Partners: CTTC, EAB, ICC, NFI, NFR, NXW, QLC, QRT, TID, TUD, UBW, VTT, WIN











Hexa-X-II D2.6 Deliverable

D2.6 Final end-to-end system evaluation results of the overall 6G system design

Hexa-X-II

hexa-x-ii.eu

2025-05-31

Deliverable properties

Document number	**D2.6**
Document title	Final end-to-end system evaluation results of the overall 6G system design
Document editor	**WIN**: Sokratis Barmpounakis, Vasiliki Lamprousi, Vasilis Tsekenis, Panagiotis Demestichas
Authors	**CTT:** Pol Alemany, Raul Muñoz, Ricard Vilalta, Behnam Ojaghi, Daniel Adanza, Lluís Gifre **EAB:** Selim Ickin **ICC:** Ioannis Tzanettis, Grigorios Kakkavas, Anastasios Zafeiropoulos **NFI:** Akshay Jain, Dani Korpi **NFR:** Sylvaine Kerboeuf **NXW:** Erin Elizabeth Seder, Pietro G. Giardina, Pietro Piscione **QLC:** Yann Lebrun, Taesang Yoo **QRT:** Mohammad Hossein Moghaddam **TID:** Riccardo Nicolicchia **TUD:** Carl Collmann, Ahmad Nimr **UBW:** Luís Santos **VTT:** Pawani Porambage, Rafael Pires; **WIN:** Sokratis Barmpounakis, Vasiliki Lamprousi, Vasilis Tsekenis, Panagiotis Demestichas

Deliverable properties

Editorial team	
Target dissemination level	PU
Version	1.0
File Name	Hexa-X-II-D2.6-slideset_FINAL

Revision	Date	Issued by	Description
0.1	01.10.2024	Hexa-X-II, T2.5, WINGS	Complete first draft towards internal review phase
0.2	28.03.2025	Hexa-X-II, T2.5, WINGS	Internal review phase document version
0.3	30.04.2025	Hexa-X-II, T2.5, WINGS	External review phase document version
1.0	30.05.2025	Hexa-X-II, T2.5, WINGS	GA approval phase document version
1.0	30.06.2025	Hexa-X-II, T2.5, WINGS	Submission to EC

Table of Content

Chapter 2: The E2E 6G System

Chapter 3: E2E system-level validation design

Chapter 4: System- and Component-PoC #C validation and evaluation results



Chapter 2
The E2E 6G System

Overarching 6G system blueprint

Application enablement platform & Application layers:

- a unified platform for developers and third-party applications to leverage 6G capabilities.

Network functions layer

- a single radio access technology, an extended core, and beyond communications capabilities

Infrastructure layer

- delivering a cloud continuum

Pervasive functionalities

- enabling AI/ML, data, security, and automation across all layers

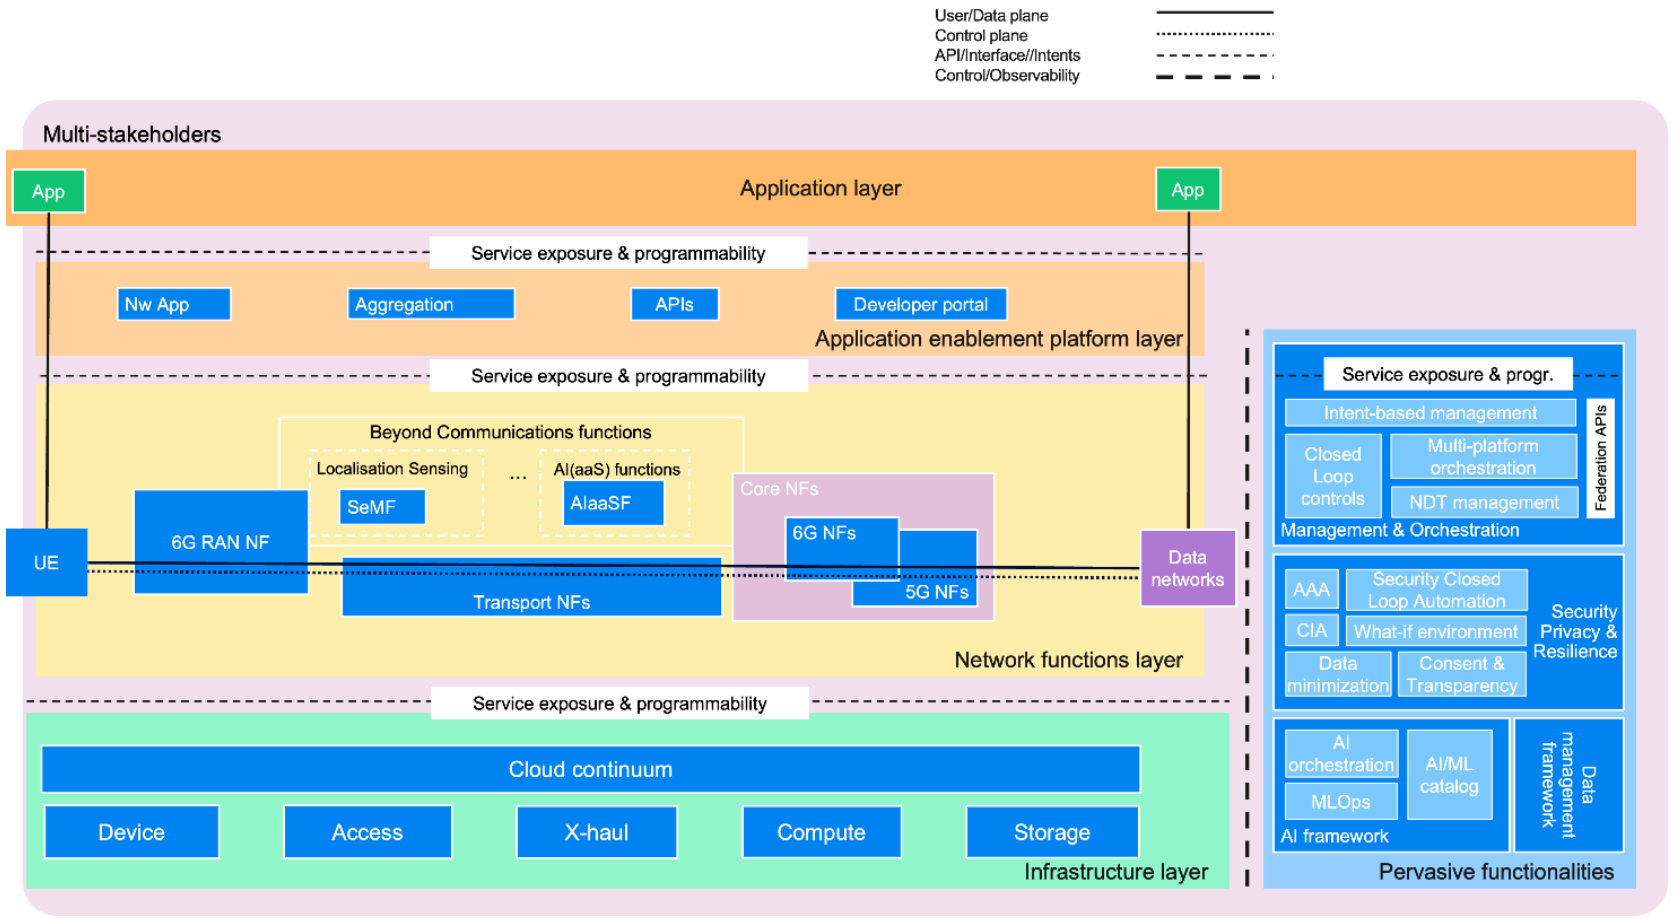

PoC components mapping to the E2E system blueprint

- Application layer components: Intuitive Trigger/UI, intelligent components (functionality allocation, swarm planning and management, etc.) real-time video streams, sensor/device status, etc.
- Exposed compute/AI and data resources via the application enablement platform layer
- Advanced 6G functionalities and devices including flexible topologies, 6G-based sensing, AI-native interfaces, energy harvesting devices, adaptive transceiver design and E2E compute offloading mechanisms in E2E extended reality setups.
- Synergetic monitoring and orchestration: Dynamic resource allocation, seamless integration, and continuous trust evaluation via trust evaluation functions.

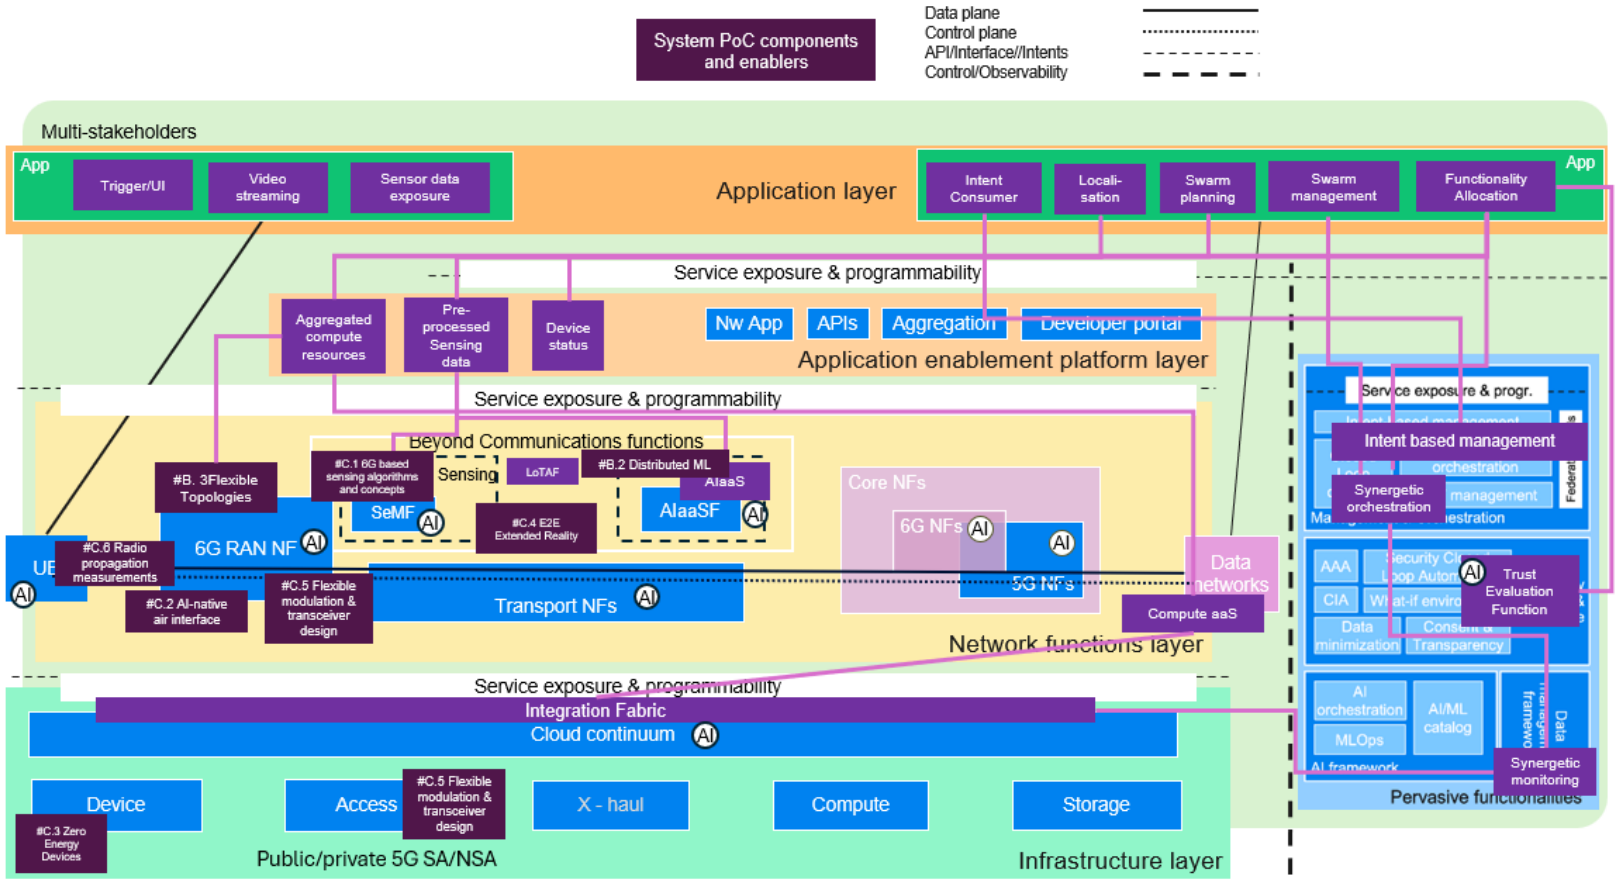



Chapter 3
E2E system-level validation design

E2E system validation

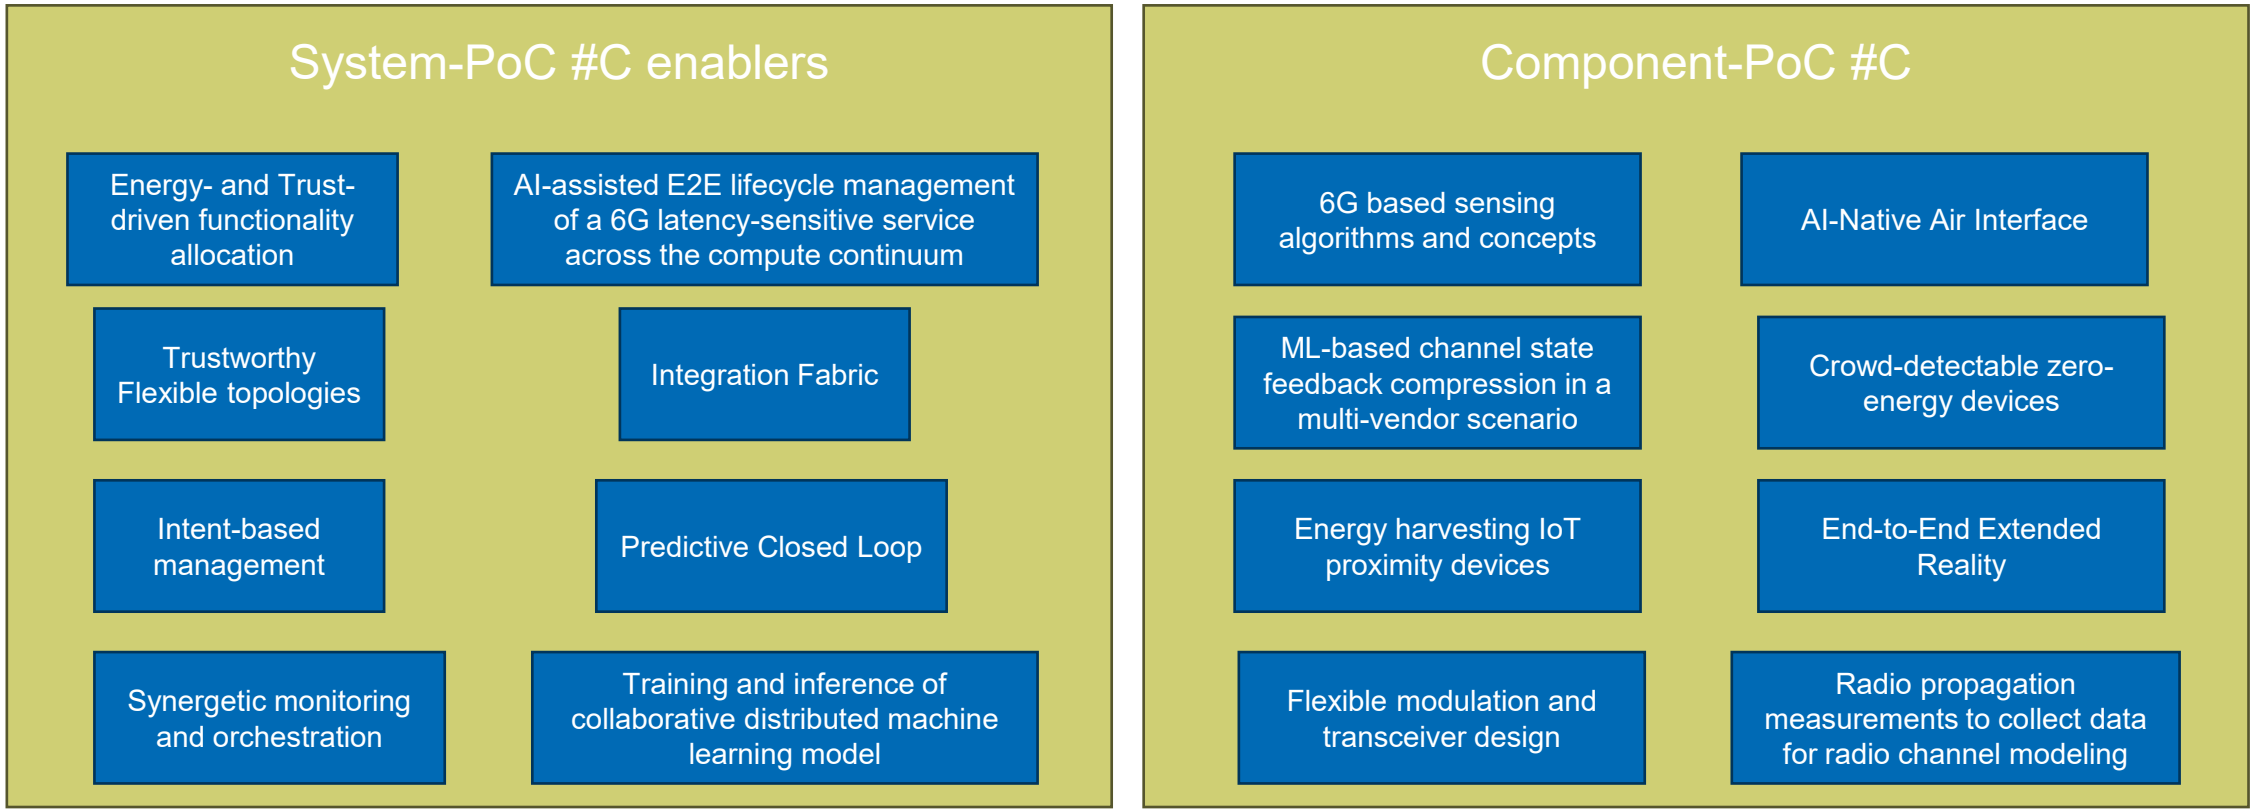

System PoC design evolution

- System-PoC C introduces energy harvesting-capable proximity-sensing devices in the warehouse inventory management scenario, designed to monitor shelf occupancy with high precision.
- Flexible topology nodes provide communication resources to the energy harvesting devices
- Equipped with energy harvesting capabilities using solar panels, represent a step forward in sustainable and autonomous operation: real-time digital twinning of the warehouse storage areas is updated via the energy harvesting devices for shelf occupancy; this information is then utilised by the evolution of the FA algorithm
- New APIs leveraging exposure of Application Enablement Platform capabilities, such as Quality-on-Demand (QoD), device status/density, as well as integrating System PoC-B capabilities, such as advanced flexible topologies functionality and compute offloading features
- Integration Fabric adoption as a core enabler in multi-domain resource orchestration scenarios
- Trust Evaluation Function- and Energy-driven Functionality allocation of physical and compute tasks
- Intent-driven orchestration leveraging intent hierarchy, i.e., business intents, as well as intent-based DSM capabilities.

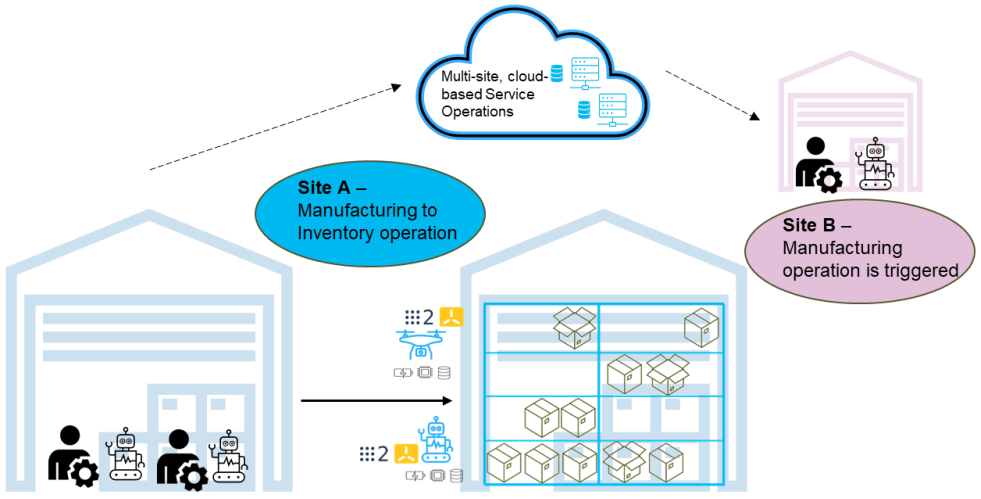

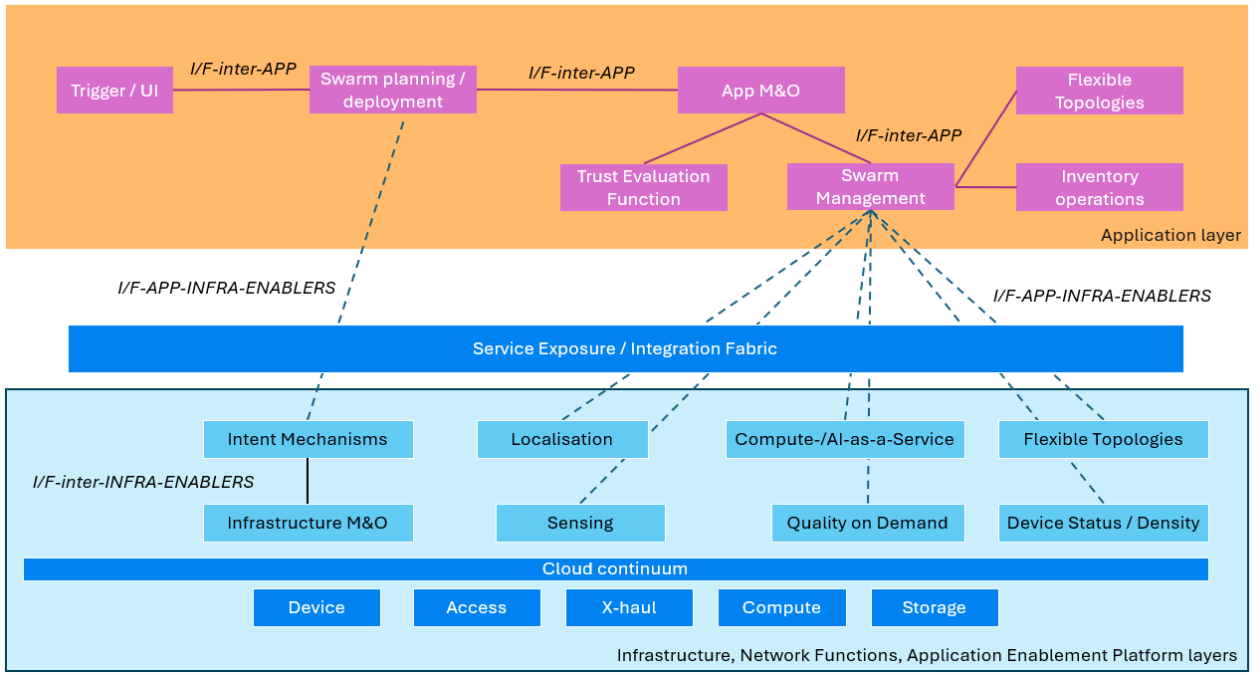

System PoC: Cobot-powered Warehouse Management in the context of multi-site synergetic monitoring and orchestration scenarios

- IBM on Transport Resources
  - An IBN solution based on the intent-based DSM (detailed in Chapter 4) for configuring the network based on business intent.
  - The implemented solution is based on the "Intent translation and provisioning" enabler described in D2.3.

