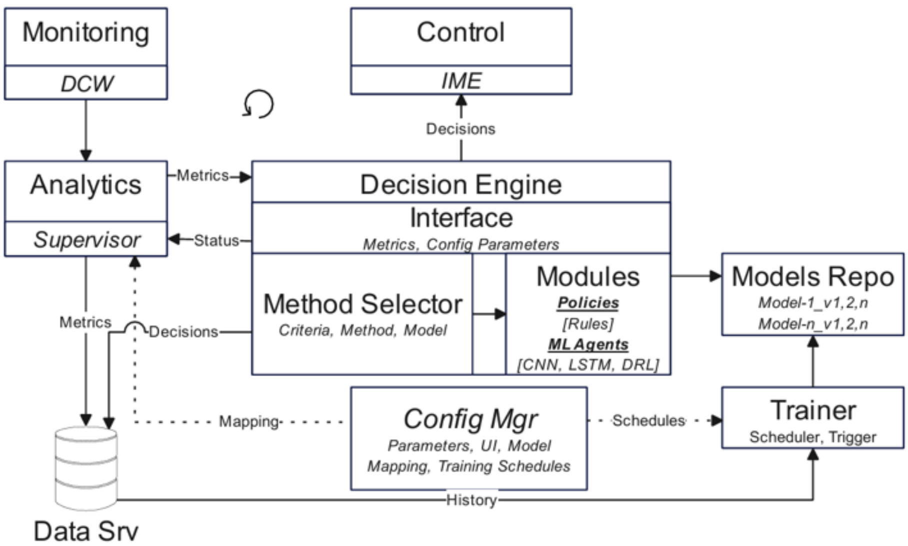

Figure 52: Modular architecture of the Decision Engine within the ZSM framework, highlighting internal submodules and runtime adaptability.

The implementation was realised as a set of containerised microservices orchestrated via Kubernetes, deployed across Edge and Cloud nodes of the Smart Highway testbed. The ZSM Analytics and Decision Engine components operate as a management layer on top of Kubernetes, while the managed vehicular services run as containerised workloads on the edge Roadside Units. Telemetry is normalised and distributed through a unified JSON schema over a Zenoh data bus. The CNN model was developed using the PyTorch framework, serialised as a .pth artefact, and integrated directly into the inference engine of the DDE. The experimental deployment followed a Simultaneous Validation configuration, with two Decision Engines running MCDM and CNN methods in parallel across Edge and Cloud, fed by a synchronised multicast telemetry stream from a single Managed Domain.

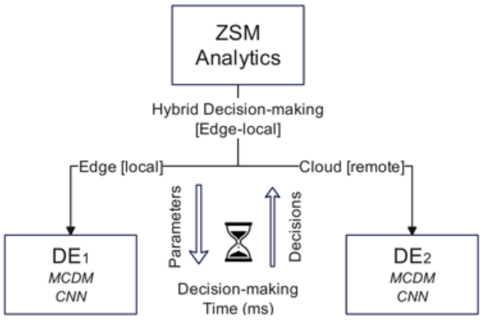

Figure 53: Hybrid Decision-Making experiment plan across Cloud and Edge

The experimental validation was conducted on the Smart Highway (SHW) testbed, covering a four-kilometre segment of the E313 highway in Antwerp, Belgium, with eight Roadside Units operating as far-edge nodes.

Two decision-making methods were evaluated in parallel on both Edge and Cloud deployments (Figure 4):

- A rule-based Multi-Criteria Decision-Making (MCDM) algorithm, evaluating end-to-end latency, CPU utilisation, and energy consumption KPIs to continuously rank and select the optimal execution node.
- A lightweight Convolutional Neural Network (CNN) model, trained on historical MCDM placement decisions from normalised telemetry snapshots, achieving 94.24% agreement with the MCDM baseline over 100 training epochs.

