Grant Agreement No: 101016941

Full title: Open cooperative 5G experimentation platforms
for the industrial sector NetApps

Start date: 01/01/2021 End date: 30/06/2024

Duration: 42 Months

Project URL: *https://www.5G-induce.eu/*

Deliverable D6.2

5G-INDUCE Showcasing of Industrial NetApp use case scenarios

Document type	Deliverable		
Title	5G-INDUCE Showcasing of Industrial NetApp use case scenarios		
Contractual due date	30/6/2024	Actual submission date	30/8/2024
Nature	Report	Dissemination Level	Public
Lead Beneficiary	OTE		
Responsible Author	Christina Lessi		
Contributions from	OTE, UBITECH, ERC, UWS, SUITE5, BEKO, CNIT, WIND3, FORD, ILINK, ININ, OCU, PPC		

Revision history

Version	Issue Date	Changes	Contributor(s)
v0.1	11/12/2023	ToC	OTE
v0.2	20/05/2024	Section 3	ERC
v0.3	11/06/2024	Sections 1,2,4	OTE, UWS, SUITE5
v0.4	28/07/2024	Sections 5,6	OTE, BEKO
v0.5	19/08/2024	Sections 4, 5	OCU
V0.6	29/08/2024	Update all sections	OTE
V1.0	30/08/2024	Final revision and formatting	CNIT

Disclaimer

The content of the publication herein is the sole responsibility of the publishers and it does not necessarily represent the views expressed by the European Commission or its services.

While the information contained in the documents is believed to be accurate, the authors(s) or any other participant in the 5G-INDUCE consortium make no warranty of any kind with regard to this material including, but not limited to the implied warranties of merchantability and fitness for a particular purpose.

Neither the 5G-INDUCE Consortium nor any of its members, their officers, employees or agents shall be responsible or liable in negligence or otherwise howsoever in respect of any inaccuracy or omission herein.

Without derogating from the generality of the foregoing neither the 5G-INDUCE Consortium nor any of its members, their officers, employees or agents shall be liable for any direct or indirect or consequential loss or damage caused by or arising from any information advice or inaccuracy or omission herein.

Copyright message

© 5G-INDUCE Consortium, 2020-2023. This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both. Reproduction is authorised provided the source is acknowledged.

Table of Contents

List of Figures

List of Tables

Glossary of terms and abbreviations used

Abbreviation / Term	Description
5G	Fifth Generation
AGV	Automated Guided Vehicles
AI	Artificial Intelligence
API	Application Programming Interface
BW	Bandwidth
CPU	Central Processing Unit
DL	Downlink
E2E	End-to-End
EM	Emergency Maintenance
ExFa	Experimentation Facility
FHD	Full High Definition
gNB	gNodeB
GPS	Global Positioning System
HDF	Heat Dissipation Failure
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
KPI	Key Performance Indicator
LPM	Logistic Process Manager
MEC	Multi Edge Computing
ML	Machine Learning
MP	Manual Preventive Maintenance
MTTR	Mean Time To Repair
NAO	Network Application Orchestrator
nApp	Network Application
NR	New Radio
NSA	Non-Standalone
OSF	Overstrain Failure
PWF	Power Failure

QR	Quick-Response
RNF	Random Failure
RSRP	Reference Signal Receiving Power
RSRQ	Reference Signal Received Quality
RTMP	Real-Time Messaging Protocol
SA	Standalone
SIM	Subscriber Identity Module
SINR	Signal Interference Noise Ratio
TBR	Time Between Repairs
TCP	Transmission Control Protocol
TWF	Tool Wear Failure
UAV	Unmanned Aerial Vehicle
UC	Use Case
UL	Uplink
VNF	Virtual Network Function
VR	Virtual Reality
VSLAM	Visual Simultaneous Localization And Mapping
XMS	Experience Manager System

Executive Summary

D6.2, as part of the 5G-INDUCE project, is the outcome of Task 6.3 carried out under WP6. It is the report on the Showcasing of Industrial nApp use case scenarios and the evaluation of the results at a technical level. This report is a demonstration of the practical application and effectiveness of 5G technology in the Industry 4.0 sector and at the same time the performance of nApps implemented with the aim of proposing specific solutions to challenges faced by the sector.

This report details the environment, processes and results of the experimentation of the 8 use cases of 5G-INDUCE. The demonstrations focused on testing services developed on 5G technologies and related to the Industry 4.0 sector. The aim was to provide useful information about the performance, challenges and opportunities that the use of 5G can offer in combination with technological solutions such as machine learning, data analysis, etc. Additionally, it describes the execution of the tests and their results, while also evaluating the contribution of the network characteristics to the extraction of these results.

At the same time, this deliverable can serve as a source for understanding the potential impact of 5G on improving operational efficiency, security and overall performance in Industry 4.0, by studying how the proposed solutions affected the operation in the use cases examined.

From the specific report of 5G-INDUCE, there are key insights that can be extracted regarding:

- The **utilization of 5G networks** and the advantages offered by its characteristics such as low latency and high throughput.
- Methods and procedures for the **effective and immediate utilization of the data collected**, so that the services offered are more efficient.
- The **implementation of edge infrastructure and its combination with core computing**, offering a choice between the two sites depending on the application’s requirements.

1 Introduction

This introductory section describes the general purpose of this deliverable and its relation to the project's workflow and other deliverables. The structure is also presented at the end of the section.

1.1 Deliverable Purpose

D6.2 reports the outcomes of T6.3. It focuses on the showcasing of nApps over the experimentation facilities, providing the testing results. Additionally, it will show the benefits of the developed nApps and the usability of the 5G platform. This report is a demonstration of the practical application and effectiveness of 5G technology in the Industry 4.0 sector

1.2 Relation to other deliverables and tasks

As a direct follow-up to the evaluation activities carried out in T6.2, T6.3 focuses on actions that highlight the benefits of the developed nApps and the usability of the 5G platform. The deliverable as a report of the activities of T6.3, takes into account activities carried out in other WPs:

- The specific use case scenarios identified in T5.1 and information about them are in D5.1.
- From T6.1, it leverage the results on nApp adaptability to 5G platform over DevOps testbed.
- the release of nApp developments and integration methodology and requirements was carried out in the framework of WP4 activities (T4.1). In the same WP, the technical description of the final releases of the use cases was specified (T4.2). The resulting outcomes were used in T6.2 to determine the test set-up.

The results obtained from T6.3, which are presented in this deliverable, will be used in WP7 as dissemination material, as well as a critical source of information on the exploitation and potential market opportunities of the proposed solutions.

1.3 Structure of Deliverable

Following the current introductory section 1, the deliverable is structured as follows:

- Section 2 describes the trial results in ExFa-SP. More specifically, information is provided, per use case, about the set-up of the tests (services and network applications, their deployment in ExFa and interaction with the devices), the service-level validation at ExFas and KPIs measurements and finally the analysis of the results.
- Section 3 describes the trial results in ExFa-GR, providing the same information as in ExFa-SP per use case.
- Section 4 describes the trial results in ExFa-IT, providing the same information as in ExFa-SP and ExFa-GR per use case.
- Section 5 captures the major conclusions, key outcomes and contributions from the research work carried out and documented in this report.

In the Annex there are the results of a demo that was executed in ExFa-SP with all Spanish UCs running simultaneously.

2 Trial results in ExFa-SP

2.1 UC1

2.1.1 ExFa setup

To evaluate this use case, three NetApp components were deployed: AGV Coordinator, Logistic Process Manager (LPM) and the VSLAM algorithm. These applications allow the autonomous fleet management indoors. In this use case a logistic operation is replicated, in that both AGVs are involved in the same process.

The general architecture is illustrated in Figure 1.

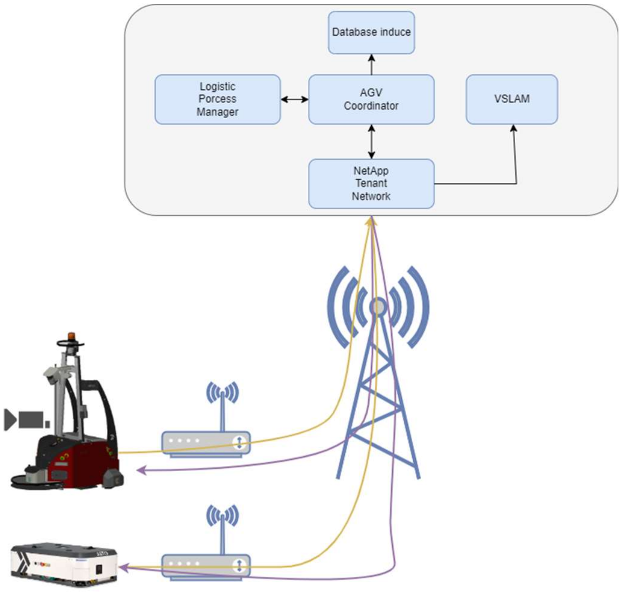

Figure 1 General Architecture of UC1.

