xG-ALOE

Alliance for Open Enterprise Networks

PLUGFEST 2025
TEST REPORT

EANTC
European Advanced Networking Test Center

Fraunhofer
HHI

Table of contents

Introduction

From September 29 to October 2, 2025, the Fraunhofer Heinrich Hertz Institute (HHI) in Berlin hosted the first xG-ALOE Plugfest, organized in collaboration with EANTC AG. This exciting event brought together industry partners, research organizations, and universities to test Open RAN implementations in a cooperative environment. Participants included AiVader, airpuls, brown-iposs, Fraunhofer FOKUS, Fraunhofer HHI, Fraunhofer IIS, Hochschule Offenburg, IS-Wireless, QUALIGON, RWTH Aachen, TU Dresden, and Universität Passau.

The xG-ALOE, initiative hosted by Fraunhofer HHI provides a joint platform for companies and research teams to develop and evaluate private 5G and 6G network technologies. The plugfest format enables participants to interconnect their systems, perform tests, and observe interoperability in practice. For smaller companies and startups, this shared environment provides a practical approach to addressing integration challenges that would be difficult to solve alone.

Eight test scenarios were executed, covering different aspects of network deployment and performance. These included evaluations of the integration of Time-Sensitive Networks with 5G networks, security mechanisms for Open RAN systems, and localization and sensing functions relevant to industrial applications. Other scenarios showcased a portable network-in-a-box solution and evaluated multi-user scalability. Two additional tests focused on establishing PDU sessions for off-the-shelf devices and another comparing optical wireless backhaul performance with traditional Ethernet utilizing Open6GCore.

Configuration mismatches, protocol inconsistencies, and hardware compatibility issues were among the challenges faced. These findings proved valuable, as participants could work together to troubleshoot, find root causes, and improve both implementations and testing methods. The collaborative environment transformed technical challenges into opportunities to enhance interoperability and clarify standardization requirements.

Recording these challenges and solutions will help participants solve similar issues more efficiently during future plugfests, with each event building on past lessons. This event was a great success in fostering collaboration and innovation in the field of network technology. The following sections describe the test setups, procedures, and detailed results.

xG-ALOE Plugfest 2025 Participants

1 Ethernet PDU Support for 5G-TSN

This scenario addresses the transport of Time-Sensitive Networking (TSN) traffic over a private 5G network using 5G Ethernet PDU sessions instead of regular IP PDU sessions. These Ethernet PDU sessions enable the transmission of plain layer-2 traffic (Ethernet frames) without an additional IP header, which is a fundamental requirement for 5G-TSN, targeting deterministic communication. The implementation represents one of the first real-world over-the-air tests combining generalized Precision Time Protocol (gPTP) and 5G over Ethernet PDU sessions.

The scenario examines two primary aspects: time synchronization using IEEE 802.1AS (gPTP) and overall system performance, assessed through end-to-end latency and jitter testing.

The setup includes prototypes of TSN translators (detailed in 3GPP 23.501) for time synchronization and TSN switches that implement traffic shaping with IEEE 802.1Qbv, which minimizes network jitter by artificially introducing latency. The testing aims to validate requirements critical for 5G-TSN Use Cases, particularly jitter reduction and edge-cloud control scenarios, such as wireless remote control of Automated Guided Vehicles (AGVs) over proprietary protocols like PROFINET.

USE CASE 1:
gPTP Time Synchronization

Traditionally, gPTP time synchronization over 5GS has been run using L2TP tunnels on top of IP PDU sessions. This Use Case demonstrates running gPTP directly over Ethernet PDU sessions using plain L2 Ethernet packets, eliminating the need for additional tunneling and simplifying the architecture. The primary objective was to verify the correct functioning of the synchronization procedure and message transport, with particular focus on L2 multicast functionality, rather than achieving high synchronization accuracy. Stable synchronization in the range of a few hundred nanoseconds can be expected.

USE CASE 2:
End-to-End Latency and Jitter Testing

This test evaluates one-way end-to-end latency in both uplink and downlink directions, encompassing the 5G UPF, 5G RAN, 5G UE, two TSN switches, and L2 switches for interconnection. Data is generated from the load generator and the measurement device and captured on the same device again. Both TSN switches can operate with or without an enabled time-aware shaper (TAS), utilizing a gate-control list derived from the 5G RAN TDD pattern (a 5 ms cycle, 4.5 ms closed, and 0.5 ms open for both uplink and downlink).

Test Topology

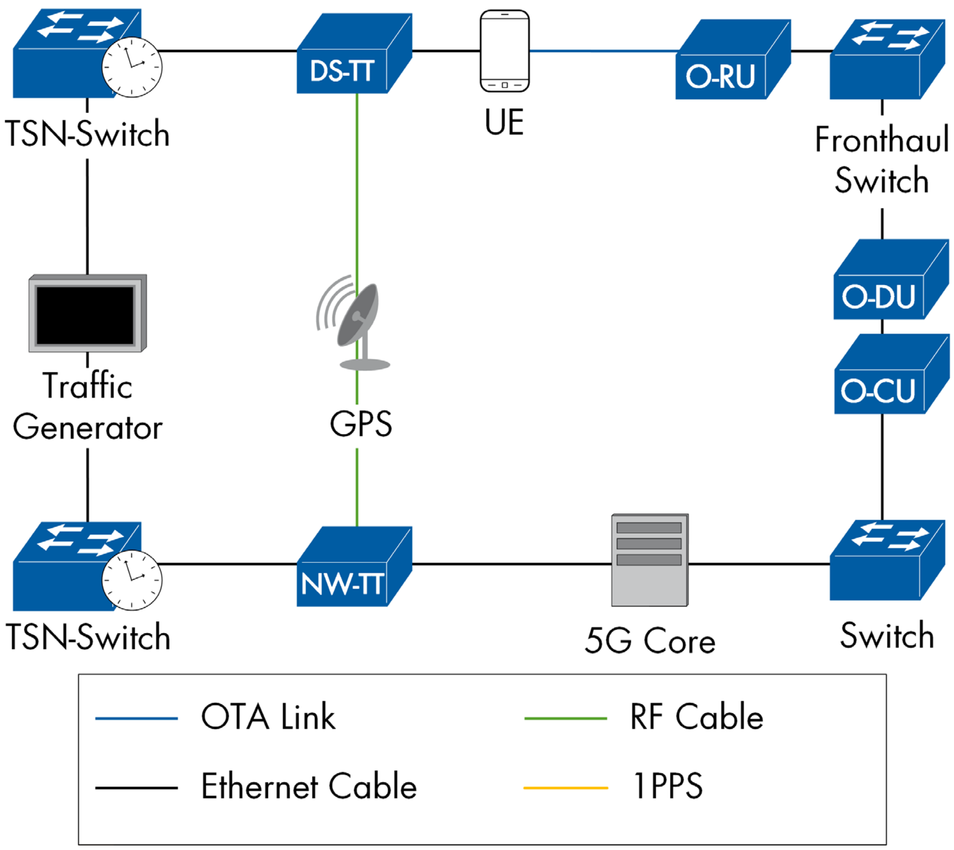

FIGURE 1: Test Setup

COMPONENT	DESCRIPTION
RAN	airpuls (O-CU/O-DU)
TSN Translator (DS-TT)	Hochschule Offenburg on TI SK-AM64B
TSN Translator (NW-TT)	Hochschule Offenburg on TI SK-AM64B
TSN Switches	Kontron KSwitch D10 MMT Series (2 units)
5G Core	OAI Core

TABLE 1: Test Setup Components

Network Configuration and Test Parameters

This section presents the key configuration parameters and test traffic settings used in the evaluation.

PARAMETER		VALUE
5G Network	PLMN	001-01
	Center Frequency / Bandwidth	n78 / 100 MHz
	TDD Configuration / Periodicity	2-1-2 (DDSUU) / 2.5 ms
	Transmit Power (RU) / UE Distance	24 dBm / ca. 1.5 m
Traffic (iMIX: 64B-58.33%, 594B-33.33%, 1518B-8.33%)	Uplink / Downlink Load	5 Mbps / 10 Mbps (1 stream each)
TSN Switches	Cycle Time / Gate Closed / Open	5 ms / 4.5 ms / 0.5 ms (UL/DL)
gPTP	Protocol / Interface	IEEE 802.1AS over Ethernet PDU / eth0
	Addressing / Time Source	Unicast (L2 multicast workaround) / GPS sim

TABLE 2: Network Configuration and Test Parameters

Test Result

USE CASE 1
gPTP Time Synchronization

Testing successfully confirmed that gPTP over Ethernet PDU Sessions is functional. The gPTP Slave successfully transitioned through the expected states: MASTER → UNCALIBRATED → SLAVE, indicating proper clock selection and synchronization establishment.

Synchronization Performance
The ptp4l output demonstrates stable synchronization across two distinct operational phases. During the initial convergence period (the first 10 seconds), the system achieved rapid stabilization within 10 iterations, significantly reducing the RMS error from 7,346 ns to 55 ns, while the path delay stabilized at 388-390 ns.
Following convergence, steady-state performance maintained typical RMS errors between 50 and 150 ns, with occasional spikes of up to 673 ns, likely due to varying wireless network conditions. The frequency offset stabilized at approximately -3,000 ppb, while the path delay consistently remained at 388-389 ns ±1 ns throughout steady-state operation.

PHASE	RMS ERROR	FREQUENCY OFFSET	PATH DELAY
Initial (start)	7,346 ns	-10,085 ppb	389 ns ±1
After convergence (10s)	55 ns	-3,088 ppb	388 ns ±0
Steady-state (typical)	50-150 ns	ca. -3,000 ppb	388-389 ns ±1
Steady-state (max spike)	673 ns	-	-

TABLE 3: Synchronization Performance

This performance demonstrates successful gPTP functional verification over Ethernet PDU Sessions, achieving rapid convergence to 55 ns RMS error within 10 seconds, meeting the sub-100-nanosecond target. During steady-state operation, typical RMS errors ranged from 50-150ns with occasional spikes to 673ns. The variation in steady-state performance can be attributed to the wireless nature of 5G over-the-air transmission, which introduces timing variability due to radio channel conditions, and the use of GPS simulation rather than authentic GPS signals. Despite these factors, the achieved synchronization accuracy remains well within functional requirements for 5G-TSN applications.

Sample ptp4l Output Log

```
ptp4l[20123.311]: port 1 (eth0): SLAVE to MASTER on
ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
ptp4l[20123.311]: selected local clock 2ccf67.fffe.403c05 as best master
ptp4l[20123.311]: port 1 (eth0): assuming the grand master role
ptp4l[20211.679]: selected best master clock 2ccf67.fffe.403c24
ptp4l[20211.679]: port 1 (eth0): MASTER to UNCALIBRATED on RS_SLAVE
ptp4l[20212.448]: port 1 (eth0): UNCALIBRATED to SLAVE on
MASTER_CLOCK_SELECTED
ptp4l[20212.698]: rms 7346 max 14368 freq -10085 +/- 9113 delay  389 +/-  1
ptp4l[20213.824]: rms 3696 max 7239 freq -12394 +/- 3677 delay  390 +/-  0
ptp4l[20214.825]: rms 3382 max 3681 freq  -3439 +/- 1304 delay  390 +/-  0
ptp4l[20215.824]: rms 2623 max 3466 freq  -1230 +/- 134 delay  389 +/-  0
ptp4l[20216.824]: rms  975 max 1536 freq  -1759 +/- 261 delay  390 +/-  0
ptp4l[20217.824]: rms  222 max  407 freq  -2587 +/- 307 delay  390 +/-  0
ptp4l[20218.824]: rms  230 max  302 freq  -3050 +/- 135 delay  390 +/-  0
ptp4l[20219.824]: rms  114 max  157 freq  -3073 +/-  69 delay  388 +/-  0
ptp4l[20220.824]: rms   84 max  147 freq  -3018 +/- 114 delay  388 +/-  0
ptp4l[20221.824]: rms   55 max   89 freq  -3088 +/-  52 delay  388 +/-  0
```

USE CASE 2:
End-to-End Latency and Jitter Testing

Traffic streams were generated with the load generator and looped back on the same device, which recorded transmission and reception timestamps. The test was conducted in two phases: with the TSN switches operating in their default configuration (TAS disabled), and with TAS enabled using the gate-control list derived from the TDD pattern of the 5G RAN.

End-to-End Performance Comparison: With and Without Time-Aware Shaper (TAS)

		WITHOUT TAS			WITH TAS ENABLED			CHANGE WITH TAS		
Direction	Metric	Avg*(ms)*	Min*(ms)*	Max*(ms)*	Avg*(ms)*	Min*(ms)*	Max*(ms)*	Avg	Min	Max
	Latency	8.94	3.52	90.96	8.59	3.54	71.18	↓ -3.9%	↑+0.6%	↓ -21.8%
Uplink	Jitter	0.98	0	36.44	0.98	0	59.2	-	-	↑+62.5%
	Latency	3.01	1.7	16.17	3.0	1.69	11.76	↓ -0.3%	↓ -0.6%	↓ -27.3%
Downlink	Jitter	0.44	0	14.54	0.44	0	9.75	-	-	↓ -32.9%

TABLE 4: End-to-End Performance Comparison: With and Without Time-Aware Shaper (TAS)

The following figure visualizes the average delay and jitter values measured and described in Table 4.