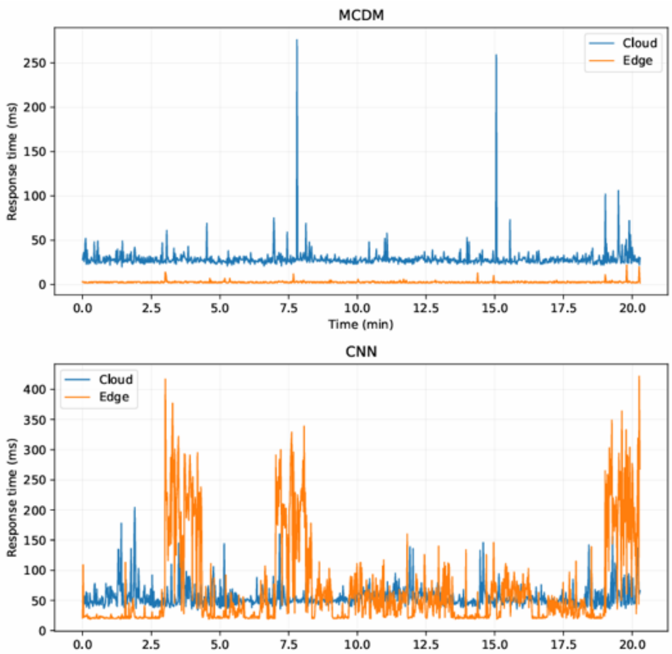

Figure 54: Decision Engine response time for MCDM and CNN under varying traffic loads.

Telemetry from one Managed Domain was multicast to both Decision Engines in a 1x2 mapping configuration, enabling strict side-by-side comparison of method behaviour and the impact of deployment location. Decision response time was measured as the round-trip time from telemetry dispatch to decision reception (Figure 5), including network transmission and compute time. Results confirmed that Edge-deployed MCDM achieved the lowest response time (avg. 2.69 ms), while CNN performed better in the Cloud (avg. 55.20 ms) than at the Edge (avg. 65.19 ms), indicating that computational constraints at the edge outweigh network latency benefits for compute-intensive ML workloads. The experiment confirmed that the ZSM Analytics module correctly routed requests to the intended Decision Engines based on tagged requirements, validating the feasibility of the Adaptive Hybrid Decision-Making mechanism.

After the integration of the ZSM framework within the German trial site is complete, different decision-making and orchestration mechanisms will be applied using T4 application components, in order to ensure optimal placement and service continuity of such applications within the distributed edge-cloud compute continuum.

Decision-Making Algorithm

The decision-making logic of T4 UC is illustrated in Figure 55. The algorithm governs the transitions among four operational states, namely *Automated Driving, Preparation for Remote Operation, Remote-Operated Driving, and Minimal Risk Condition (MRC).* Its purpose is to ensure that the vehicle always remains in the safest possible operating mode according to the current autonomy capability, tele-operation feasibility, and contextual risk.

To support these transitions, the algorithm relies on three indicators that capture the fitness of automation, the feasibility of tele-operation, and the level of contextual risk. The *Autonomous Fitness Score (AFS)* quantifies how well the automated-driving system can safely operate under the current conditions by combining factors such as perception confidence, compliance with the Operational Design Domain (ODD), planner stability, and overall system health. The *Tele-operation Fitness Score (TFS)* evaluates whether safe remote control can be established and maintained, considering the quality of the communication link, the availability and readiness of the remote operator, and the status of the supporting tele-operation infrastructure. The *Context Risk Index (CRI)* expresses the level of environmental and situational risk derived from onboard sensing and V2X-enhanced context awareness, e.g., in the presence of vulnerable road users, construction zones, occlusions, or other hazards that increase operational complexity. Together, these three indicators provide the decision-making function with a structured basis for determining whether the vehicle should remain in automated driving, prepare for remote operation, switch to tele-operated driving, or transition to the Minimal Risk Condition.

Under normal conditions, the vehicle remains in Automated Driving. In this state, the decision-making function continuously evaluates the Autonomous Fitness Score (AFS), the Tele-operation Fitness Score (TFS), and the Context Risk Index (CRI). As long as the AFS remains above the upper threshold, or the TFS remains above the minimum level required for safe fallback support, the vehicle continues automated operation. However, if the AFS drops below the down-transition threshold or the CRI exceeds the high-risk threshold, the vehicle enters the Preparation for Remote Operation state. This intermediate state is used to prepare a possible handover by activating the tele-operation support chain and assessing whether remote intervention can be safely established.