The implementation of the use case began with a virtual map of the workplace to trace a route with the outdoor AGV, in order to replicate a logistic process in an autonomous way. The outcome of this phase is shown in Figure 2.

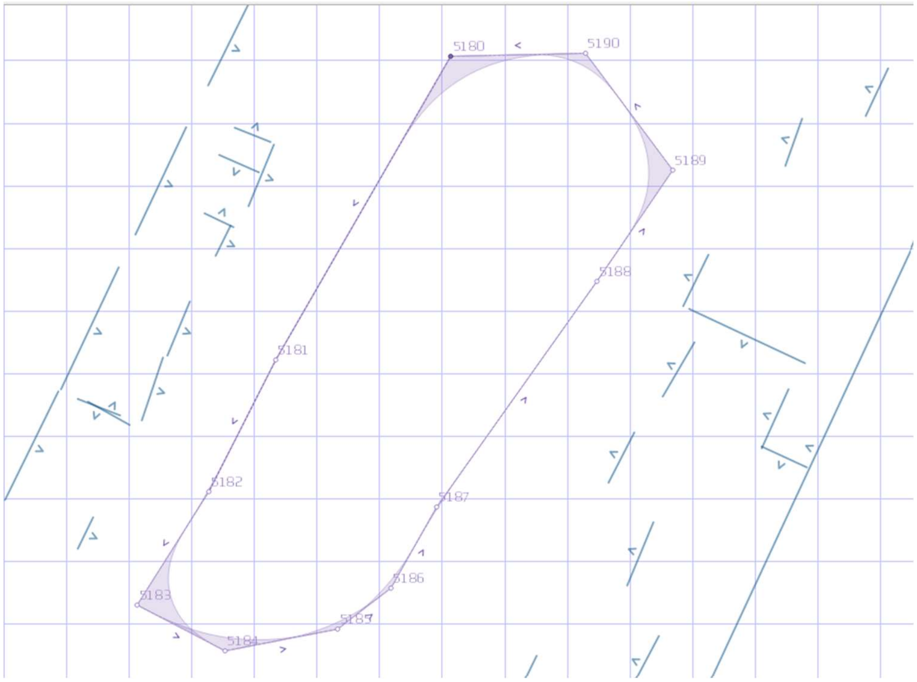

Figure 2 Virtual map.

Different tests were performed, until a good repeatable one was obtained in the stop of carrying about the trolley. After that, we proceeded to make a map with the indoor AGV with QR codes. Overlaying both maps and routes the result has the shape outlined in Figure 3 Both maps.

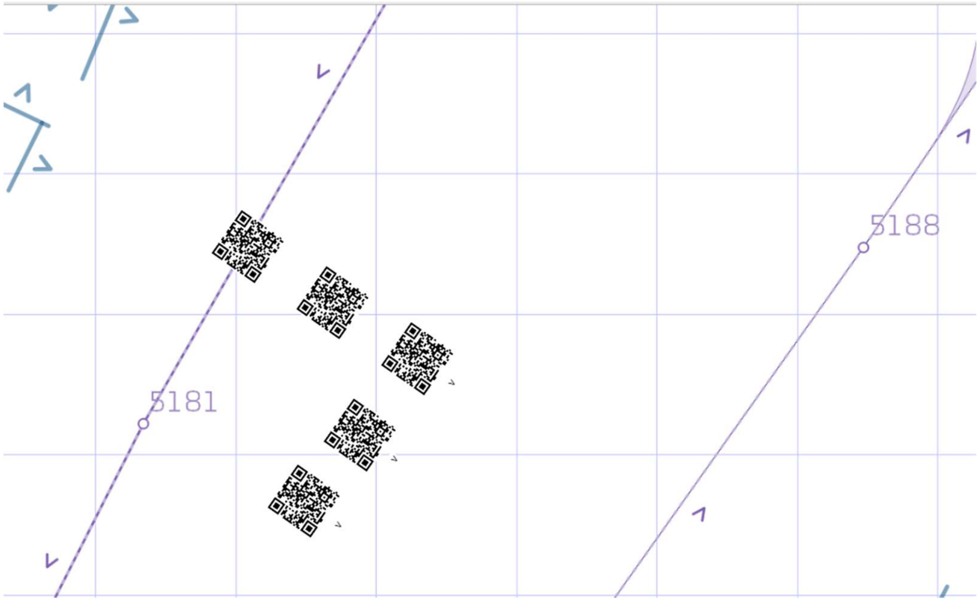

Figure 3 Both maps.

Before validating the nApps, it was necessary to ensure that the logistic process is correctly operating in the local net. As long as the outdoor AGV does not move out of the route, it has a good repeatable outcome, so that the indoor AGV can move in and out from the trolley. Besides, on the movement they do not get in each other's way by seeing each other, as this would paralyze the process.

Once checked the operation, before establishing connection and deployment in the Ford server, tests were performed locally with a Wi-Fi net and the dockers from every nApp were launched. Once they were launched, the existence and healthy status of communication between the dockers coordinator and LPM were checked. The communication was fine, since the values were updated and in the part of the VSLAM nApp the frames of the camera were received.

Later on, the network was changed, replacing the Wi-Fi with 5G, including the Fivecomm modems in every AGV and adapting the port forwarding to enable the downlink communication.

The final architecture of the NetApp's, for the coordinator and logistic process manager, is illustrated in Figure 4 Architecture coordinator and LPM in ford server.

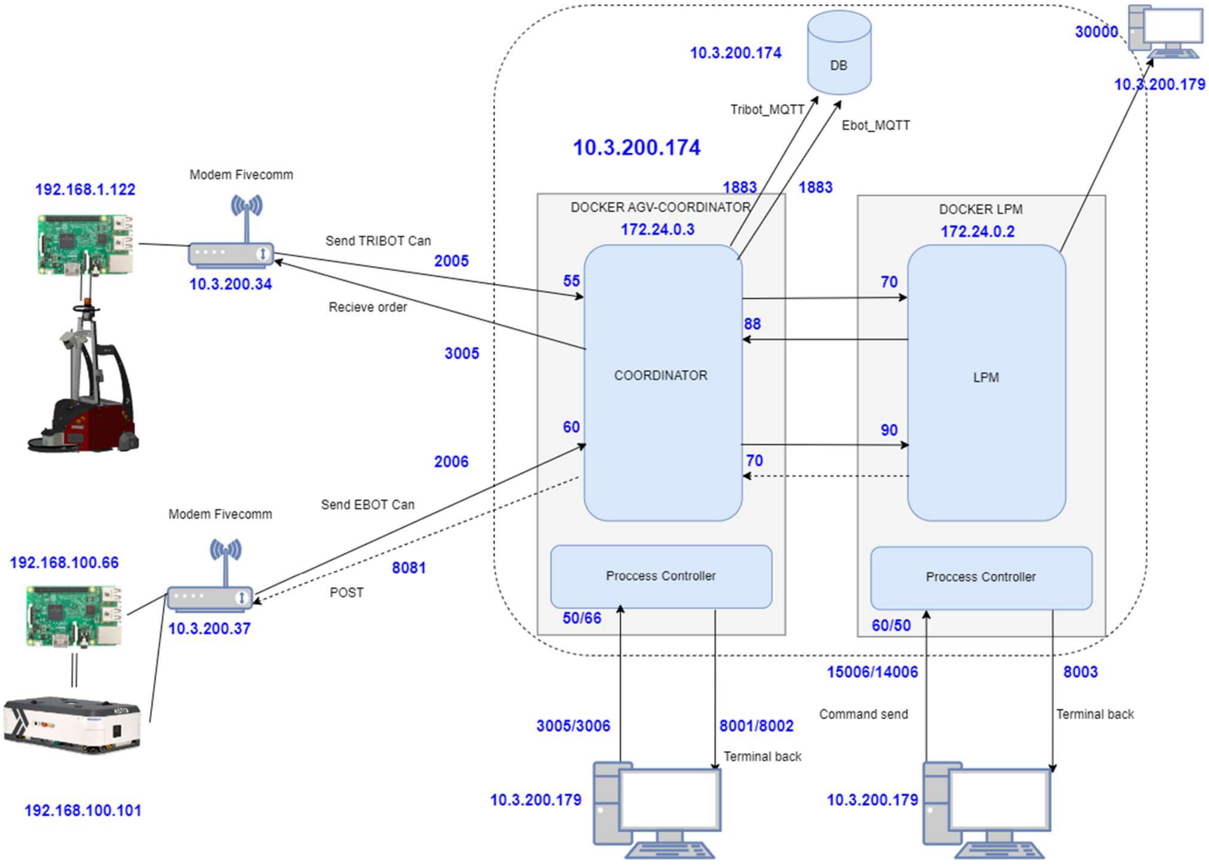

Figure 4 Architecture coordinator and LPM in ford server.