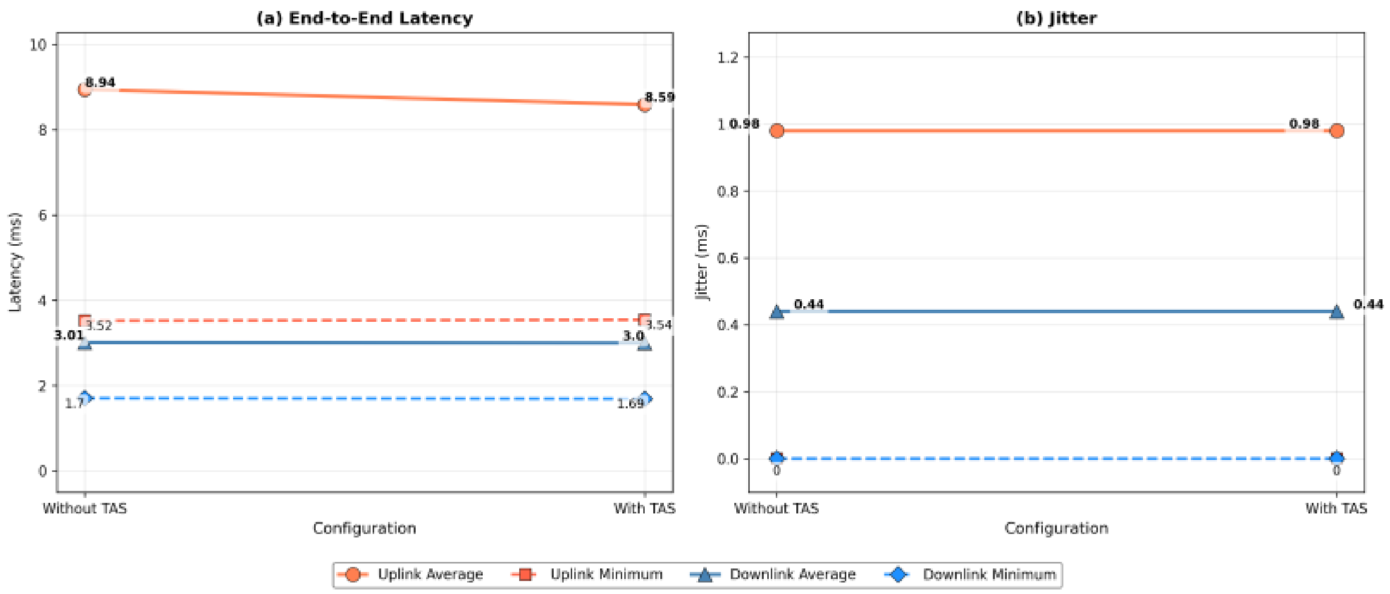

FIGURE 2: End-to-End Performance Comparison: With and Without Time-Aware Shaper (TAS)

Analysis

The results show minimal differences between TAS-enabled and disabled configurations, with average performance metrics remaining nearly identical. Average jitter values for both uplink (0.98 ms) and downlink (0.44 ms) were the same regardless of TAS status. Average latency showed only slight variations, with the uplink marginally decreased 8.94 ms to 8.59 ms, and the downlink remaining essentially constant at 3.01 ms. Maximum latency values showed more noticeable improvements with TAS enabled—uplink improved from 90.96 ms to 71.18 ms (21.8% reduction) and downlink from 16.17 ms to 11.76 ms (27.3% reduction). However, maximum jitter results were mixed, with the uplink increasing from 36.44 ms to 59.2 ms, while the downlink improved from 14.54 ms to 9.75 ms.

Several factors may contribute to the minimal TAS impact observed in these results. This testing phase focused on baseline performance characterization using single-stream traffic per direction. TAS benefits are typically more notable in scenarios involving multiple traffic classes with different quality of service requirements and competing flows, where traffic competition necessitates the enforcement of priorities and the isolation of traffic. Under the current test conditions—with no competing traffic streams and a single flow per direction—there was limited opportunity for TAS to demonstrate these capabilities. Additionally, the measurement configuration accumulated results over one-second intervals, which may have dimmed short-term timing variations and outliers that could better illustrate the effectiveness of TAS.

Future testing iterations should incorporate more complex scenarios, including multiple simultaneous traffic streams with different priority levels, mixed time-sensitive and best-effort traffic, and network load conditions that create contention. Higher temporal resolution measurements would also help capture fine-grained timing behavior. These enhancements would better demonstrate TAS's effectiveness in managing deterministic traffic requirements typical of industrial 5G-TSN applications.

Technical Challenges and Limitation

The testing described in this scenario was performed with the modified OpenAirInterface (OAI) 5G Core. While the original test plan aimed to repeat all measurements with Amarisoft Core to compare open-source versus commercial core network performance characteristics, resource and scheduling constraints prevented completion of this comparative phase.

During the gPTP time synchronization testing, we encountered two technical issues that required workarounds to maintain functionality. First, the UE did not forward L2 multicast packets, which are essential for standard gPTP protocol operation. To overcome this limitation, unicast addresses were configured for gPTP communication instead of the standard multicast addresses. While this workaround enabled successful functional testing, it represents a deviation from standard TSN practices that must be resolved for production deployments.

Second, GPS signal availability was insufficient for proper TSN translator operation at the test location. A GPS simulation solution was used as an alternative, providing a sufficient timing reference for functional validation purposes. However, this approach introduces a slight degradation in accuracy compared to authentic GPS synchronization. Both issues require comprehensive investigation in later testing phases to ensure full protocol compliance and evaluate their potential impact on synchronization accuracy and overall system performance in production environments.

CONCLUSION 1

Ethernet PDU Support for 5G-TSN

The test results showed that 5G networks can support time synchronization with an accuracy of around 50-150 nanoseconds in steady state, which meets the typical requirements of industrial automation. The gPTP protocol converged from an initial 7,346 ns to 55 ns within about 10 seconds. Time-Aware Shaping had a minimal effect on average latency, with the uplink decreasing slightly from 8.94 ms to 8.59 ms, and the downlink remaining essentially unchanged at ca. 3.0 ms. There was no change in average jitter values, as the uplink (0.98 ms) and downlink (0.44 ms) remained constant, suggesting its benefits may become more apparent in scenarios with multiple competing traffic flows. However, maximum latency improved by 21.8% uplink and 27.3% downlink.

2 Security and Attack Detection

This scenario focuses on strengthening the security of Open RAN systems by enabling the detection of malicious activities within the RAN through an Intrusion Detection System (IDS). As 5G networks transition to fully IP-based architectures, attack surfaces are expanding significantly, making IDS a vital first line of defense against potential threats. By design, an IDS functions as an early warning system that monitors network activities for malicious behavior—it does not block threats but alerts administrators to potential attacks that may have bypassed other security measures. The IDS leverages an xApp deployed on the Near-Real-Time RAN Intelligent Controller (Near-RT RIC) to collect and analyse Key Performance Metrics (KPMs) from the network. By continuously monitoring these metrics using anomaly-based detection, the system can identify unusual patterns or deviations from normal behavior that may indicate potential attacks, making it effective against unknown or emerging threats independently of the overlying communication protocol.

In this test case, a protocol-independent detection mechanism is integrated into the O-RAN RIC as an xApp, enabling closed-loop security monitoring and response. Under benign operation, the xApp subscribes to the O-RAN environment via the E2 interface, gathers operational metrics and control events, and creates a machine learning model of the system under observation. During regular operation, this model is then used to compare real-time captured KPMs and identify deviations from expected behavior, raising alerts when anomalies indicative of intrusions are detected. This setup illustrates how an intelligent, data-driven security application can be seamlessly deployed within the O-RAN ecosystem.

A critical aspect of this approach is the low-latency reporting and processing of KPM data. Rapid detection and response are essential to ensure that security threats are recognized in near real time and at the network's edge, minimizing their spread and impact on network performance and stability. This scenario demonstrates how integrating a protocol-independent IDS with the Near-RT RIC can significantly enhance the resilience and trustworthiness of Open RAN deployments.

Test Topology

The testbed comprises several components connected in an end-to-end architecture. The user equipment consists of four User Equipment (UEs). The network flow operates as follows: the UE connects over the 5G air interface to the RAN components, which interface via Ethernet with the Open5GS Core. Test devices (traffic generators) are connected at both ends of the network, specifically behind the N6 interface of the core network and to the LAN behind the UE. Key Performance Measurements (KPMs) provide low-level data collected by network components. They provide raw insights into UE activity, resource utilization, overall cell conditions, and other key metrics. Its underlying idea is to detect certain abnormal system states and misbehaving entities using the key performance metrics that can be retrieved via the E2 interface.

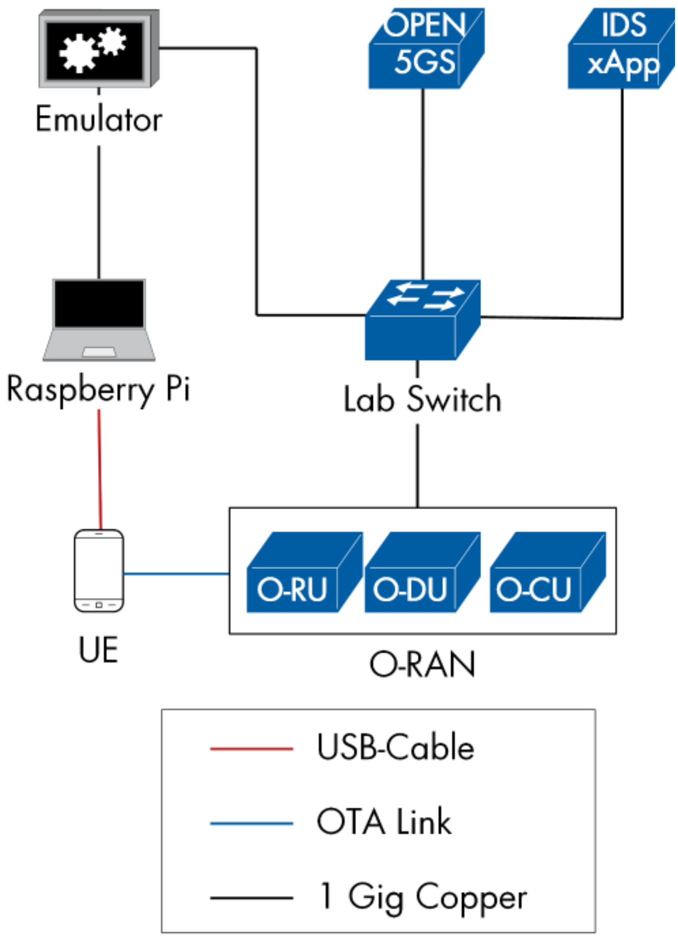

FIGURE 3: Test Setup

COMPONENT	DESCRIPTION
RAN Intelligent Controller	SC RIC - (Universität Passau)
Intrusion Detection System	xAPP (RWTH Aachen University)

TABLE 5: Test Setup Components

Test Case Description

The test follows a four-step methodology consisting of baseline verification, regular operation validation, attack simulation, and recovery assessment. Initially, the system under test's operational status is confirmed to ensure all services are functioning correctly. Legitimate traffic is then established and monitored to document normal traffic rates and verify error-free operation without security alerts. Once stable, attack traffic will be generated while legitimate traffic continues running. During this attack phase, the system is expected to maintain legitimate traffic stability while generating security alarms, with documented rates of attack traffic. Finally, the attack traffic is stopped, and the system is verified to return to stable operation with alarm generation ceasing. This controlled approach allows systematic evaluation of the SUT's security detection capabilities and operational resilience under attack conditions.

Technical Challenges and Limitations

Integration challenges related to multi-vendor interoperability between the system under test and third-party RAN components prevented successful test execution within the available timeframe. These issues require further investigation and resolution in subsequent testing phases to enable intrusion detection testing.

3 Localization and Sensing

This scenario examines the integration of the brown-iposs Protocol Value Grabber and the Fraunhofer IIS Positioning Engine through the Qualigon KAFKA data broker. The brown-iposs Protocol Value Grabber can extract the data from public 4G networks over the air via its built-in antenna. For the Fraunhofer IIS Positioning Engine to utilize these measurement data, both components must be connected via the Qualigon KAFKA data broker. Additionally, the measurement data from the Protocol Value Engine requires formatting and restructuring to ensure seamless integration. The integration of all three components is the subject of this scenario.

The integration encompasses: The deployment of the brown-iposs Protocol Value Grabber and ensuring that it extracts data from public 4G networks. The extraction feature for 5G networks of the Protocol Value Grabber is still in development; however, the integration chain would still work with the new feature to extract data from 5G networks. The connection between the Protocol Value Grabber and the KAFKA data broker, and the publication of raw measured data to the data broker. The preprocessing of the raw measured data by a Data Conditioning Unit and publishing the optimized data to the data broker again. The reception of the optimized data by the Positioning Engine in the form of a Positioning Protocol Skeleton.

Apart from the integration, the continuous operation of the integrated solution, the interpretation of the data by the Positioning Engine, and the long-term logging of raw measured and optimized data were under evaluation.