The transition from Preparation for Remote Operation to Remote-Operated Driving occurs when the tele-operation conditions are adequate, i.e., when the TFS is above the minimum threshold, and the remote operator confirms takeover. Once this condition is satisfied, the vehicle switches to Remote-Operated Driving, where control is taken over by the remote operator through the tele-operation system. While in this state, the algorithm continues to monitor the TFS and the AFS. If the TFS remains above the recovery threshold, remote-operated driving can be maintained safely. If, in parallel, the AFS rises above the upper threshold for a dwell time D and the remote operator releases control, the vehicle can transition back to Automated Driving.

The algorithm also defines the logic for entering the Minimal Risk Condition (MRC) whenever safe continuation of the journey is no longer possible. A direct transition from Automated Driving to MRC is triggered when both the autonomous and tele-operation capabilities degrade to emergency level, i.e., when the AFS falls below the emergency threshold and, at the same time, the TFS is below its own emergency threshold. In this case, neither autonomous control nor remote operation can be relied upon, and the safest action is to execute the predefined minimal-risk manoeuvre. Similarly, if the vehicle is already in Remote-Operated Driving and the TFS drops below the emergency threshold, the system transitions from remote operation to MRC. This reflects the situation in which the tele-operation link or the remote-control conditions have become insufficient for safe control. In both cases, the MRC corresponds to a safe fallback behaviour, such as controlled deceleration, stopping in a safe position, or another predefined low-risk manoeuvre.

Overall, the algorithm implements a structured and safety-oriented state machine in which transitions are driven by the joint evaluation of autonomy fitness, tele-operation fitness, and environmental risk. This enables the vehicle to move progressively between automated driving, remote assistance, and safe fallback operation according to the current operational conditions.