- Management of Transport Domain
  - The SDN controller for managing and controlling the network across the different domains.
  - TeraFlowSDN as the selected solution to deliver an SDN Controller to the multiple sites connectivity.

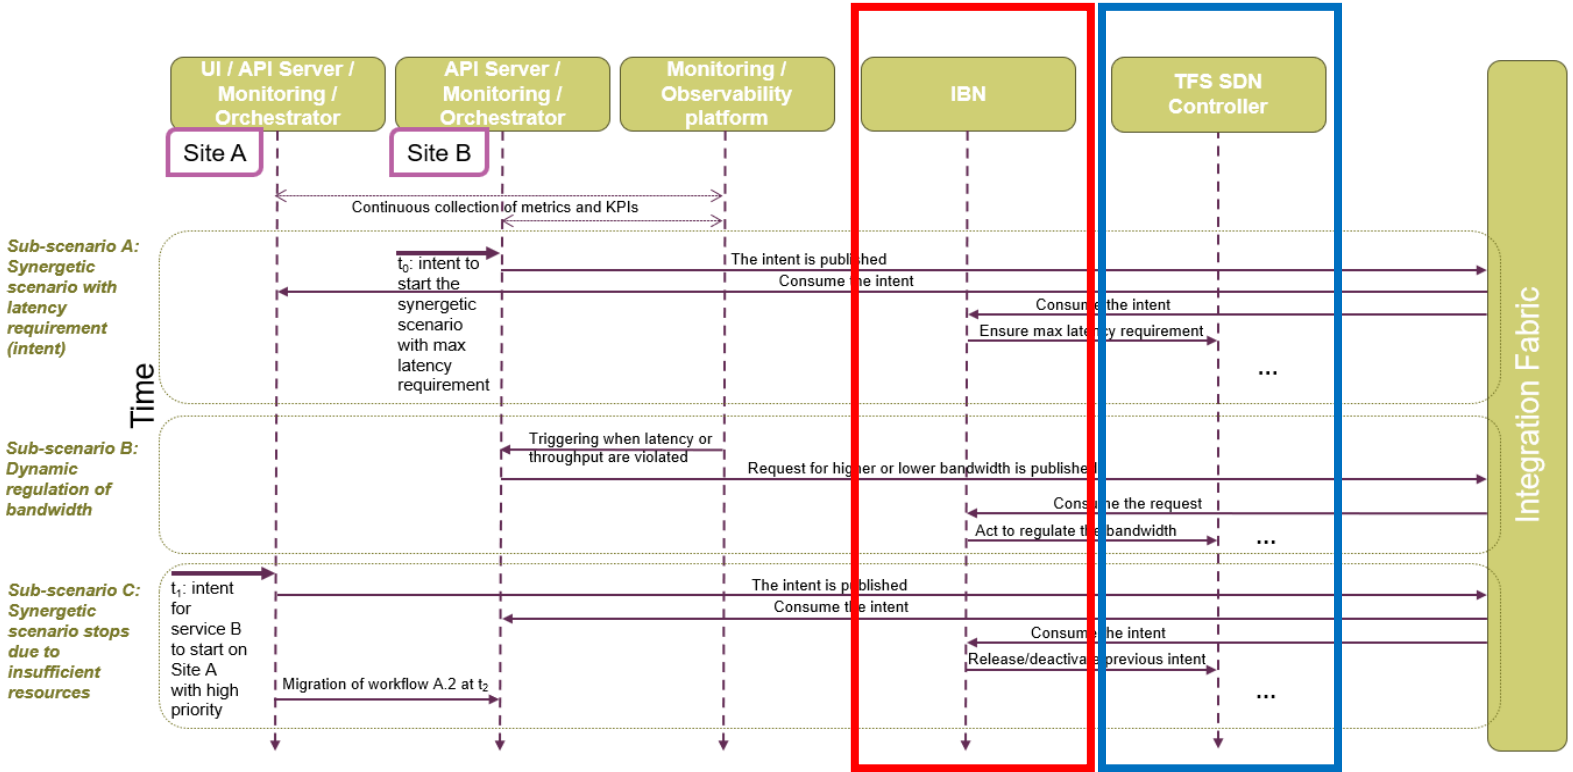

System PoC: Integration Fabric role

- Provides a standardized (e.g., REST-based) interface that abstracts complex underlying system capabilities, enabling easy access to diverse resources and services.
- Seamlessly connects different components at service level in multiple domains and sites, supporting real-time resource allocation and management.
- Facilitates high-level, goal-oriented control that enables and improve intent-driven service management and closed-loop automation.
- Serves as the central integration point for various system enablers, enabling real-world interoperability testing and API refinement.

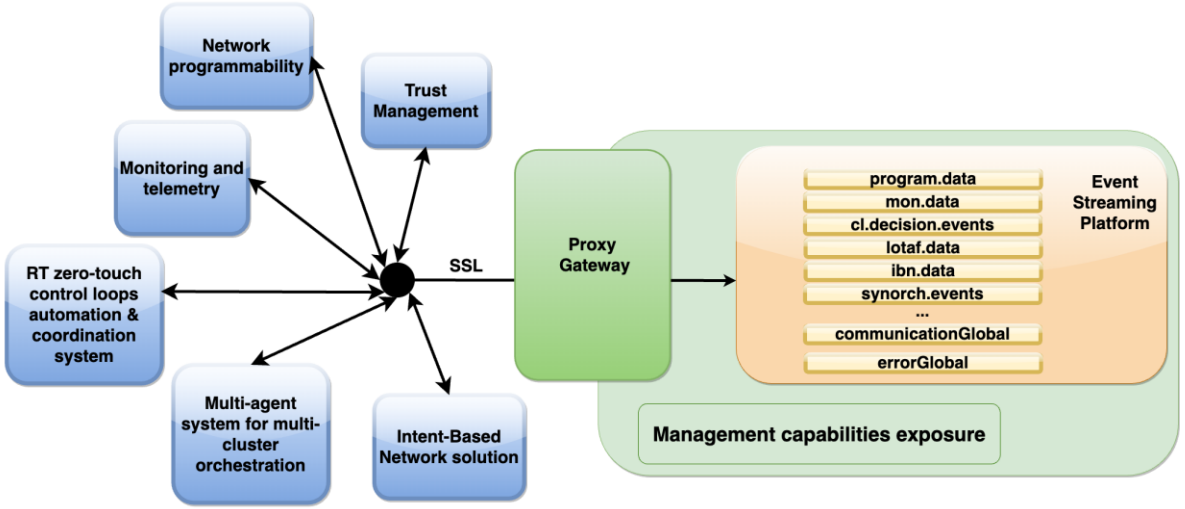

Sustainability-oriented orchestration in 6G

- **Warehouse Inventory Management scenario powered by collaborative robots (Cobots)**
  - Utilizes advanced computer vision, sensor data fusion, and dynamic translation systems to efficiently manage inventory using autonomous mobile robots (AMRs) and unmanned aerial vehicles (UAVs) for real-time item detection, counting, and navigation.
- **Functionality Allocation (FA) & Trust Management**
  - The FA mechanism optimizes workload and task management towards energy efficiency and maximum trustworthiness, integrating a Trust Evaluation Function (TEF) to assess trust levels of system's compute nodes based on real-time and historic data.
- **Swarm planning and FA mechanism enhancements**
  - Responsible for resource reservation and management of AMRs/UAVs, ensuring smooth inventory operations while monitoring runtime performance and managing events like battery depletion or hardware failures.
- **Trust system enhancement with Level of Trust Assessment Function (LoTAF)**
  - The LoTAF mechanism enhances trust assessment by evaluating additional real-time features such as CPU, memory, network, and disk data, with trust indexes integrated into the system for improved reliability and coordination.
- **Energy efficiency and device integration**
  - New solar energy harvesting devices with proximity sensors improve AMR/UAV energy consumption and operational efficiency, integrated with FA and swarm planning services to optimize resource usage.

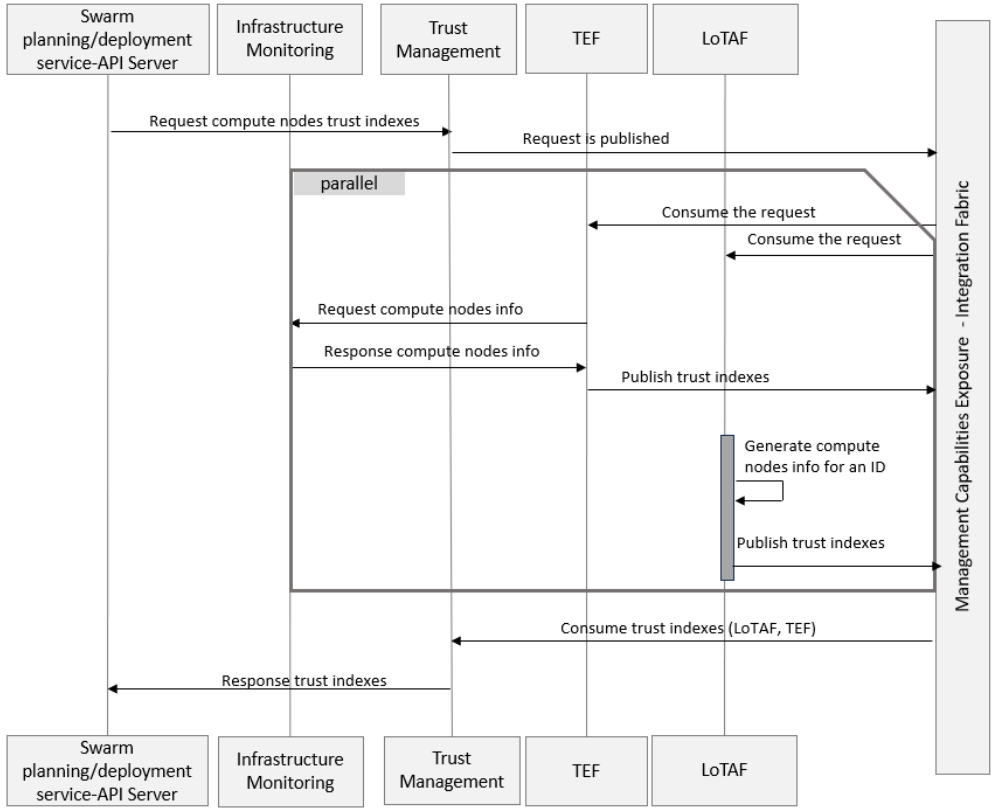

Message Sequence Diagram for Trust Management Processing: Receiving TEF and LoTAF Indexes and Feeding the swarm planning/deployment service.

AI-assisted E2E lifecycle management of a 6G latency-sensitive service across the compute continuum

- **Smart manufacturing application**: A microservice-based system manages a stationary robotic arm using two service chains—one for remotely controlling a surveillance vehicle and another for detecting human presence via ML to safely pause/resume arm operations.
- **Teleoperation workflow**: Involves live video encoding, a media server for stream handling, and a user interface—designed to support high-bandwidth, continuous video streaming (5–30 Mbps) with low-latency requirements for real-time teleoperation.
- **Object detection and safety alerts**: Continuously monitors the robotic workspace with a dynamic sampling rate. If a person is detected, the system escalates to warning or danger states, triggering the robotic arm to halt and resume based on ongoing detection results.
- **Distributed orchestration across clusters**: In a multi-cluster environment, agents manage service scaling and migration based on workload, proximity, and system performance. These agents may operate independently or collaboratively, depending on deployment strategy.
- **Reinforcement learning-based control**: Initially managed centrally with RL, the system evolved into a decentralized Multi-Agent RL (MARL) setup where agents manage different and heterogeneous orchestration tasks, improving flexibility and scalability.
- **Real-time decision making**: Orchestration strategies rely on continuous monitoring of latency, CPU, memory, and request load to make optimal placement and scaling decisions, ensuring service-level objectives are met efficiently.

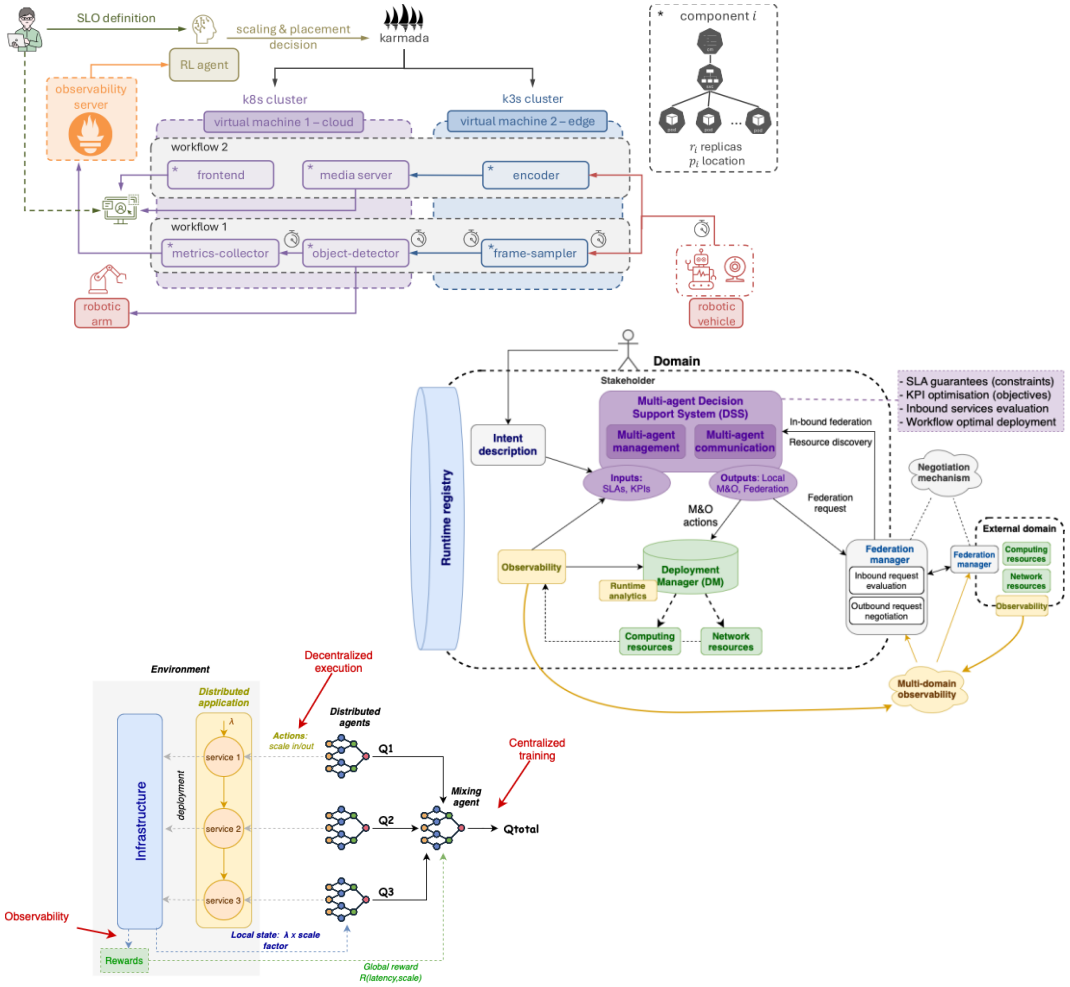

The figures originate from deliverable D6.5 [HEX225-D65] as this work was carried out in close alignment with the orchestration mechanisms developed under WP6.

Training and inference of collaborative distributed machine learning model on a dynamically changing heterogeneous 6G architecture

- Component PoC B.2 showcases the potential of split learning for secure, efficient, and adaptive AI-driven applications in 6G networks.
- Collaborative model training across application and networks using split learning, a technique for vertical federated learning without sharing raw data between entities.
  - Distributed input nodes processing domain-specific data
  - a generalization intermediate node at the network
  - output nodes that can be deployed at the external applications
    - delay estimation
    - object bitrate estimation
- Dynamically offloading NN model layers based on resource availability (compute, memory, energy)
  - Ensures efficient use of resources while maintaining model accuracy.
  - Energy-aware model layer offloading to balance performance and energy consumption.

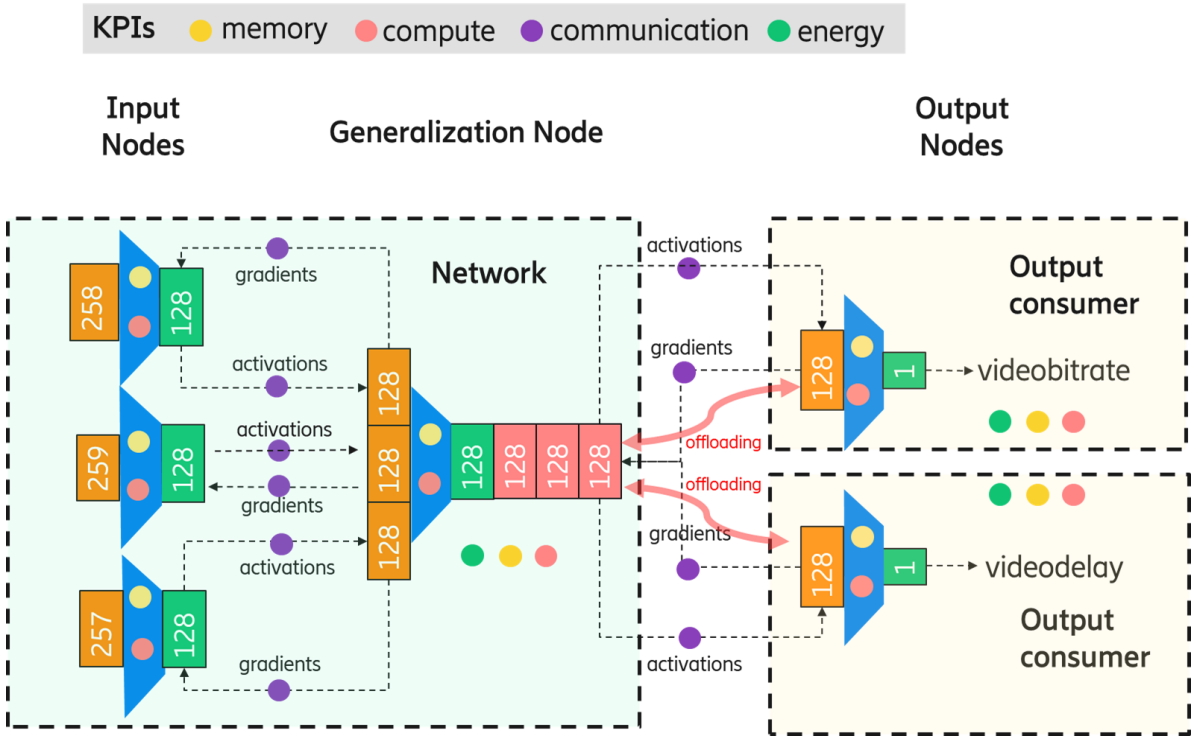

See slide 32 for results

Trustworthy flexible topologies in 6G, leveraging on “beyond communication” aspects

- Dynamic, trustworthy network architectures designed to adapt in real-time to varying service requirements and environmental conditions.
- Beyond Communication Services (BCS): Integration of advanced data processing and insight extraction into the network fabric, enhancing network functionality beyond traditional communication.
- Trustworthiness extends beyond security to include factors like capacity, energy efficiency, throughput, and cost, with a Trust Manager and Trust Evaluation Function assessing critical parameters.
- Adhoc network controller: Orchestrates network node selection, continuously monitoring the environment and generating a composite trust metric to optimise node selection and function allocation.
- Dual topology instantiation: Combines systematic methods and heuristic optimisation for efficient node selection, using algorithms like Minimum Spanning Tree (MST) for UAV placement and interconnectivity.
- AI/ML algorithms are used to optimise UAV flight paths, reducing redundant movements and improving energy conservation.
- CAMARA-compliant APIs: Enable standardized exposure of network metrics and seamless interoperability between network functions, supporting real-time coordination of connectivity and resource allocation.
- Flexible Topology Network (FTN): Orchestrates operations like authentication, resource allocation, and mobility management, guided by performance metrics such as latency, throughput, and energy consumption.
- UAV-assisted topologies provide flexibility and resilience in scenarios like disaster recovery, ensuring optimal network performance even in compromised environments.

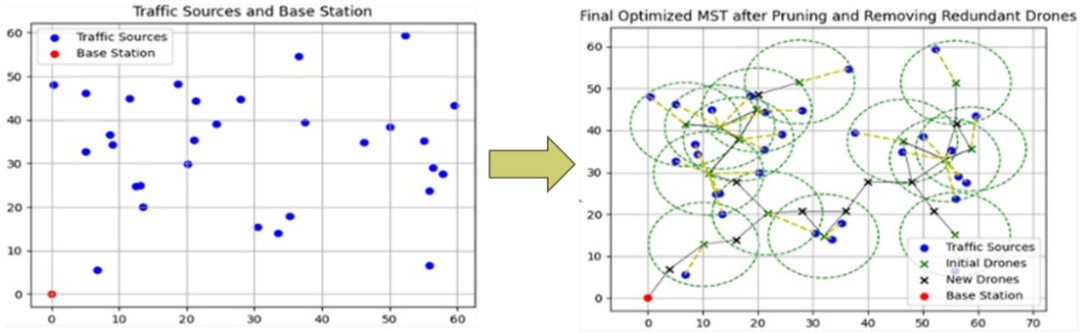

Flexible UAV placement process transforming an initial distribution of traffic sources into an optimized, low-cost network topology.

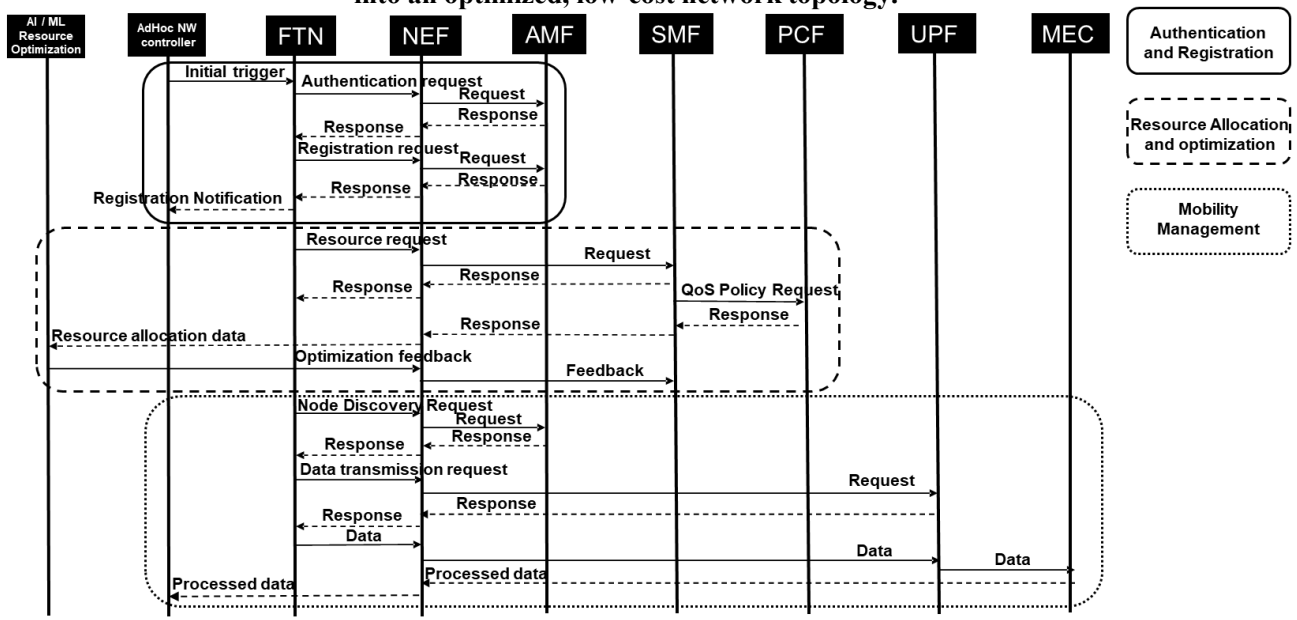

High-Level Message Flow for Flexible Topology Network Orchestration Using Exposed Network Functions.