Test Topology

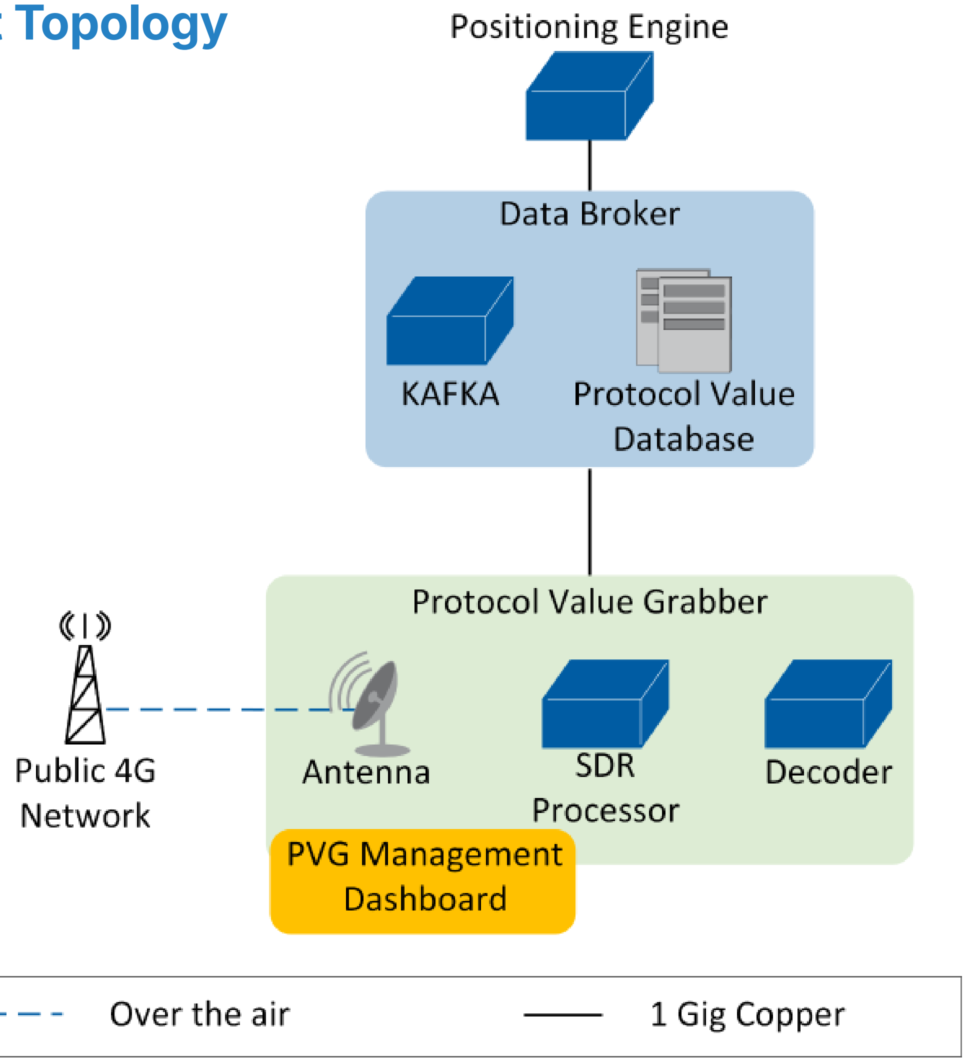

FIGURE 4: Test Setup

COMPONENT	DESCRIPTION
Protocol Value Grabber and PVG Management Dashboard	brown-iposs
Data Broker Industrial PC	Qualigon - Apache Kafka
Positioning Engine	Fraunhofer IIS – including Protocol Skeleton, API Stub and Positioning Stack

TABLE 6: Test Setup Component

Network Configuration and Test Parameters

PARAMETER		VALUE
Protocol Value Grabber	configured EARFCNs	75, 101, 125, 252, 275, 475, 1300, 1444, 1600, 1650, 1801, 2850, 3050, 3350, 3500, 3600, 3749, 6200, 6300, 6400, 9260, 9360, 9460, 10020

TABLE 7: Test Parameters

Test Results

TEST CASE 1:
Protocol Value Grabber functionality

The Protocol Value Grabber was deployed at the Plugfest venue. Upon connecting a power source and a management connection, the Protocol Value Grabber turned on and was accessible via the PVG Management Dashboard. Next, the GPS connection of the Protocol Value Grabber was established through a GPS repeater. Thereafter, the Protocol Value Grabber began measuring for 4G networks in the area around the venue. Sixteen separate 4G cells were detected across 103.000 measurement points. The sixteen detected cells were part of the public 4G networks from three separate public operators.

The measurements of the Protocol Value Grabber continued for two days without issue, confirming the system's stability. Throughout the two days, the system was moved inside the venue to provide different location points. The Protocol Value Grabber published the taken measurements to the KAFKA Data Broker.

Since all criteria for the successful operation of the Protocol Value Grabber, namely automatic start-up, measurements, detection of one or more 4G networks, and forwarding of measurements, could be confirmed, this test is evaluated as passed.

```
=== PCI Statistics ===
 pci  count    mean  mcc  mnc   tac                          operator
  54     56 -108.68  262    1  1494          Telekom Deutschland GmbH
  71     13 -109.48  262    1  1494          Telekom Deutschland GmbH
  72   2997 -112.34  262    3 22020 Telefónica Germany GmbH & Co. oHG
 203     90 -111.43  262    1  1494          Telekom Deutschland GmbH
 216      5 -107.17  262    1  1494          Telekom Deutschland GmbH
 228     37 -107.49  262    2 49101                  Vodafone D2 GmbH
 252    357 -108.87  262    1  1494          Telekom Deutschland GmbH
 264      6 -101.07  262    2 49101                  Vodafone D2 GmbH
 268   8667 -105.79  262    2 49101                  Vodafone D2 GmbH
 276     20 -107.48  262    3 22020 Telefónica Germany GmbH & Co. oHG
 278      1 -107.59  262    3 22020 Telefónica Germany GmbH & Co. oHG
 355  23364  -99.71  262    2 49101                  Vodafone D2 GmbH
 372  20820 -104.38  262    1  1494          Telekom Deutschland GmbH
 388     36  -96.49  262    2 49101                  Vodafone D2 GmbH
 478  23016 -108.24  262    1  1494          Telekom Deutschland GmbH
 499  24352  -95.72  262    2 49101                  Vodafone D2 GmbH
```

FIGURE 5: Example Measurements

TEST CASE 2:
KAFKA data broker functionality

After the initial setup of the KAFKA data broker at the Plugfest venue, it started automatically. After the startup, the KAFKA data broker was accessible to the Protocol Value Grabber and the Positioning Engine, as well as other devices, such as third-party KAFKA clients, in the test environment. The data produced by the Protocol Value Grabber and the Data Conditioning Unit was handled by the KAFKA data broker, which was verified through the aforementioned third-party KAFKA clients.

During the Plugfest period, the KAFKA data broker remained continuously accessible to the Protocol Value Grabber, the Positioning Engine, and the Data Conditioning Unit. Furthermore, it also stored historical data.

As all expected results for the KAFKA data broker functionality, namely automatic startup, continuous accessibility and operation, and storage of historical data, could be verified, this test is evaluated as passed.

TEST CASE 3: Interoperability of the Setup

After the on-site team established accessibility between the Systems under Test, the Protocol Value Grabber could provide raw measurement data through the KAFKA Data Broker to the Positioning Engine. The Positioning Engine can process this raw measurement data; however, for this test setup, an extra component, the Python-based Data Conditioning Unit, preprocessed the measurement data.

The Data Conditioning Unit processed the measurement data into a JSON format that the Positioning Engine can understand. The final mapping of raw measurement data to the output JSON format was agreed upon and implemented during the Plugfest timeframe and represents a new accomplishment in the setup implementation. The most critical value for fingerprint-based positioning is the signal strength value, which was measured at 4gDbm. There are many more fields in the API of the Positioning Engine, which are populated by the mapping, such as frequency and channel (number), or configuration, technology, and vendor-specific values. In each mapping, there may be an n:m ratio between raw data fields and preprocessed data fields. There are additionally mapping functions implemented to normalize and combine field values.

The Positioning Engine utilizes widely accepted UUIDs for identifying nodes within the system. A node is a participant (such as a UE or a base station) in the positioning technology, and its virtual instance. It is crucial to provide pseudo-unique identifiers. They can be generated, but in the case of referencing measurements, it is also necessary


```json
{
    "timestamp": 1759235168,
    "rss": [
        {
            "id": "bb880f01-099c-5806-949f-85a7b12bd4b6",
            "rssValue": -92.1,
            "frequency": 783.0,
            "channelType": "4G"
        },
        {
            "id": "2a1ea4a6-78b2-56b1-8668-37b17ab506f2",
            "rssValue": -109.4,
            "frequency": 796.0,
            "channelType": "4G"
        },
        {
            "id": "049b627e-4cc1-524f-a329-d4068b6373a7",
            "rssValue": -91.1,
            "frequency": 806.0,
            "channelType": "4G"
        },
        {
            "id": "38c14dae-9f8b-52c1-9aff-7bda60aa8e85",
            "rssValue": -94.3,
            "frequency": 816.0,
            "channelType": "4G"
        },
```


FIGURE 6: Grouping

to make them reproducible. For the unique station identification based on Protocol Value Grabbers measurements, the on-site team decided on and implemented a combination of MCC, MNC, 4G-Tac, and 4G-Ci for the feed of the UUID generation. With this newly established function, a sufficient identification of stations is possible.

The Data Conditioning Unit groups data, primarily using time frames, and structures it into datagrams in a format usable by the Positioning Engine. Therefore, timestamps are crucial information. Currently, only group timestamps are preserved during the grouping process, which is absolutely suitable for the application. This procedure means that functionally, there is no loss of relevant information, but not each measurement‘s component preserves its own timestamp.

To summarize, the Data Processing Unit preprocesses the raw measurement data for the Positioning Engine by reformatting it into a JSON format, assigning each node a UUID, and grouping the measurement data into time frame-based groups. The needed processing time was at least a magnitude lower than the update rate. The Data Conditioning Unit then published the datagrams, i.e., the preprocessed measurement data, back to the KAFKA Data Broker, which forwarded them to the Positioning Protocol Skeleton.

After the on-site team aligned the protocol, the Positioning Protocol Skeleton had a complete understanding of the datagrams. The Positioning Protocol Skeleton logged incoming datagrams continuously. Thus, a logging system based on the consumption of KAFKA-provided data was implemented and operationalized. Additionally, the KAFKA Data Broker stored the datagrams itself. To selectively access stored data, the on-site team used KAFKA's functionality to observe windows and offsets.


```
{"version":"1.0.1","tim":1759230726,"loc":{"ts":1759230726,"lat":52.518785,"lng":13.322385,"acc":0,"alt":56.400000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230729,"loc":{"ts":1759230729,"lat":52.518785,"lng":13.322388,"acc":0,"alt":56.300000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230733,"loc":{"ts":1759230733,"lat":52.518788,"lng":13.322388,"acc":0,"alt":56.200000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230736,"loc":{"ts":1759230736,"lat":52.518790,"lng":13.322387,"acc":0,"alt":56.200000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230737,"loc":{"ts":1759230737,"lat":52.518788,"lng":13.322388,"acc":0,"alt":56.400000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230743,"loc":{"ts":1759230743,"lat":52.518780,"lng":13.322385,"acc":0,"alt":56.500000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230744,"loc":{"ts":1759230744,"lat":52.518780,"lng":13.322385,"acc":0,"alt":56.500000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230752,"loc":{"ts":1759230752,"lat":52.518778,"lng":13.322388,"acc":0,"alt":56.300000,"spd":0,"typ":0}
{"version":"1.0.1","tim":1759230759,"loc":{"ts":1759230759,"lat":52.518780,"lng":13.322390,"acc":0,"alt":56.500000,"spd":0,"typ":0}
```


FIGURE 7: Incoming Datagrams

The received datagrams could be visualized by the Positioning Engine using plotting. Therefore, the on-site team could provide an analysis of the integration characteristics: Although the stream of measurement data was continuous, it is noticeable that the update rate of several seconds only allows applications with slow movements at the moment. Furthermore, even if the absolute signal strengths are low, they are usable. But, below the -110 dBm level, there is too much uncertainty. Most of the time, however, stable signal levels are observed in static measurements, with several levels that need to be considered. Outliers beyond those stable signal levels didn't cause significant problems to the integration, except for one outlier, which shifted the levels implausibly. An investigation by the on-site team indicates that the measurement hardware is the source of this anomaly. Furthermore, the on-site team discovered that the location of the Protocol Value Grabber's Antenna affects signal level, stability, and noise, as well as gaps and level-dependent effects substantially.

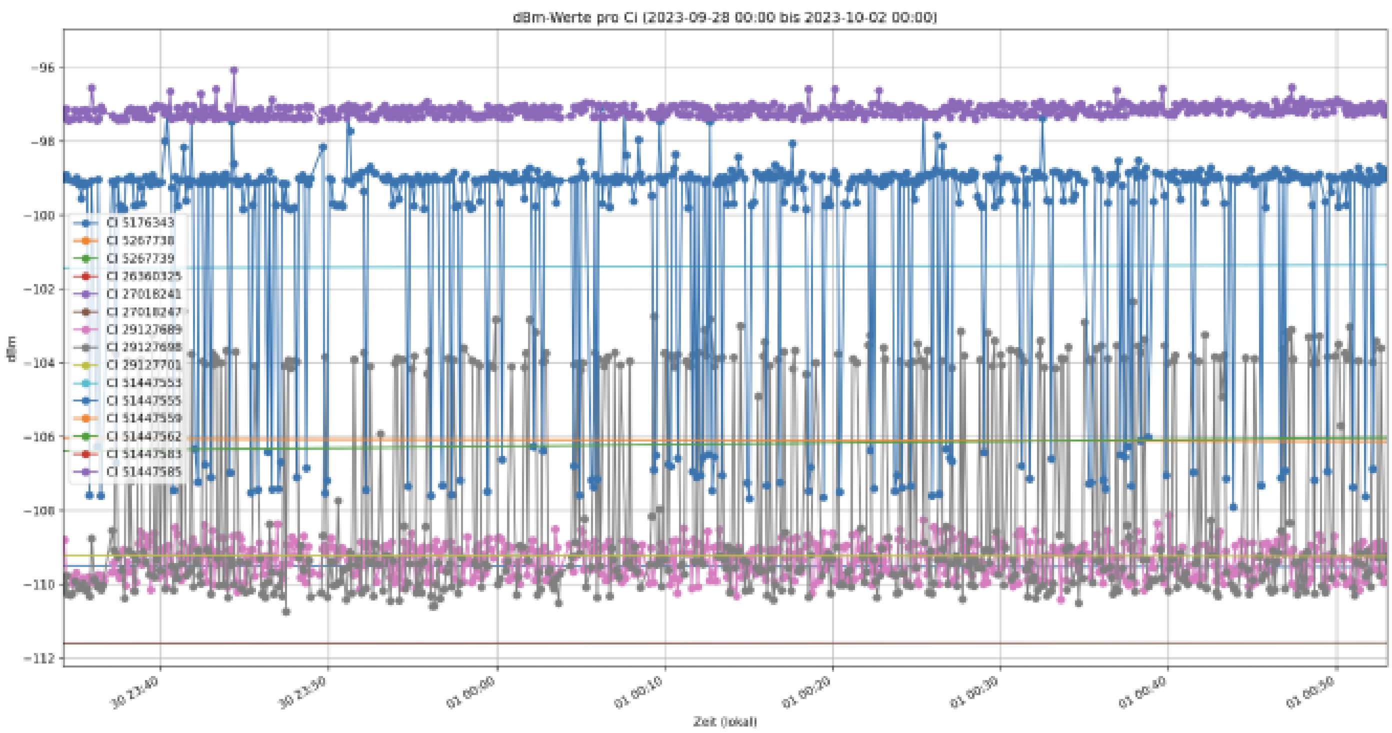

FIGURE 8: Visualized Datagrams

CONCLUSION 3

Localization & Sensing

The integration of the brown-iposs Protocol Value Grabber, Qualigon KAFKA data broker, and Fraunhofer IIS Positioning Engine was completed during the Plugfest. The Protocol Value Grabber operated continuously for two days, detecting 16 cells from three operators. The Data Conditioning Unit, integrated for the first time, successfully transformed raw measurement data into JSON format. The Protocol Value Grabber currently supports only 4G networks, as 5G features are under development. Due to time constraints, the alternative positioning engine core was tested only at the interface level. The current update rate limits applicability to slow-moving scenarios, and measurements below -110 dBm show high uncertainty. Future work will focus on adding 5G support and completing positioning accuracy validation.