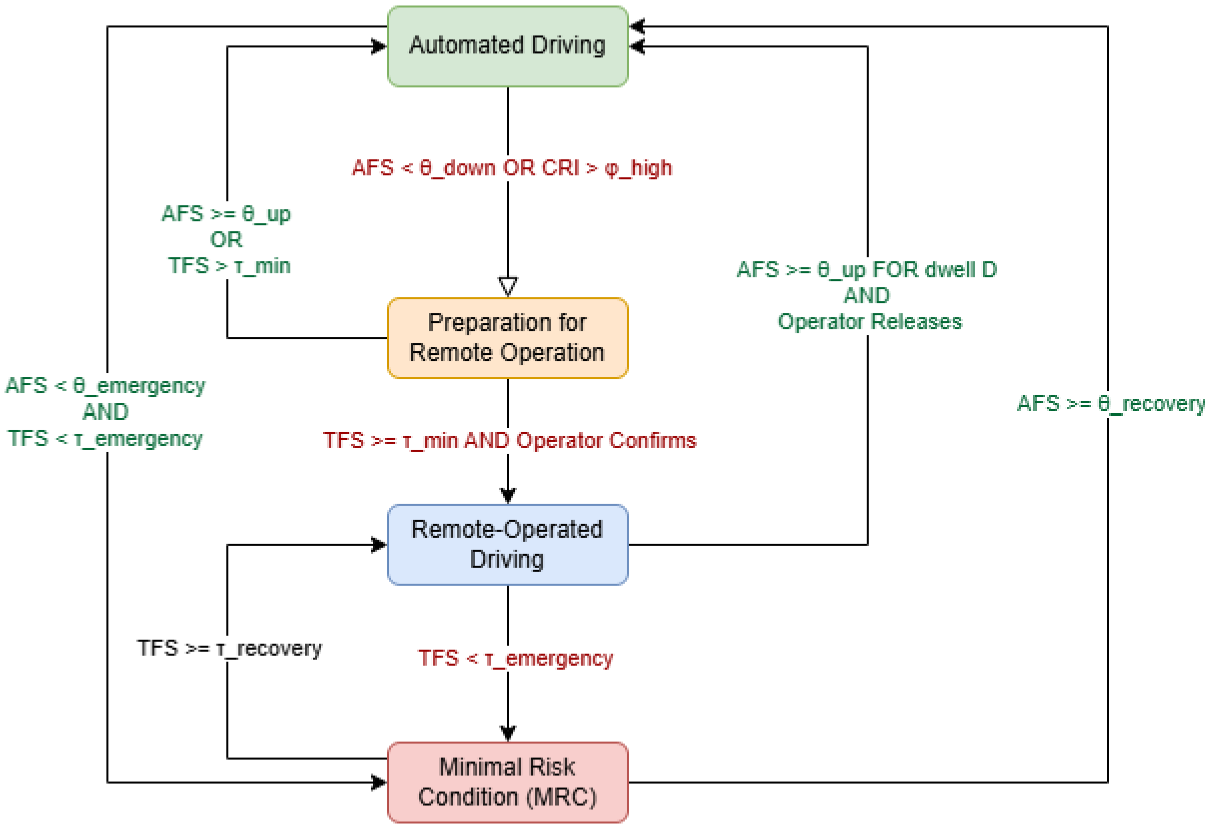

Figure 55: The T4 UC Decision-Making Algorithm

4.4 Transport 5 Use Case

This section presents the current integration status and early-phase activities for the T5 Use Case at the Container Terminal of Thessaloniki Port Authority (ThPA). As of the reporting period (M15), the T5 UC is in an early integration phase, with preparatory work focused on network architecture design, equipment provisioning, and initial platform configuration. The first pre-trial (T5_GR_MVS_1) is planned for Q2 2026.

The integration effort is jointly led by CERTH and ACROMOVE, with CERTH responsible for the Digital Twin platform, AI/ML models, IoT sensor integration, and use case orchestration, and ACROMOVE responsible for the B5G/6G network infrastructure (private 5G RAN, Core, and edge compute) through its ServerPack™Edge 5G and ServerPack™Edge Ultra platforms.

4.4.1 T5 Integration status

The following activities have been completed or are actively in progress as of M13:

Network infrastructure: The ACROMOVE ServerPack™Edge 5G (SPE5G) unit has been provisioned and pre-configured with the 5G SA Core Network functions. The network architecture for the ThPA Container Terminal has been designed, including Small Cell placement planning to ensure coverage of the operational area (quay, yard, and truck loading zones). Access to the ThPA network infrastructure has been coordinated. Virtual environment deployments for the 5G Core and ZSM components are in preparation.

Edge compute platform: The ACROMOVE ServerPack™Edge Ultra (SPE Ultra) has been provisioned for edge AI/ML workloads. Initial containerised deployments of the Digital Twin platform components are being prepared.

IoT and sensor integration: The IoT sensor inventory has been defined, including AXIS PTZ cameras for crane and yard monitoring, air quality monitoring stations for environmental KPIs (PM10, dust emissions), and GNSS/RTK-enabled devices for hybrid localisation of straddle carriers. Integration of the IoT Exposure Function APIs with the Digital Twin data ingestion pipeline is in the design phase.

ZSM and network management: The Zero-touch Service Management (ZSM) architecture has been designed for automated network provisioning and self-healing. RAN and RIC deployment planning has been completed. The ZSM Monitoring and Analytics framework and the telemetry collection pipeline are being set up for integration with the SPE5G network monitoring capabilities.

Enablers from WP3: The communication enablers (automated gNB configuration, network performance monitoring and control, identification and selection of backhauling) are being adapted for integration with the ACROMOVE private 5G infrastructure. The Digital Twin framework and the AI-as-a-Service components (ML models catalogue, MLOps pipelines) are under development and will be deployed on the SPE Ultra edge platform.

4.4.2 T5 Preliminary test results

As the T5 UC is in the early integration phase, no formal trial execution or KPI measurement campaigns have been conducted yet. The first pre-trial (T5_GR_MVS_1) is scheduled for Q2 2026 and will focus on the initial deployment of the network infrastructure and a dry-run of B5G connectivity.

However, the following preparatory demonstrations and validations have been performed to confirm the readiness of individual technology components:

Private 5G network validation: The ACROMOVE SPE5G platform has been validated in laboratory conditions, confirming the ability to deploy a fully functional private 5G SA network (including Core, RAN, and up to 4 Small Cells) within the target 15-minute deployment window. PTP synchronisation and basic throughput/latency benchmarks have been verified internally by ACROMOVE. These baseline results will inform the on-site network design at ThPA.