Synergetic monitoring and orchestration (1/2)

- Dynamic, real-time management of services across multiple sites, optimising resource allocation and performance to meet strict SLAs and adapt to changing network conditions.
- Two deployment options: Services can be hosted on local infrastructure (full control but maintenance costs) or offloaded to a Cloud Provider (CP) with service tier guarantees and associated costs.
- Service deployment and network guarantees: The Service Provider (SP) manages where to deploy services, network guarantees from Network Providers (NP), and processing guarantees from the Cloud Provider (CP).
- Key services: Service A (latency-sensitive manufacturing service) deployed with workflows (A.1: video streaming and A.2: object detection) and Service B (warehouse inventory management service).
- Workflow migration: Workflow A.2 migrates from Site A to Site B due to insufficient resources, with performance metrics monitored before and after migration.
- Sub-scenario Breakdown:
  - Sub-scenario A: Latency requirement management through intent-based requests and TeraflowSDN (TFS) network adjustments.
  - Sub-scenario B: Dynamic regulation of bandwidth when latency/throughput violations occur.
  - Sub-scenario C: Workflow stop due to insufficient resources, with a TFS request for available resources and an update to the status through Integration Fabric (IF).

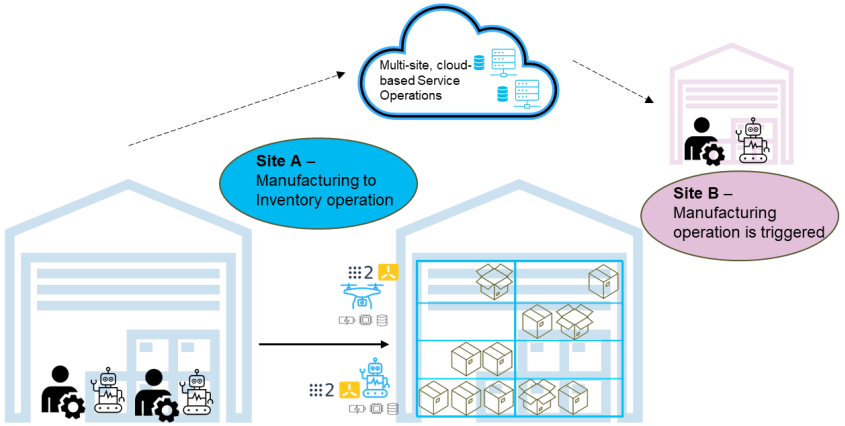

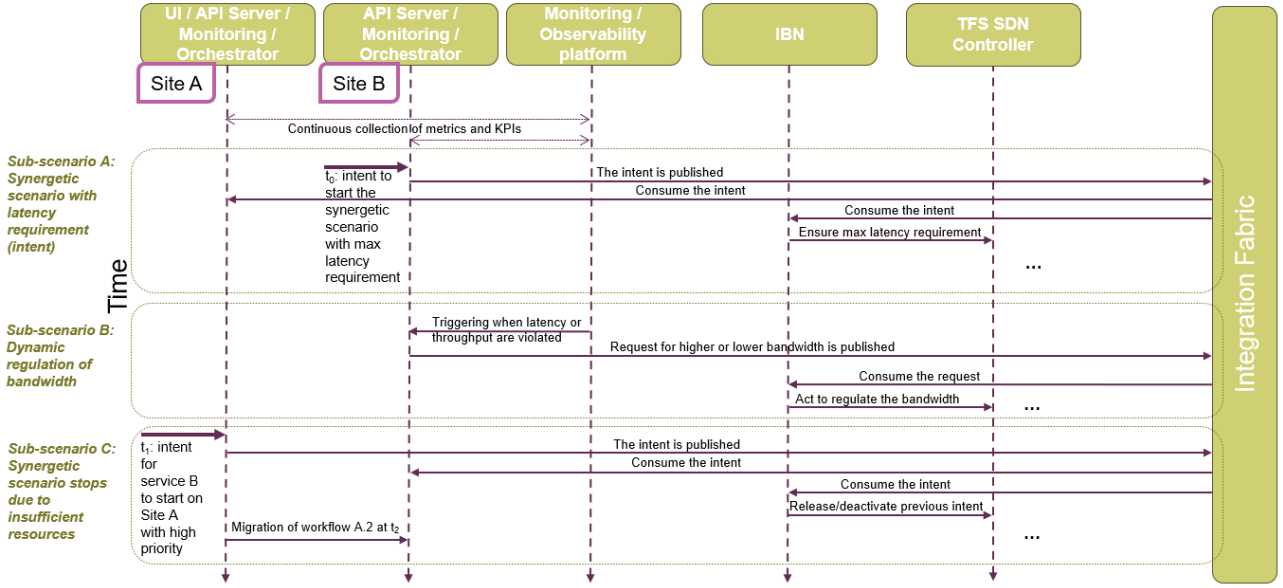

Synergetic monitoring and orchestration (2/2)

Infrastructure: Two sites - Site A & Site B
Service components:

- *Frame Sampler* – deployed at Site B (near robot capturing frames)
- *Object Detector* – dynamically migrated between Site A and B based on resource availability

Observability:

- Distributed data collection via Monitoring & Telemetry Framework
- Data stored at observability server

Migration Logic:

- High-level intents translated into migration directives
- To Site A: Site A starts service, Site B drops it after evaluation of active intents
- To Site B: Triggered when Site A resources are reassigned to other apps

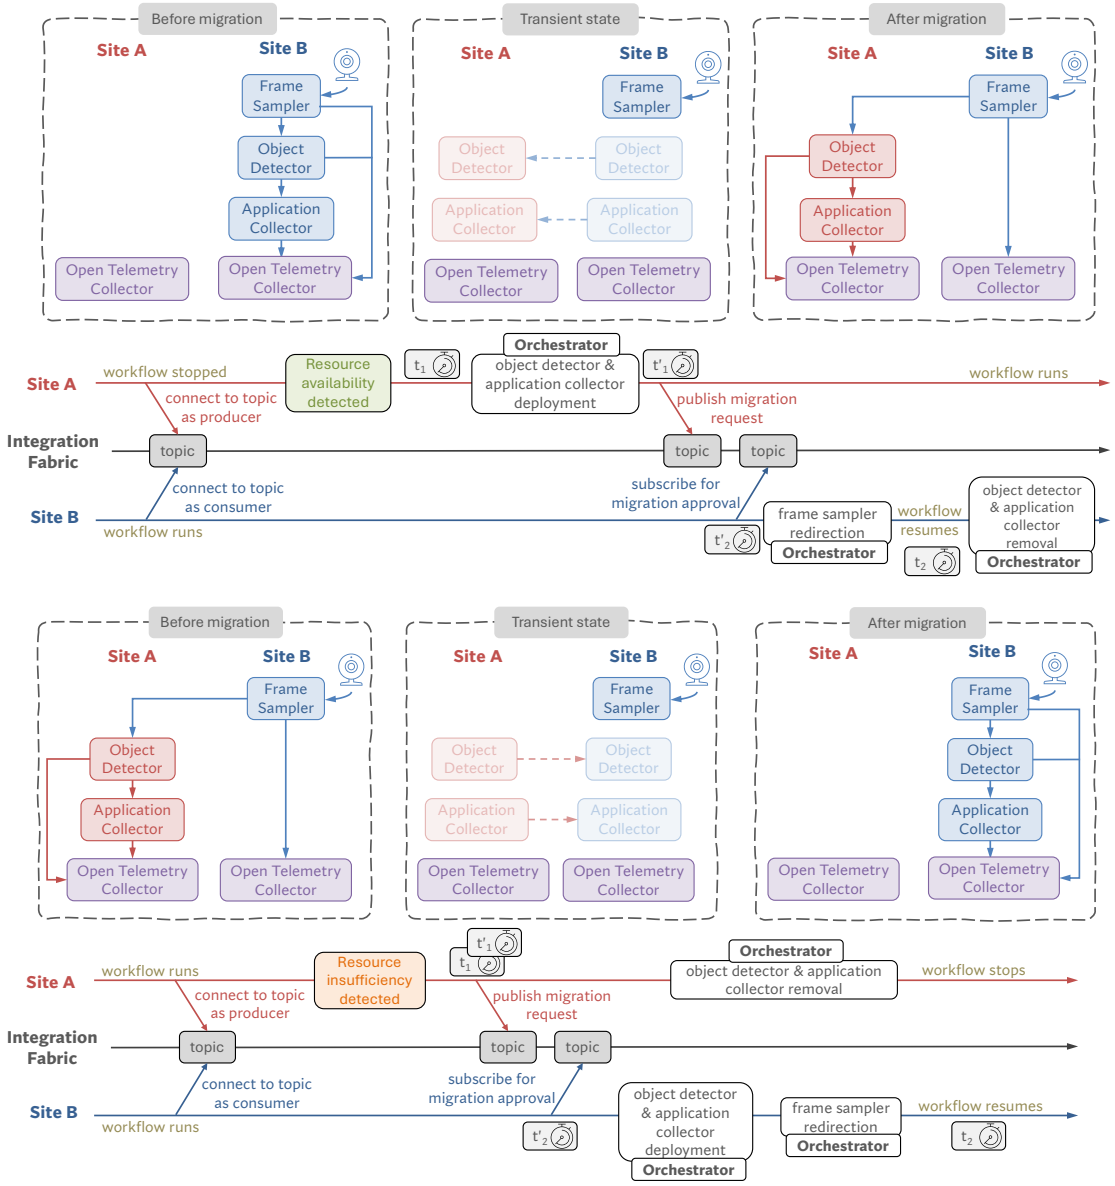

Intent-based Cobot Service Provisioning and Migration with Integrated Closed Loop

- Intent-Based service provisioning: Expanded in the cobot platform, featuring a digital twin for enhanced surveillance and human interaction.
- 5G-powered digital twin utilising real-time SLAM point cloud data to create a virtual representation of physical environments, enabling remote interaction and control.
- Cobots equipped with LiDAR and video cameras transmit data over 5G, while onboard edge computing and Kubernetes manage tasks.
- XR visualization: SLAM data is rendered into a digital twin accessible via XR headset, providing immersive feedback for monitoring and control.
- Predictive Closed-Loop (CL): Enhances reactive CL by incorporating AI reasoning, advanced monitoring, and AI-enabled analysis for automatic cobot service migration.
- Closed-Loop functions: Includes monitoring, analysis, decision, and execution functions, coordinated through MQTT messaging for service migration and reconfiguration.
- AI-enabled analysis: Adaptive Neuro-Fuzzy Inference System (ANFIS) used to predict the optimal time for cobot service migration based on battery levels, positions, and states.
- Enhanced monitoring: AI-based enhancements to monitoring, analysis, and decision functions enable proactive, data-driven cobot service migration.

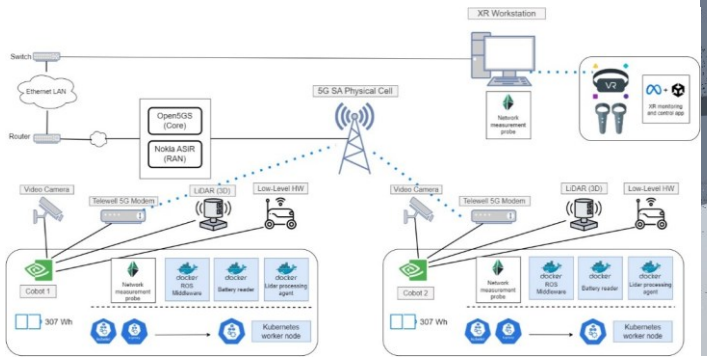

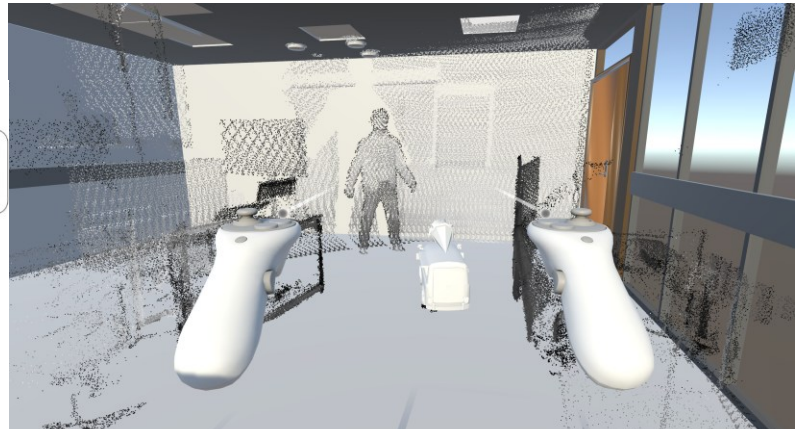

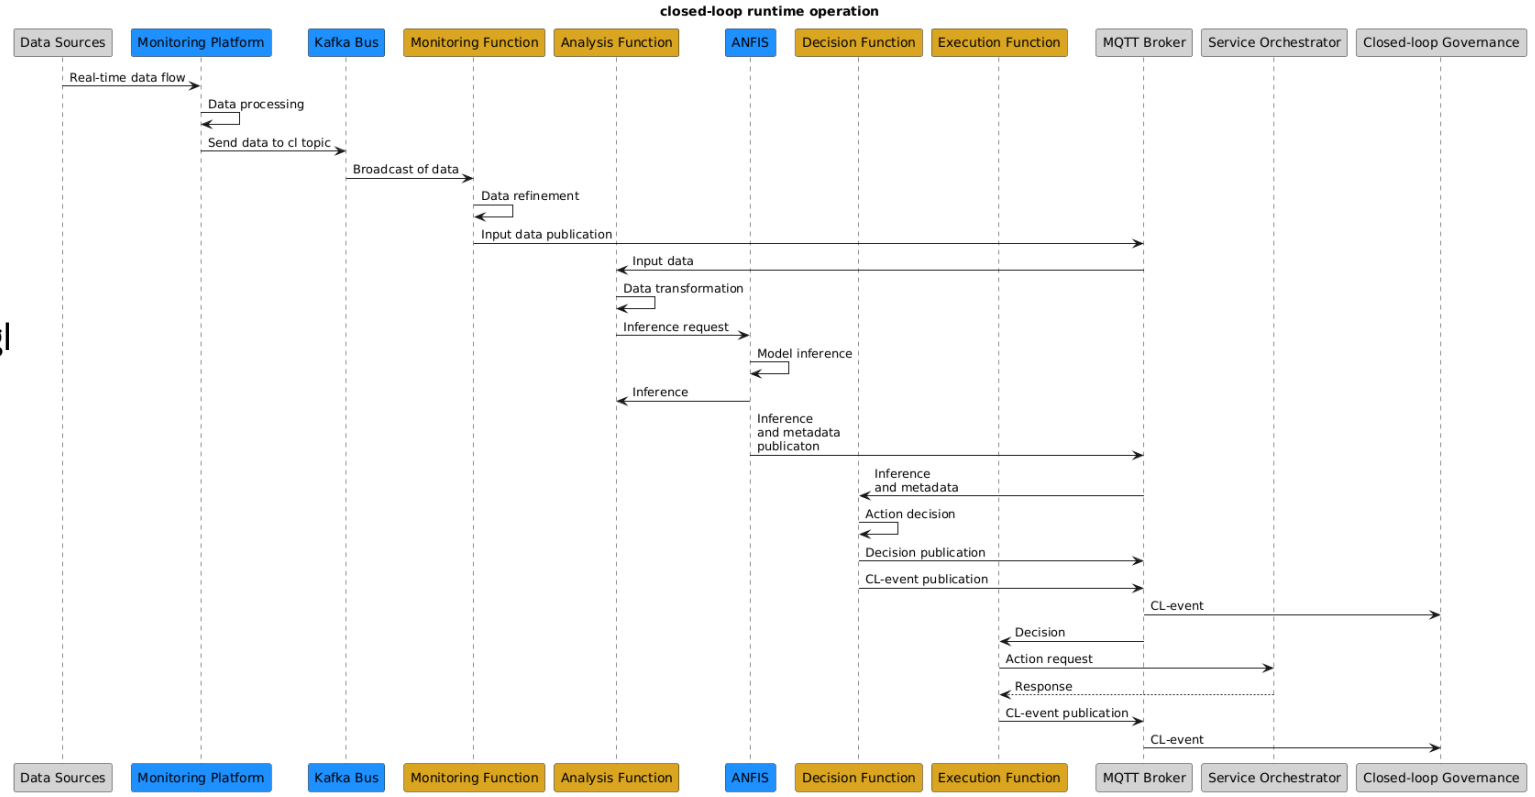

Components of the IBN-IME architecture

- IBN-IME architecture derived from DSM-IME ([HEX225-D25]) enables intent-based network management.

- Key components:
  - WebUI: User-facing interface for intent requests.
  - NBI: Filters inputs; directs requests to Interpreter (for NLP parsing) or Manager (for formatted intents).
  - Interpreter: Converts human-readable intents into 3GPP format using NLP.
  - Manager: Oversees intent lifecycle.
  - Intent Fulfilment: Conducts feasibility checks and triggers actions.
  - Intent Reporting: Updates status and logs data [28.312].
  - Intent DB: Stores all intent objects.
  - SBI-Connectors: Adapt data to appropriate formats for execution.

- Validated in two System-PoC scenarios:
  - *Zero-touch Cobot-based Video Surveillance (System-PoC #B and #C)*
  - *Synergetic Monitoring and Orchestration (System-PoC #C)*

- Demonstrates adaptability and consistent execution performance across use cases.

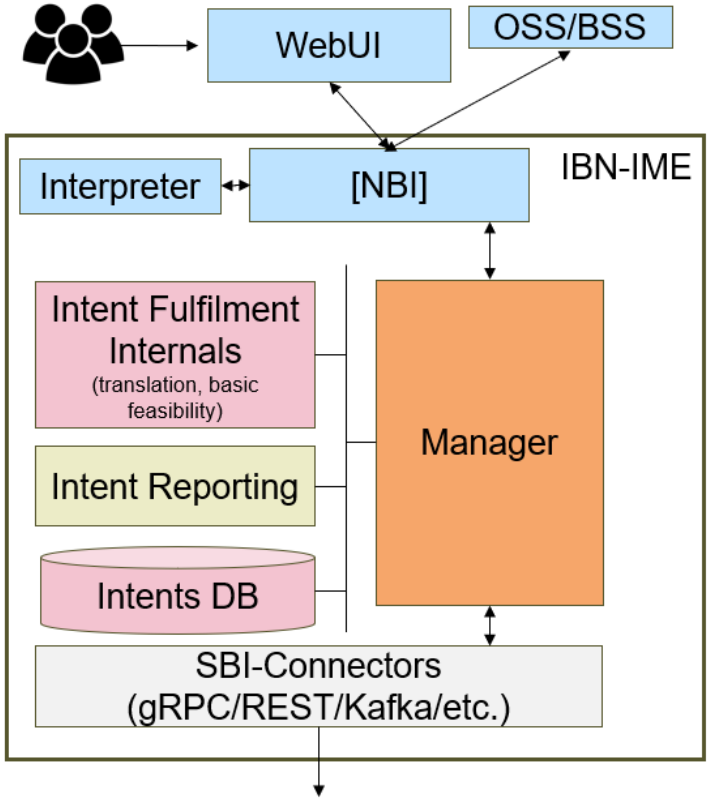

Intent-based service provisioning with Integration fabric mediation

Subscription to the Integration Fabric - Executed once

Service provisioning stage:

- Executed as many times as the number of service requested to be provisioned

- INB-IME fabric notified each time Service state is updated

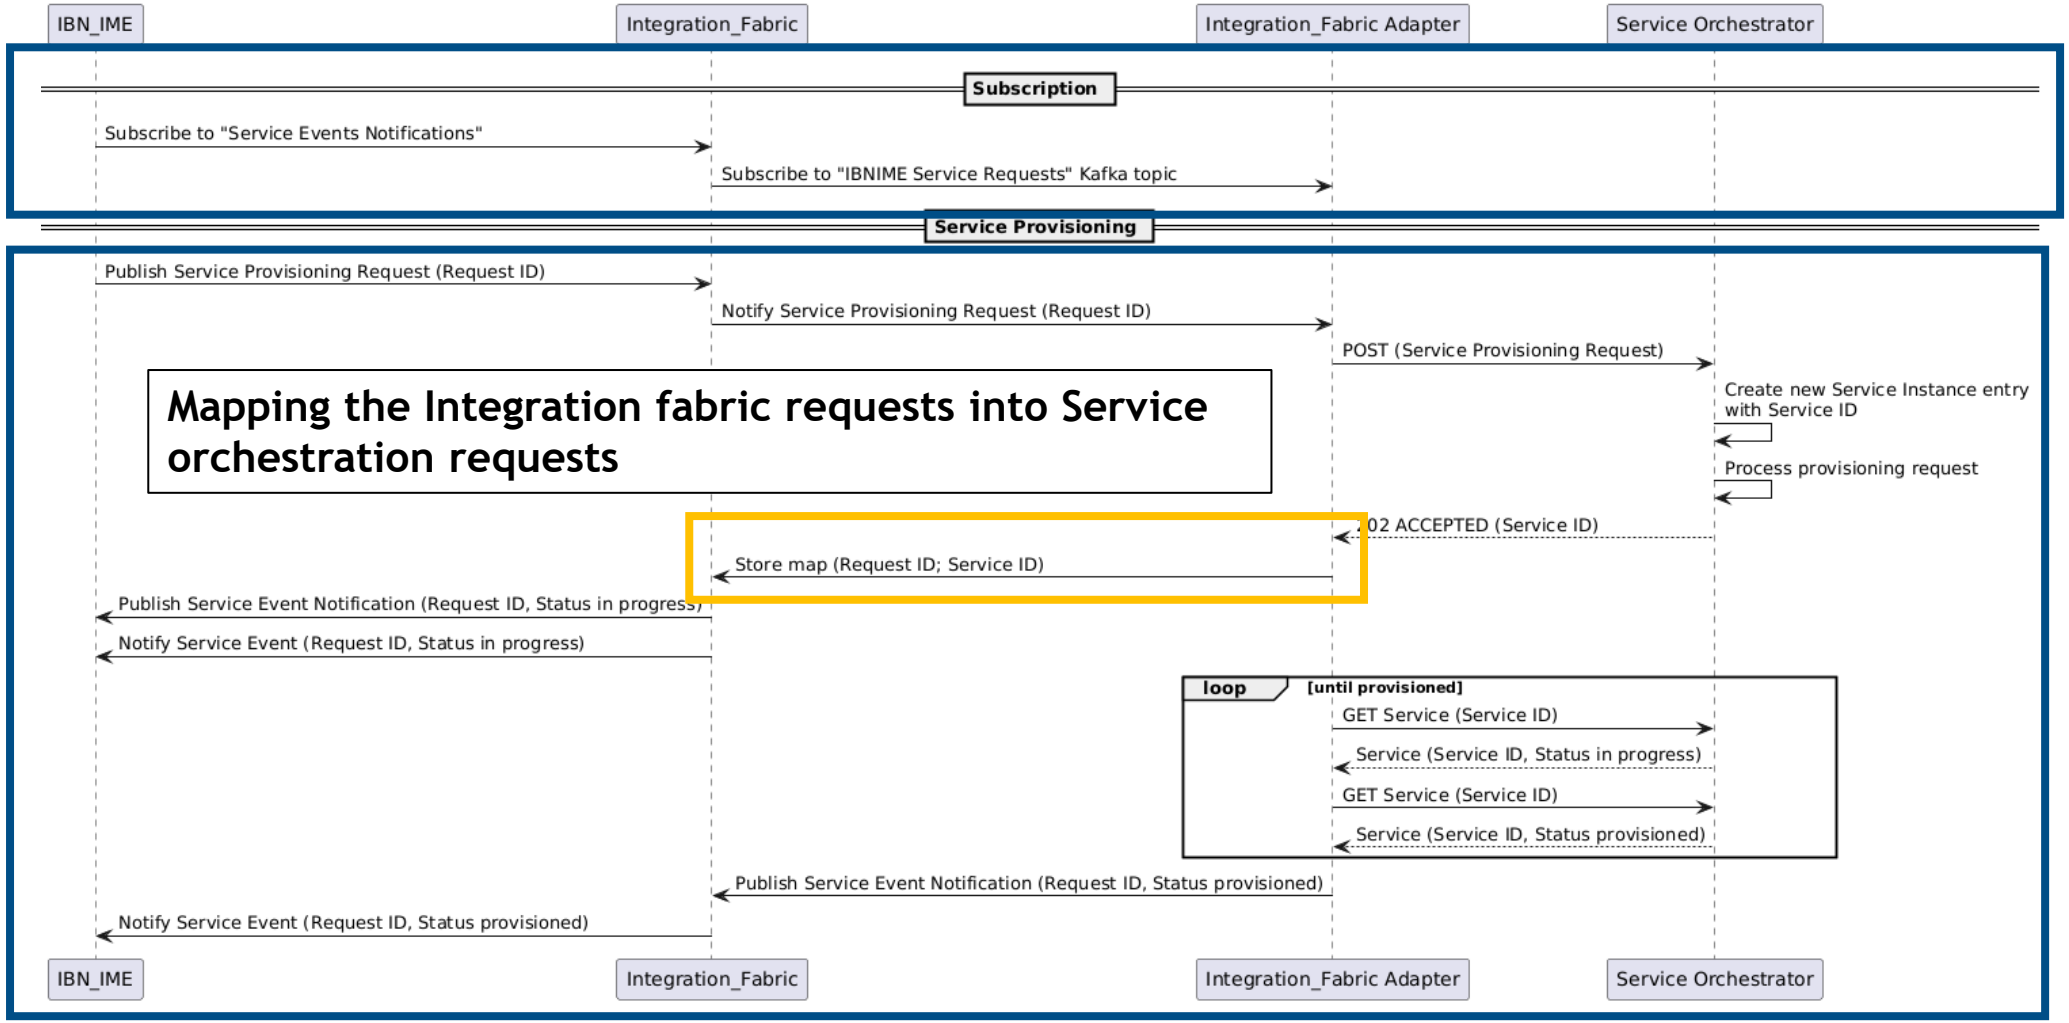

6G Radio & 6G devices Component PoCs

- **6G based sensing algorithms and concepts with real-time performance:** Demonstrates real-time 6G sensing algorithms to detect and localize both connected cobots and non-connected human workers, improving safety and efficiency in industrial environments.
- **AI-Native Air Interface:** Validates an AI-native air interface using end-to-end learned transmission without DMRS pilots, showing significant throughput gains and improved spectral efficiency in various mobility scenarios.
- **ML-based channel state feedback compression in a multi-vendor scenario:** Evaluates ML-based CSI feedback compression in a multi-vendor setup to enhance beamforming accuracy and reduce overhead without sharing proprietary models.
- **Flexible modulation and transceiver design:** Develops a dynamically reconfigurable transceiver platform that adapts radio parameters in real time, supporting flexible modulation and waveform selection for diverse 6G use cases.
- **Radio propagation measurements to collect data for radio channel modelling:** Conducts sub-THz radio propagation measurements in factory settings to build accurate channel models and coverage maps for network planning and performance prediction in industrial environments.
- **Crowd-detectable zero-energy devices:** Showcases zero-energy backscatter devices integrated into cellular networks for cobot positioning and logistics applications, enabling energy-efficient localization and asset tracking.
- **Energy harvesting IoT proximity devices:** Designed for cobot-powered warehouse inventory management, these devices use a solar panel and proximity sensor to detect shelf occupancy and trigger alerts, optimizing energy use and integrating real-time data processing for efficient operations.
- **End-to-End Extended Reality:** Implements scalable compute offloading for cobots and VR-controlled robots to reduce power consumption while ensuring real-time sensor data processing for navigation and safety.

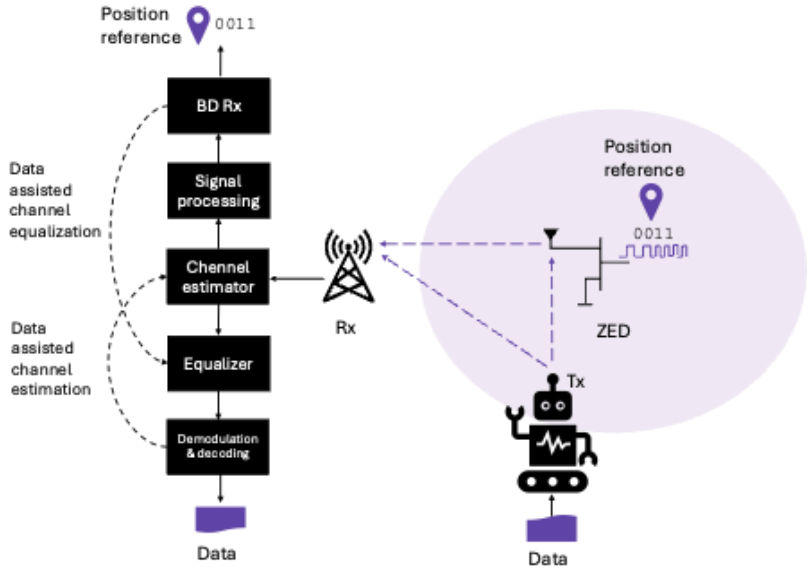

ZED assisted cobot positioning

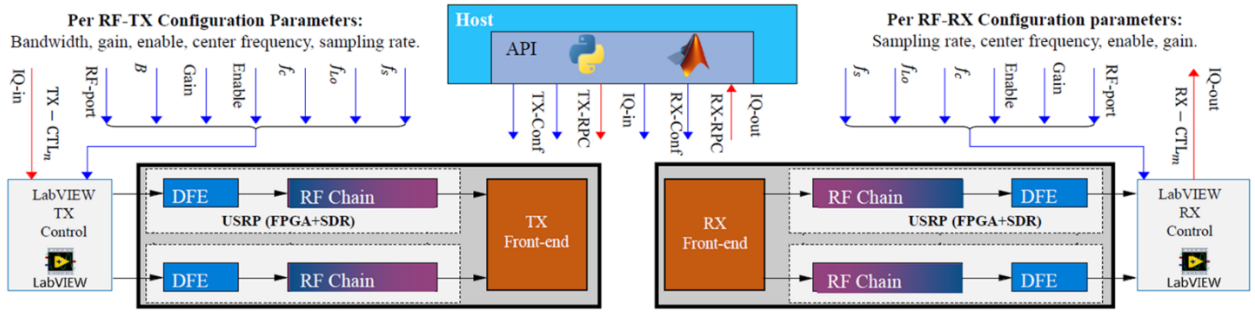

Flexible transceiver system schematic

6G Radio Components - 6G based sensing algorithms and concepts with real-time performance

- A bi-static system setup implemented in this POC where 5G NR waveforms were used for sensing human targets in an indoor scenario (office area).
- These functionalities are tested using our POC:
  - Communication, using adjusted 5G NR waveforms
  - Sensing targets in an indoor scenario, using same adjusted 5G NR waveform
  - Gesture recognition, using adjusted 5G NR waveform and training different targets using CNN in collaboration with IMEC

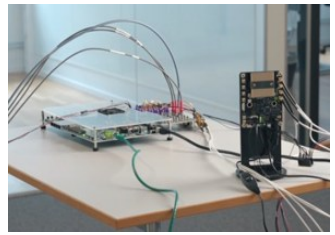

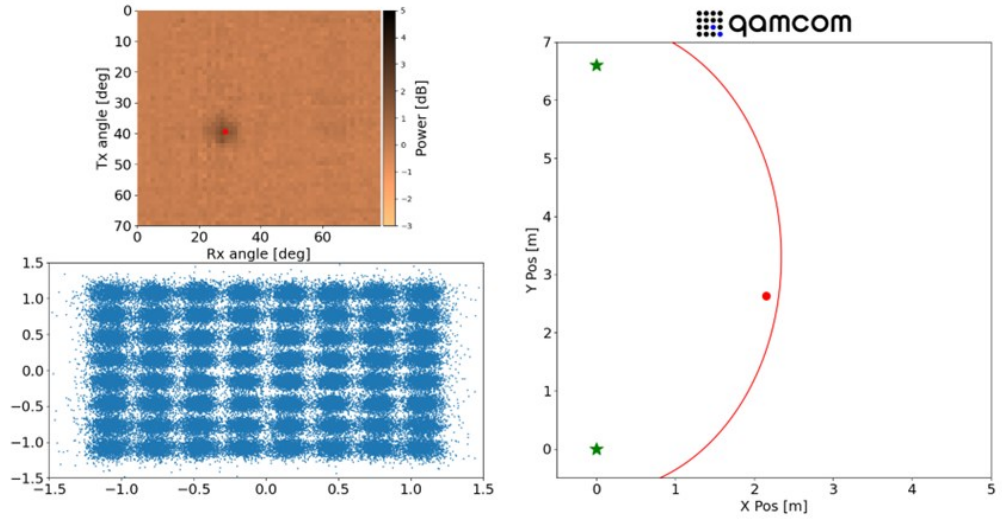

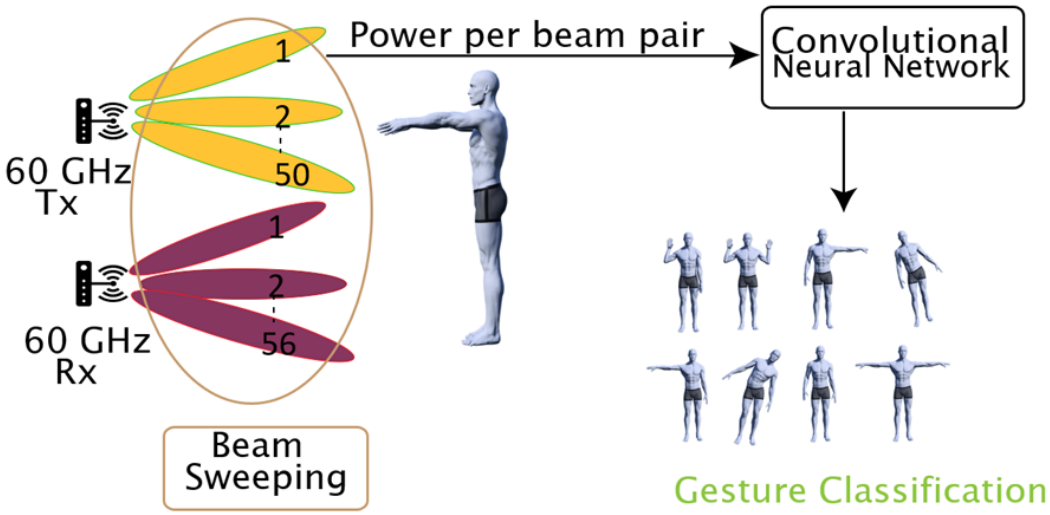

6G Radio Components - AI-Native Air Interface

- Objective of this PoC is to demonstrate the practical feasibility and performance gain of an end-to-end learned transmission link
- The PoC consists of an OFDM-transmitter utilizing a machine learned constellation, which is transmitting data to an ML-based receiver (DeepRx)
- Experiments were performed with channel emulator and indoor over-the-air transmissions
- Performance gains were demonstrated with the PoC measurements, ranging roughly between 15% and 25%

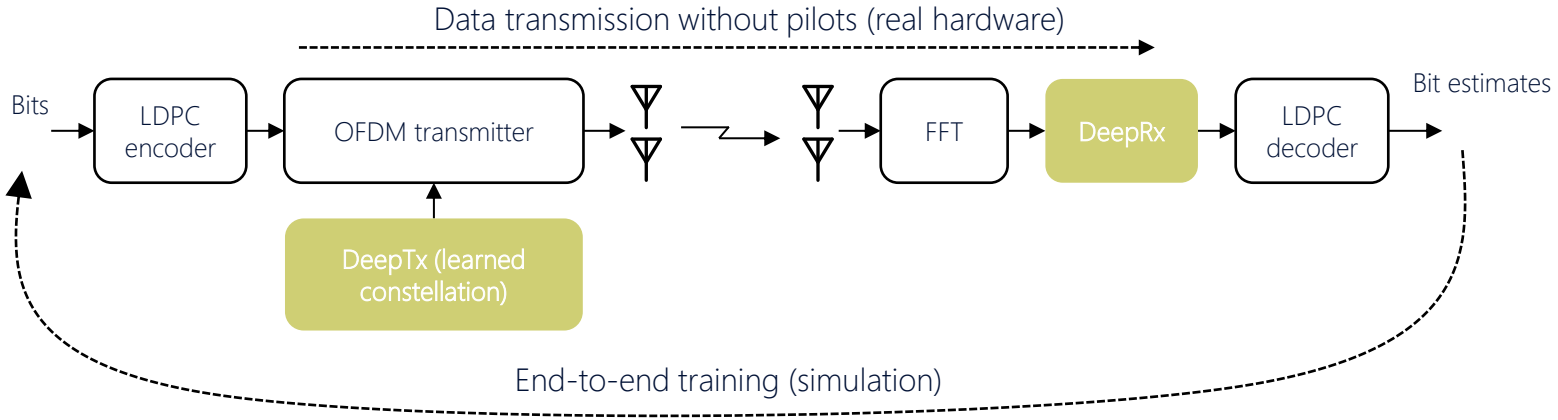

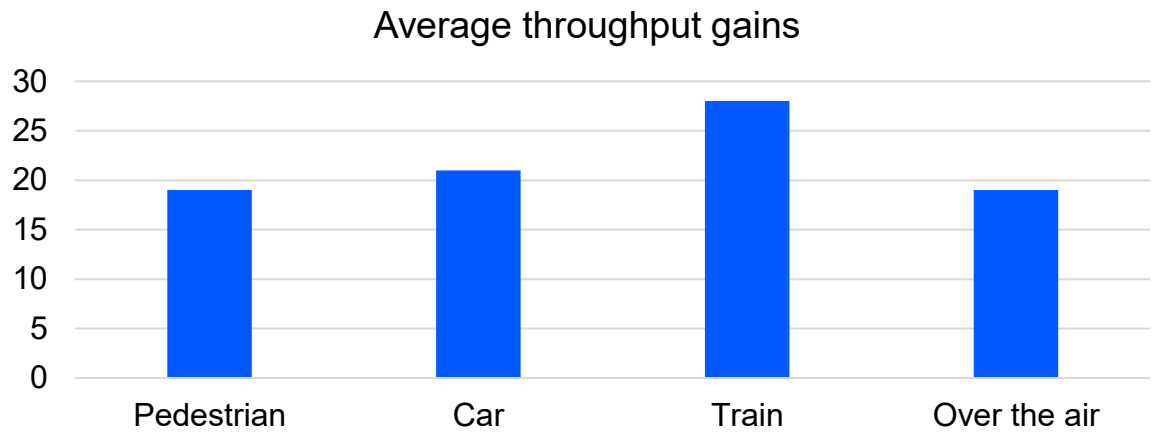

6G Radio Components - ML-based channel state feedback compression in a multi-vendor scenario

- The goal of this PoC is the demonstration that sequential-based training can provide improved channel state information at the network
- The AI model in the NW device is trained first and then the model in the UE is trained. Both devices share a common dataset to make sure that both models are trained on the same data. It does not need devices to share AI model
- OTA data collected has been collected over a 5G test network operating at 3.5GHz, with a Nokia gNB with 64 antennas and 8 antennas ports and Qualcomm mobile test device with 4 antennas
- OTA data were collected in different environment, in both indoor and outdoor and varying mobility, static, pedestrian and vehicular.
- Synthetic models have also been trained and evaluated on outdoor and indoor OTA environment

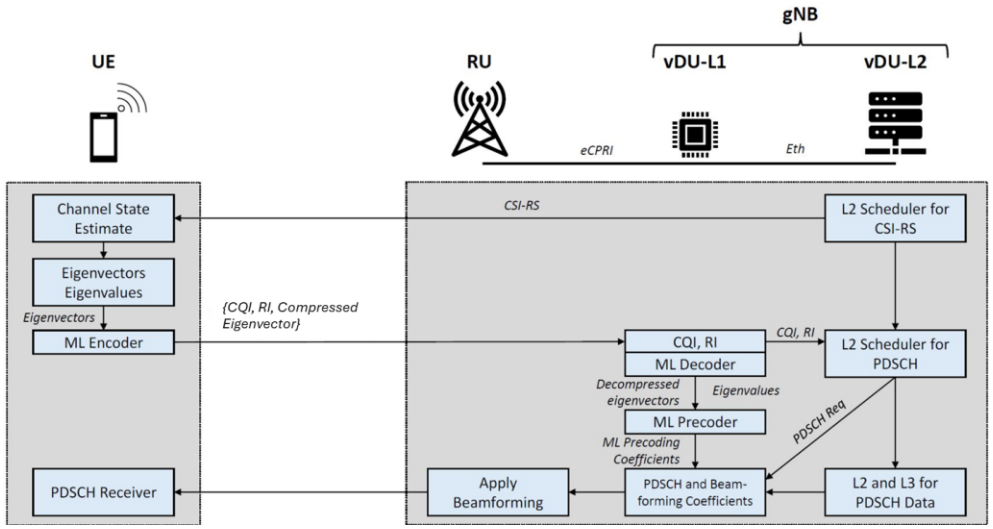

Figure: Block diagram of the implement OTA ML-based CSF testbed.

6G Radio Components - Flexible modulation and transceiver design

- Built on software-defined radio (SDR) for dynamic reconfiguration, enabling experimental testing of mobile communication algorithms beyond static setups
- Supports real-time RF parameter adjustments (e.g., frequency, bandwidth, modulation) via Python/MATLAB API, with independent TX/RX control for beamforming and carrier aggregation.
- Operational prototype includes phase calibration for angle estimation; dynamic parameter adjustment via RPCs remains under refinement
- Scalable platform for mobile communication research, detailed in, facilitating performance evaluation in complex scenarios

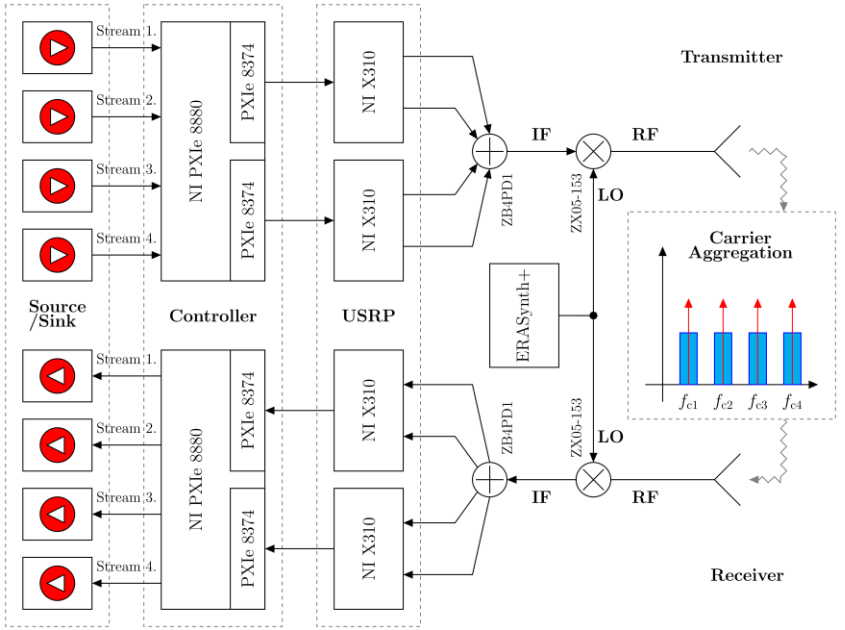

Scematic of flexible transceiver system used for carrier aggregation



Flex. Transceiver, video streaming at 10 GHz, demo at 6G Summit dresden 2025

6G Radio Components - Radio propagation measurements to collect data for radio channel modelling

- Radio propagation measurements: Collects radio-channel data in cobot-equipped factory environments with multipath richness and blockages.
- Sub-THz measurements: Path-loss and delay-spread profiles for factory halls and warehouses support throughput and latency forecasts.
- Virtual sensing: Data supports cobot applications needing environmental awareness and adaptive control.
- Radio coverage mapping: Identifies signal strength and dead zones for optimal access point placement.
- Integration with E2E 6G: Empirical data enhances adaptive radio control and predictive link management in the infrastructure layer.
- APIs and coverage maps: Expose site-specific models for dynamic service orchestration.
- Pervasive capabilities: Supports AI orchestration, intent-based management, and closed-loop resilience.
- Challenges: Includes capturing multipath dynamics, ensuring calibration, and generalizing profiles.
- Next steps: Broaden measurements, automate coverage map generation, and refine parameter-extraction methods.

6G Devices Components in PoCs - Crowd-detectable zero-energy devices

- Crowd-detectable Zero-Energy Devices: Backscatter-based IoT devices powered by ambient light, modulating signals from mobile systems without hardware modifications.
- Signal modulation: Devices use frequency shift keying (FSK) for encoding data, distinguishing signals from natural multipath components in the Doppler frequency domain.
- PoC implementation: Based on LTE, using software-defined radios (SDR) to detect and filter BD signals during downlink and uplink communication.
- Enables proximity-based indoor localization and object identification for cobots, similar to BLE beacons but with zero energy consumption.
- Provides low-data-rate AIoT services using standard cellular signals, eliminating the need for separate carrier waves, improving spectral efficiency.
- Uses bi-static configuration, enhancing spectral and energy efficiency, and aligns with 3GPP AIoT standards for cellular systems.
- Challenges identified:
  - Limited BD reading range, which could be improved with Forward Error Correction (FEC).
  - Need for firmware updates in standard UE modems to support BD signal demodulation.
  - Differences in reference signal availability between LTE and 5G NR, especially for BD illumination.
  - Synchronization issues and potential interference from unsynchronized BD signals.
- Next steps: Improve detection range with FEC, address synchronization challenges, explore solutions for 5G NR reference signals, and integrate BD technology with core network functions.

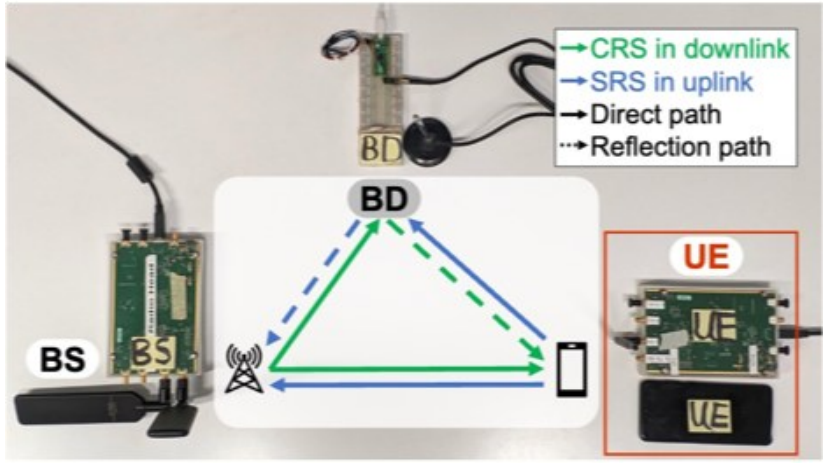

Operation principle of the crowd-detectable zero-energy devices integrated to LTE network.