Both AGVs send their internal variables by 5G through the Fivecomm modems to carry on the information to dockers in the Ford server. Specifically, the coordinator is responsible to translate the data and to send the interesting variables to LPM. Furthermore, to update the database, which collects all the updated values of the internal variables, after updating of the state machine in the LPM, the container sends the necessary information to the coordinator and this one sends the orders to modems, so that the information arrives at AGVs. Besides, as one can check in the graph Figure 4 a process controller has been added to control the state of the nApp, and even one to look over the current status of the state machine from outside with values from the coordinator and the state.

The architecture of VSLAM is depicted in Figure 5 VSLAM architecture in Ford server.

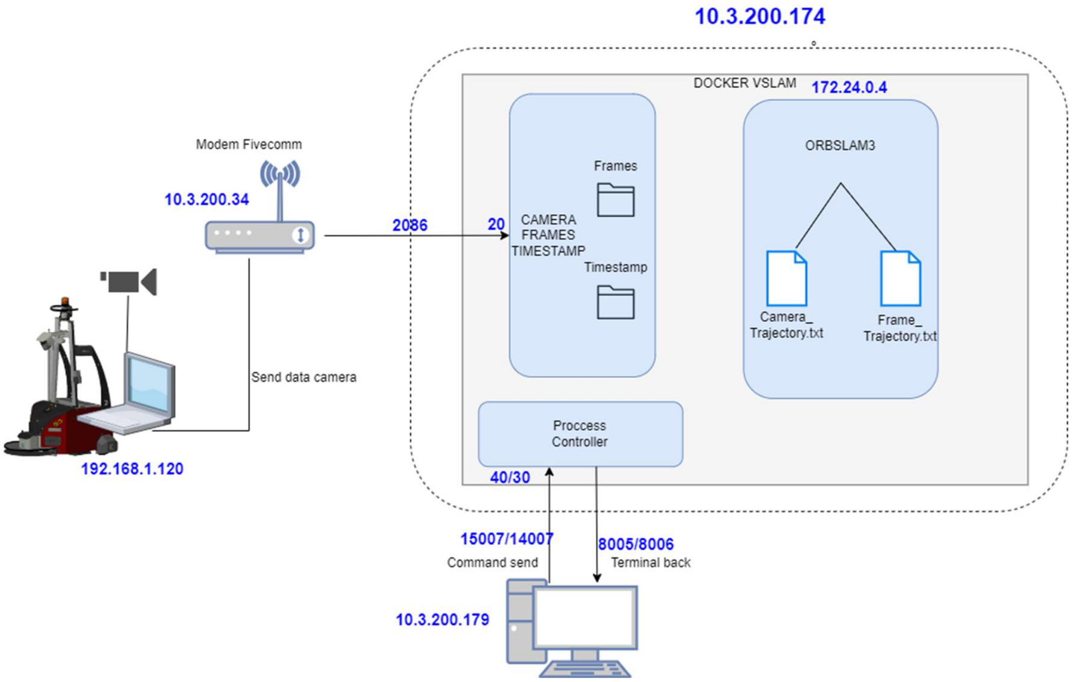

Figure 5 VSLAM architecture in Ford server.

A depth camera sends frames through 5G to the VSLAM docker. The frames and timestamps are kept inside the docker. The kept frames are used to launch the VSLAM algorithm.

2.1.2 Service-level validation at ExFas and KPIs measurements

For the validation of the use case, two scenarios have been proposed, whose target is to validate the functionality of the AGV Coordinator, Logistic Process Manager and the VSLAM from a user and from the nApps' technical point of view, respectively. On the one hand, a virtual map is done with outdoor AGV and the autonomous navigation is validated for the traced route. Later, the indoor AGV is localized in the initial position in its map. Once ubicated both AGVs in their initial positions, the three nApps can be launched. When the NApps are running, it is necessary to check that there is communication among AGVs and dockers and between dockers. The coordinator, the LPM and VSLAM module are working correctly.

After these checks, the correct working of the logistic operation is controlled; when 1 round is finished, the process is stopped and a map can be done with the received images in the Docker VSLAM.

The scenario KPIs' exposed are the following:

- SVM-02-nApp specific functional test, visually checking that the logistic operation is performed exactly and that a VSLAM map is generated with the received frames.

The outcomes are shown in the Table 1 SVM-02-UC1.

Table 1 SVM-02-UC1

KPIs-UC1 Scenario 1	SVM-02
Visual check	OK
MAKE VIRTUAL MAP	OK

The operation of the logistic process can be watched at the following link:

1 full round

- SVM-07-Perceceived application operation related latency. It consists in measuring the delay end to end, from issuing the order to AGV up to the start of the motion from a user view. In our case, all the state changes were in less than 2 seconds (Table 2 SVM-07-UC1).

Table 2 SVM-07-UC1

KPIs-UC1 Scenario 1		SVM-07
Time state changes		<2 sec

On the other hand, the second scenario its target is to validate the nApps from a technical view. In this case the measured KPIs are the following:

- SVM-01- Service deployment. The overall time from the instantiation of the application through the VAO until the VSLAM point cloud sender is launched in the AGV (Table 3 SVM-01-UC1).

Table 3 SVM-01-UC1

KPIs-UC1 Scenario 2	SVM-01
Time	15 sec

- SVM-04 – Data capacity. In this KPI, the consumed bandwidth is measured in the transmission of the frames from the depth camera to the VSLAM docker (Table 4 SVM04-UC1).

Table 4 SVM04-UC1

KPIs-UC1 Scenario 2	SVM-04
Max Bandwidth	21.4Mbps

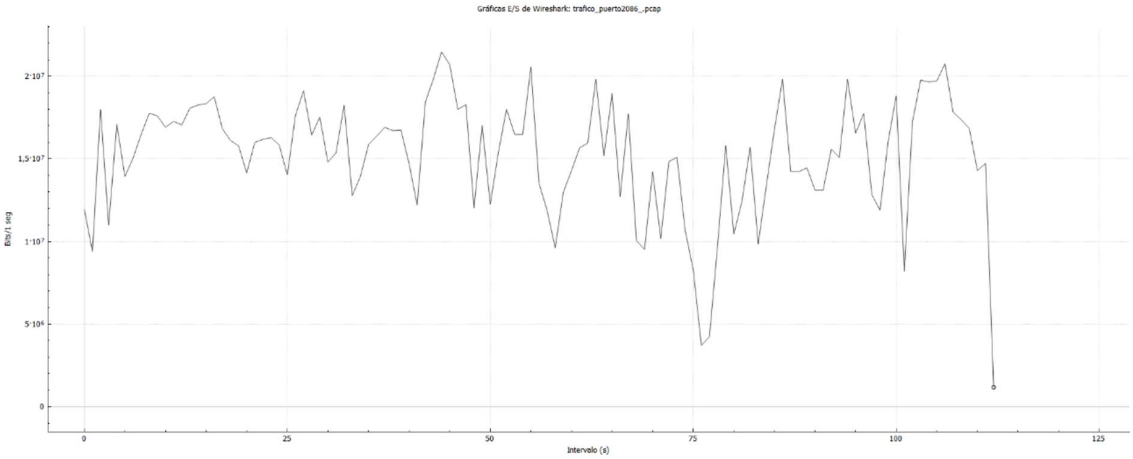

Figure 6 Bandwidth in port 2086 TCP/IP.

Figure 6 shows the consumed bandwidth in the Ford server with IP 10.3.200.179 and the port 2086. The frames were sent by TCP/IP, so it is one of the filters in the file with extension pcap. The graph shows all the frames sent during one route in the final demonstration.

- SVM-07 -Perceived AGV operation latency. Evaluation of the end-user perceived operational latency when the AGV is controlled remotely from the Edge (Table 5 SVM-07-UC1).

Table 5 SVM-07-UC1

KPIs-UC1 Scenario 2	SVM-07
State 1→2	0.09 sec
State 2→3	1.76 sec
State 3→5	0.05 sec
State 5→6	1.88 sec
State 6→1	1.57 sec

In the following link one can see these times in some time snippets focusing on these moments:

Video change state

2.1.3 Analysis of the results

The benefits that 5G contributes us in our case is the low latency. This is achieved by reducing in some cases the expected values terrifically, even when the net is handling a big quantity of data due to both AGVs sending their internal variables. The internal variables include the AGV speeds, like 1 Gbps for the outdoor and 250 Mbps for the indoor. These speeds are provided by the CAN bus. We have to add the sending of images. In our cases the camera sent to 29 fps and the receptor received up to 21 fps, with no packets loss, favoring the VSLAM algorithm for map creation.