4 Multi-RU feature within the 5G NR

This scenario validates the Multi-RU feature within the 5G NR system configuration. The test environment was deployed using two O-RUs, configured together to operate as a single cell. From the User Equipment (UE) perspective — represented by the Amarisoft UE Simbox — the PCI was expected to remain constant across the entire service area, ensuring seamless radio continuity.

Within this configuration, the Reference Signal Received Power (RSRP) distribution should exhibit the highest values in the proximity of RU 1 and RU 2, with the lowest RSRP observed in the mid-point region between the two RUs. Consequently, no handover or cell reselection events should occur during UE mobility across the coverage area, as both RUs serve the same cell entity. Additionally, data throughput and link-reliability performance were evaluated, particularly at the border area between the two RUs.

Test Topology

The testbed consists of a 5G NR network architecture with two coordinated radio units operating as a single cell. The UE Simulator connects to both RUs over the 5G air interface at different positions within the coverage area. The RAN components are connected to the 5G core network, which manages traffic routing and session establishment.

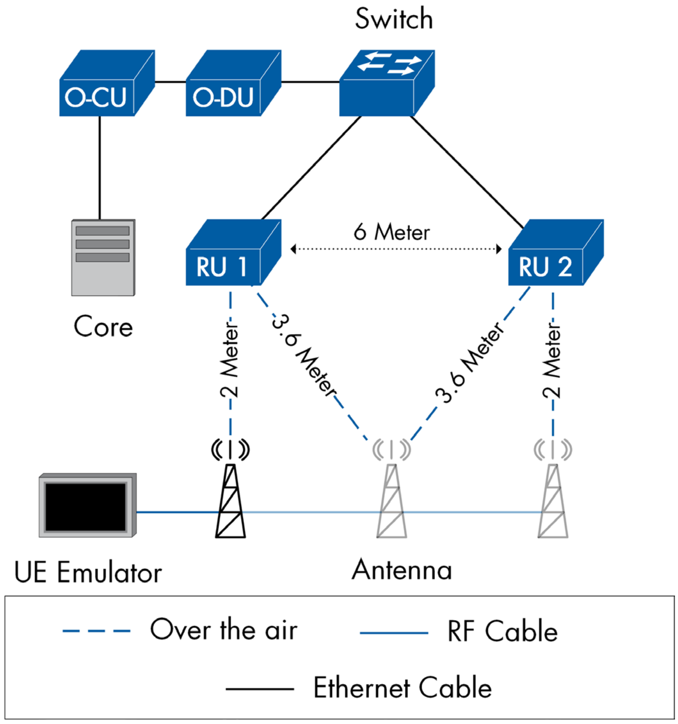

FIGURE 9: Test Setup

COMPONENT	DESCRIPTION
RAN	IS-Wireless RAN (O-DU, O-CU)
UE Simulator	Amarisoft UE Simbox

TABLE 8: Test Setup Components

Network Configuration and Test Parameters

This section presents the configuration parameters and test settings used in the evaluation.

PARAMETER		VALUE
5G NETWORK	PLMN	999-82
	Center Frequency	n78
	Bandwidth	40 MHz
	TDD Configuration	7-2 (DDSUU)
	TDD Pattern Periodicity	2.5 ms
	Physical Cell ID (PCI)	30
	Transmit Power RU	24 dBm
	Number of RUs	2 (RU1 and RU2)
TEST POSITIONS	Position 1	In front of RU1 (2 m)
	Position 2	Between RU1 and RU2 (6 m separation)
	Position 3	In front of RU2 (2 m)
IPERF3 CONFIGURATION	Port	5204
	Protocol	UDP
	Target Bitrate	100 Mbps
	Packet Size	1320 bytes
	Direction	Downlink
	Interval	2 sec

TABLE 9: Network Configuration and Test Parameters

Test Results

USE CASE 1:
RSRP/RSRQ/SNR Performance

The antenna of the UE Simulator was positioned at three locations to validate Multi-RU operation: directly in front of RU1 (2m), at the midpoint between both RUs, and in front of RU2 (2m). At each position, RSRP, RSRQ, SNR, and throughput measurements were recorded while verifying that PCI remained constant (PCI 30) throughout.

All measurements confirmed consistent PCI 30 visibility across the entire coverage area with no handover events.

RSRP and SNR values demonstrated expected distribution patterns:

- **Position 1 (In front of RU1):**
  RSRP: -58.6 dBm, SNR: 36.0 dB
- **Position 2 (In front of RU2) :**
  RSRP: -56.2 dBm, SNR: 36.5 dB

The UE successfully connected to cell PCI 30 and achieved high SNR at positions 1 and 2. At position 3 (between the two RUs), the SNR was approximately 4 dB lower than at positions 1 and 2, and throughput was also lower than directly in front of an RU.

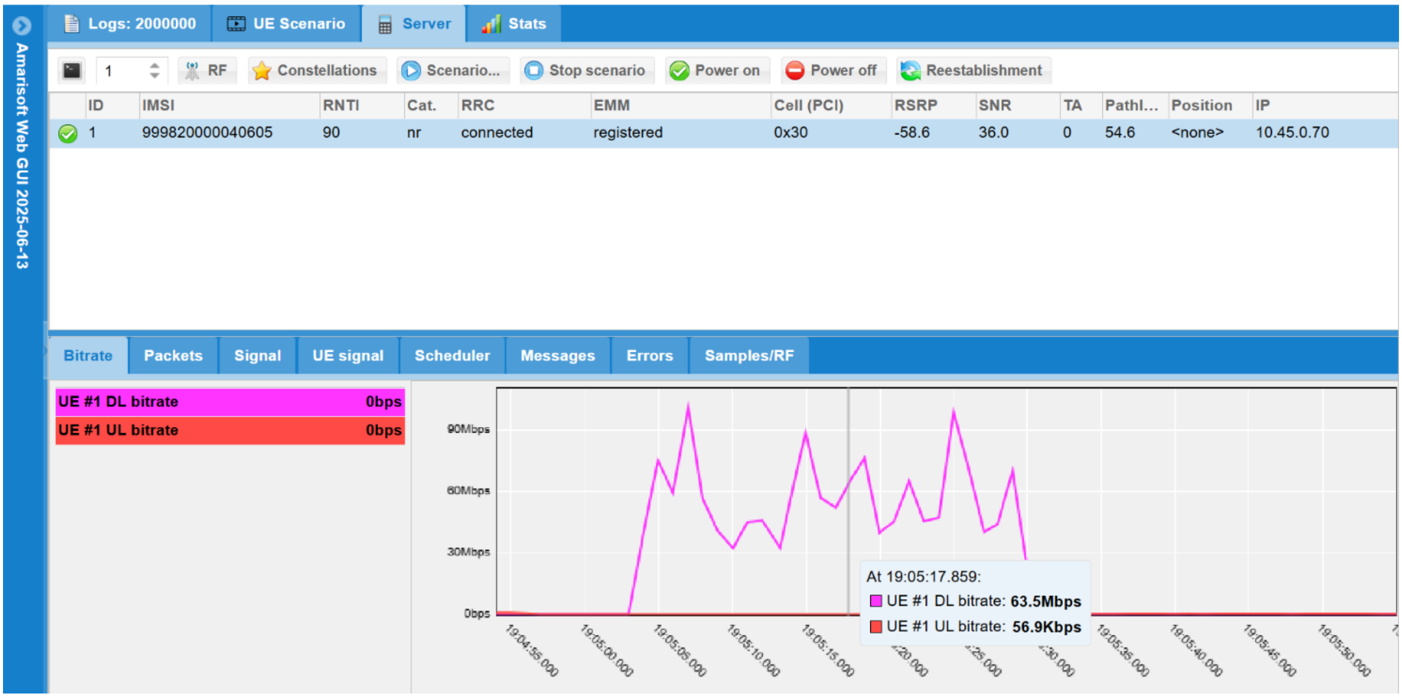

FIGURE 10: RSRP, SNR and DL-Troughput at UE-Position 1

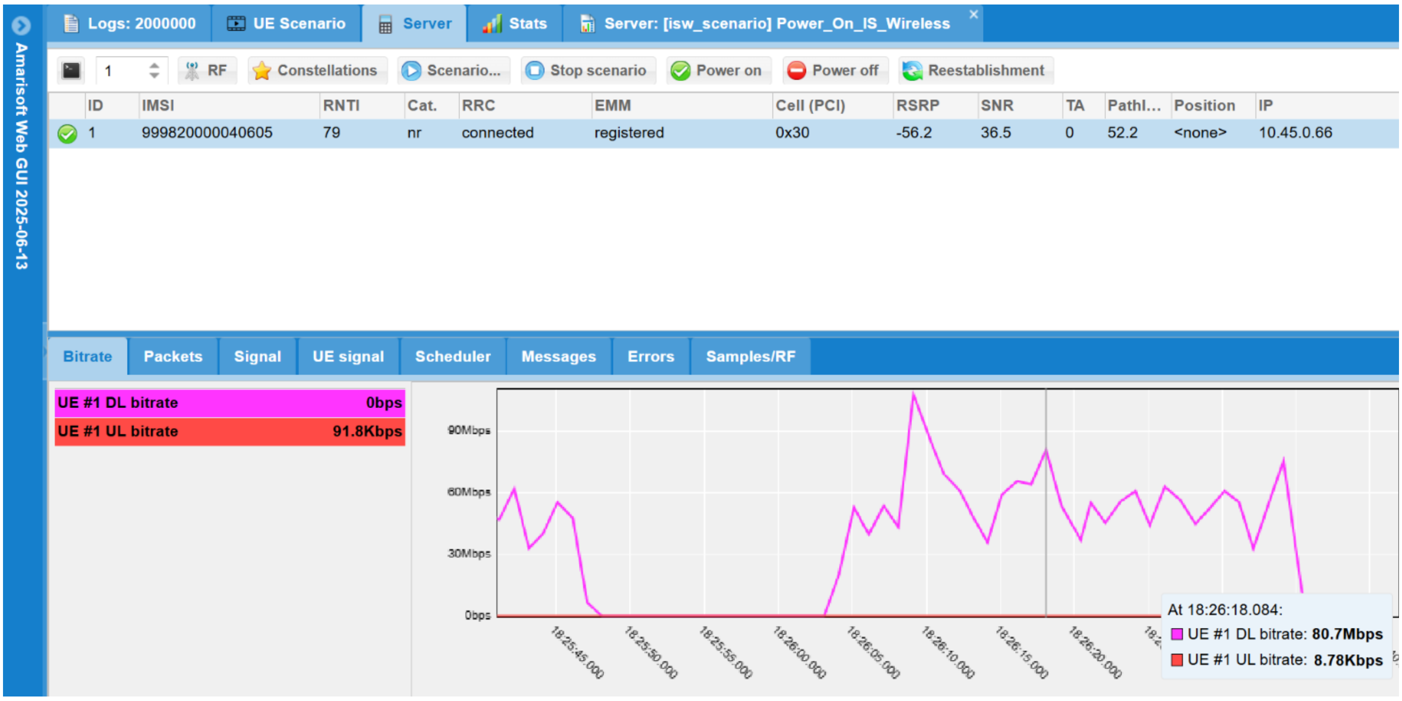

FIGURE 11: RSRP, SNR and DL-Troughput with AMARI Simbox GUI at UE-Position 2

USE CASE 2:
Link Reliability and Throughput Performance

To evaluate the impact of the Link Reliability feature on network performance, tests were conducted at the midpoint position between RU1 and RU2 under varying transmission power conditions. The UE Simulator was tested with the Link Reliability feature both enabled and disabled, comparing performance at high transmit power and reduced transmit power. At each configuration, downlink and uplink throughput, RSRP, RSRQ, and SNR measurements were recorded to evaluate the feature's effectiveness in maintaining connection stability and performance under challenging RF conditions.

Downlink throughput measurements were conducted at the midpoint position under both high and low power configurations, with the Link Reliability feature toggled on and off.