6G Devices Components in PoCs - Energy harvesting IoT proximity devices

- Developed for the cobot-powered inventory management warehouse PoC to enhance real-time inventory operations using occupancy detection.
- Equipped with a proximity sensor to detect the occupancy of warehouse shelves and trigger alerts for cobot inspection when a box or item is detected.
- Includes a 9V, 3W solar panel, generating up to 300 mA under ideal sunlight, with real-world conditions providing less than 150 mA.
- Features a 2400 mAh LiPo battery for energy storage and a Sunny Buddy solar charger for managing battery charging.
- Incorporates a 4.2V to 3.3V DC-DC converter, drawing approximately 9mA during both sleep and normal operation modes.
- Ultrasonic proximity sensor: Utilizes a DFRobot A02YYUW ultrasonic sensor with a range from 3 cm to 400 cm for accurate occupancy detection.
- Microcontroller and communication module: Controlled by an STM32 microcontroller and communicates wirelessly via an ESP32 C6 Wi-Fi module.
- Protected by an IP67-rated enclosure, ensuring durability in harsh warehouse conditions (dust, water, etc.).

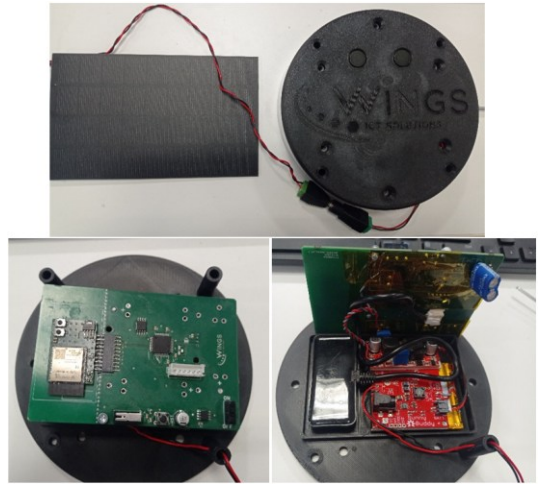

Warehouse inventory management scenario new EH devices for advanced inventory operations based on real-time occupancy information.

Component	Power Analysis	Time Active per Cycle	Notes
Solar Panel (9 V, 3 W)	Produces up to 300 mA (theoretical)	Continuous (daytime)	Real output <150 mA typically
Solar Panel Under Lamp	Produces ~27-30 mA +	Continuous (lamp on)	Lamp above the panel at 10 cm
Sunny Buddy Charger	Consumes 100 µA -	Continuous	Manages LiPo charging
2400 mAh LiPo Battery	N/A	Continuous	Stores energy
DC-DC Converter (4.2 V→3.3 V)	Consumes ~9 mA in sleep mode -	Continuous	Dominates sleep consumption
STM32 MCU	Consumes ~30 mA Consumes (Active), - <0.5 mA (Sleep) -	2 s every 30 s cycle 28 s every 30 s cycle	Data acquisition & control
ESP Wi-Fi Module	Consumes ~50 mA - (Transmission)	~ 45 s-1 m every change of state	High consumption for sending data
Ultrasonic Sensor (A02YYUW)	Consumes ~30 mA with STM during measurement	2 s every 30 s	Distance measurement

6G Devices Components in PoCs - End-to-End Extended Reality

- This PoC demonstrates the benefits of exploiting L4S and adaptive rate control at the UE side for improving user experience with XR devices in a distributed spatial compute framework
- It shows that distributed spatial compute, by offloading heavy processing to remote devices, can enables high quality rendering and reduced power constraints at the XR device
- This PoC also demonstrate that exploiting network congestion information and adaptive rate control at the UE side can provide enhanced user experience.

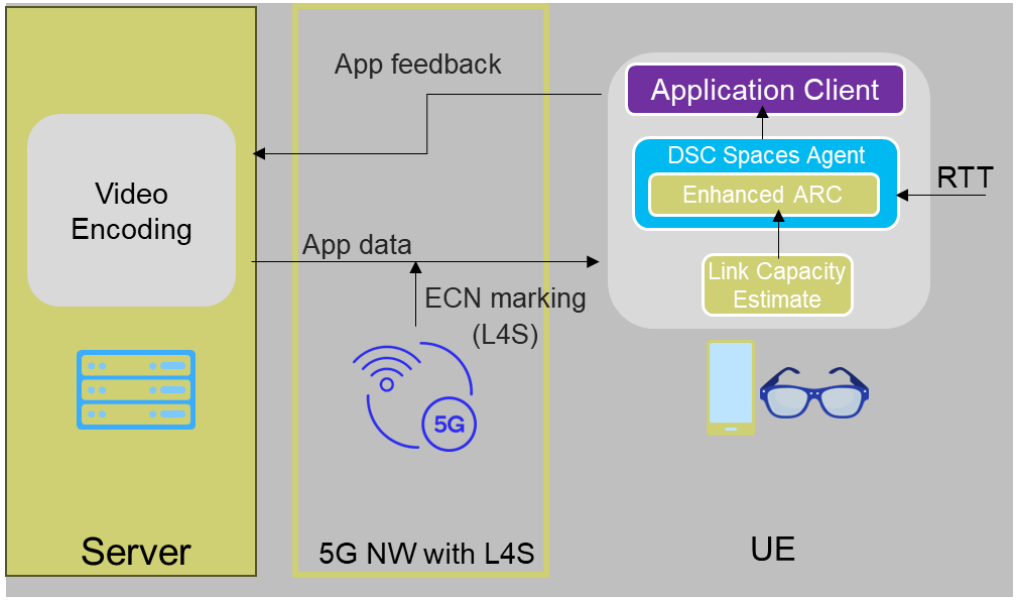

The testbed comprises of a user with an AR glass device and a phone running the client application, and with a 5G connection to a gNB and a PC running the server-side application.



Chapter 4
System- and Component-PoC #C validation and evaluation results

System-PoC #A key results

- FA Mechanism is a key component of System-PoC #A designed to optimize computational workload placement across robotic units, edge servers, and cloud resources, focusing on energy efficiency while maintaining performance.
  - It analyzes input parameters such as workload requirements, compute node capabilities, and edge device data to generate an optimized placement plan, enforced by an orchestrator.
  - The system triggers the FA mechanism dynamically based on conditions like increased latency, high power consumption, or hardware malfunctions in robotic units.
  - The FA mechanism employs a metaheuristic algorithm that is scalable and suitable for large-scale experiments.
- Validation results:
  - In tests involving seven compute nodes (3 robotic units, 2 edge servers, 2 cloud servers) and workloads ranging from 4 to 28 tasks, the FA mechanism outperformed baseline approaches (Round-Robin Placement and Feasible Random Placement) in energy efficiency.
  - The FA mechanism achieved up to 28.6% power consumption reductions, with larger gains observed as the workload increased.
- The Intent-Based Network-Intent Management Entity (IBM-IME) solution was integrated into two PoC scenarios: “Synergetic Monitoring and Orchestration” and “Intent-based Service Provisioning with Integrated Closed Loop Deployment and Reactivity for a Cobot Service Migration”.
  - The IBM-IME solution incorporated Generative AI (GenAI) to improve the efficiency of the entire DSM-IME system, enhancing its performance across multiple domains.

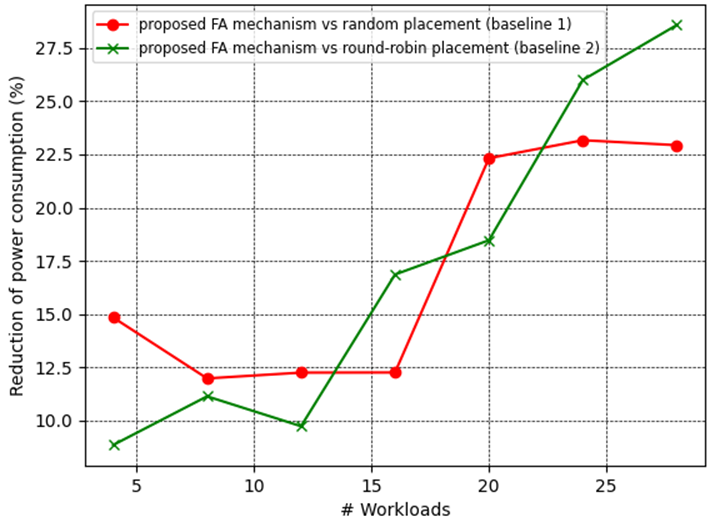

Reduction of power consumption with increasing number of compute workloads of the FA mechanism compared with two baseline algorithms

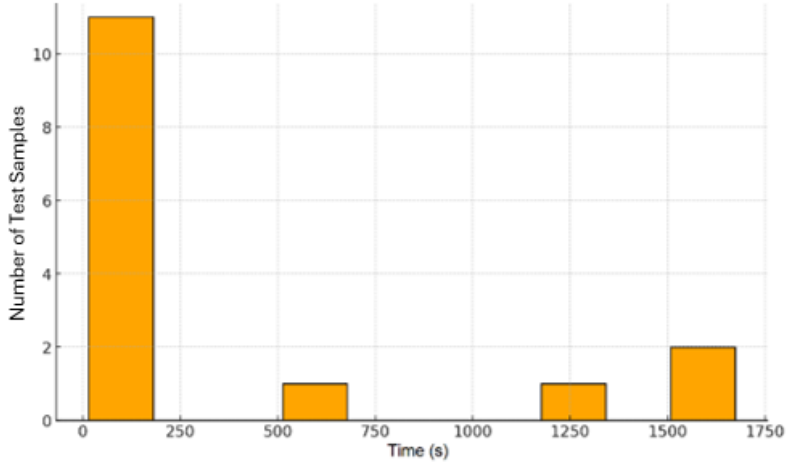

Latency to deploy the intent-based service request

System-PoC #B key results - Component-PoC#B.1

- Evaluated orchestration mechanisms for latency-sensitive applications across multiple clusters using a Karmada Multi-cluster testbed and AI-assisted service management.
- Experiment setup: Tested horizontal scaling across Cloud and Edge clusters; Cloud reduced compute latency by 10% but added 15ms communication delay, while Edge suffered major latency increases due to CPU constraints.
- Orchestration mechanism: Implemented a centralized, RL-based autoscaling and migration system focused on minimizing resource use while meeting E2E latency SLA requirements.
- Key insight: Effective orchestration across heterogeneous clusters improves service efficiency by leveraging trade-offs in location-based performance characteristics.

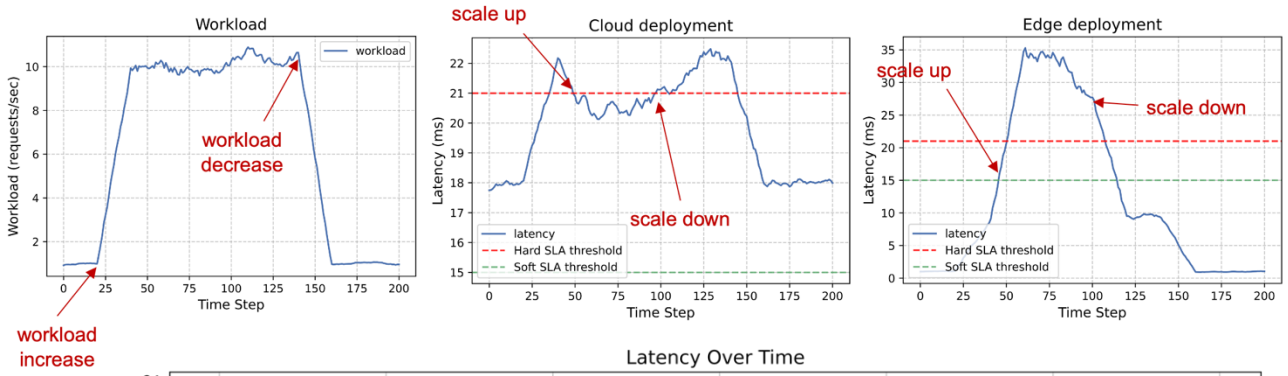

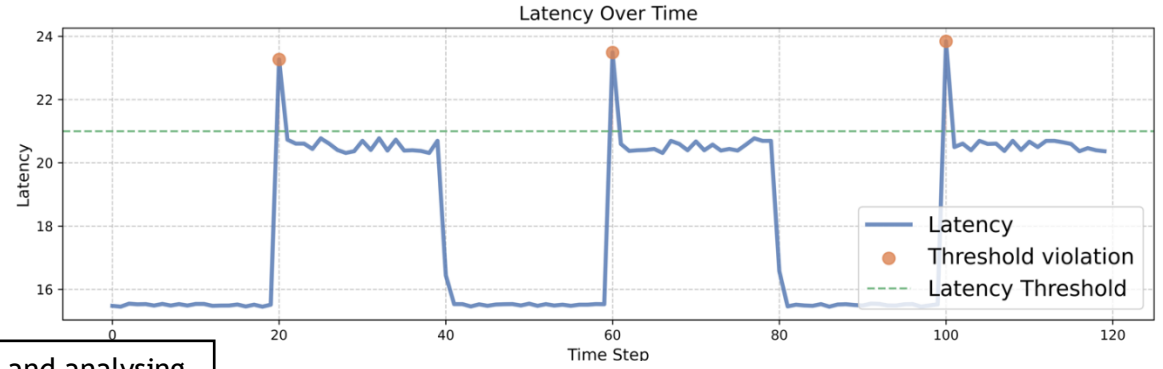

With the multi-stage zero-touch closed-loop mechanism, and relying on the battery cobot's data analysis the cobot video streaming surveillance service is migrated from robo-pi1 (cobot1) to robo-pi2 (cobot2)

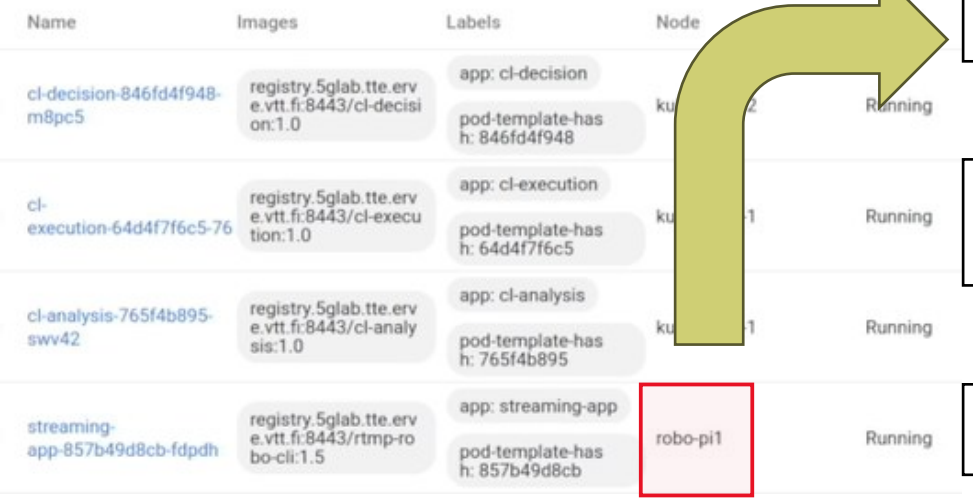

Analysis Stage: Collecting and analysing battery metrics

Decision Stage: based on the metrics collected, it performs the migration decision

Execution Stage: it executes the service migration request

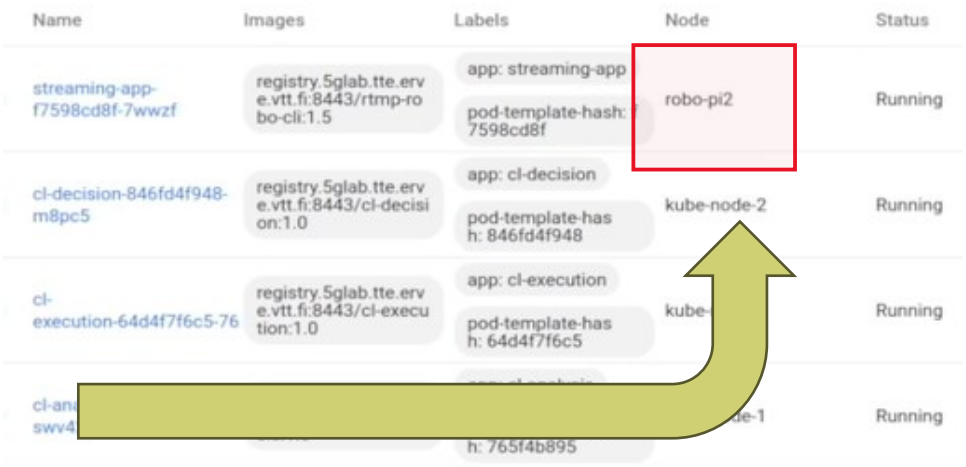

System-PoC #B key results - Component PoC#B.2

- Input, generalization, and output nodes deployed as isolated Kubernetes pods to emulate local datasets and workloads(1,1)
- Increasing accuracy for video bitrate and delay estimation simultaneously via multi-task learning (1,2)
- Offloading: model layers were dynamically offloaded between the generalization and output (1,2)
  - computational overhead decreased at the output nodes, reducing energy consumption from 22 to 5 J as given in (2,2),
  - a temporary impact on the accuracy, at t~20 s was observed due to the aggregation of model layers, which then started increasing again over the training rounds and stabilized.(5,2)
  - The offload from generalization node to the output nodes does not have impact on the accuracy (at t~40 s) as the overall neural network weights are kept unchanged and are just relocated
  - Energy reduction is mainly expected to be at the device during the offloading (3,2)
  - The model layers that are offloaded to the network increases the energy consumption at the network side (3,1)
- Scalability when offloading from many devices/applications to network: Impact on network energy consumption at the network can be reduced with the model layer aggregation of ML layers in the network (5,2)

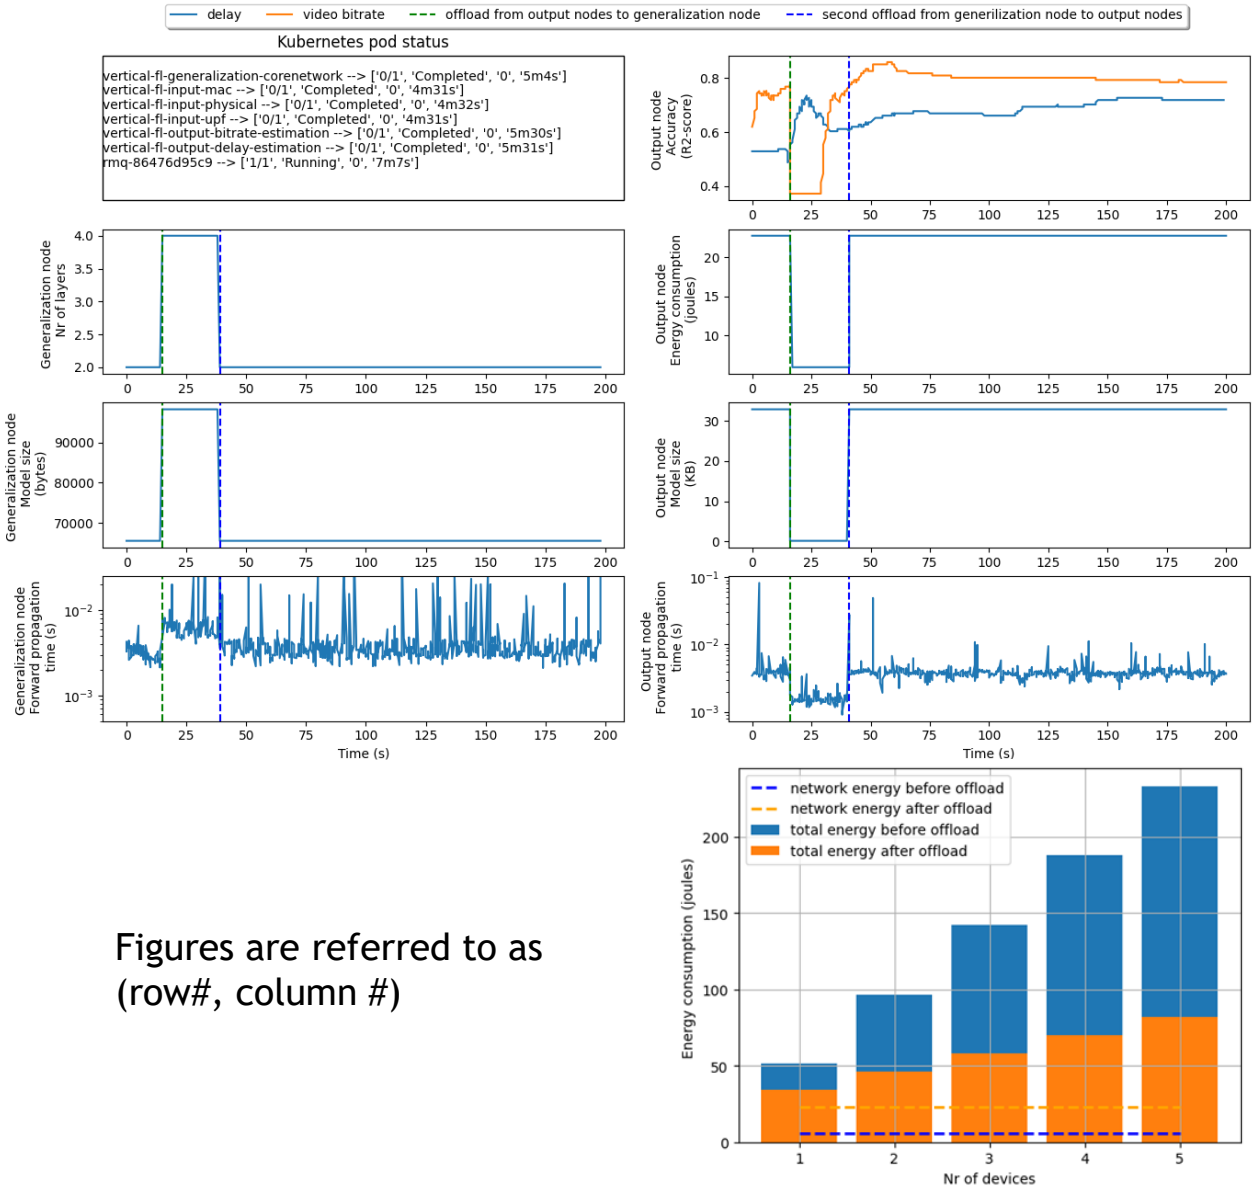

Figures are referred to as (row#, column #)

System-PoC #B key results - Component PoC#B.3

- System-PoC #B Warehouse inventory management showcased:
  - Advanced management & orchestration features,
  - Business intent-driven resource allocation,
  - Novel architectural enablers like trustworthy flexible topologies and beyond communication enablers.
- Trustworthy Flexible Topologies: Enabled dynamic use of temporary communication and computational resources, with a UAV deployed when the ground cobot was out of coverage.
- Beyond communications: Incorporated advanced cloud-based services (compute-as-a-service, AI-as-a-service) for real-time resource availability, used by the FA algorithm to optimize placement.
- Integration Fabric: Ensured seamless interaction and coordination between systems and services, including foundational devices and computational resources.
- Pervasive functionalities: Included advanced management and orchestration capabilities (intent-based management, closed-loop controls, multi-platform orchestration) and trust evaluation mechanisms in the security framework.
- Software-Defined Networking (SDN): Introduced SDN capabilities to manage the transport domain, enhancing system flexibility and efficiency.
- Key results:
  - Gains in trustworthiness and energy consumption through the FA algorithm.
  - Flexible topology node (FTN) energy consumption.
  - Comparison of on-device (UAV/AMR) vs. offloading performance in terms of inference times and energy consumption.