2.2 UC2

2.2.1 ExFa setup

Previous activities to those shown in this deliverable are described in detail in WP4 deliverables. D4.2 explained all tasks related to pre-DevOps and pre-ExFa (5TONIC premises) validation testing, which were considered as phase 1, phase 2 and phase 3 in our initial planning. In this deliverable D6.2, we now explain the design, development, and integration process within the environment of the Spanish ExFa in Ford, according to phases 4, 5 and 6. The fourth phase consisted of the integration and manual onboarding of the network applications in the final facilities where the tests were done, i.e., the Ford factory. Phase 4 entails the deployment of applications and demos for the collective operation of the three Spanish ExFa use cases. Phase 5 consisted of the end-to-end demonstration of the use case, together with UC1 and UC3. Finally, in phase 6 we onboarded all components in the DevOps testbed thanks to the NAO and demonstrated the final use case through different tests by using the 5G-INDUCE platform. Note that we did a complete demonstration of the use case and got results against the values set in D5.1 in all phases 4, 5 and 6. Such results are shown in Section 3.2.2.

Phase 4: Manual onboarding and first tests at Ford

The fourth phase was about the necessary development and the onboarding of the network application in the final facilities, located on the Ford factory shop floor. In this phase, we made use of the final AGV for outdoor scenarios, the ASTI TriBOT. In addition, the deployment of the network application and its performance was validated on two different locations, an Edge server located in Ford and another one deployed at UPV, with a direct fiber connection to Ford. More details about this setup is provided in previous deliverables.

In this phase, all VNFs were deployed and validated on the Ford server. In this case, a new version of the gesture recognition VNF was created to work without an Internet connection, since it was needed for licensing checking, but was not available in Ford at that moment. Integration and validation of the Fivecomm 5G modems at Ford also took place, including their operability with the AGV. We then checked the interoperability of the AGV and the 5G network, and did a first demonstration, where the AGV TriBOT for outdoor scenarios was controlled by gestures in a controlled environment (close to the server, not the final location). We produced a video for dissemination purposes, which is available in the project YouTube channel. In addition, a first set of 5G latency and throughput measurements was obtained.

Figure 7 shows the AGV connected to the 5G network at Ford and being controlled by the remote operator via hand gestures.

Figure 7 Phase 4 integration work at Ford.

Phase 5: first validation of the use case with UC1 and UC3 (manual deployment)

In the fifth phase, the three use cases of the Spanish ExFa were validated together. For these tests, the 5G modems from Fivecomm were configured for a simultaneous use in the three use cases. Additionally, the operation of use case 2 was validated at the final location of the Ford factory, with instructions given to the AGV TriBOT during the demonstration of UC1, while the demonstration of UC3 was also running.

In this phase, the containers for the different nApp systems were manually deployed at the Ford server, without using the 5G-INDUCE platform orchestrator. This allowed validating the correct operation of the application and the performance offered by the network before the final onboarding.

Figure 8 illustrates the final environment where the demo for the three use cases was conducted simultaneously.

Figure 8 Phase 5 integration work at Ford.

Phase 6: platform onboarding and final demonstration

In the final phase, the deployment was done similarly to phase 5. However, this time, the use case network application deployment was done through the NAO, including the allocation of necessary resources and execution of the components. In this phase, some problems arose in the communication between the two VNFs of gesture recognition and AGV control. However, such problems did not affect the validation of the application's functionality nor the individual validation of each VNF, since each one was deployed correctly. Once all VNFs were deployed and executed, the use case was validated simultaneously together with UC1 and UC3 in the final scenario envisaged at Ford. All use cases were deployed by the NAO and executed correctly.

Figure 9 shows the AGV located in this scenario, with the VNFs deployed by the NAO directly at Ford. Note that, as done in previous phases, both VNFs were deployed and tested in the two available servers located directly in Ford and in UPV, connected through fiber. Results were gathered, which are shown in the following sections. Among others, the deployment time of different VNFs was validated, the consumed bitrate was measured both in uplink and downlink on both servers, and end-to-end latency for video was measured from the moment it is obtained until it returns to the processed user.

Figure 9 Final demo of UC2 at Ford.

Additional work at Fivecomm offices (post-demo)

Part of the work performed in use case 2 was also used in use case 3. The data gathered from the 5G modems was sent to a data base, accessible from the VR application to be shown to the worker wearing the glasses. Since this step could not be completed in time at Ford, further development was done later to finalize the task.

At Fivecomm offices, we set the modems, so the coverage data collected is stored in an Influx database, located at the Edge server in UPV. To accomplish this, a script was executed within the modem itself, obtaining parameters such as RSRP, RSRQ, SINR, GPS coordinates, as well as minimum, average, and maximum latency through the database API. These parameters can be retrieved from the modem using AT commands and then transmitted via the HTTP API integrated into the Influx database. Figure 10 displays the table generated after collecting these measurements, indicating a name for the device and another for the measurement conducted, along with the generated timestamp. This database can be accessed by UC3 owners or any other application developer when showing 5G parameters to the user in real-time.

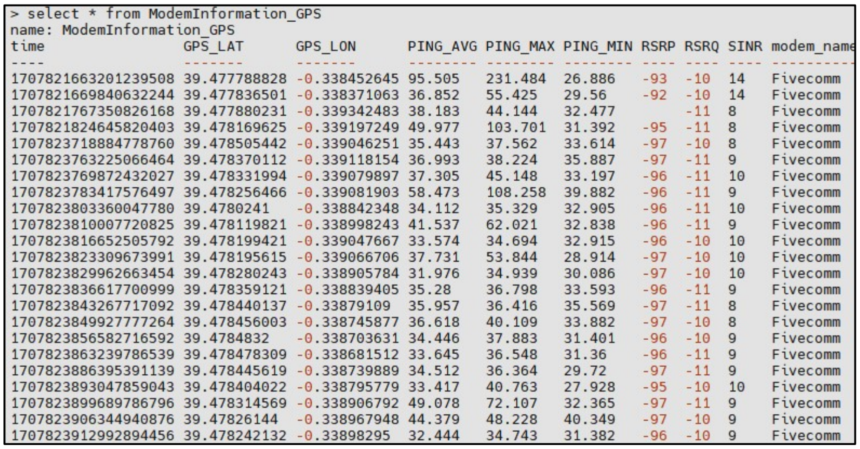

```
> select * from ModemInformation_GPS
name: ModemInformation_GPS
time                GPS_LAT      GPS_LON      PING_AVG PING_MAX PING_MIN RSRP RSRQ SINR modem_name
----                -------      -------      -------- -------- -------- ---- ---- ---- ----------
1707821663201239508 39.477788828 -0.338452645 95.505   231.484  26.886   -93  -10  14   Fivecomm
1707821669840632244 39.477836501 -0.338371063 36.852   55.425   29.56    -92  -10  14   Fivecomm
1707821767350826168 39.477880231 -0.339342483 38.183   44.144   32.477        -11  8    Fivecomm
1707821824645820403 39.478169625 -0.339197249 49.977   103.701  31.392   -95  -11  8    Fivecomm
1707823718884778760 39.478505442 -0.339046251 35.443   37.562   33.614   -97  -10  8    Fivecomm
1707823763225066464 39.478370112 -0.339118154 36.993   38.224   35.887   -97  -11  9    Fivecomm
1707823769872432027 39.478331994 -0.339079897 37.305   45.148   33.197   -96  -11  10   Fivecomm
1707823783417576497 39.478256466 -0.339081903 58.473   108.258  39.882   -96  -11  9    Fivecomm
1707823803360047780 39.4780241   -0.338842348 34.112   35.329   32.905   -96  -11  10   Fivecomm
1707823810007720825 39.478119821 -0.338998243 41.537   62.021   32.838   -96  -11  9    Fivecomm
1707823816652505792 39.478199421 -0.339047667 33.574   34.694   32.915   -96  -10  10   Fivecomm
1707823823309673991 39.478195615 -0.339066706 37.731   53.844   28.914   -97  -10  10   Fivecomm
1707823829962663454 39.478280243 -0.338905784 31.976   34.939   30.086   -97  -10  10   Fivecomm
1707823836617700999 39.478359121 -0.338839405 35.28    36.798   33.593   -96  -11  9    Fivecomm
1707823843267717092 39.478440137 -0.33879109  35.957   36.416   35.569   -97  -11  8    Fivecomm
1707823849927777264 39.478456003 -0.338745877 36.618   40.109   33.882   -97  -10  8    Fivecomm
1707823856582716592 39.4784832   -0.338703631 34.446   37.883   31.401   -96  -10  9    Fivecomm
1707823863239786539 39.478478309 -0.338681512 33.645   36.548   31.36    -96  -11  9    Fivecomm
1707823886395391139 39.478445619 -0.338739889 34.512   36.364   29.72    -97  -11  9    Fivecomm
1707823893047859043 39.478404022 -0.338795779 33.417   40.763   27.928   -95  -10  10   Fivecomm
1707823899689786796 39.478314569 -0.338906792 49.078   72.107   32.365   -97  -11  9    Fivecomm
1707823906344940876 39.47826144  -0.338967948 44.379   48.228   40.349   -97  -10  9    Fivecomm
1707823912992894456 39.478242132 -0.33898295  32.444   34.743   31.382   -96  -10  9    Fivecomm
```

Figure 10 Database deployed at UPV server, getting info from Fivecomm 5G modem.