USE CASE 2:
Link Reliability and Throughput Performance

TEST NO.	LINK RELIABILITY ON (MBPS)	LINK RELIABILITY OFF (MBPS)	REMARKS
1	38	45	Stable
2	44	45	Stable
3	41	16	UE re-attach observed
4	42	36	Slight throughput variation
5	42	34	Moderate fluctuation
Average	41.4 Mbps	35.2 Mbps	Stability is better with the feature ON

TABLE 10: High Power Configuration Results

Low Power Configuration (RSRP: -81 dBm, SNR: 21 dB):

TEST NO.	LINK RELIABILITY ON (MBPS)	LINK RELIABILITY OFF (MBPS)	REMARKS
1	25	21	Stable
2	23	25	Stable
3	21	26	Stable
Average	23.0 Mbps	24.0 Mbps	Similar results: low SNR dominates performance

TABLE 11: Low Power Configuration Results

Analysis

The test results demonstrate the successful implementation of the Multi-RU feature, with both radio units operating as a unified cell entity. The consistent PCI visibility and absence of handover events confirm that the system is configured correctly and that coordination between RU1 and RU2 is effective. RSRP distribution patterns validated the expected signal strength topology, with peak values near each RU and minimum values at the mid-point.

At higher power levels (RSRP: -68 dBm, SNR: 30 dB), enabling Link Reliability improved system stability and consistency in throughput. The average throughput increased by approximately 18%, and no reattachments or service drops were recorded. Disabling the feature led to sporadic throughput degradation and a UE re-attach event during one iteration. The comparison of measurement data clearly shows that the bit rate fluctuates less when the reliability function is ON, with throughput up to 10 Mbps higher (≈ 6 Mbps higher on average) than when the reliability function is OFF.

At lower power levels (RSRP: -81 dBm, SNR: 21 dB), throughput performance was primarily constrained by radio conditions. The Link Reliability feature still maintained stable operation but yielded only minor improvements (approximately 4%). At reduced power and weaker radio conditions, both configurations achieved similar average throughput, with the Link Reliability feature providing slightly smoother performance.

Overall, the Link Reliability mechanism effectively enhances connection robustness and throughput stability in moderate-to-strong radio conditions, while providing marginal benefit in low-SNR environments where link quality is already the limiting factor.

Technical Challenges and Limitations

During the plugfest, initial connectivity issues prevented the UE Emulator from establishing a connection to the IS-Wireless network. A concurrent network outage limited remote troubleshooting capabilities, and with constrained vendor support, immediate resolution was not possible during the event.

After switching to a 100 MHz Bandwidth on the RAN side and configuring the Amarisoft UE Simbox differently, a connection could be established. Subsequently, switching back to 40 MHz, the connection still could be established.

CONCLUSION 4

Multi-RU feature within the 5G NR

The Multi-RU feature operated two radio units as a unified cell with consistent PCI visibility and no handover events. At higher power levels (RSRP -68 dBm, SNR 30 dB), Link Reliability improved average throughput by approximately 18% and eliminated UE re-attach events. At lower power levels (RSRP -81 dBm, SNR 21 dB), the feature provided only marginal improvements (ca. 4%) as radio conditions became the dominant limiting factor.

The results confirm that Multi-RU with Link Reliability can extend coverage while maintaining service quality, particularly in border regions between radio units.

An additional challenge was that the short distance between the RUs which prevented testing with RSRP values lower than -100 dBm

5 Network-in-a-box

This Scenario audits the interoperability and functional integration of the AiVader Network-in-a-Box, the brown-iposs „Classification and Root-Cause Analyzing Tool Agent, and the Fraunhofer IIS Positioning Engine through the Qualigon KAFKA data broker. Furthermore, a validation of the positioning and location services of the System-under-test was conducted. The AiVader Network-in-a-Box provides a private 5G network, including the core and an SMO. In the Network-in-a-Box, the brown-iposs CARAT-Agent is installed. It extracts the RRC MeasurementReport messages sent by the UE at the F1 interface of the RAN. The extracted messages are then dispatched over the KAFKA data broker to the Fraunhofer IIS Positioning Engine. The Positioning Engine then converted the raw messages into a usable format, on which it based its position estimation and visualization.

The integration encompasses: The deployment of the AiVader Network-in-a-Box. The deployment of the brown-iposs CARAT-Agent on the Network-in-a-Box and ensuring that it can extract data at the F1 interface. The setup of the KAFKA data broker and the Fraunhofer IIS Positioning Engine. The publication of extracted data from the CARAT-Agent, through the KAFKA data broker, to the Positioning Engine. And the subsequent preprocessing of the raw data by the Positioning Engine.

The validation of the positioning and localization services offered by the System-under-test, which had the precondition of integration being completed, was tested in the following manner. The emulated UE from the UE Emulator or the COTSUE connected to the Network-in-a-Box and, as configured by the RAN, supplied MeasurementReport messages, which the CARAT-Agent extracted. The MeasurementReport messages continue to be sent while the UE, real or emulated, is in motion. The Positioning Engine then supplies an estimation of the UE position, which is compared to the actual position of the UE. The aim is to validate the accuracy of the UE localization.

Test Topology

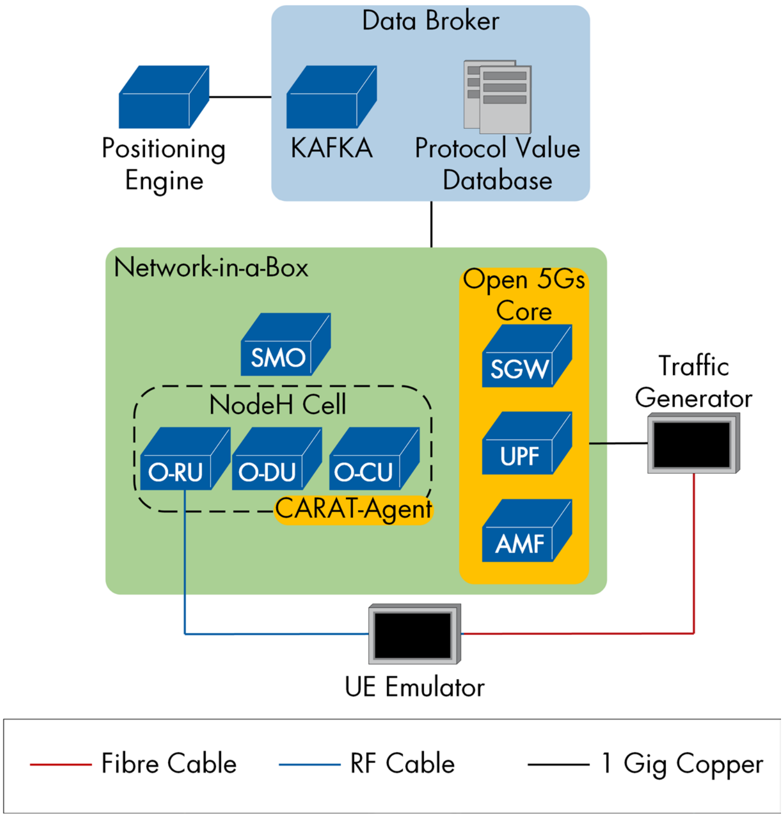

FIGURE 12: Test Setup

COMPONENT	DESCRIPTION
COTS UE	POCO M4 Pro 5G
RAN (O-RU, O-DU, O-CU)	Node-H - Askey Hardware
Network-in-a-Box Server	AiVader - AMD Ryzen 9 16-Cores
5G Core	Open5GS
CARAT-Agent	brown-iposs
Data Broker Industrial PC	Qualigon - Apache Kafka
Positioning Engine	Fraunhofer IIS – including Protocol Skeleton, API Stub and Positioning Stack

TABLE 12: Test Setup Components

Network Configuration and Test Parameters

PARAMETER		VALUE
5G NETWORK	MCC	001
	MNC	01
	TAC	1
	APN	internet
	SST	1
	SD	0x111111
5G NR	ARFCN	650000 (3.75 GHz)
	Bandwidth	100 MHz
	SCS	30 KHz
	Transmit Power	10 dBm
	TDD	10:7/2
	PCI	1

TABLE 13: Test Parameters

Test Results

TEST CASE 1
Network-in-a-box integration test

This test was done with the UE Emulator and the COTS UE. The relative services and components were UP and running. The cell was activated. Subsequently, the emulated UE and the COTS UE were attached to the RAN and registered to the network. The emulated UE was in excellent radio conditions, as the RSRP value of -71 dBm is over the limit of -75 dBm. However, the COTS UE did not achieve excellent radio conditions, as evident by the RSRP value of -83 dBm. The bidirectional UDP traffic from the UE Emulator and the video streaming on the COTS UE were then started. For the emulated UE from the SMO, an UL data rate (based on Total Bytes transmitted) of 72.344 Mbps and a DL data rate (based on Total Bytes received) of 5.475 Kbps were observed.

For the COTS UE from the SMO, an UL data rate (based on Total Bytes transmitted) of 21 Kbps and a DL data rate (based on Total Bytes received) of 54 Kbps were observed. Since both UEs could attach and register to the network, and the data rate was displayed for both in the SMO, this test is evaluated as passed.

TEST CASE 2
CARAT-Agent integration test

This test was done with the UE Emulator and the COTS UE directly, as an extension after test case 1. As described in the Test Case 1 result section above, the relative services were UP and running, and the cell was activated. Then both UEs attach and registered to the network, with the emulated UE achieving excellent radio conditions and the COTS UE being slightly worse. Then, the bidirectional UDP traffic and video streaming were started on the emulated and COTS UE, respectively. Subsequently, for both UEs, the SMO did show the data rate. Then, with a packet capture on the F1AP interface of the Network-in-a-Box RAN, the forwarding of Measurement Report messages by the CARAT-Agent was verified, as shown in Figure 13 for the COTS UE MeasurementReport messages.

Protocol	rsrp	Info
F1AP/NR RRC	-83dBm <= SS-RSRP < -82dBm (74)	SACK (Ack=7, Arwnd=67108864) , ULRRCMessageTransfer, Measurement Report MAC=0xa70d0c2c
F1AP/NR RRC	-83dBm <= SS-RSRP < -82dBm (74)	ULRRCMessageTransfer, Measurement Report MAC=0x80fc4575
F1AP/NR RRC	-83dBm <= SS-RSRP < -82dBm (74)	ULRRCMessageTransfer, Measurement Report MAC=0x2961a017
F1AP/NR RRC	-83dBm <= SS-RSRP < -82dBm (74)	ULRRCMessageTransfer, Measurement Report MAC=0x4ceecb52
F1AP/NR RRC	-83dBm <= SS-RSRP < -82dBm (74)	ULRRCMessageTransfer, Measurement Report MAC=0x39073873

FIGURE 13: Packet Capture on the F1 interface filtered for MeasurementReport messages

The MeasurementReport messages were then already converted into JSON format and published on the KAFKA data broker, as shown in Figure 14.

Since all auditing goals done in test case 1, as well as the forwarding of MeasurementReport messages by the KAFKA data broker from the F1 interface, could be verified, this test is evaluated as passed.


```
"f1ap, RRC\", \"Procedure_Code\": \"id-ULRRCMessageTransfer (13)\", \"Msg_Type\": \"ULRRCMessageTransfer, measurementReport
f1ap\", \"Procedure_Code\": \"Unknown (57)\", \"Msg_Type\": null, \"Msg\": {\"per.choice_index\": \"0\", \"f1ap.F1AP_PDU\":
f1ap, RRC\", \"Procedure_Code\": \"id-ULRRCMessageTransfer (13)\", \"Msg_Type\": \"ULRRCMessageTransfer, measurementReport (
"f1ap, RRC\", \"Procedure_Code\": \"id-ULRRCMessageTransfer (13)\", \"Msg_Type\": \"ULRRCMessageTransfer, measurementReport
"f1ap, RRC\", \"Procedure_Code\": \"id-ULRRCMessageTransfer (13)\", \"Msg_Type\": \"ULRRCMessageTransfer, measurementReport
"f1ap, RRC\", \"Procedure_Code\": \"id-ULRRCMessageTransfer (13)\", \"Msg_Type\": \"ULRRCMessageTransfer, measurementReport
```


FIGURE 14:
MeasurementReport messages converted into JSON format and published on the KAFKA data broker

TEST CASE 3

CARAT-Agent and Positioning engine integration and localization feature test.

This test was done with the UE Emulator and the COTS UE.

For the emulated UE of the UE Emulator, the relative services were UP and running, and the cell was active; thus, it could attach and register to the network as in previous tests. As configured, the emulated UE then started UL UDP traffic and started emulating movement in a diamond-shaped pattern. On the KAFKA data broker, a continuous data stream was recorded representing the raw incoming MeasurementReport messages. This data stream represents the predefined trajectories. The Positioning Engine then received and interpreted the raw MeasurementReport messages. The on-site team implemented several optimizations for this process during the Plugfest, and further possibilities for improvement were identified.

For the COTS UE, the relative services were UP and running, and the cell was active, as well. Therefore, it too could attach and register to the network as in previous tests. The COTS UE was placed in the four corners of Plugfest's venue. The locations are shown on an outdoor global map in Figure 15.

FIGURE 15: COTS UE Position

To ensure the consistency of the coordinate systems, reference locations are given based on the mapping data. At each of the four locations, a continuous data stream of sufficient length was recorded from the KAFKA data broker, representing the raw MeasurementReport messages of the COTS UE scraped by the CARAT-Agent. For extensive testing, a wider and even more complex area needs to be calibrated. However, for the basic characterization of the recorded data, it serves as a representative starting point

The analysis is continuing after the Plugfest timeframe. So far, the only analysis results are that the gathered data from the tests can enable the virtual calibration of the Plugfest venue area and the proven integration chain. As the analysis is ongoing, no results are available on the accuracy of the positioning and localization services offered by the integrated solution. To prove the accuracy, a comparison is needed between the position data produced by the system under test and the emulated or actual position of the emulated or COTS UE, respectively. Since this is still ongoing, this test is currently evaluated as inconclusive.