In addition, during the days of the Project's 4th plenary meeting that was held in Thessaloniki, an initial test implementation of the ACROMOVE ServerPack™Edge 5G (SPE5G) was made to further inform the design and plan of the upcoming trial tests (Figure 56).

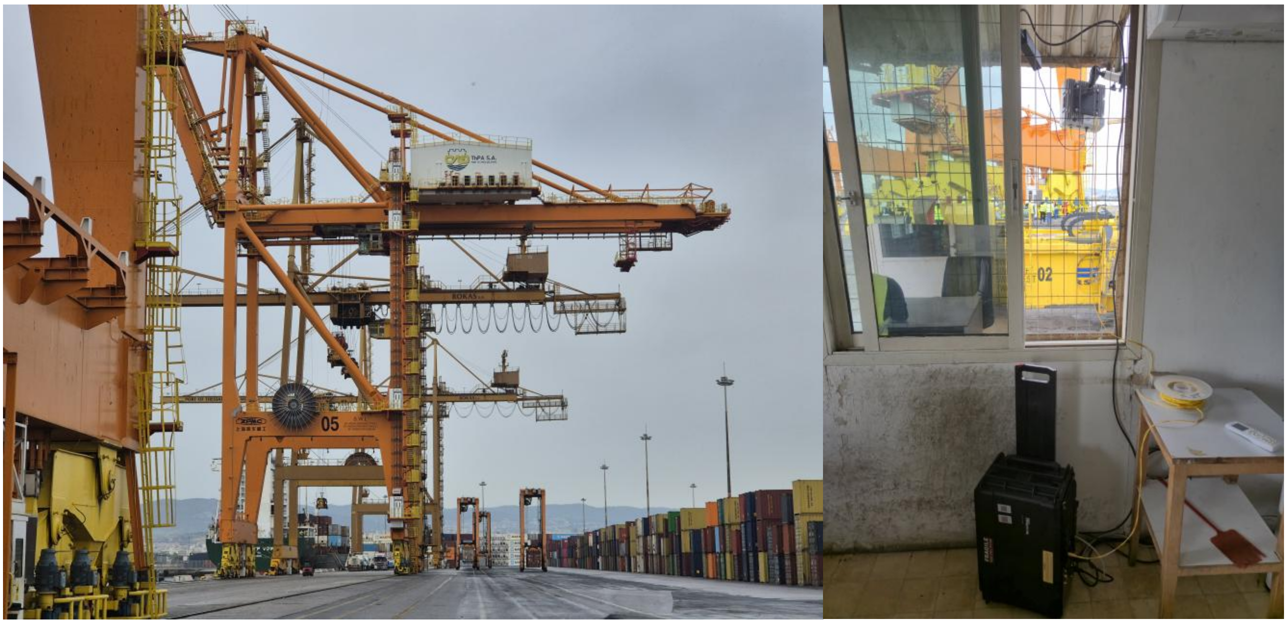

Figure 56: STS cranes and the demo installation of ACROMOVE ServerPack™Edge 5G (SPE5G) in the port of Thessaloniki

Digital Twin platform development: The AnyLogic-based simulation engine for port process modelling is under active development. Initial simulation scenarios covering vessel unloading, straddle carrier yard routing, and truck loading sequencing have been prototyped. The containerised deployment architecture for the Digital Twin services has been defined and tested in a development environment.

AI functionality for incident detection in moving vehicles: Partners have worked on the development of AI functionality for incident detection in real-time through cameras. The final functionality will be trialled in the Port of Thessaloniki. In the following screenshots, we see some pre-trial results/screenshots from cameras in a street of Athens, which can detect persons not wearing seatbelts and other security incidents in the street by using AI image recognition. Images are captured by cameras, and GDPR is respected through blurring/masking of sensitive data (e.g. people's identification or vehicles' identification are not recognised by the system).

Figure 57: T5 AI functionality for incident detection (preliminary/indicative results pre-trial)

Anonymised statistics are provided in user-friendly dashboards, as shown below as well. Such statistics include the type of vehicle, type of detected incident, total number of incidents, etc.