2.2.2 Service-level validation at ExFas and KPIs measurements

The following section provides the results obtained during the tests performed in the 5G-INDUCE project. This includes not only those obtained in Ford (phases 4, 5 and 6), but also those gathered in previous stages described in D4.2 (pre-ExFa measurements).

Pre-ExFa KPIs (obtained in Valencia – phase 2, 5TONIC – phase 3)

As described in D4.2, the gathering of a set of parameters was planned prior to the use of the Spanish ExFa at Ford. The measurements taken in phase 2 in Valencia, and phase 3 in 5TONIC premises in Madrid, are here shown.

It is also important to note that for these tests, the AGV EBOT was used, instead of the TriBOT intended for the final demos. More details about the use of this AGV are provided in D4.2.

- SVM-02: nApp specific functional tests

Table 6 SVM-02-UC2

KPIs-UC2	SVM-02
Successful operation	OK

This test was done in phases 2 and 3. This test included: i) successful video stream from AGV to server; ii) video reception from AGV; and iii) AGV control from the Edge. Firstly, the proper functioning of video transmission was validated. It was verified that there were no issues on video reception on the server from both the user and the robot. Additionally, it was confirmed that the commands generated by the processed video reached the AGV correctly.

- SVM-07: Perceived operation latency

Table 7 SVM-07-UC2

KPIs-UC2	SVM-07
< 900 ms	169 (video) / 339 ms (end-to-end)

Done as part of phase 3 in 5TONIC. This is the measurement of the end-to-end delay from the moment a gesture is made until the AGV starts moving. In this case, two types of tests were conducted to measure the E2E delay. Firstly, the latency was measured from the moment the video is captured until it is received from the user. Secondly, using a video recorded in super-slow motion, we measured the time from when the user completes a STOP gesture until the AGV actually stops.

Figure 11 shows the results for the first test, where just the video part is included. It can be observed that the average latency obtained is 169 ms.

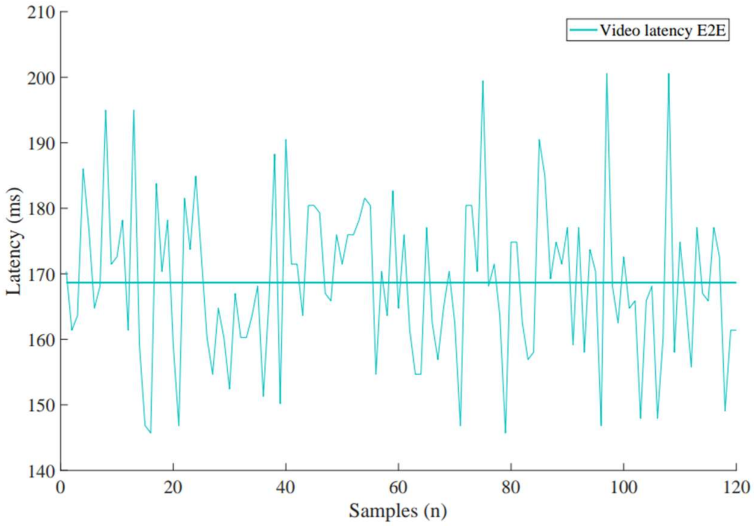

Figure 11 Figure 13. E2E Latency measurements (video part).

For the second set of tests, measurements were taken from various recorded videos as shown in Table 8, where it can be observed that an average end-to-end delay of 339 ms was obtained.

Table 8 Pre-ExFa tests: E2E perceived operation latency.

Samples	Delay
Sample 1	357 ms
Sample 2	352 ms
Sample 3	338 ms
Sample 4	309 ms
Sample 5	341 ms
Average	**339 ms**

- SVM-06: HR video quality

Table 9 SVM-06-UC2

KPIs-UC2	SVM-06
15 Mbit/s (UL), 10 Mbit/s (DL)	40 Mbit/s (4G) and 75 Mbit/s (5G) for UL/DL (UL-limited)

Done already in phase 2, in Valencia, using the available networks at UPV campus. This measurement is about the video in real time for gesture recognition. To validate that the desired bandwidth could be transmitted, measurements of the available throughput were taken. As shown in Figure 12, for both cases, the specified KPIs for the video are met in both the uplink and downlink.

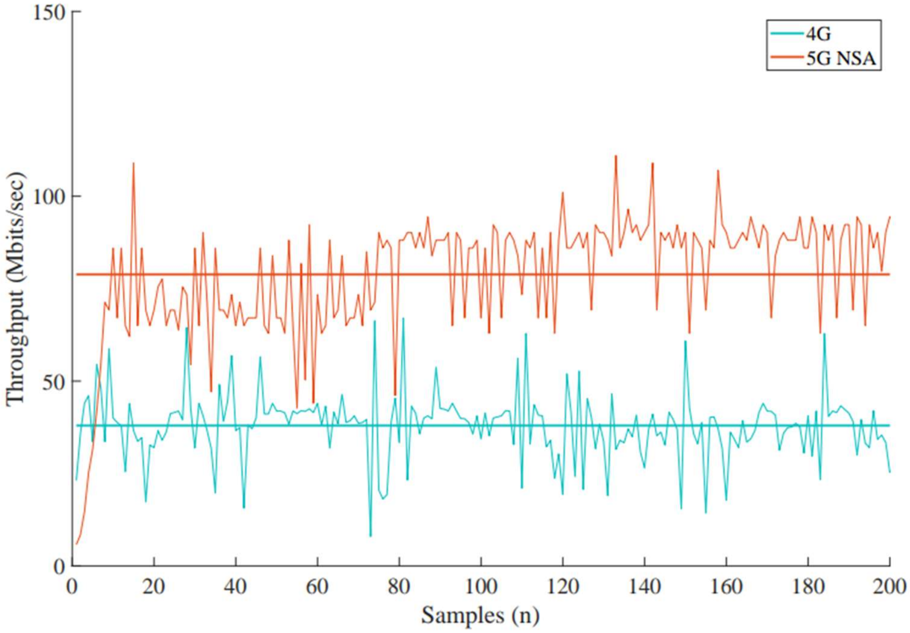

Figure 12 Network throughput measurement values obtained.

ExFa KPIs (Ford – phases 4, 5 and 6)

The second part of the tests were directly performed at Ford, during the last phases of the project. As done in the previous part, the obtained values were compared against the requirements set in D5.1.

- SVM-01: Measurement of the service deployment time

Table 10 SVM-01-UC2

KPIs-UC2	SVM-01
< 60 s	~10 s

For all deployments of the different VNFs for the use case, the specified KPI of not exceeding 60 seconds was always met. If this were to happen, the video would not be received in the gesture recognition VNF, and the container would need to be restarted. The approximate time measured for the deployment of NetApp was about 10 seconds.

- SVM-09: end-to-end latency

Table 11 SVM-09-UC2

KPIs-UC2	SVM-09
< 800 ms	224 ms

Measurement of the end-to-end delay from the moment a gesture is made until the AGV start moving. In this case, we focused on the glass-to-glass video latency. As done in the pre-ExFa measurements, we measured this value by showing a clock on the final display where the gestures are shown.

The measurements were performed for three potential scenarios. The first scenario involved the manual deployment of network applications on the Ford server. In the second one, a similar manual deployment was done at UPV server. Lastly, the deployment of network applications on the Ford server was validated using the NAO orchestrator. Figure 13 shows one of the measurements taken for the last scenario.

Figure 13 End-to-end video metrics for deployment on both the server and the NAO.

Figure 14 shows the latency results obtained for these three cases. The measurements were taken from the moment the camera captured an image until the processed frame was transmitted back to the tablet through the gesture recognition NetApp. The obtained values were, on average, 246 ms, 394 ms and 224 ms, for the three aforementioned scenarios respectively. The most favourable result was achieved through deployment via the NAO at Ford (224 ms), where resource management significantly improved the performance, leading to the lowest latency. Conversely, the least favourable result was observed with the deployment on the UPV server.