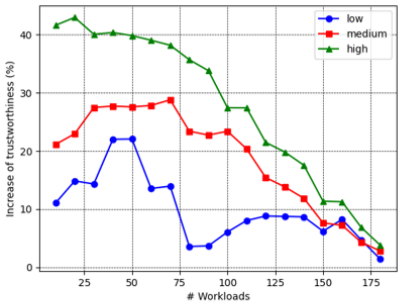

Conf. #	Setup
1	Transmission Load: 2Mbps, CPU Utilization: 22%, CPU cores: 7, Battery Autonomy: 0.32h
2	Transmission Load: 100Mbps, CPU Utilization: 22%, CPU cores: 7, Battery Autonomy: 028h
3	Transmission Load: 0Mbps, CPU Utilization: 60%, CPU cores: 7, Battery Autonomy: 0.30h

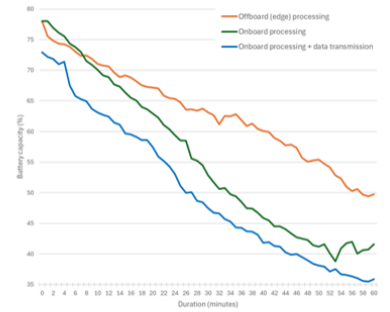

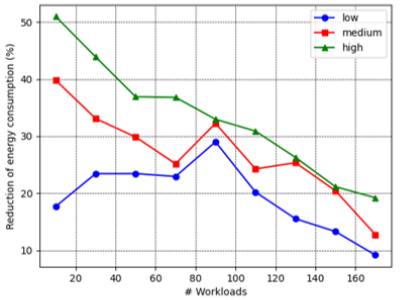

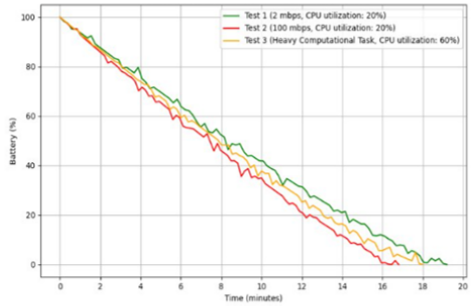



(i) (ii) (iii)

(i) Trustworthiness measurements (top) and energy consumption (bottom) of the FA algorithm; (ii) UAV (flexible topology node) battery depletion under different configurations; (iii – top) Power consumption comparison of on-device (AMR/UAV) vs edge computing setup; (iii-bottom) Performance benefits of offloading computation via 5G on CPU utilisation

System-PoC #C - Integration fabric general analysis on performance and impact

- Scalable, secure communication:
  Integration Fabric enables efficient, scalable, and secure messaging in distributed systems, with moderate overhead that needs careful management as system complexity increases.
- Modern, event-driven architecture:
  It shifts from traditional client-server or RPC models to an event-driven approach, making integration easier and allowing broader, more flexible connectivity between services.
- Democratized, encrypted data exchange:
  By using topic-based messaging, it democratizes data communication, improves interoperability, and ensures security through encrypted channels and access controls.
- Loosely coupled, reliable systems:
  The architecture promotes loosely coupled components that autonomously trigger or respond to events, reducing central dependencies and handling dynamic user activity with high reliability.
- Consistent performance & scalability:
  The system maintains strong throughput and stable latency—even under heavy load—by efficiently batching messages and aligning producer-consumer speeds, though some latency spikes during peak events indicate areas for further optimization.

Route	Avg Latency	Median Latency	P95 Latency	P99 Latency	Throughput	Error Rate	Availability
register	1.864 s	1.838 s	1.883 s	2.582 s	0.54 req/s	0.00%	100.00%
registerNewTopic	1.663 s	1.637 s	1.846 s	1.889 s	0.60 req/s	0.00%	100.00%
deleteTopic	3.776 s	3.653 s	4.745 s	4.823 s	0.26 req/s	0.00%	100.00%
unregister	3.476 s	3.471 s	3.753 s	3.785 s	0.29 req/s	0.00%	100.00%

Entry	Message Metrics	Latency Avg	Latency P50	Latency P95	Errors	Error Rate	Message Rates	Duration
Producer	Sent: 2783838 messages Delivered: 2394616 messages	0.2348 s	0.0434 s	0.7323 s	0	0.00 err/sec	Send Rate: 15651.67 msg/sec Delivery Rate: 13463.87 msg/sec	152.22 s
Consumer	Received: 2936826 messages	0.5438 s	0.5005 s	1.1895 s	0	0.00 err/sec	Receive Rate: 14839.73 msg/sec	197.99 s

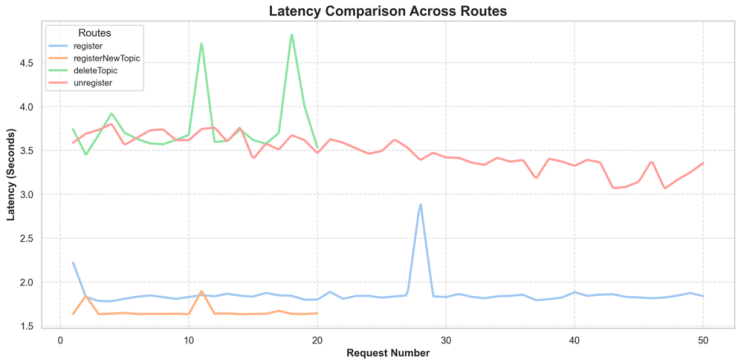

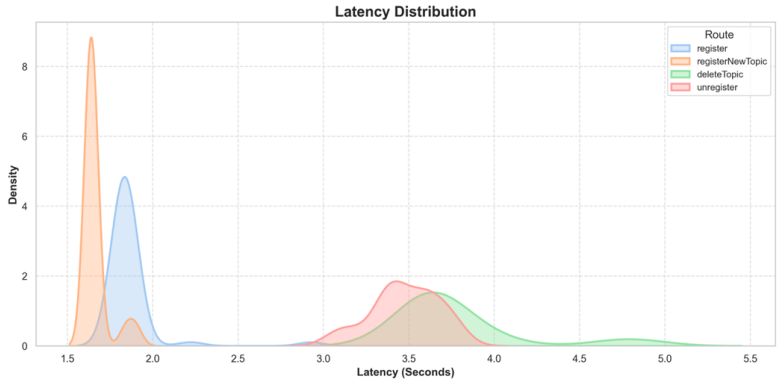

System-PoC #C - Cobot-powered Warehouse Inventory Management scenario leveraging BCS exposure & Trustworthy Flexible Topologies

- Warehouse inventory management scenario: Combines cobots, UAVs, and dynamic network components for real-time operational efficiency and adaptive inventory management.
- Uses trustworthy flexible topologies (Flextop) and BCS exposure for seamless coverage and connectivity, even in the event of disruptions.
- Task planner generates optimised task lists, and Fleet Manager coordinates AMRs and UAVs to ensure efficient deployment and real-time task management.
- CAMARA-compliant APIs: Exposed Network Functions (ENFs) like Sensing/Localisation, Device Status/Location, and Quality-on-Demand enable real-time decision-making and data offloading to edge/cloud servers.
- Performance evaluation:
  - Novel functional capabilities successfully validated leveraging exposure APIs
  - Flextop UAVs and ground AMRs evaluated in different configurations; Flextop UAV-extended links showed broader coverage with slightly higher latency and fluctuating performance.
  - Connection re-establishment time was considerably lower in the 5G SA mode comparing to non-3GPP technologies solution

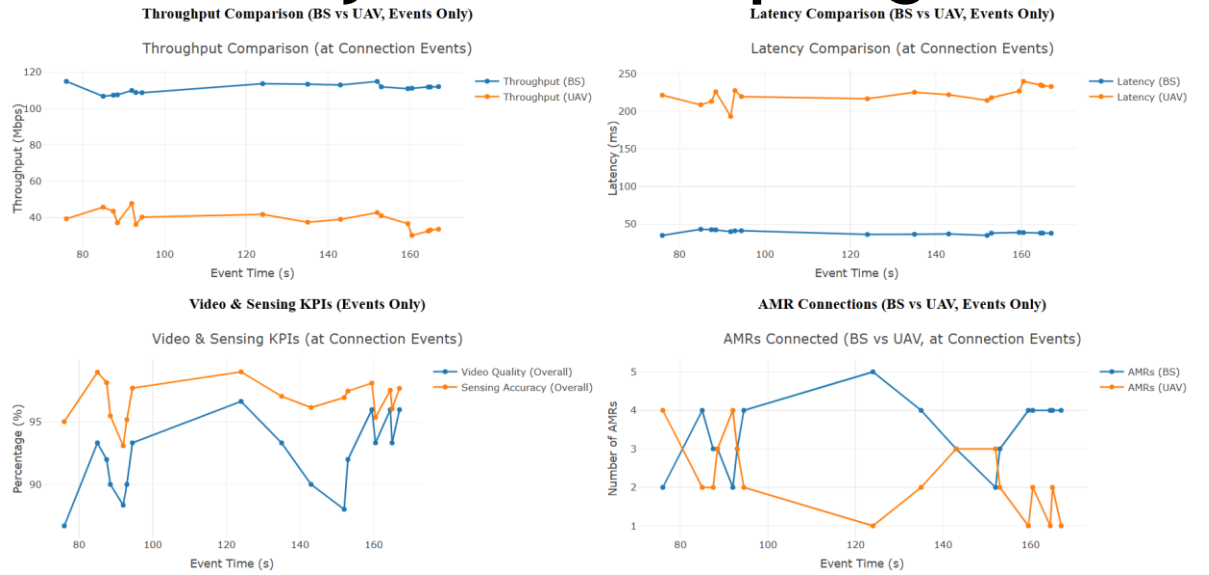

Performance Evaluation of Warehouse Inventory Management Architecture with Ground AMRs and Flextop UAVs

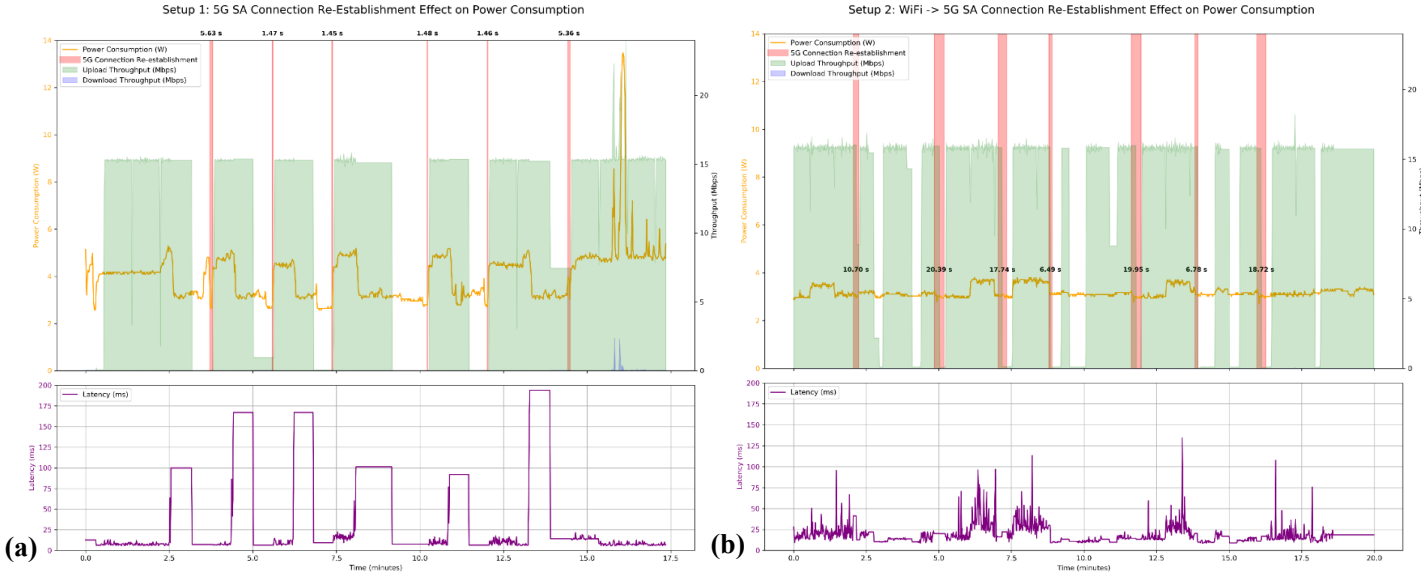

Power and latency during: (a) periodic coverage loss in direct 5G SA connection and (b) Wi-Fi handovers with 5G SA backhaul.

System-PoC #C - Intent-based Service Provisioning with Integrated Closed Loop Deployment and Reactivity for a Cobot Service Migration

- Autonomous Mobile Vehicles (AMV) applications require detailed environmental information often provided by LiDAR sensors.
- Results show Lidar Uplink Traffic measurements using a 5G SA network. The measurement ran for 5 minutes while autonomous robot patrolled the premises.
- Measurements show the performance in an indoor environment. The results benchmark the network requirements for live lidar data transfer
- Predictive, AI-based and zero-touch Closed-loop for the automated migration of the cobot surveillance service

Parameter	Value
Median Load	59 Mbps
Mean Load	58 Mbps
Median Delay	35 ms
Mean Delay	38 ms

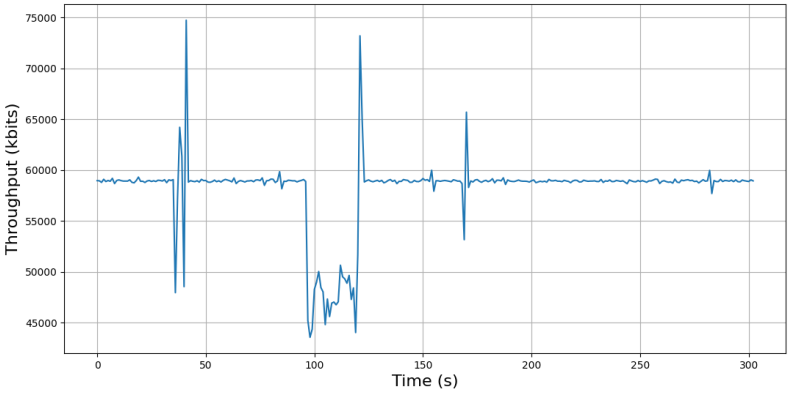

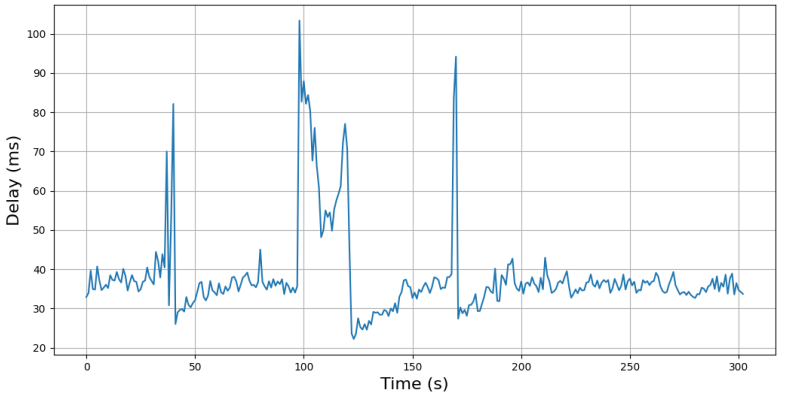

Analysis stage: analyse collected data and infer cobot battery draining

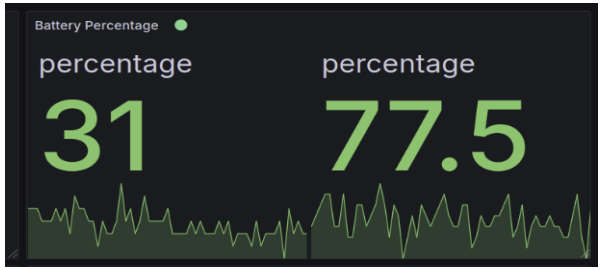

Decision stage: decide to which cobot migrate the surveillance service to (robo-pi2)

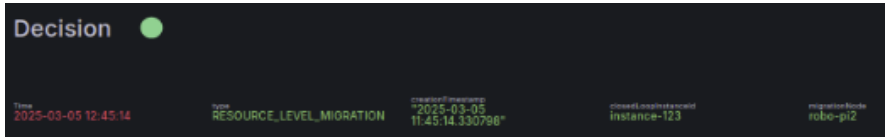

Execution stage: performing the video surveillance service migrration process to a charged cobot



System-PoC #C - Synergetic monitoring and orchestration (1/3)

- Planning mechanism manages: i) Site placement, ii) Service scaling factor (replicas), iii) Cloud Processing guarantees, iv) Network Latency guarantees and v) Power consumption
- Cost factors: Power consumption cost, Cloud and Network SLA tiers cost
- Cost function:

$$C_{tot} = (1 - x) \cdot C_{pow} + x \cdot (C_{cloud}(PL) + C_{net}(NL))$$

- Experiments run to obtain observations as latency SLAs grow larger
- Local deployment is preferred when able to support SLA locally
- Cloud / Network lower (less costly) tiers are preferred
- Total cost reduces

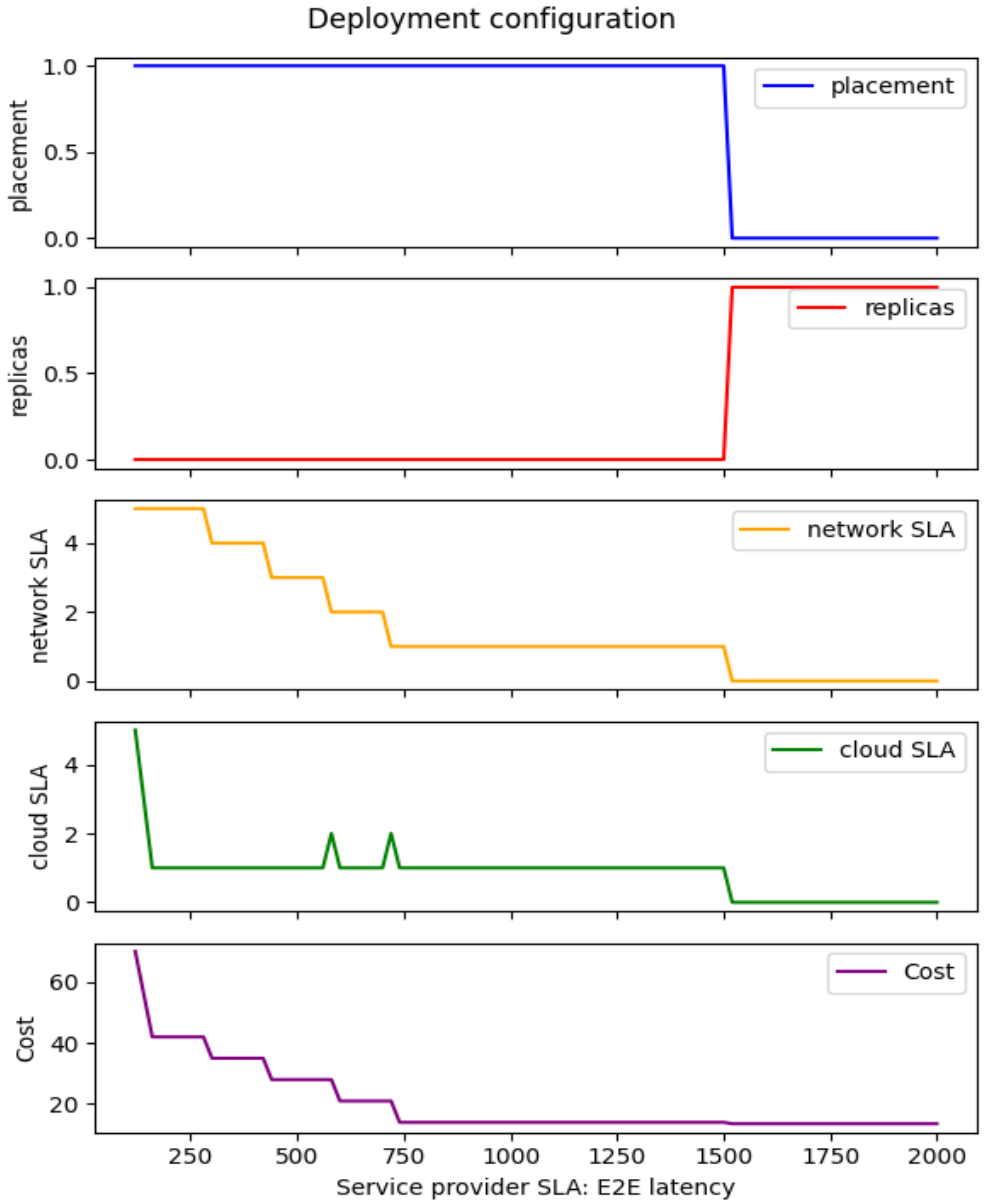

System-PoC #C - Synergetic monitoring and orchestration (2/3)

Constraints: E2E application latency
- Threshold-based (1s)
- Rolling window guarantees (80% of the time)

Observations
- Latency constraints are generally met, except during peak hours
- During remote server unavailability, service is migrated locally
- Local server overload leads to higher latency
- Additional resources needed to stay below latency threshold

SLA Compliance
- Over 83% satisfaction in 24-hour rolling window
- SLA remains valid despite occasional breaches

Cost Insights
- Cost drops notably when SLO is not satisfied
- Power cost during local deployment

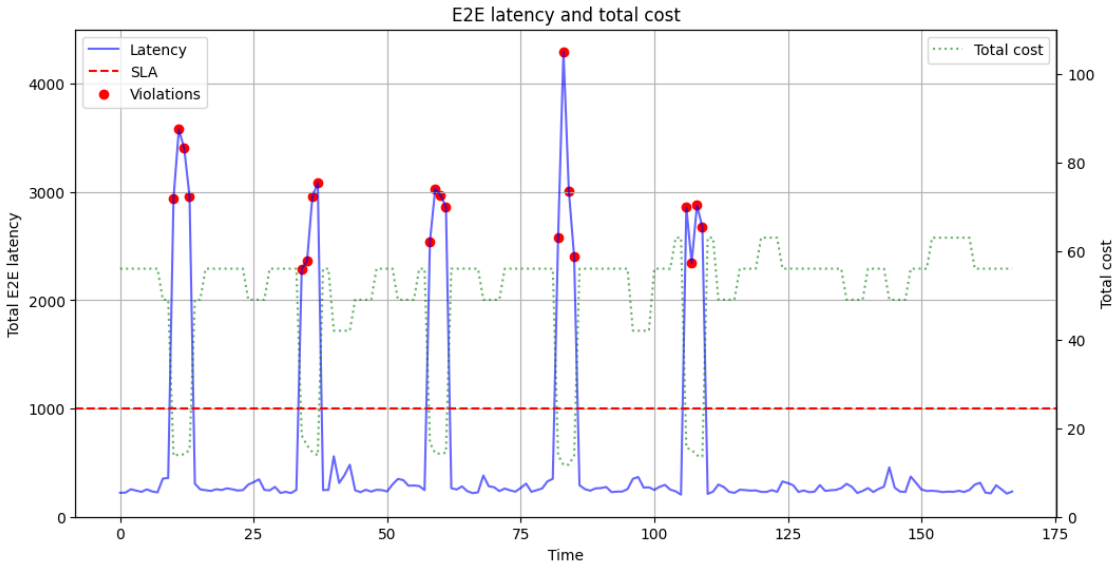

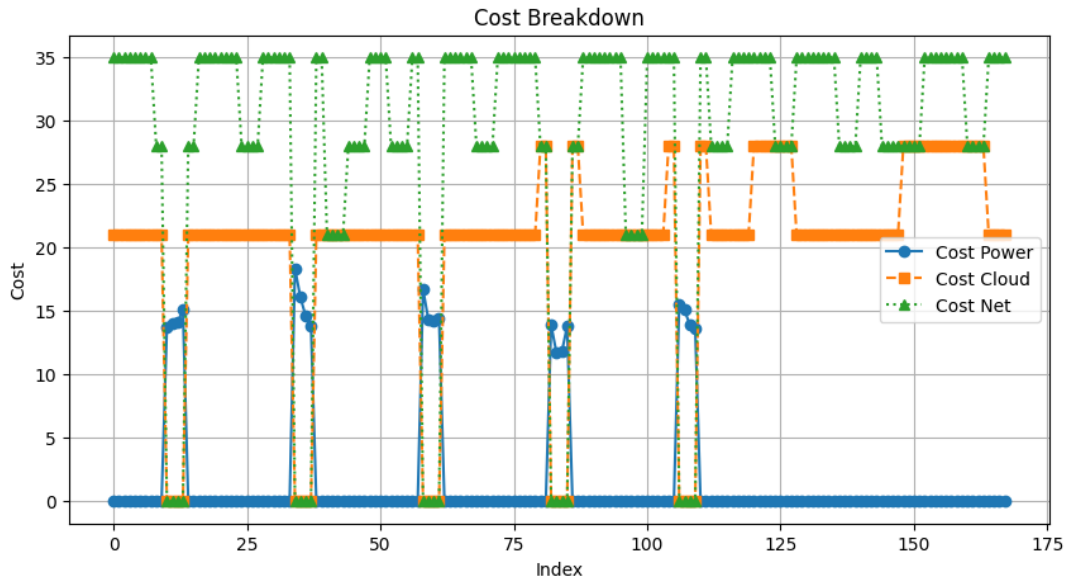

System-PoC #C - Synergetic monitoring and orchestration (3/3)

a) IBN networking - cumulative probability vs execution time: Execution time ranges approximately from 2.6s to 4.2s, with the majority of them being complete between 2.8s and 3.8s. The cumulative probability reaches near 1.0 at around 4.1s, meaning nearly all executions are completed by this time.

b) IBN networking - execution time breakdown: Feasibility check is the most time-consuming component, accounting for 46% of the total execution time. Then, Kafka messages consume 40% of the time while the TFS calls account for 13% of the time. Finally, the Database storage takes 1% and the task to interpret intent it is so low that it can be considered as a negligible processing time.

c) Service orchestration: Integration Fabric messaging performance ranges on average between 100–300ms with the worst-case being ~800ms. Service start & drop metrics for the two computing sites demonstrate:

- Site A start/stop: 550ms / 0ms (offloading to cloud: no drop -> workload redirection)
- Site B start/stop: 100ms/100ms (lightweight server)

d) Cloud placement overhead: At Site A operations downtime increases by ~150%, which can impact SLA compliance when downtime is critical.