Technical Challenges and Limitations

The emulated UE stopped sending MeasurementReport messages containing the necessary RF data to the RAN after receiving an RRCReconfig message from the RAN, which disabled the MeasurementReports. This issue occurred in every test run, approximately 10 to 15 minutes after the start. The RAN sent this RRCReconfig message because the UE had requested less data to be sent, due to resource conflicts in the Uu connection. The setup couldn't be configured in time, so that this issue does not occur. Therefore, only test runs with an approximate 10- to 15-minute runtime could be achieved.

CONCLUSION 5

Network-in-a-box

The integration worked continuously. The AiVader Network-in-a-Box continually supplied a private 5G network, from which the CARAT-Agent could then extract the RRC Measurement Report messages at the F1 interface. Those raw messages were then forwarded to the Fraunhofer IIS Positioning Engine, via the KAFKA data broker, where they were then analysed. Both emulated and commercial UEs connected successfully to the private 5G network. During the Plugfest, several protocol and data format optimizations were implemented to improve the data flow between components, and measurement data was collected from four reference locations within the venue. However, the positioning accuracy could not be validated because the venue area has not yet been virtually calibrated—a process needed to establish the reference signal map against which the system's calculated positions can be compared. Additionally, emulated UEs stopped sending measurement reports after 10-15 minutes due to RAN-initiated RRCReconfig messages. Thus, providing only limited data for analysis. Future work will focus on completing the virtual calibration of the venue, resolving this timing constraint, and verifying positioning accuracy.

6 PDU Session Setup for COTS UE with Open6GCore and Multi-User Scalability

As the development of 6G networks advances, early validation of experimental core implementations is crucial for ensuring the readiness of real-world deployments. Open6GCore is a modular, cloud-native platform designed for flexible deployment and integration with a wide range of RAN components and user equipment. It has a different internal architecture than the 5G core, but remains 5G compatible by maintaining the interfaces towards the RAN. In this test, it was used as a functional 5G core, rather than evaluating 6G-specific capabilities. This evaluation presents a comprehensive two-phase assessment: first, validating Open6GCore with commercial user equipment under varying traffic conditions; second, assessing multi-user scalability using the Amarisoft UE Simbox to verify concurrent connection stability and sustained traffic performance.

The initial phase focuses on validating interoperability with commercial 5G user equipment and a multi-vendor O-RAN setup under real application traffic, measuring latency and jitter across different bandwidth and frame size configurations to establish a performance baseline.

Testing with multiple physical user equipment devices presents practical challenges due to equipment availability and the complexity of coordination. To address these constraints, the second phase utilizes the Amarisoft UE Simbox as both a testing tool and device under test, evaluating its capability to emulate realistic multi-user behavior in an over-the-air environment with 16 concurrent emulated users.

Test Topology

The testbed comprises an end-to-end 5G architecture where user equipment connects over the air interface to RAN components (O-RU, O-DU, O-CU), which interface via Ethernet with the core network. Traffic generators are positioned behind the N6 interface of the core and behind the UE. Phase 1 utilizes a Teltonika RUTX50 Industrial 5G Router as the commercial UE, while Phase 2 employs the Amarisoft UE Simbox to emulate 16 concurrent users.

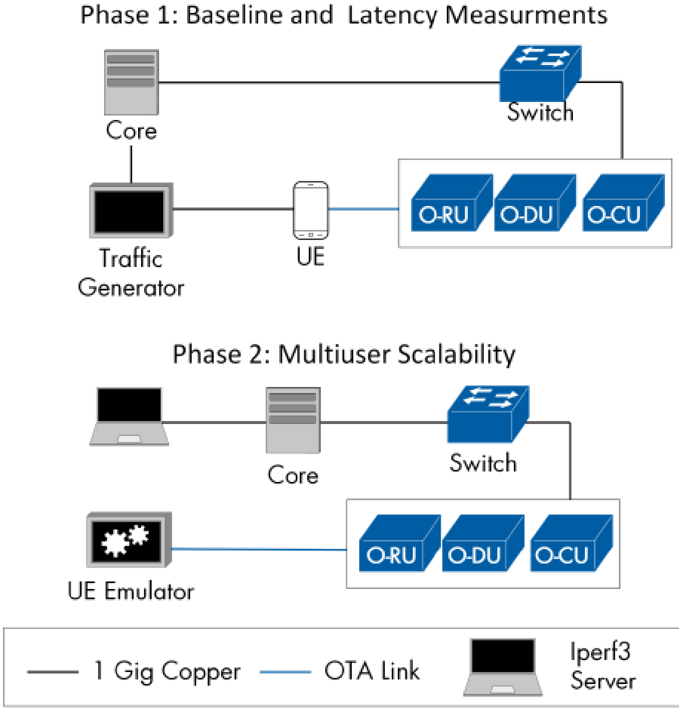

FIGURE 16: Test Setup

COMPONENT	DESCRIPTION
RAN	airpuls RAN (O-DU, O-CU)
Core	Fraunhofer FOKUS Open6GCore (6G core implementation operating in 5G-compatible mode)
UE Emulator (Phase 2)	Amarisoft UE Simbox

TABLE 14: Test Setup Components

Network Configuration and Test Parameters

The evaluation uses a common 5G radio configuration for both single-user and multi-user phases, differing only in test traffic endpoints:

PARAMETER	VALUE
PLMN	999-99
Center Frequency / Bandwidth	n78 / 100 MHz
TDD Configuration / Periodicity	2-1-2 (DDSUU) / 2.5 ms
Transmit Power (RU) / UE Distance	24 dBm / ca. 2 m
Phase 1 Traffic Endpoints	Commercial UE (Teltonika RUTX50) ↔ Load Generator
Phase 2 Traffic Endpoints	16 Emulated UEs (Amarisoft UE Simbox) ↔ iperf3 Servers
Phase 1 Throughput Settings	1, 10, 50 Mbit/s; 128, 512, 1500 Byte frames; UL/DL
Phase 2 Throughput Settings	12 Mbit/s per UE; 1320 Byte packets; Downlink; UDP
KPIs Measured	Latency (average/max), Jitter (average/max)

TABLE 15: Network Configuration and Test Parameters

Test Results

Phase 1:
Single-User Baseline Testing

Open6GCore successfully established and maintained PDU sessions across all tested conditions with no observed session drops or protocol instability. Table 16 summarizes the latency and jitter metrics for all tested configurations:

# RUN	DIRECTION	BANDWIDTH (MBIT/S)	FRAME SIZE (BYTE)	AVG LATENCY (MS)	MAX LATENCY (MS)	AVG JITTER (MS)	MAX JITTER (MS)
1	Downlink	1	128	11.1	133.62	1.84	90.03
	Uplink	1	128	47.74	102.77	2.21	93.55
2	Downlink	1	512	9.13	106.82	2.6	92.04
	Uplink	1	512	45.89	99.97	7.3	90.94
3	Downlink	1	1500	5.57	110.75	0.76	92.94
	Uplink	1	1500	48.01	100.59	18.05	83.79
4	Downlink	10	128	8.73	15.15	0.16	8.27
	Uplink	10	128	24.17	103.9	0.22	97.35
5	Downlink	10	512	8.29	90.28	0.73	55.29
	Uplink	10	512	43.09	104.2	0.76	94.64
6	Downlink	10	1500	8.54	78.88	1.88	94.18
	Uplink	10	1500	45.33	103.32	2.06	97.85
7	Downlink	50	1500	10.94	170.77	0.39	17.97
	Uplink	50	1500	38.37	106.13	0.34	96.74

TABLE 16: Latency and Jitter Results

The results confirm that Open6GCore can establish and maintain PDU sessions across all test conditions, with no observed session drops or protocol instability. This demonstrates successful multi-vendor integration between the RAN components and Open6GCore. Figures 17 and 18 visualize the same data, highlighting latency and jitter trends across the different test configurations.

Additionally, a preliminary iperf3 test was conducted to determine the single-user maximum downlink UDP throughput, which measured approximately 200 Mbit/s. This value serves as a baseline reference for the subsequent multi-user scalability evaluation in Phase 2.

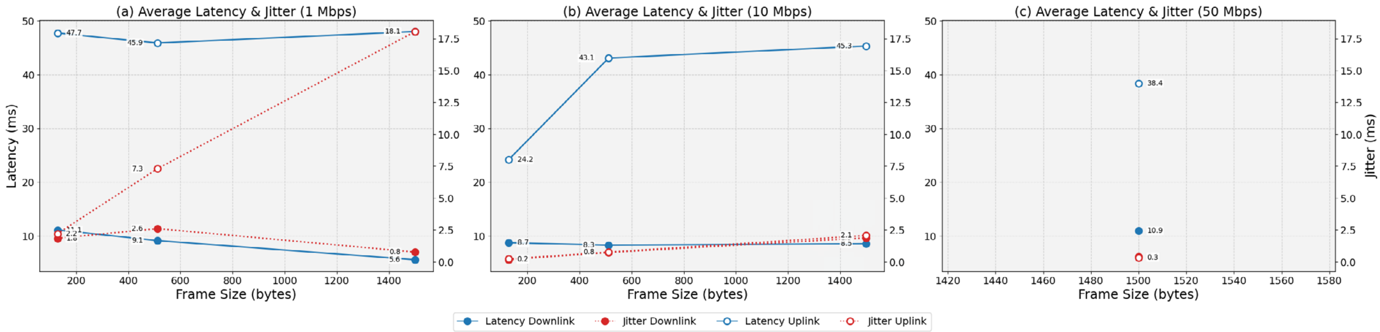

FIGURE 17: Average latency and jitter for each test condition, for both downlink and uplink

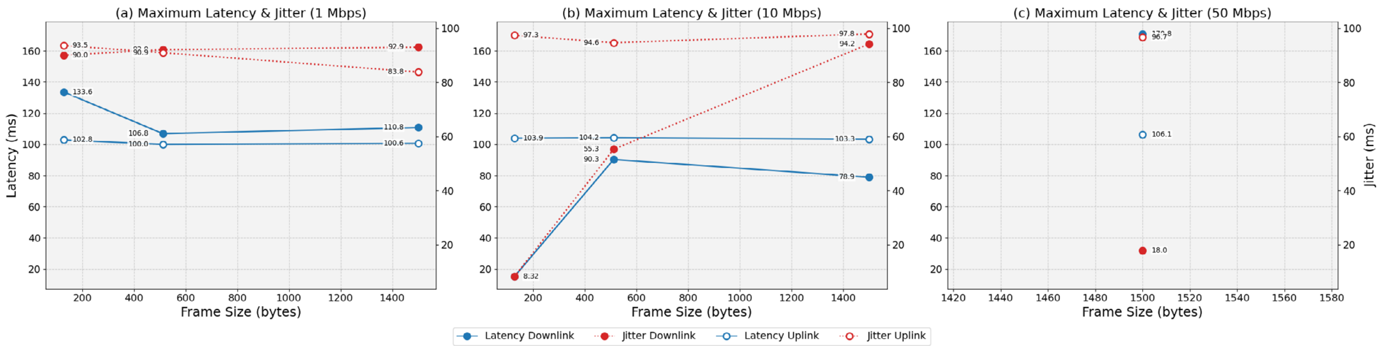

FIGURE 18: Maximum latency and jitter for each test condition, for both downlink and uplink

PHASE 2:
Multi-User Scalability Testing

All 16 emulated user equipment devices successfully connected to the airpuls Open RAN system over the air and received IP address assignments from both FOKUS Open6GCore and Open5GS core networks. The throughput measurements presented below were conducted with the Open6GCore, while the Open5GS 5G core was tested only for device registration capability.

The aggregate throughput across all 16 users achieved approximately 200 Mbit/s, matching the single-user baseline capacity established in Phase 1. Individual user data rates averaged approximately 12 Mbit/s per UE as configured in the iperf3 parameters, with observable fluctuations in the total downlink data rate attributed to over-the-air channel variability. Performance metrics captured from iperf3 log files confirmed consistent per-user throughput distribution across all emulated devices.

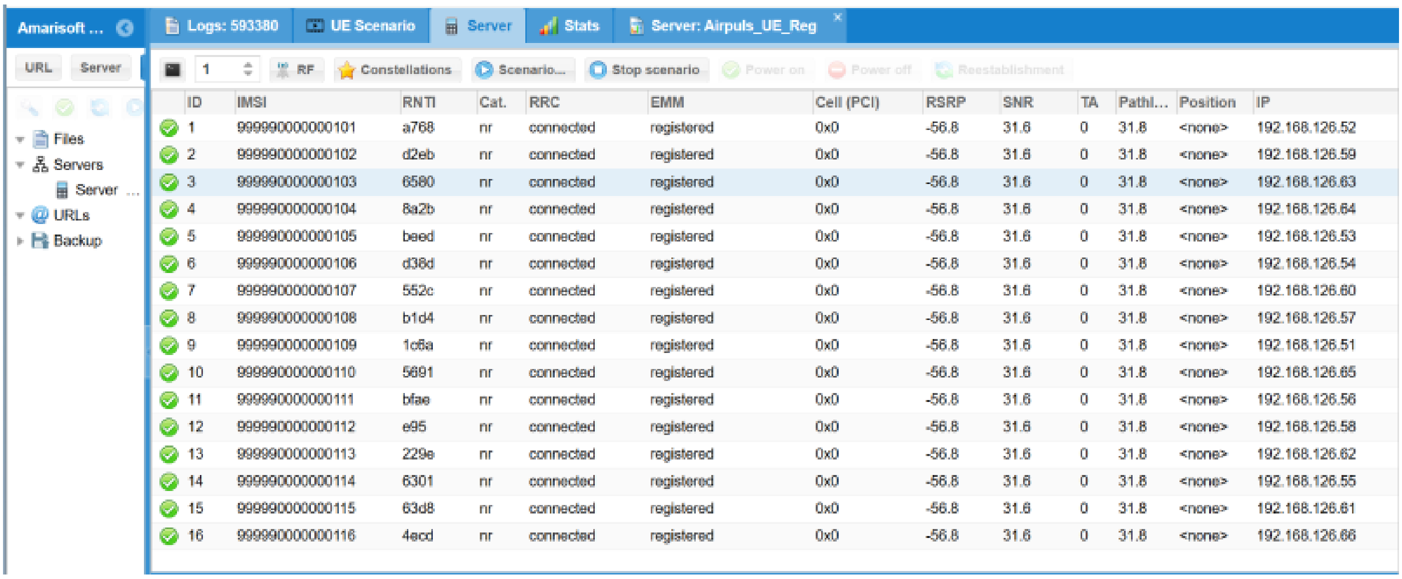

ID	IMSI	RNTI	Cat.	RRC	EMM	Cell (PCI)	RSRP	SNR	TA	Pathl...	Position	IP
1	999990000000101	a768	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.52
2	999990000000102	d2eb	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.59
3	999990000000103	6580	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.63
4	999990000000104	8a2b	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.64
5	999990000000105	beed	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.53
6	999990000000106	d38d	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.54
7	999990000000107	552c	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.60
8	999990000000108	b1d4	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.57
9	999990000000109	1c6a	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.51
10	999990000000110	5691	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.65
11	999990000000111	bfae	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.56
12	999990000000112	e95	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.58
13	999990000000113	229e	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.62
14	999990000000114	6301	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.55
15	999990000000115	63d8	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.61
16	999990000000116	4ecd	nr	connected	registered	0x0	-56.8	31.6	0	31.8	<none>	192.168.126.66

FIGURE 19: Amarisoft UE Simbox Dashboard - Displaying all 16 UEs successfully connected to the network

Figure 20 shows the total downlink data rate for the 16 UEs in the Amarisoft web GUI. The data rate fluctuates slightly due to the inconsistent OTA channel.