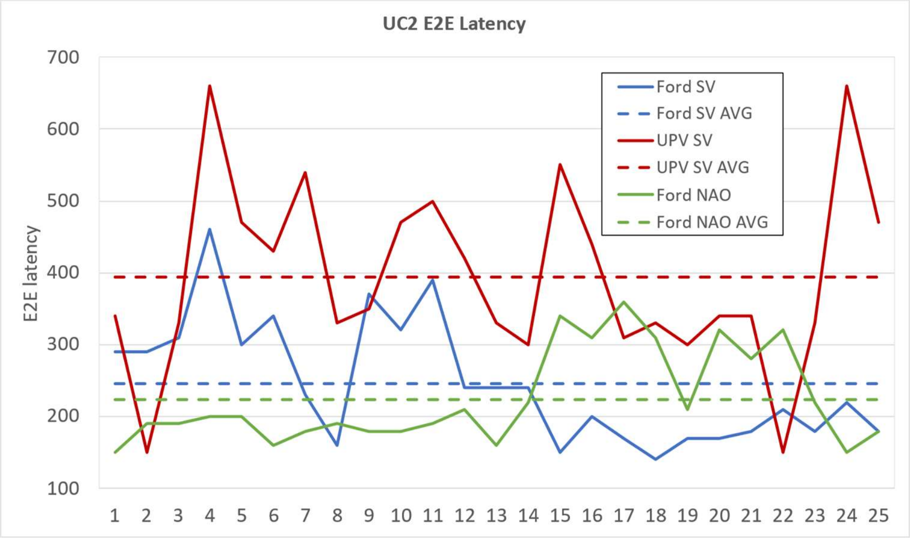

Figure 14 E2E Latency measurements.

As depicted in Figure 15, the server had fewer resources available for NetApp execution, resulting in higher latency values. Overall, it can be observed that even for the worst-case scenario, the specified KPI of 800ms for the end-to-end latency of the NetApp is always met.

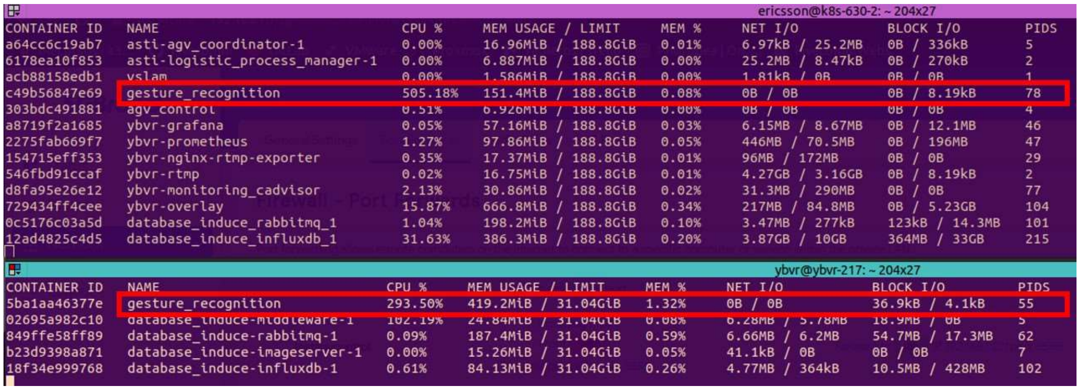

```
ericsson@k8s-630-2: ~ 204x27
CONTAINER ID   NAME                               CPU %     MEM USAGE / LIMIT     MEM %     NET I/O           BLOCK I/O         PIDS
a64cc6c19ab7   asti-agv_coordinator-1             0.00%     16.96MiB / 188.8GiB   0.01%     6.97kB / 25.2MB   0B / 336kB        5
6178ea10f853   asti-logistic_process_manager-1    0.00%     6.887MiB / 188.8GiB   0.00%     25.2MB / 8.47kB   0B / 270kB        2
acb88158edb1   vslam                              0.00%     1.586MiB / 188.8GiB   0.00%     1.81kB / 0B       0B / 0B           1
c49b56847e69   gesture_recognition                505.18%   151.4MiB / 188.8GiB   0.08%     0B / 0B           0B / 8.19kB       78
303bdc491881   agv_control                        0.51%     6.926MiB / 188.8GiB   0.00%     0B / 0B           0B / 0B           4
a8719f2a1685   ybvr-grafana                       0.05%     57.16MiB / 188.8GiB   0.03%     6.15MB / 8.67MB   0B / 12.1MB       46
2275fab669f7   ybvr-prometheus                    1.27%     97.86MiB / 188.8GiB   0.05%     446MB / 70.5MB    0B / 196MB        47
154715eff353   ybvr-nginx-rtmp-exporter           0.35%     17.37MiB / 188.8GiB   0.01%     96MB / 172MB      0B / 0B           29
546fbd91ccaf   ybvr-rtmp                          0.02%     16.75MiB / 188.8GiB   0.01%     4.27GB / 3.16GB   0B / 8.19kB       2
d8fa95e26e12   ybvr-monitoring_cadvisor           2.13%     30.86MiB / 188.8GiB   0.02%     31.3MB / 290MB    0B / 0B           77
729434ff4cee   ybvr-overlay                       15.87%    666.8MiB / 188.8GiB   0.34%     217MB / 84.8MB    0B / 5.23GB       104
0c5176c03a5d   database_induce_rabbitmq_1         1.04%     198.2MiB / 188.8GiB   0.10%     3.47MB / 277kB    123kB / 14.3MB    101
12ad4825c4d5   database_induce_influxdb_1         31.63%    386.3MiB / 188.8GiB   0.20%     3.87GB / 10GB     364MB / 33GB      215

ybvr@ybvr-217: ~ 204x27
CONTAINER ID   NAME                              CPU %     MEM USAGE / LIMIT     MEM %     NET I/O           BLOCK I/O         PIDS
5ba1aa46377e   gesture_recognition               293.50%   419.2MiB / 31.04GiB   1.32%     0B / 0B           36.9kB / 4.1kB    55
02695a982c10   database_induce-middleware-1      102.19%   24.84MiB / 31.04GiB   0.08%     6.28MB / 5.78MB   18.9MB / 0B       5
849ffe58ff89   database_induce-rabbitmq-1        0.09%     187.4MiB / 31.04GiB   0.59%     6.66MB / 6.2MB    54.7MB / 17.3MB   62
b23d9398a871   database_induce-imageserver-1     0.00%     15.26MiB / 31.04GiB   0.05%     41.1kB / 0B       0B / 0B           7
18f34e999768   database_induce-influxdb-1        0.61%     84.13MiB / 31.04GiB   0.26%     4.77MB / 364kB    10.5MB / 428MB    102
```

Figure 15 *Ford and UPV servers resources.*

- SVM-06: High-resolution video in real time for gesture recognition

Table 12 SVM-06-UC2

KPIs-UC2	**SVM-06**
15 Mbit/s (UL), 10 Mbit/s (DL)	>15 Mbit/s (UL), >10 Mbit/s (DL) – not all bandwidth needed

Throughout the tests conducted on both servers, measurements of the consumed bandwidth in both the uplink and downlink were obtained specifically for UC2. The results of these measurements are presented in Figures 16 and 17. It is evident that, in all cases, the limiting bandwidth in the uplink does not exceed 8 Mbps.

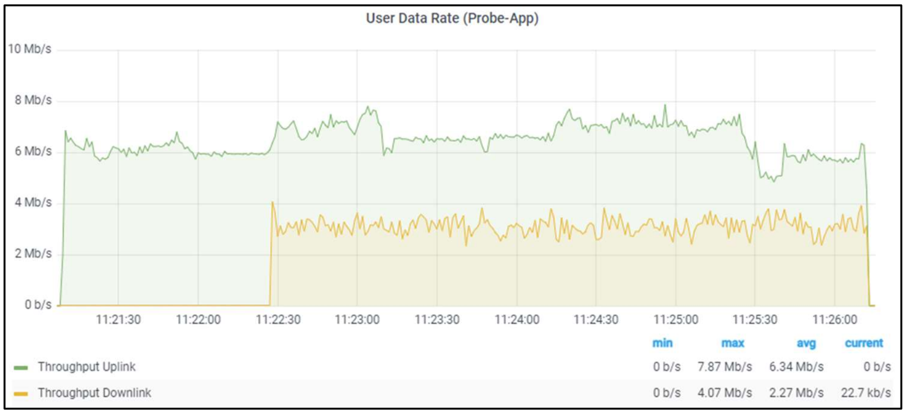

	min	max	avg	current
Throughput Uplink	0 b/s	7.87 Mb/s	6.34 Mb/s	0 b/s
Throughput Downlink	0 b/s	4.07 Mb/s	2.27 Mb/s	22.7 kb/s

Figure 16 *UC2 BW with app deployed on UPV server.*

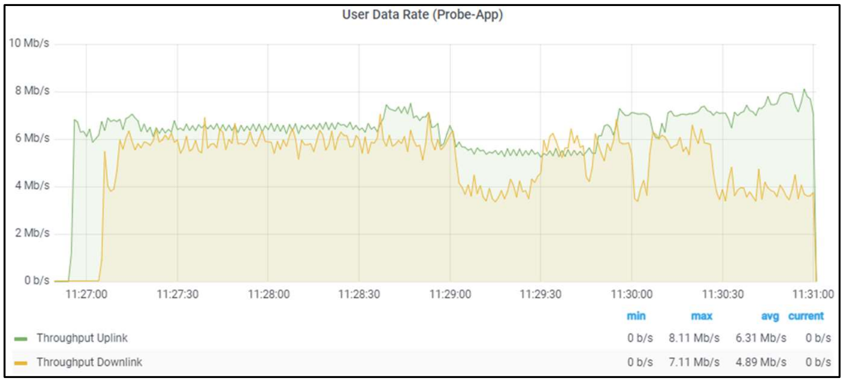

	min	max	avg	current
Throughput Uplink	0 b/s	8.11 Mb/s	6.31 Mb/s	0 b/s
Throughput Downlink	0 b/s	7.11 Mb/s	4.89 Mb/s	0 b/s

Figure 17 *UC2 BW with app deployed on Ford server.*

At Ford's premises, it was validated that transmitting the video at approximately 6Mbps of bitrate was sufficient for the proper operation of the gesture recognition VNF. Although it would be possible, due to network capabilities, to transmit at a bitrate of 15Mbps as indicated in the KPIs.