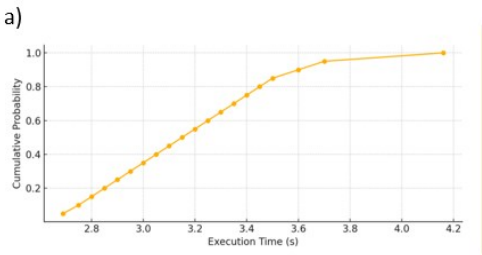

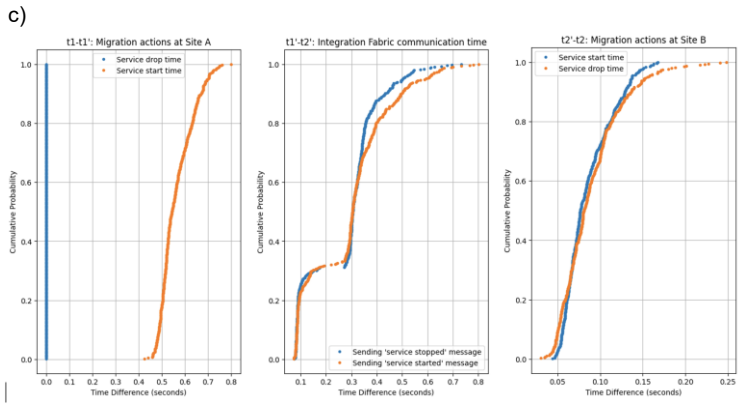

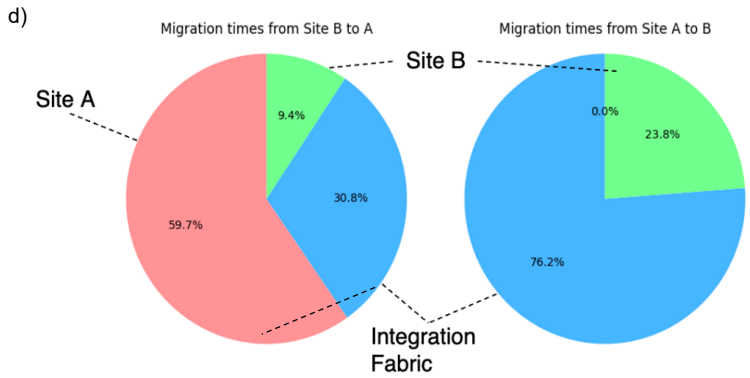

System-PoC #C - Evaluation of the Integration fabric Adapter implementation

IBN-IME requests to create a new service (nginx as example)

CREATED notification sent to Integration Fabric

```
{"reqId":"9724fc1e-328c-4ed0-93ce-97374e2e30eb", "status":"CREATED"}
Message: b'68c2f71c-287d-49c4-a3ef-4dc3131c7c54' successfully produced to Topic: test.nextworks
2025-03-13 10:26:32.710040 Delta time:  0:00:00.136248
```

INSTANTIATING notification to Integration Fabric

```
Response after instantiat
2025-03-13 10:26:32.76784
```

```
Instance stauts changed in: INSTANTIATING
2025-03-13 10:26:32.826376 Delta time:  0:00:00.058534
Sent notification through kafka
{"reqId":"9724fc1e-328c-4ed0-93ce-97374e2e30eb", "status":"INSTANTIATING"}
Message: b'eaddc184-6e17-4f7d-8f28-0f7cb716361e' successfully produced to Topic: test.nextworks
2025-03-13 10:26:32.868912 Delta time:  0:00:00.042536
```

Service is instantiated by Service Orchestrator

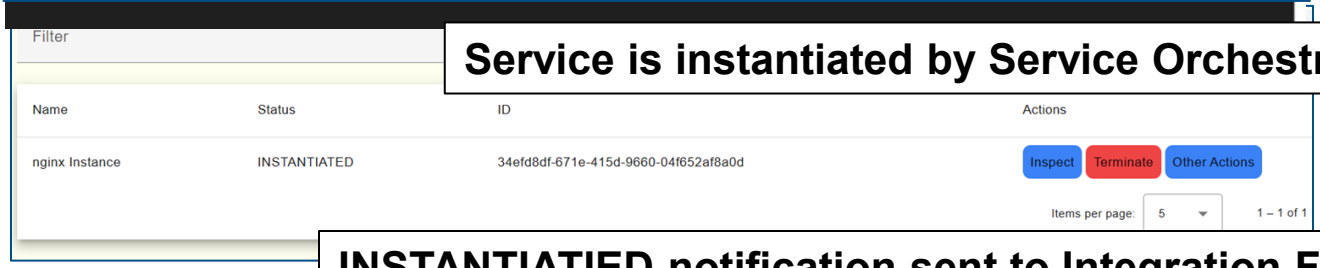

INSTANTIATIED notification sent to Integration Fabric

Evaluation of the delay overhead introduced by Integration Fabric Adapter:

- Measured the time for publishing each notification
- Run multiple time the Intent-based service provisioning workflow

Results:

- On average, a notification publishing on Integration fabric: requires 0.133 seconds
- The whole service workflow instantiation introduces an overhead of less than one seconds

System-PoC #C - Sustainability and trustworthy-oriented orchestration in 6G

- Physical task planning FA mechanism:
  - Simulated environment with 30 heterogeneous virtual AMRs and UAVs executing vision-based tasks in a 20×20×4 m warehouse.
  - Tested against a nearest-neighbour baseline under varying task loads.
  - The proposed task planning FA algorithm achieved up to 35.9% energy savings and 60% reduction in task completion time compared to the nearest-neighbour heuristic.
  - Energy efficiency and scalability improved when sufficient robotic resources were available, but gains decreased as robot availability dropped due to task saturation.
- Trust Management System:
  - Assesses the trust indexes of compute nodes, interacting with both TEF and LoTAF.
  - Integrated through the Management Capability Exposure (MCE) integration fabric for improved trustworthiness estimation.
  - Time-related results from the integration, measured over 26 workflow repetitions.

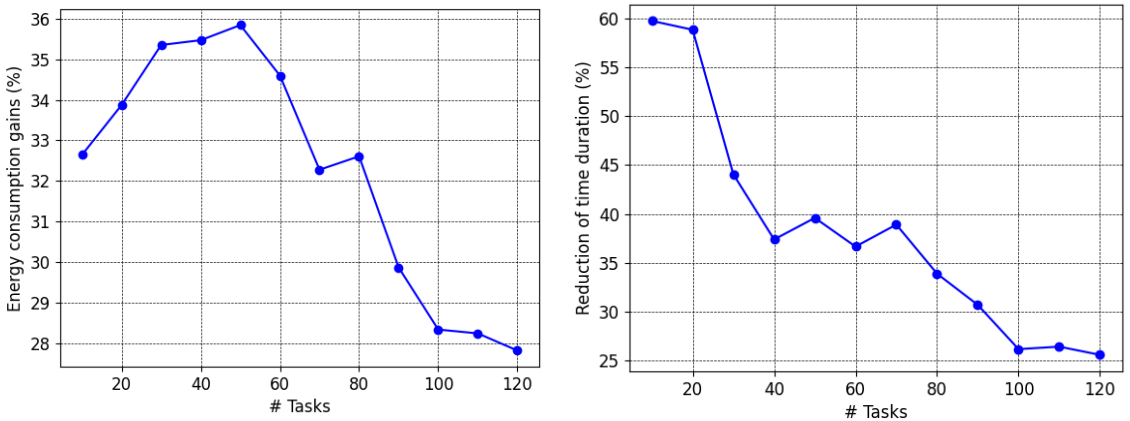

Energy consumption gains and reduction of time duration using the physical task planning FA algorithm compared to the nearest-neighbour heuristic.

Publish trust management request to MCE	~ 0 s	Publish TEF output to MCE	~ 0 s
Consume request from TEF	0.0402 s	Publish LoTAF output to MCE	0.0746 s
Consume request from LoTAF	0.0321 s	Consume TEF and LoTAF output	2.0250 s
TEF execution time (including infrastructure monitoring communication)	0.3247 s	Execution time of trust management calculations	0.0001 s
LoTAF execution time	0.0002 s	E2E latency (from publishing trust management request till its final calculations)	2.7369 s

Time-related results of trust management (utilising MCE).

Component-PoC #C - 6G based sensing algorithms and concepts with real-time performance

- Reached more than 90 percentile within 1 meter for human target detection in an indoor scenario
- Extensive developments both on the hardwar and software made to make the process as close as possible to the realtime experience. As our final test, we could reach 10 FSP for target detection

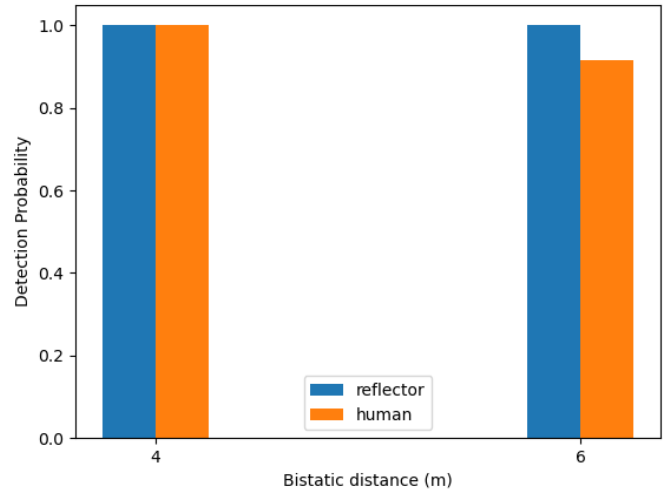

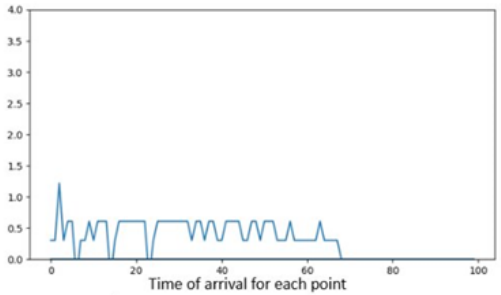

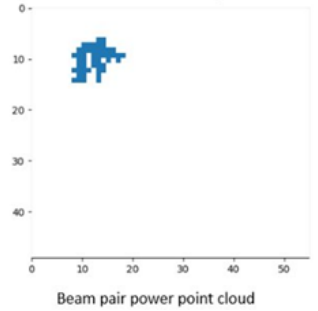

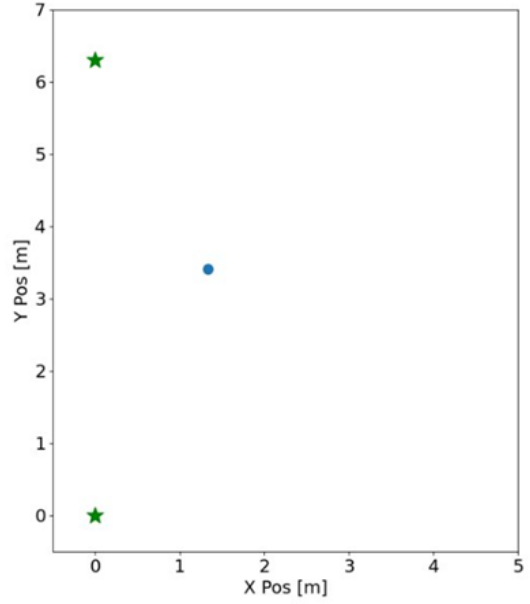

Component-PoC #C - AI-Native Air Interface

- PoC testing setups: Tested in two setups: through a channel emulator and over-the-air (OTA) inside the laboratory.
- Channel emulator experiments: Conducted with three mobility levels—pedestrian, car, and train scenarios—compared with non-ML algorithms.
- Throughput gains:
  - Highest gain of 30% achieved in the train scenario (highest velocity).
  - Lower velocity scenarios (pedestrian/car) showed gains of 15-20%.
  - Gains attributed to reduced overhead of the AI-based approach, trained to communicate without DMRS pilots.
- Spectral efficiency: The AI-based approach achieves similar bit error rate performance to conventional methods with DMRS, resulting in higher spectral efficiency.
- OTA measurements: Achieved 19.2% throughput gain, similar to the low-velocity pedestrian scenario, demonstrating performance with real physical radio channels.

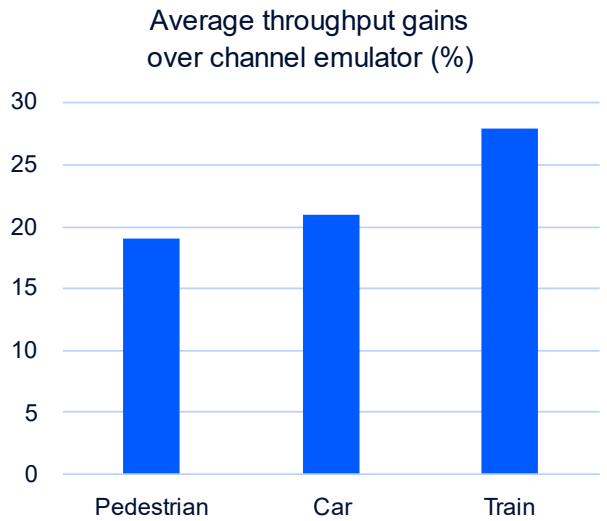

Throughput gains measured over channel emulator.

Component-PoC #C - ML-based channel state feedback compression in a multi-vendor scenario

- Evaluation of sequential training during OTA trials confirmed a 35-90% DL throughput gain with OTA model vs. 3GPP Type-I CSF
- Synthetic models, trained on simulator-generated cluster delay line (CDL) channels covering mixed LoS/NLoS indoor and outdoor cases, show a 25-65% gain compared to Type-I CSF
- Synthetic models performed within 2–15% of OTA-trained AI models indoors and matched performance outdoors
- Conclusion: sequential-based training is a promising solution for ML-based CSF

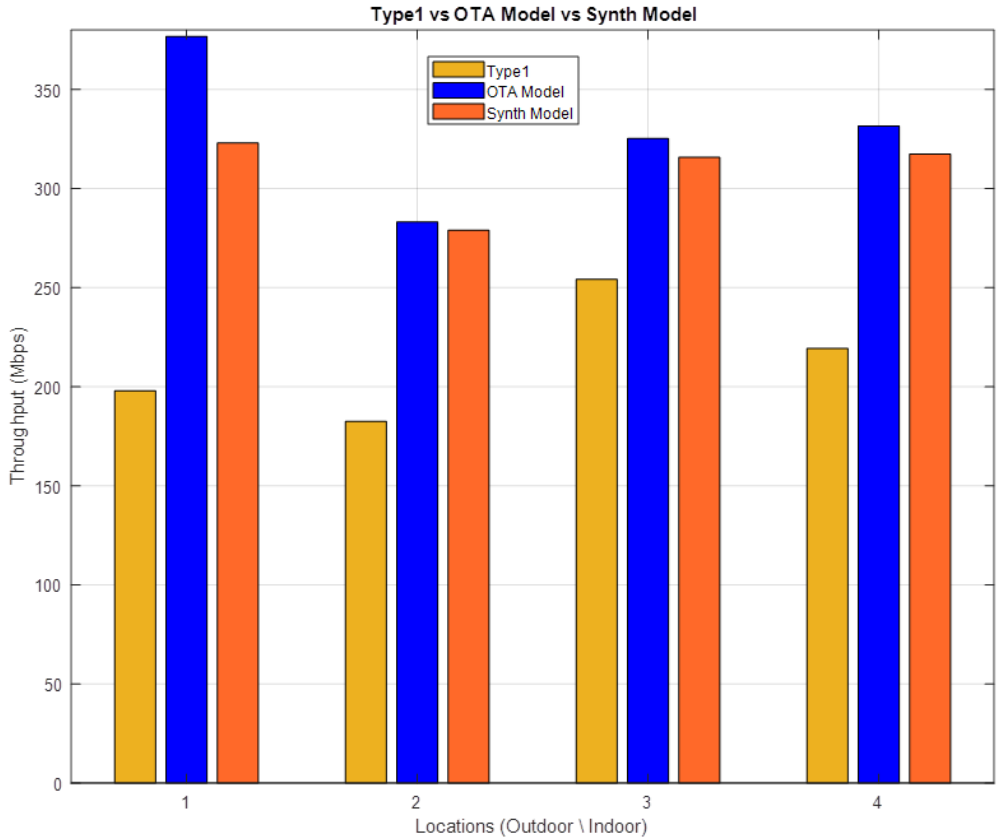

Figure: comparison of the downlink throughput between the Type 1, ML-based model with common training and synthetized-based ML model CSF, at 4 different locations.

Component-PoC #C - Crowd-detectable zero-energy devices

- Evaluation of Zero-Energy Devices: Measured bit error rate (BER) performance in a laboratory environment, compared with the theoretical BER curve of a non-coherent FSK receiver.
- Field trials: Conducted two field trials in Paris (Orange Gardens) and Espoo (Aalto University) using a commercial LTE network and SDR-based UE.
- Performance results:
  - BD signals were readable by a mobile phone even with the base station outside, and the BD/UE inside, with a reading distance of 0.5 meters.
  - Measured BER performance and compared to theoretical non-coherent FSK receiver results.
- Field trial setup: LTE base station illuminated the BD, and SDR implementation of the UE acted as the receiver for backscatter-modulated messages.

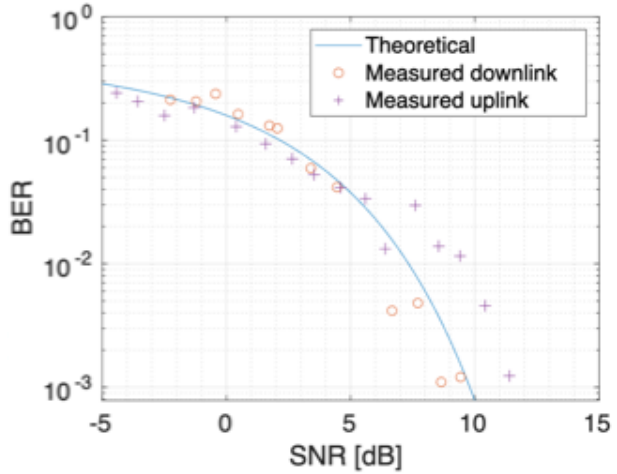

Field trial in Paris

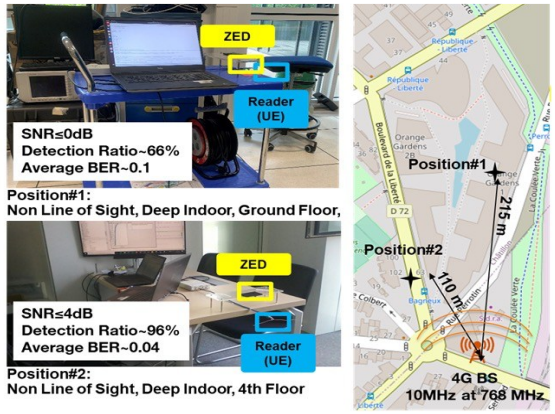

	Position 1	Position 2
Average 4G SNR	≤0 dB	≤4 dB
Observation duration	4171s	4171s
Transmitted frames	852	867
Detected frames	561	835
Detection Ratio	65.8%	96.27%
Average Data BER	0.1008	0.0439

Field trial in Espoo

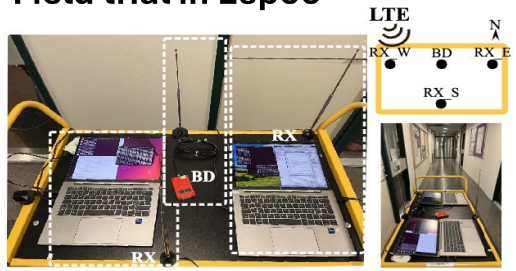

TABLE I
MEDIAN AND 90TH PERCENTILE BER VALUES BY ENVIRONMENT AND FLOOR (TWO-DIGIT ACCURACY)

Environment	Floor	Median	90th Percentile
Bridge	1st	1.1×10^{-2}	1.2×10^{-1}
Bridge	2nd	2.3×10^{-2}	9.7×10^{-2}
Bridge	3rd	8.4×10^{-3}	7.9×10^{-3}
Corridor	1st	4.7×10^{-2}	1.3×10^{-1}
Corridor	2nd	2.1×10^{-2}	7.4×10^{-2}
Corridor	3rd	2.1×10^{-2}	7.8×10^{-2}

Component-PoC #C - Energy harvesting IoT proximity devices

Evaluation of the energy harvesting IoT proximity device in two scenarios:

- Scenario 1:
  - Sensor cycle: 30s (2s active at 30mA, 28s sleep at 9mA).
  - Wi-Fi transmission every 10 cycles for 45s at 50mA.
  - Lamp provides continuous current of 27-30mA, 10cm from the panel.
  - Observations:
    - Total current drawn by the device in 10 cycles and transmission: 5370 mA·s over 345s.
    - Average current consumption in this scenario: 15.6mA.
  - Outcome: The lamp's 27-30mA current provides enough energy to slowly recharge the battery, preventing it from draining.
- Scenario 2:
  - Sensor cycle: 30s (2s active at 30mA, 28s sleep at 9mA).
  - Wi-Fi transmission every 2 sensor cycles, lasting 30s at 50mA.
  - Lamp provides continuous current of 27-30mA, 10cm from the panel.
  - Observations:
    - Device’s current draw: 35.4mA, while the lamp provides 27mA.
    - Network current drain: 8.4mA.
    - Battery life with current drain: 2400mAh / 8.4mA = 285.7 hours.
  - Outcome: Removing the DC-DC converter (which consumes 9mA) could improve battery autonomy.