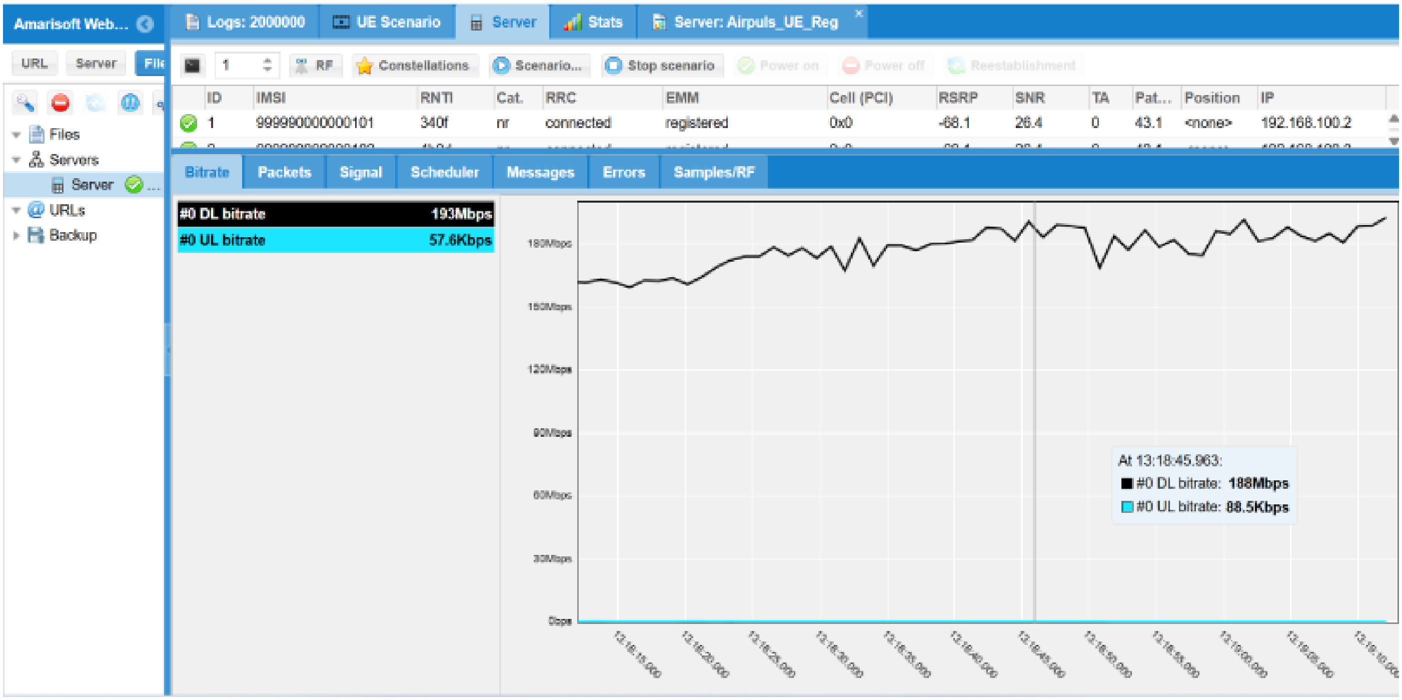

FIGURE 20: Aggregate Throughput Graph - Total downlink data rate over time for all 16 users

Analysis

In Phase 1, downlink latency is consistently lower than uplink latency, with average values typically ranging from 7 to 14 ms and remaining relatively stable even as frame size or bandwidth increases. Uplink latency is higher, averaging around 30–54 ms depending on load and frame size.

Jitter values remain generally low for downlink, while uplink jitter increases with larger frame sizes at low bandwidths (notably at 1 Mbit/s and 1500-byte frames, where maximum jitter exceeds 90 ms). For the highest bandwidth (50 Mbit/s), both average and maximum jitter values are moderate, indicating stable transport when the radio is fully utilized.
The graphical results reinforce these trends: both maximum and average latency/jitter are stable across frame sizes for the downlink, while uplink exhibits greater sensitivity, especially at low rates and large frames. The observed outlier is the higher maximum latency for the downlink at 50 Mbit/s and 1500-byte frames (peak ca. 174 ms), likely due to burst effects at maximum throughput. Overall, the core and RAN together handle all tested conditions without loss of functionality or extreme delay spikes. This functional stability provides confidence for further 6G research using this platform.

In Phase 2, the test results demonstrate successful integration of the UE emulator with the 5G system and the multi-user scalability of the 5G Open RAN system. The achievement of aggregate throughput matching the single-user baseline of 200 Mbit/s indicates effective radio resource scheduling and allocation across multiple simultaneous users, with individual user data rates averaging the configured 12 Mbit/s target. Observed fluctuations in total downlink data rate represent expected behavior in over-the-air radio environments due to multipath propagation and channel variability. Throughput measurements were conducted exclusively with the Open6GCore; the Open5GS core was tested only for device registration.

Technical Challenges and Limitations

Phase 1: Testing with small frame sizes (128 bytes and 512 bytes) at 50 Mbit/s resulted in high packet drop rates due to the high packet-per-second load that small frames generate at higher bandwidths. Therefore, valid measurements could only be collected at 50 Mbit/s for the 1500-byte frame size. This outcome does not affect the integrity of results for other frame sizes and bandwidth combinations.
Phase 2: No significant technical challenges or limitations were encountered during this test. All 16 emulated UEs successfully connected to the network and achieved the expected traffic performance targets without requiring extensive troubleshooting beyond the initial single-user baseline configuration.

CONCLUSION 6

PDU Session Setup for COTS UEs with Open6GCore & Multi-User Scalability

The Open6GCore successfully established and maintained PDU sessions across all tested conditions without any session drops or protocol instability, showing reliable multi-vendor integration between airpuls RAN components and Open6GCore. Downlink latency consistently remained lower than uplink, averaging 7–14 ms versus 30–54 ms. Meanwhile, jitter was low for downlink but increased in uplink under low bandwidth and large frame sizes (e.g., over 90 ms at 1 Mbps with 1500-byte frames). Building on this stable baseline, the network successfully connected all 16 emulated user equipment devices and maintained steady connections, achieving an aggregate throughput of approximately 200 Mbps—matching the single-user reference—with individual UEs averaging 12 Mbps each. Observed rate fluctuations due to over-the-air variability highlight the real-world effects of channel variations. Finally, while throughput measurements were performed exclusively with Open6GCore, successful UE registration with Open5GS confirmed multi-vendor interoperability across different core implementations.

7 Compare LiFi Backhaul to Ethernet Backhaul

The integration of alternative backhaul technologies into future 6G architectures is crucial for enhancing network flexibility and capacity. This test investigates the use of a LiFi optical wireless link as a backhaul medium between the RAN and core network, replacing the typically used 1GbE Ethernet connection. The Open6GCore has a different internal architecture than the 5G core, but remains 5G compatible by maintaining the interfaces towards the RAN. In this test, it was used as a functional 5G core, rather than evaluating 6G-specific capabilities. The primary objective is to compare the end-to-end latency and jitter characteristics of standard Ethernet backhaul with those of a LiFi-based one under identical radio and core configurations.

The wireless backhaul link is intended to behave transparently like a cable to the 5G system. This adds great flexibility to deployments, particularly for nomadic scenarios such as in agriculture, where cells are placed temporarily where coverage is needed but wired connections are difficult to achieve. Potential Use Cases also include nomadic cells on robots or drones.

Test Topology

The setup uses the same 6G testbed as the previous test scenario 6, comprising the Open6GCore, RAN components, and a COTS industrial 5G router, with traffic generated via the load generator. The topology is based on the previous setup, with the addition of LiFi transceivers placed between the RAN and the core network. The LiFi pair provides a transparent Layer 2 bridge, replacing the Ethernet backhaul. The rest of the setup remains unchanged: the UE connects to the RAN over the 5G air interface, while measurement endpoints generate and receive IP traffic in both directions.

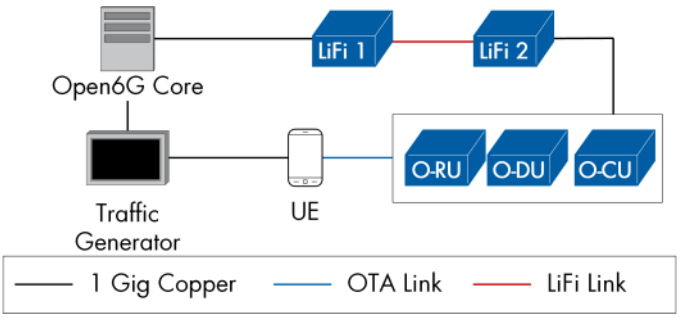

FIGURE 21: Test Setup

The two LiFi transceivers were placed at approximately 4 meters distance and carefully aligned. To isolate the observed effects, the LiFi link was later tested independently to evaluate its inherent transport characteristics.

COMPONENT	DESCRIPTION
RAN	airpuls RAN (O-DU, O-CU)
Core	Fraunhofer FOKUS Open6GCore (6G core implementation operating in 5G-compatible mode)
LiFi Backhaul	Fraunhofer HHI fiber-like LiFi Transceivers

TABLE 17: Test Setup Components

Network Configuration and Test Parameters

This section presents the key configuration parameters and test traffic settings used in the evaluation.

PARAMETER		VALUE
5G NETWORK	PLMN	999-99
	Frequency Band / Channel Bandwidth	n78 / 100 MHz
	TDD Configuration / Periodicity	2-1-2 (DDSUU) / 2.5 ms
	Transmit Power (RU) / UE Distance	24 dBm / ca. 2 m
LIFI BACKHAUL LINK	Center Wavelength	850 nm
	Nominal Throughput	1 Gbps @ 150 m, 250 Mbit/s @ 300 m
	Average Latency	2 ms
	Protocol	ITU-T G.9991 (Wave 2 G.hn) with ARQ support
	Type	Half-Duplex
	Transceiver Distance	ca. 4 m
TEST TRAFFIC	Bandwidth Settings	1 Mbit/s, 10 Mbit/s, 50 Mbit/s
	Frame Sizes	128 bytes, 512 bytes, 1500 bytes
	Traffic Direction	Uplink and Downlink
	Test Duration	2 minutes per run
	KPIs Measured	Average Latency, Maximum Latency, Average Jitter, Maximum Jitter

TABLE 18: Network Configuration and Test Parameters

Test Results

The performance evaluation compared end-to-end latency and jitter between the establishment of a PDU Session for a COTS UE with Open6GCore scenario, and the LiFi backhaul configuration. Additionally, standalone LiFi testing was conducted to isolate the link's inherent transport characteristics.

Latency and Jitter Performance Results with LiFi Backhaul

# RUN	DIRECTION	BANDWIDTH (Mbit/s)	FRAME SIZE (Byte)	AVG LATENCY (ms)	MAX LATENCY (ms)	AVG JITTER (ms)	MAX JITTER (ms)
1	Downlink	1	128	9.38	19.33	1.87	34.07
	Uplink	1	128	54.34	111.85	2.23	100.47
2	Downlink	1	512	9.22	40.55	3.03	10.75
	Uplink	1	512	49.09	116.48	7.38	91
3	Downlink	1	1500	7.74	39.85	2.58	27.88
	Uplink	1	1500	51.75	130.29	18.28	90.97
4	Downlink	10	128	9.52	46.84	0.16	33.37
	Uplink	10	128	30.5	129.03	0.2	97.89
5	Downlink	10	512	10.33	40.82	0.73	29.27
	Uplink	10	512	49.18	111	0.75	94.83
6	Downlink	10	1500	10.17	40.65	1.92	33.36
	Uplink	10	1500	49.34	130.11	2.08	102.67
7	Downlink	50	1500	13.99	173.86	0.39	16.46
	Uplink	50	1500	42.95	130.11	0.21	99.98

TABLE 19: Measured Latency and Jitter Values Over Core, LiFi System, and RAN Components

Latency and Jitter Performance Results - LiFi Link Only

# RUN	DIRECTION	BANDWIDTH (Mbit/s)	FRAME SIZE (Byte)	AVG LATENCY (ms)	MAX LATENCY (ms)	AVG JITTER (ms)	MAX JITTER (ms)
1	Downlink	1	128	2.23	35.48	1.1	35.31
	Uplink	1	128	2.07	35.71	1.09	35.32
2	Downlink	1	512	2.12	35.25	-	-
	Uplink	1	512	2.03	35.17		
3	Downlink	1	1500	2.15	35.25		
	Uplink	1	1500	2.14	34.49		
4	Downlink	10	128	0.94	35.91		
	Uplink	10	128	0.93	35.37		
5	Downlink	10	512	2.36	35.91		
	Uplink	10	512	2.34	35.56		
6	Downlink	10	1500	2.35	35.6		
	Uplink	10	1500	2.42	35.75		
7	Downlink	50	128	1.48	36.91	0.022	0.033
	Uplink	50	128	1.53	36.94	0.022	0.032
8	Downlink	50	512	3.44	35.54		
	Uplink	50	512	3.43	35.4		
9	Downlink	50	1500	2.19	37		
	Uplink	50	1500	2.33	37.05		

TABLE 20: Measured Latency and Jitter Values Over The LiFi System

Figures 22 and 23 visualize average and maximum latency for all tested conditions. For illustration and comparison, the graphs also include results from the previous scenario (6G setup with Ethernet backhaul), allowing direct visual estimation of the performance impact introduced by integrating the LiFi backhaul.