- SVM-07: AGV control from the Edge

Table 13 SVM-07-UC2

KPIs-UC2	SVM-07
100%	100%

In all tests, the proper operation and control of the AGV for all scenarios was validated, ensuring it always received the necessary commands. If the AGV did not receive an instruction at the required frequency, it would activate the brakes automatically.

2.2.3 Analysis of the results

After analysing and validating the results, it can be observed that 5G networks represent a minor factor when calculating the entire latency budget in an end-to-end configuration. In this case, the latency introduced by the network is minor, considering the E2E latency values obtained. We observed that, since the use case is about autonomous driving, the value of video E2E latency in SVM-09 is the real E2E latency that needs to be considered. It has also been observed that the use of 5G networks and the 5G-INDUCE platform in particular are extremely beneficial when sending high data rates with very low latency. The high uplink throughput allows multiple good quality video streams to be sent, in this case, for video transmission in real time. This is essential when considering human-machine applications with real-time feedback.

2.3 UC3

2.3.1 ExFa setup

SERVICES AND NETWORK APPLICATIONS

This Network Application processes immersive video feeds to build 360° images. The feeds are delivered to designated nodes defined by the Network Application processing requirements and with the goal of achieving the targeted latency at a given high availability. The Network Applications servers will be located at the optimal MEC location, elements located in the factory must be reachable by 5G premises or telco nodes, and the client devices must have connectivity to the cloud YBVR platform.

The diagram of the elements, layers and flows in the UC3 are detailed in Figure 18 .

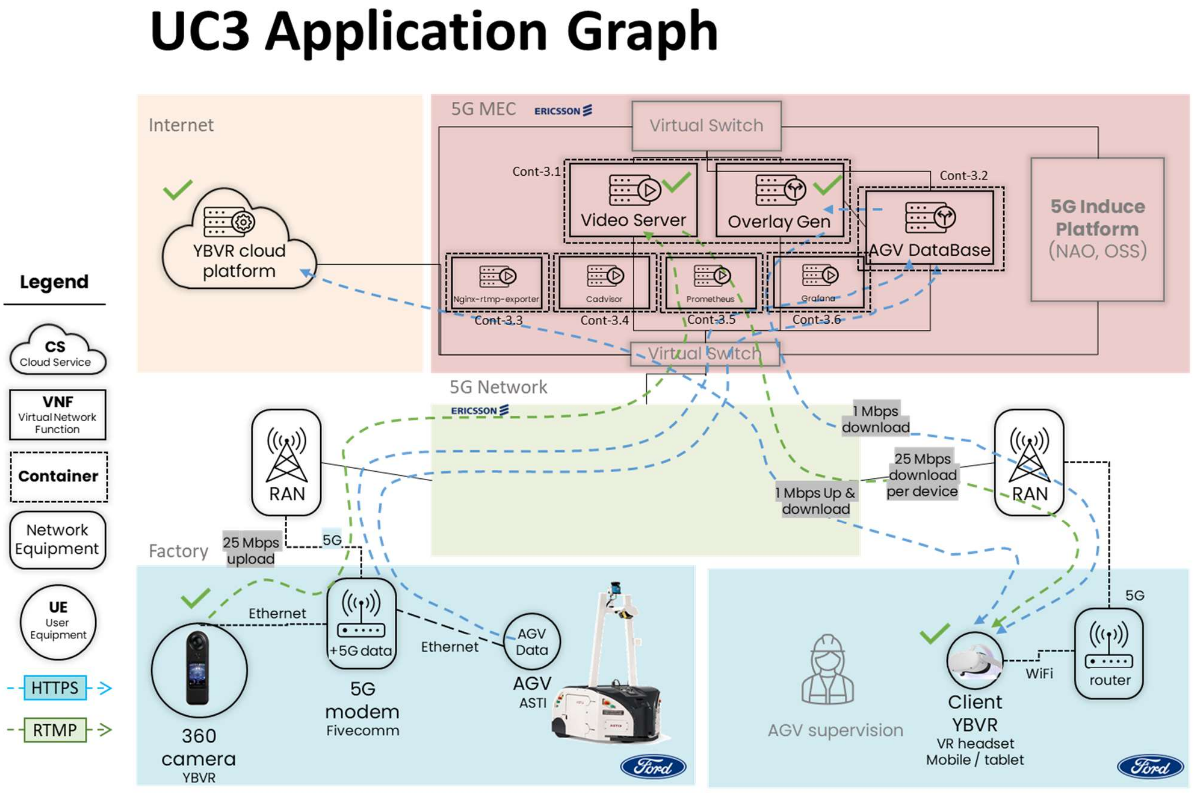

Figure 18 UC3 diagram.

Use case features:

Provided data:

- (Input) camera installed over the AGV and sensoring or navigation data from other Network Applications (modem and AGV)
- (Output) immersive image visualized on a VR head mounted display or mobile device.

Control features: Through smart phone Apps and computer devices, connection with other Network Applications

1. to provide information on actions taken
2. to receive information about alerts
3. to determine priority of data sent
4. to manage the vehicle to avoid a collision.

DEPLOYMENT IN EX-FA

This refers to the deployment of the experimentation facilities including the provisioning of the infrastructure, the deployment of the orchestration platform and porting mechanism.

In the Use Case 3, the final test was deployed in Ford Ex-Fa site facility, in Almussafes, Valencia, Spain, with the 5G network infrastructure deployed by Ericsson Spain, the AGVs provided by ABB and the 5G modems provided by FiveComm.

- SVM-01: Service deployment

Service deployment was carried out with the support of UBITECH on the MEC deployed in the Ford's factory premises.

Table 14 SVM-01 UC3

KPIs-UC3	SVM-01
Service Deployment Time	15sec

- SVM-02: Functional testing – Connectivity check

With the Network application deployed, it was checked out that all the containers were running, and the ports assigned to the services that must be accessible from outside the server (Figure 19):

```
root@k8s-630-2:/etc/kubernetes# k get svc -A | grep uc3
embb4704      forduc1te-induceuc3ybvrcadvisor3385-jhd4g23ig0-service          ClusterIP   10.96.229.124    <none>        8080/TCP
                                                                                                                       20h
embb4704      forduc1te-induceuc3ybvrexporter3376-u5fgby0rmy-service          ClusterIP   10.98.60.186     <none>        9728/TCP
                                                                                                                       20h
embb4704      forduc1te-induceuc3ybvrgrafana3403-2akepjuapb-service           ClusterIP   10.107.114.205   <none>        3000/TCP
                                                                                                                       20h
embb4704      forduc1te-induceuc3ybvroverlay3412-ihe2apvp4e-service           NodePort    10.98.22.128     <none>        8443:31699/TCP,3000:31723/TCP,80:31238/TCP
                                                                                                                       20h
embb4704      forduc1te-induceuc3ybvrprometheus3394-6nycstxrp1-service        NodePort    10.107.42.39     <none>        9090:30433/TCP
                                                                                                                       20h
embb4704      forduc1te-induceuc3ybvrrtmp3361-s3fz56gx43-service              NodePort    10.105.198.194   <none>        80:31843/TCP,1935:32763/TCP,443:30741/TCP
                                                                                                                       20h
root@k8s-630-2:/etc/kubernetes# k get po -A | grep uc3
embb4704         forduc1te-induceuc3ybvrcadvisor3385-jhd4g23ig0-deployment-4kkhj   1/1     Running   0           12h
embb4704         forduc1te-induceuc3ybvrexporter3376-u5fgby0rmy-deployment-bxfs4   1/1     Running   0           12h
embb4704         forduc1te-induceuc3ybvrgrafana3403-2akepjuapb-deployment-7nfg2n   1/1     Running   0           12h
embb4704         forduc1te-induceuc3ybvroverlay3412-ihe2apvp4e-deployment-8l492h   1/1     Running   0           8m10s
embb4704         forduc1te-induceuc3ybvrprometheus3394-6nycstxrp1-deploymen45cpn   1/1     Running   0           12h
embb4704         forduc1te-induceuc3ybvrrtmp3361-s3fz56gx43-deployment-6cb9xrgms   1/1     Running   0           12h
```

Figure 19 Pods and services from UC3.