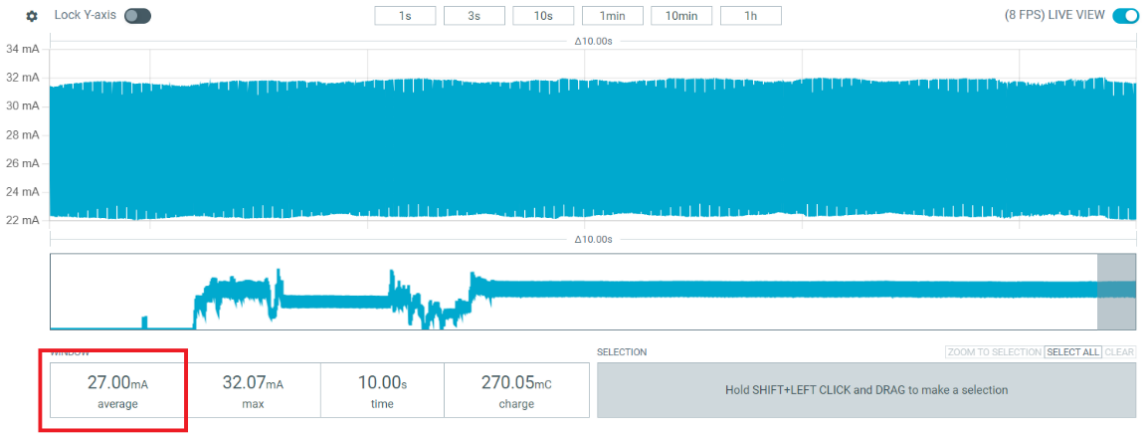

Solar current production

Component-PoC #C - End-to-end Extended Reality

- The proposed testbed exploits the low-latency low-loss (L4S) feature at the XR device to dynamically adapt the application data rate to the network load
- For the OTA demonstration, 3 users in the cell are pushing full buffer traffic for 60 seconds to generate network load then stops.
- Results show that leveraging L4S load information combined with adaptive rate adaptation at the XR device can reduce the number of impacted frames, improving the user quality of experience

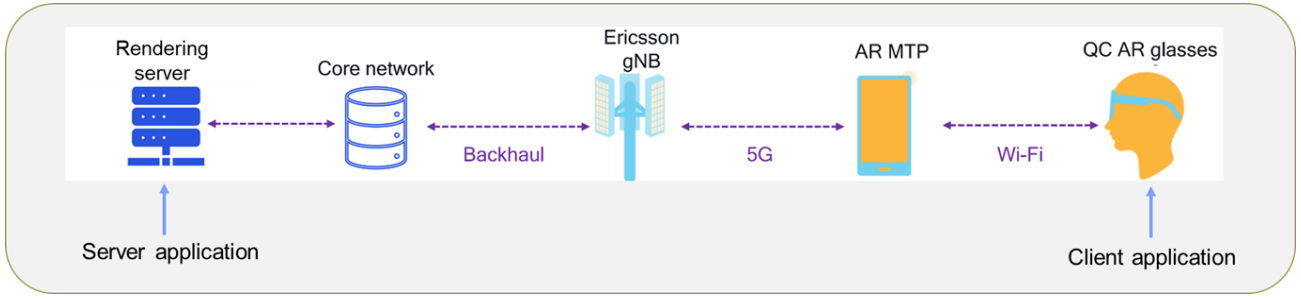

In the testbed, the AR glasses are paired with a QLC smartphone running the client application. The phone is connected to the Ericsson Network through a 5G connection and then reaches the rendering server.

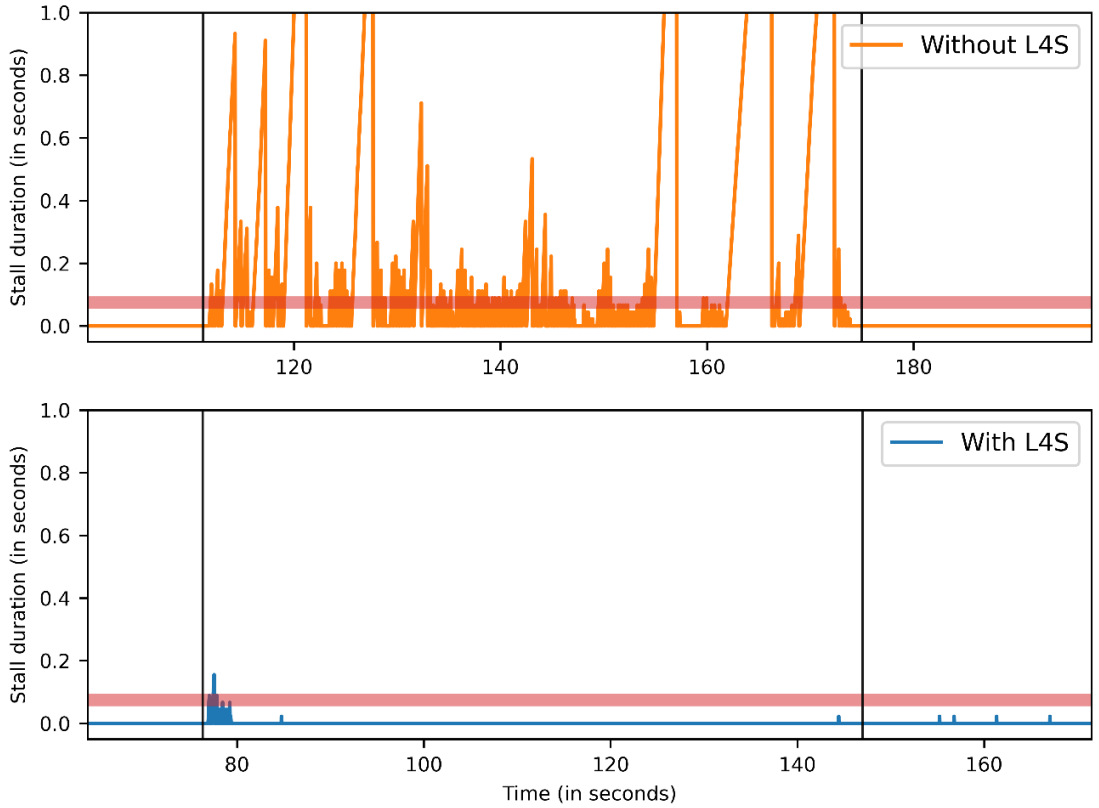

Figure: Comparison of the number of consecutive duplicated frames with and without L4S enabled in a scenario with varying NW load. The red line indicates a threshold, set to a stall duration of [50, 100]ms, for which the user experience is impacted. The black line vertical lines indicate the start and end of the user loading.

Component-PoC #C - Flexible modulation and transceiver design

- The flexible transceiver system's reconfigurable hardware supports a wide range of experimental setups
- Dynamic switching between RF frontends and real-time parameter adjustments enables efficient testing of new concepts in mobile communications and radar applications
- The modular design allows seamless integration of additional RF frontends, minimizing effort for future upgrades
- This adaptability expands the system's applicability to emerging research domains

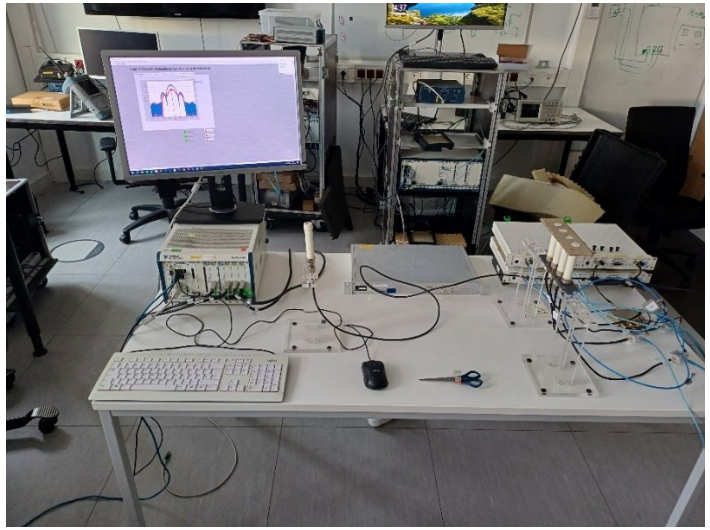

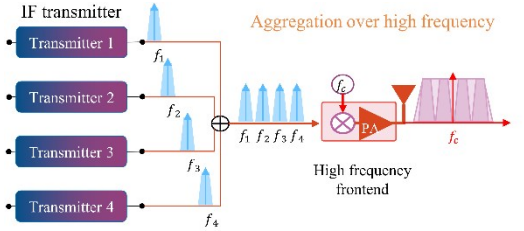

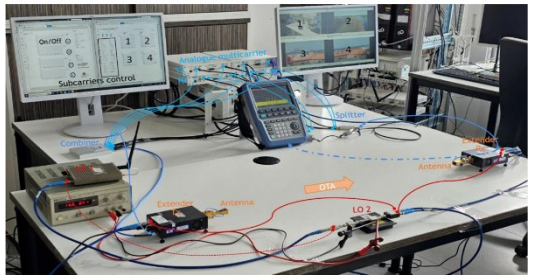

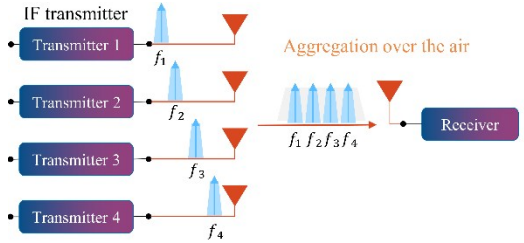

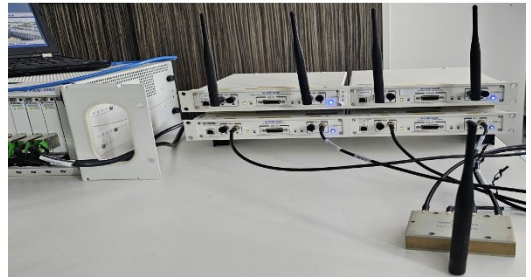

Component-PoC #C - Radio propagation measurements to collect data for radio channel modelling

- Propagation campaigns: Conducted at 15 GHz, 142 GHz, 234 GHz, and 318 GHz using USRP-based and VNA-based setups across indoor and outdoor scenarios.
- Measurement environments:
  - 15 GHz: Residential streets, university campus, parking lot (USRP-based linear route measurements).
  - 142 GHz: Entrance hall, suburban, residential, city center (VNA-based, single-directional measurements).
  - 234 & 318 GHz: Factory halls and warehouse layouts (VNA-based, bi-directional measurements).
- Measurement method: Azimuth and elevation scans at each Tx–Rx location pair to capture full multipath angular structure.
- Preliminary findings:
  - 234 GHz: Bi-directional power angular spectrum in factory hall.
  - 15 GHz: Coverage map showing path gain along routes in residential (UMi) and campus/parking lot (UMa) scenarios.
- Further details: Described in Section 3.3 [HEX225-D45] for sub-THz measurements and Section 3.7 [HEX225-D25] for 15 GHz outdoor campaign.

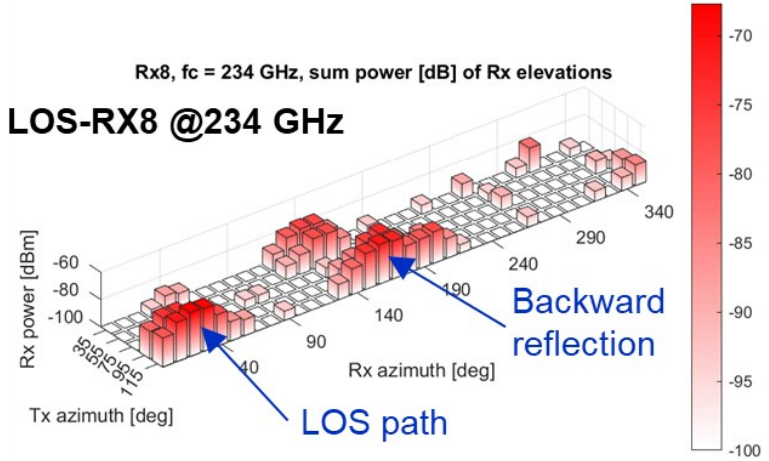

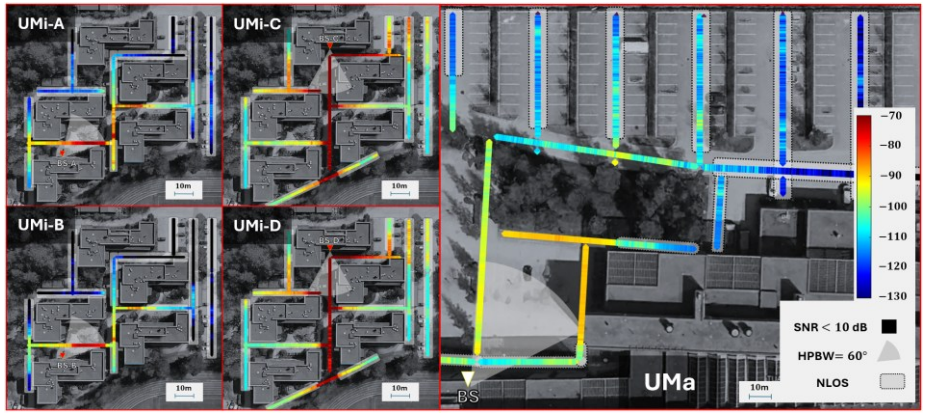

Impact - Open APIs in System-PoC / Component-PoCs

Title	SW Origin	Licence	Available in	PoC / Demo	Integration/usage context
DLT Service Federation Open API	New - Created in project	CC Attr. 4.0 International	[DLTUC324]	PoC B.1	The interfaces are used in the context of the Federated orchestration system
Monitoring jobs configuration API	Existing - Enhanced in project	Apache 2.0	[MonNXW24]	PoC B.1	Used for configuring the monitoring platform for retrieving the cobots data (battery, position and status)
Query historical monitoring data API	Existing	MIT	[InfluxDB]	PoC B.1	Provides time-series lookups on previously collected system/cobot metrics. Used by orchestration or AI modules for improved resource placement and scheduling decisions.
Multi-cluster extreme-edge resource orchestration API	New - Created in project	Apache 2.0	[MulNXW24]	PoC B.1	Enables orchestrators to manage services across multiple Kubernetes clusters, including edge clusters. In the PoC, it supports dynamic scaling and AI workflow migration between edge and cloud to meet low-latency requirements.
Closed Loop Governance – Catalogue API	New - Created in project	Apache 2.0	[CloNXW24]	PoC A, PoC B.1, PoC C	Allows for managing the applications that support the different stages of the CLs. In the PoC, it is used to onboard the specific CL functions required for the automatic control of the cobots' batteries and related remediation
Closed Loop Governance – Lifecycle Management API	New - Created in project	Apache 2.0	[LCMNXW24]	PoC A, PoC B.1, PoC C	Allows the consumer to request lifecycle operations on a given CL. In the PoC, it is consumed by the security orchestrator to request the provisioning of a CL specific for the cobot's service automation based on the battery level
Integration Fabric API	New– Created in project	CC Attr. 4.0 International	[NIntegr24]	PoC C	Provides three API endpoints used in the context of PoC for: • Retrieve keystore for the SSL connectivity with the Kafka Cluster • Listing the communication channel, onboarded entities • Onboard and offboard component and create and delete communication channels
Level of Trust Assessment Function API	New – Created in project	GPL 3.0	[VLoTAF24]	PoC C	Provides real-time trust metrics from compute nodes. These trust indexes guide the trust-aware orchestration functions.
Trust Evaluation Function API	New – Created in project	GPL 3.0	[LTEF25]	PoC B/C	Allows retrieving trust indexes after analyzing infrastructure data. Integrated in the PoC B and C.
IBN-IME API	New – Created in project	Apache 2.0	[AMV+25]	PoC B/C	A subset of the OpenAPI calls were used in both PoC B and PoC C, in the context of the intent management procedures.

Impact - Open-source tools

- Hexa-X-II PoC activities have integrated open-source tools for simulations, orchestration, security, distributed computing, and AI-driven management in 6G scenarios.
- OpenAirInterface (OAI):
  - OAI 5G virtual RAN for latency comparison between split-RAN and monolithic architectures.
  - OAI RF simulator used for F1 interface delay analysis.
  - Simulation Environment: Intel i7-11850H, 32 GB RAM, Ubuntu 20.06, n78 band (3.5 GHz), 40 MHz bandwidth.
- Kubernetes: Managed distributed workloads, AI-driven apps, and latency-sensitive orchestration across edge and cloud clusters.
- Prometheus: Monitored performance metrics (CPU, memory, latency), with Karmada and Zipkin for service propagation and tracing.
- Within Management Capabilities Exposure (MCE) Framework: Apache Kafka (secure data exposure, mTLS), OpenAPI (v3.1.0) for service integration.
- Qujata security: Post-quantum cryptography (Kyber, Frodo, HQC) for quantum-resistant encryption, hybrid PQC evaluations.
- TeraFlow SDN (TFS): Multi-domain slicing, transport reconfiguration; CockroachDB for state storage; OpenConfig for network management.
- Real-time monitoring: gNMI telemetry for packet-optical observability; Sionna for AI-driven spectrum optimization and RAN evaluation.

Impact - How the end-to-end system evaluation results of the overall 6G system design address the project objectives

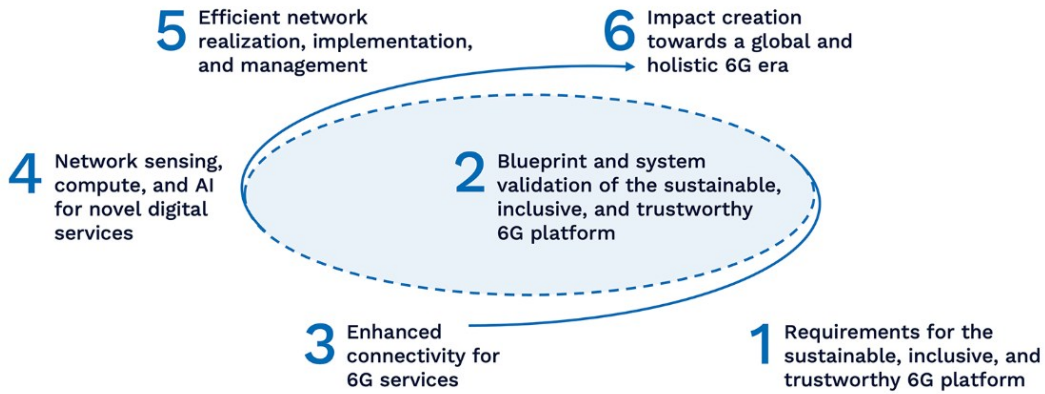

#1	#2	#3	#4	#5	#6
Introduction of energy harvesting proximity sensors and TinyML-based power optimization in warehouse shelves, supporting autonomous operation without power cables or frequent battery replacement. Trust-aware orchestration mechanisms, combining TEF and LoTAF, enhance trust and transparency across compute nodes. System design addresses vertical needs (e.g., warehouse logistics, surveillance) with privacy-, trust- and energy-conscious orchestration.	System-PoC #C showcases realistic multi-domain integration (cobots, drones, sensors, compute nodes). End-to-end evaluation across sensing, compute, trust, and orchestration layers confirms feasibility of the Hexa-X-II blueprint. Validation involves both technical KPIs (latency, energy, trust) and system integration workflows (e.g., closed loop provisioning, digital twin updates).	PoCs validate AI-native air interfaces, CSI feedback compression, and flexible transceivers enabling adaptive radio interfaces. UAVs acting as Flexible topology nodes extend coverage for ground robots and sensors in dynamic or degraded connectivity conditions. End-to-end XR and video streaming over 5G/6G testbeds demonstrates low-latency service continuity.	Implementation of 6G-native sensing algorithms and AI-powered safety detection (e.g., human-cobot collision avoidance). Exposure of Compute-as-a-Service and AI-as-a-Service over the application enablement platform, supporting advanced digital service orchestration. Energy harvesting IoT devices act as network-integrated proximity sensors, feeding data to real-time digital twins.	Integration Fabric supports standardized, real-time cross-domain orchestration, enabling scalable, zero-touch service management. Intent-driven orchestration with business intent abstraction, automated service provisioning, and AI-based workload migration across edge/cloud. Dynamic, AI-assisted functionality allocation via metaheuristics ensures workload balance and optimized energy/trust performance.	Exposure of CAMARA-compliant APIs facilitates cross-domain application development and ecosystem growth. Concrete vertical integration (warehousing, robotics, XR, safety) and validates a future-ready 6G architecture supporting multiple stakeholders. Emphasis on sustainability, energy-efficiency, and openness enhances relevance across sectors and geographies.



References

References

[AMV+25]	Alemany, P., Muñoz, R., Vilalta, R., & Adanza Dopazo, D. (2025). Designed API for the IBN-IME solution used in Hexa-X-II T2.3/T2.5 work. Zenodo. https://doi.org/10.5281/zenodo.15125456	[LTEF25]	Lamprousi, V. (2025). Trust Evaluation Function API. Zenodo. https://doi.org/10.5281/zenodo.14720466
[CloNXW24]	Nextworks (Italy). (2024). Closed Loop Governance - Catalogue API. Zenodo. https://doi.org/10.5281/zenodo.14193484	[MonNXW24]	Nextworks (Italy). (2024). Monitoring jobs configuration. Zenodo. https://doi.org/10.5281/zenodo.14192885
[DLTUC324]	Universidad Carlos III de Madrid (UC3). (2024). DLT Service Federation API. Zenodo. https://doi.org/10.5281/zenodo.14033410	[MulNXW24]	Nextworks (Italy). (2024). Multi-cluster extreme-edge resource orchestration API. Zenodo. https://doi.org/10.5281/zenodo.14193659
[InfluxDB]	InfluxDB Cloud API Service (2.x) Available: https://docs.influxdata.com/influxdb/cloud/api/v2/	[NIntegr24]	Nicolicchia, R. (2024). Integration Fabric API (v1.0.0). Zenodo. https://doi.org/10.5281/zenodo.14222126
[LCMNXW24]	Nextworks (Italy). (2024). Closed Loop Governance LCM API. Zenodo. https://doi.org/10.5281/zenodo.14193635	[VLoTAF24]	Jorquera Valero, J. M., & CyberDataLab. (2024). Level of Trust Assessment Function API. Zenodo. https://doi.org/10.5281/zenodo.14191242



HEXA-X-II.EU //

Co-funded by the European Union

6GSNS

Hexa-X-II 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 101095759.