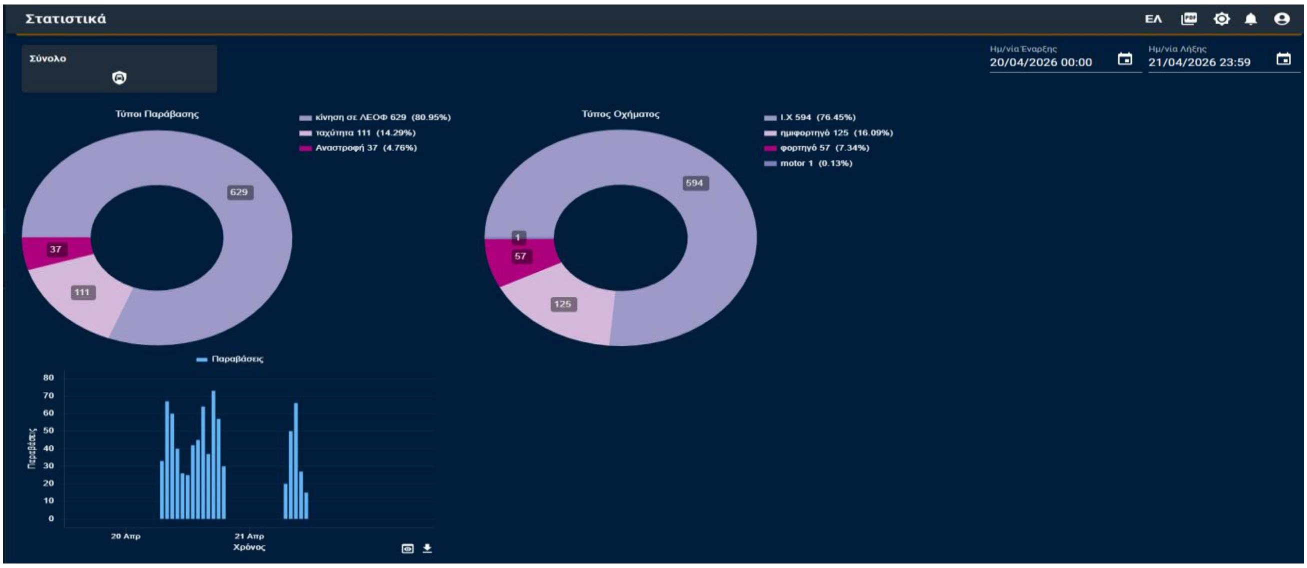

Figure 58: T5 AI functionality for incident detection – Statistics dashboard (preliminary/indicative results pre-trial)