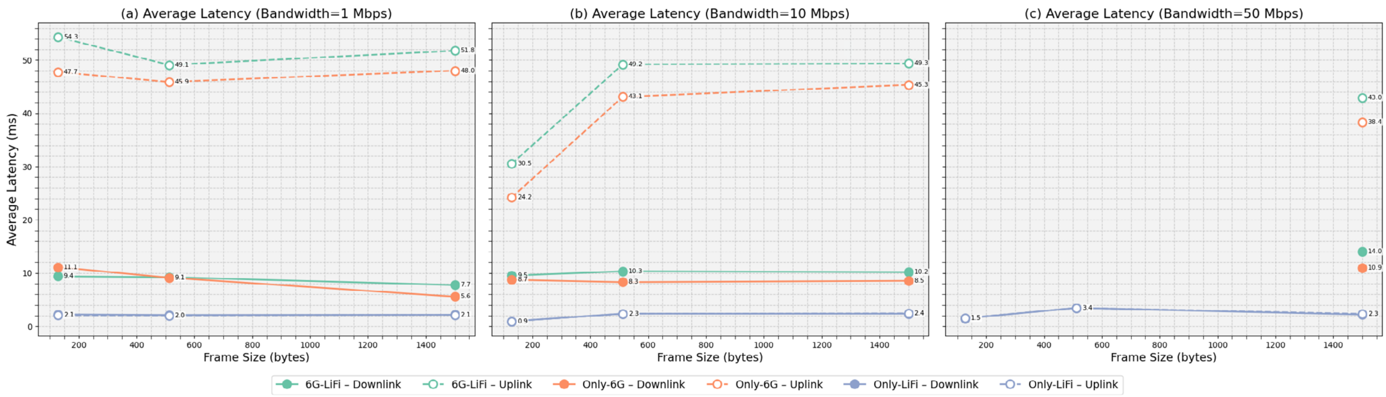

FIGURE 22: Average latency comparison

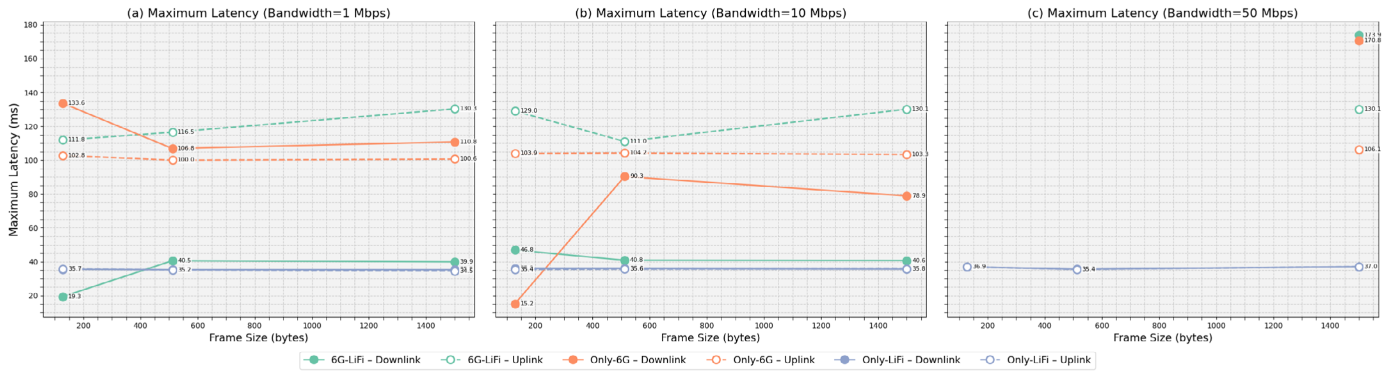

FIGURE 23: Maximum latency comparison

Similar to the previous two graphs, Figures 24 and 25 visualize average and maximum jitter for all tested conditions. including also the results from the previous scenario (6G setup with Ethernet backhaul).

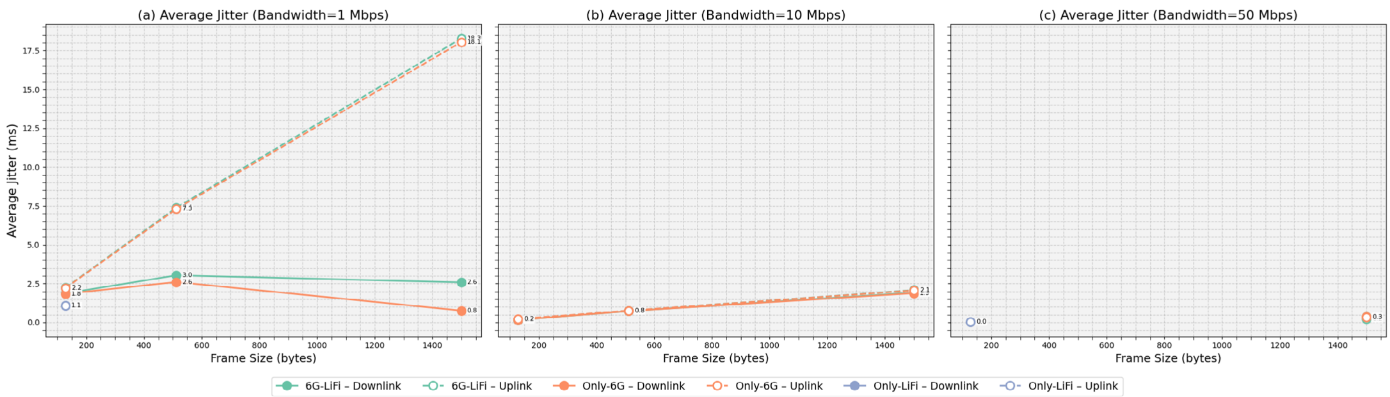

FIGURE 24: Average jitter comparison

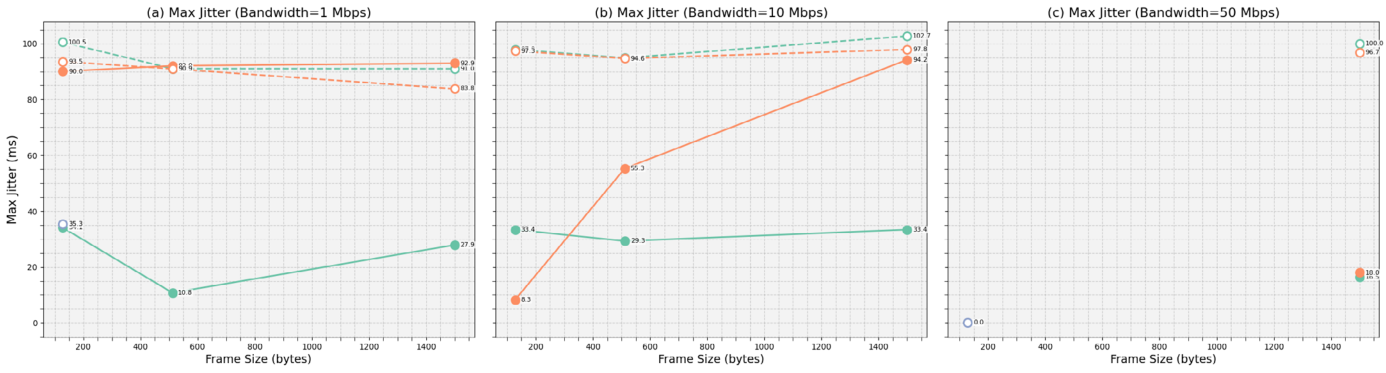

FIGURE 25: Maximum jitter comparison

Analysis

The measurements show that the 6G+LiFi setup maintained stable end-to-end connectivity with no PDU session drops or packet loss across all tested bandwidths and frame sizes. However, a clear and notable latency asymmetry was observed between the uplink and downlink directions, despite using a perfectly symmetric traffic load.

This asymmetry was particularly pronounced in the uplink direction. While downlink average latency remained consistent with the Ethernet baseline measured in the previous test scenarios, with an average increase of approximately 2-3 ms, uplink latency showed an increase, depending on the specific test configuration, of about 3-6 ms. This asymmetry was observed across all bandwidth settings and frame sizes tested.

Interestingly, the LiFi-only test results demonstrate that this asymmetry does not originate from the optical link itself. When the optical link was evaluated separately between two Ethernet endpoints, both the uplink and downlink paths exhibited nearly identical latency values, with an average one-way latency of approximately 2-3 ms and minimal jitter (around 1 ms). This proves that the LiFi system itself operates symmetrically and that the observed asymmetry in the integrated setup appears when the LiFi backhaul is combined with the 5G RAN and core system.

At higher loads (e.g., 50 Mbit/s), the overall latency peaks of the 6G+LiFi setup (ca. 170-174 ms) matched those of the wired baseline. This indicates that the optical backhaul does not amplify worst-case delays compared to the Ethernet configuration. Jitter values remained low and stable across most test conditions, with minor peaks likely linked to momentary synchronization events. The additional latency induced by the wireless backhaul link was smaller than expected and remains tolerable for most enterprise Use Cases.

Technical Challenges and Limitations

Testing with small frame sizes (128 bytes and 512 bytes) at 50 Mbit/s resulted in high packet drop rates due to the high packet-per-second load that small frames generate at higher bandwidths. Under such conditions, neither the testbed hardware nor the protocol stack is typically optimized for sustained line-rate throughput with minimum-sized packets. Therefore, valid measurements could only be collected at 50 Mbit/s/for the 1500-byte frame size. This outcome is consistent with known system constraints and does not affect the integrity of results for other frame sizes and bandwidth combinations.

The LiFi transceivers were positioned at approximately 4 meters distance for this evaluation. Future work should include testing at longer distances to assess performance degradation and validate the manufacturer's specified range capabilities. Additionally, the tests were conducted under static conditions with a single UE. Dynamic mobility tests and multi-user conditions would provide further insight into the LiFi backhaul's scalability and robustness under real-world scenarios.

CONCLUSION 7

Compare LiFi Backhaul to Ethernet Backhaul

The LiFi backhaul integration showed full functional compatibility with the Open6GCore and RAN, maintaining stable PDU sessions without session loss or connection drops. Latency and jitter values remained within the expected range for 5G systems, with the LiFi backhaul introducing only minor asymmetric latency between uplink and downlink, confirming minimal overhead. Standalone LiFi testing confirmed the optical link's inherent symmetry and low latency. Overall, the wireless optical backhaul functioned transparently relative to Ethernet, supporting its use as a flexible alternative for nomadic or cable-restricted deployments.

Final Summary & Conclusion of the xG-ALOE Plugfest 2025

The first xG-ALOE Plugfest successfully demonstrated the value of collaborative, multi-vendor testing for advancing Open RAN, 5G, and early 6G technologies. Across eight diverse scenarios, participants validated key aspects of network integration, performance, security, and flexibility, while also identifying practical challenges and areas for future improvement.

The event confirmed that 5G networks can reliably support advanced Use Cases such as time-sensitive networking, multi-user scalability, and robust positioning services. Novel components, including protocol value grabbers, positioning engines, and network-in-a-box solutions, were functionally integrated, with data flows and interoperability chains operating as intended. The tests also highlighted the importance of precise time synchronization, the effectiveness of link reliability mechanisms, and the ability of Open RAN systems to maintain stable connections and throughput under realistic multi-user loads.

Security remains a critical focus, as shown by the integration of intrusion detection systems within the RAN. While some scenarios encountered interoperability or configuration challenges, these were valuable learning opportunities, emphasizing the need for continued collaboration and refinement of interfaces and protocols.

The evaluation of alternative backhaul technologies, such as optical wireless links, demonstrated that wireless backhaul can provide performance comparable to fiber-based backhaul, with only minor integration-induced asymmetries. This opens new possibilities for flexible and nomadic network deployments, especially in environments where cabling is impractical.

Testing with open source and experimental core network implementations (Open5GS and Open6GCore) showed end-to-end Use Cases and established a stable baseline for future 6G-oriented research. The plugfest environment enabled rapid troubleshooting, knowledge sharing, and iterative improvement, turning technical challenges into actionable insights for all participants.

In summary, the plugfest fostered a spirit of innovation and cooperation, accelerating progress toward open, flexible, and interoperable network solutions. The lessons learned and solutions developed during this event will serve as a foundation for future plugfests and ongoing advancements in the field.

xG-ALOE an initiative powered by
Fraunhofer Heinrich-Hertz-Institut, HHI

Einsteinufer 37, D–10587 Berlin

https://xg-aloe.de/
info@xG-ALOE.de

BERLIN 2025