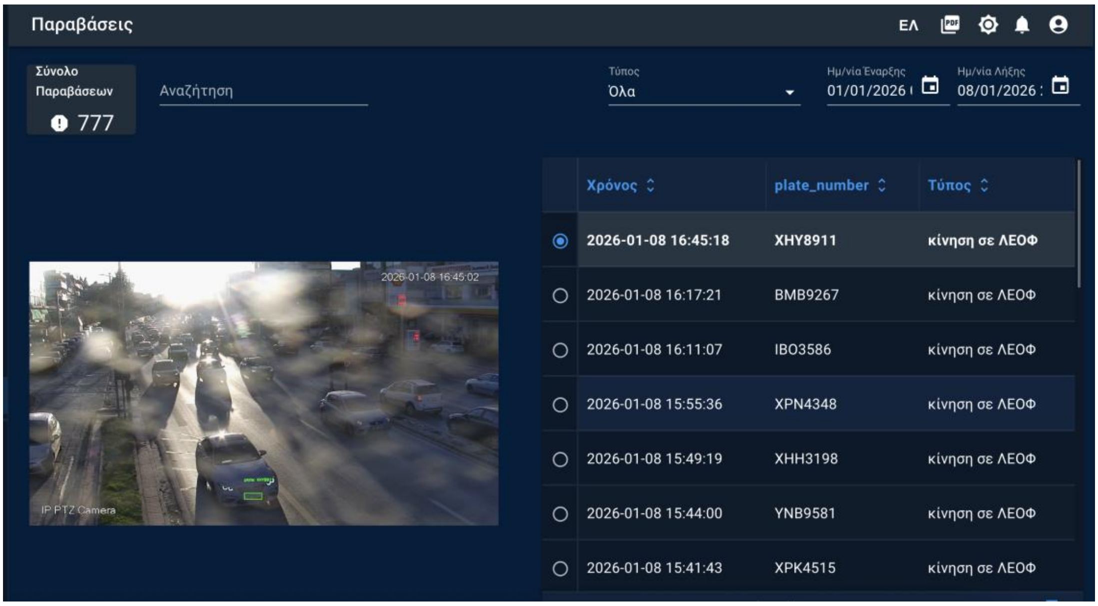

Figure 59: T5 AI functionality for incident detection – Camera View (preliminary/indicative results pre-trial)

The confusion matrix below shows an example of the identification of true values for seatbelts of vehicle drivers/passengers. In the pre-trials, the model is still under training to increase the accuracy of correct detections of issues (e.g., automated detection of non-wear of seatbelts in moving vehicles).

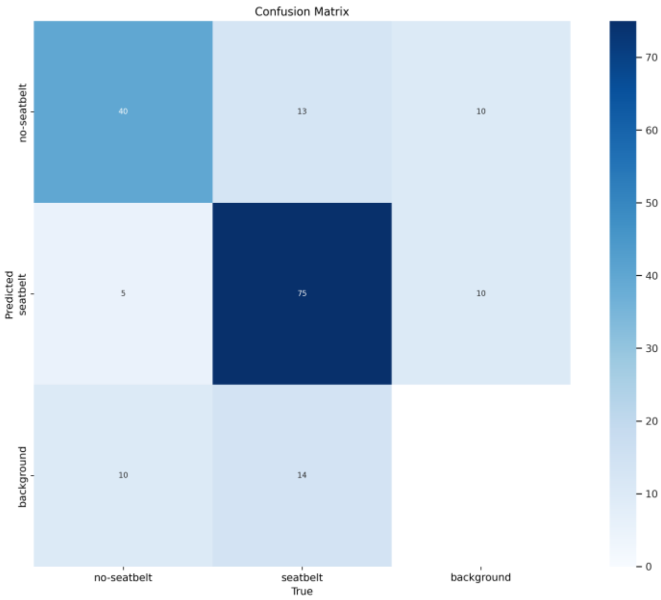

Figure 60: T5 AI functionality for incident detection – Confusion Matrix (preliminary/indicative results pre-trial)

Next steps: The immediate next steps include on-site deployment of the SPE5G at ThPA premises (Q2 2026), execution of T5_GR_MVS_1 (network bring-up, gNB configuration, ZSM dry-run), followed by progressive integration of IoT sensors, the Digital Twin platform, and the AI/ML inference pipeline on the SPE Ultra edge compute. Subsequent pre-trials (T5_GR_MVS_2 in Q4 2026, T5_GR_MVS_3 in Q2 2027) will incrementally validate yard coordination, safety monitoring, and environmental measurement scenarios, building towards the final integrated trial (T5_GR_final_1) in Q3 2027.

5 Conclusion and Next Steps

This deliverable provided the results of the activities of Task 6.1, in which the Transport Use Cases of the AMAZING-6G project have been detailed by defining the system architecture design and by identifying the needed network infrastructure components and the B5G/6G enablers. The intermediate status of development, integration, and testing activities is also provided in this deliverable based on the achievements of T6.2 and T6.3.

The design of the Transport UCs has been finalised, while implementation activities are proceeding, focusing on the development of all services required for the demonstration of the UC. The integration has started in all the trial sites involved in the Transport UCs, and some preliminary testing of single components has been performed to validate the expected functionalities in an incremental way.

The next steps will target the completion of the development activities to allow the integration of the technological enablers with the vertical services of each UC. The testing activities will proceed to validate the development components, and, after having completed the integration, to test the overall integrated system.

The content of this deliverable provides the status of the developments to be taken as reference for planning the prosecution of the implementation activities in WP6, respecting the expected timeline. This deliverable is also preparatory for setting the ground for the evaluation activities to be performed in WP7.