The camera must be configured for streaming to the video server IP with the port assigned to the ybvr-rtmp service (32763) and connected to the 5G modem via ethernet. If the RTMP server is reachable, the streaming can start. The video metrics are shown in real-time in Grafana, see Figure 20 :

Figure 20 Video server reception bitrate.

Interaction with the devices

At this point, the video and overlay URLs are configured in YBVR XMS (Experience Manager System), which is the YBVR content management system (Figure 21, Figure 22)Figure 21 .

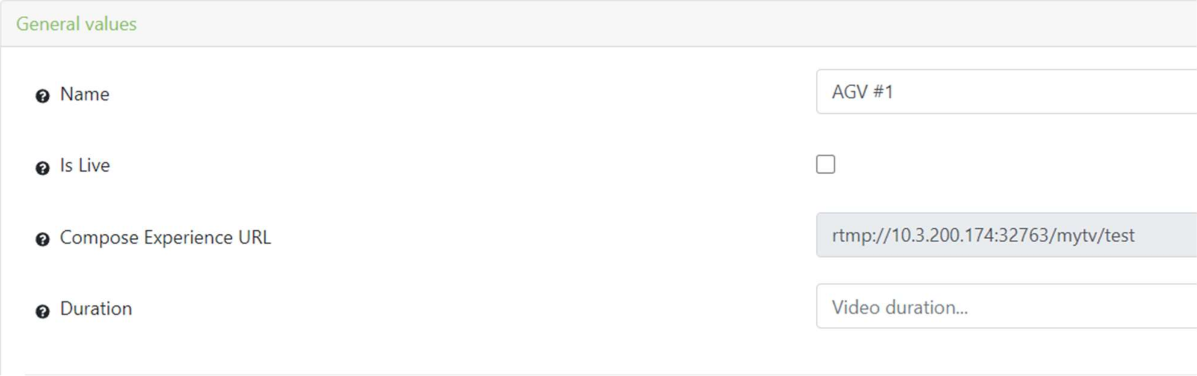

General values

Name: AGV #1

Is Live: ☐

Compose Experience URL: rtmp://10.3.200.174:32763/mytv/test

Duration: Video duration...

Figure 21 Configuration of the video URL inside the XMS.

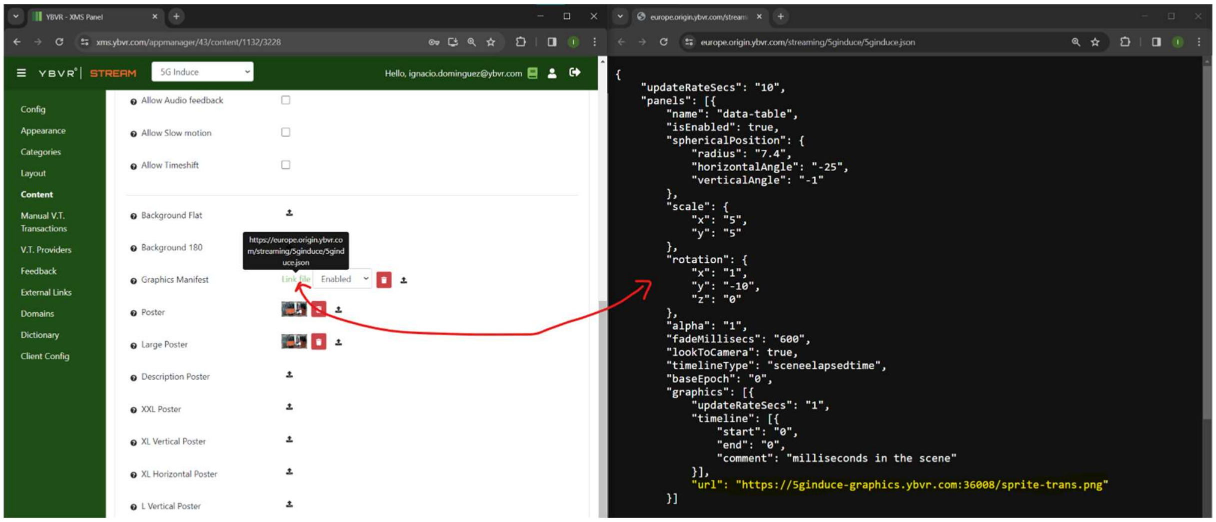

Allow Audio feedback ☐

Allow Slow motion ☐

Allow Timeshift ☐

Graphics Manifest: https://europe.origin.ybvr.com/streaming/5ginduce/5ginduce.json — Enabled

```
{
    "updateRateSecs": "10",
    "panels": [{
        "name": "data-table",
        "isEnabled": true,
        "sphericalPosition": {
            "radius": "7.4",
            "horizontalAngle": "-25",
            "verticalAngle": "-1"
        },
        "scale": {
            "x": "5",
            "y": "5"
        },
        "rotation": {
            "x": "1",
            "y": "-10",
            "z": "0"
        },
        "alpha": "1",
        "fadeMillisecs": "600",
        "lookToCamera": true,
        "timelineType": "sceneelapsedtime",
        "baseEpoch": "0",
        "graphics": [{
            "updateRateSecs": "1",
            "timeline": [{
                "start": "0",
                "end": "0",
                "comment": "milliseconds in the scene"
            }],
            "url": "https://5ginduce-graphics.ybvr.com:36008/sprite-trans.png"
        }]
```

Figure 22 Configuration of the graphic overlay in the XMS.

This XMS generates a content file accessible via the Internet for the player to know the posters (thumbnails of the videos), name of the content, the video url and the overlay url, etc.

When the player application is opened, it requests the content file to the XMS via Internet and the video appears as available in the player APP menu (Figure 23Figure 23):

Figure 23 5G-Induce Oculus application menu.

When the user clicks on the content, the application opens the current live video including the real-time data overlay (Figure 24).

Figure 24 Video + Overlay in Oculus screenshot.

Table 15 SVM-02 UC3

KPIs-UC3	SVM-02
Functionality test	Passed

2.3.2 Service-level validation at ExFas and KPIs measurements

- SVM-09: Glass2glass Latency

The glass 2 glass latency is defined as the time that elapses since the camera takes a frame until that frame is represented in the client device (glasses or tablet). The only way to measure this is visually, for this series of pictures were taken with the camera pointing at a stopwatch, sending the RTMP feed to the Network Application deployed on the server, and the tablet playing the RTMP feed from the server (Figure 25).


00:47,24
00:46,02

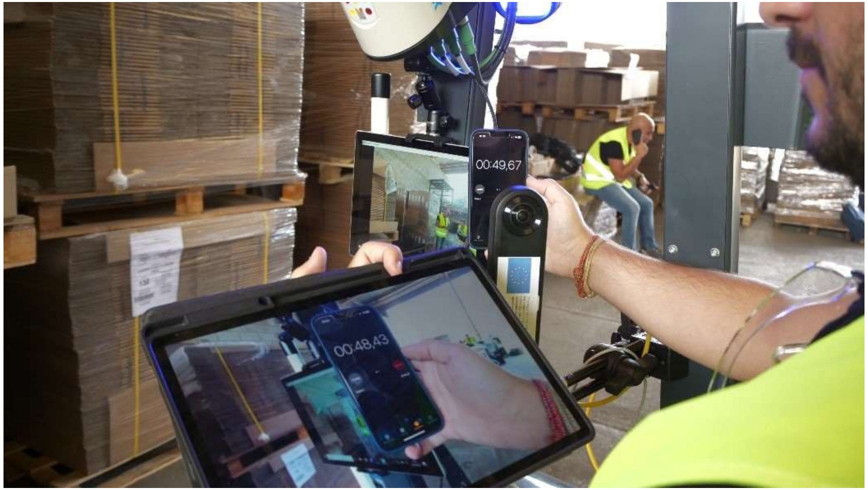
00:49,67
00:48,43


01:36,25
01:35,01

Figure 25 Glass to glass measure.

The following measures (Table 16) were got from the pictures to obtain the average:

Table 16 Glass to glass latency

Time in camera	Time in tablet	Difference
00:21.7	00:20.6	00:01.2
00:24.9	00:23.8	00:01.2
00:47.2	00:46.0	00:01.2
00:49.7	00:48.4	00:01.2
01:36.2	01:35.0	00:01.2
01:46.5	01:45.3	00:01.2
	Average Latency	00:01.2

The average delay between the video taken from the camera and the representation in the tablet was 1.2 seconds, which fulfils the KPI target of less than 2 seconds.

Table 17 SVM-09 UC3

KPIs-UC3	SVM-09
Glass2Glass latency	1.2 sec

- SVM-10: Data Refresh

Doing a database request during the tests in Valencia the following results were obtained:

In the case of the modem database (Figure 26):

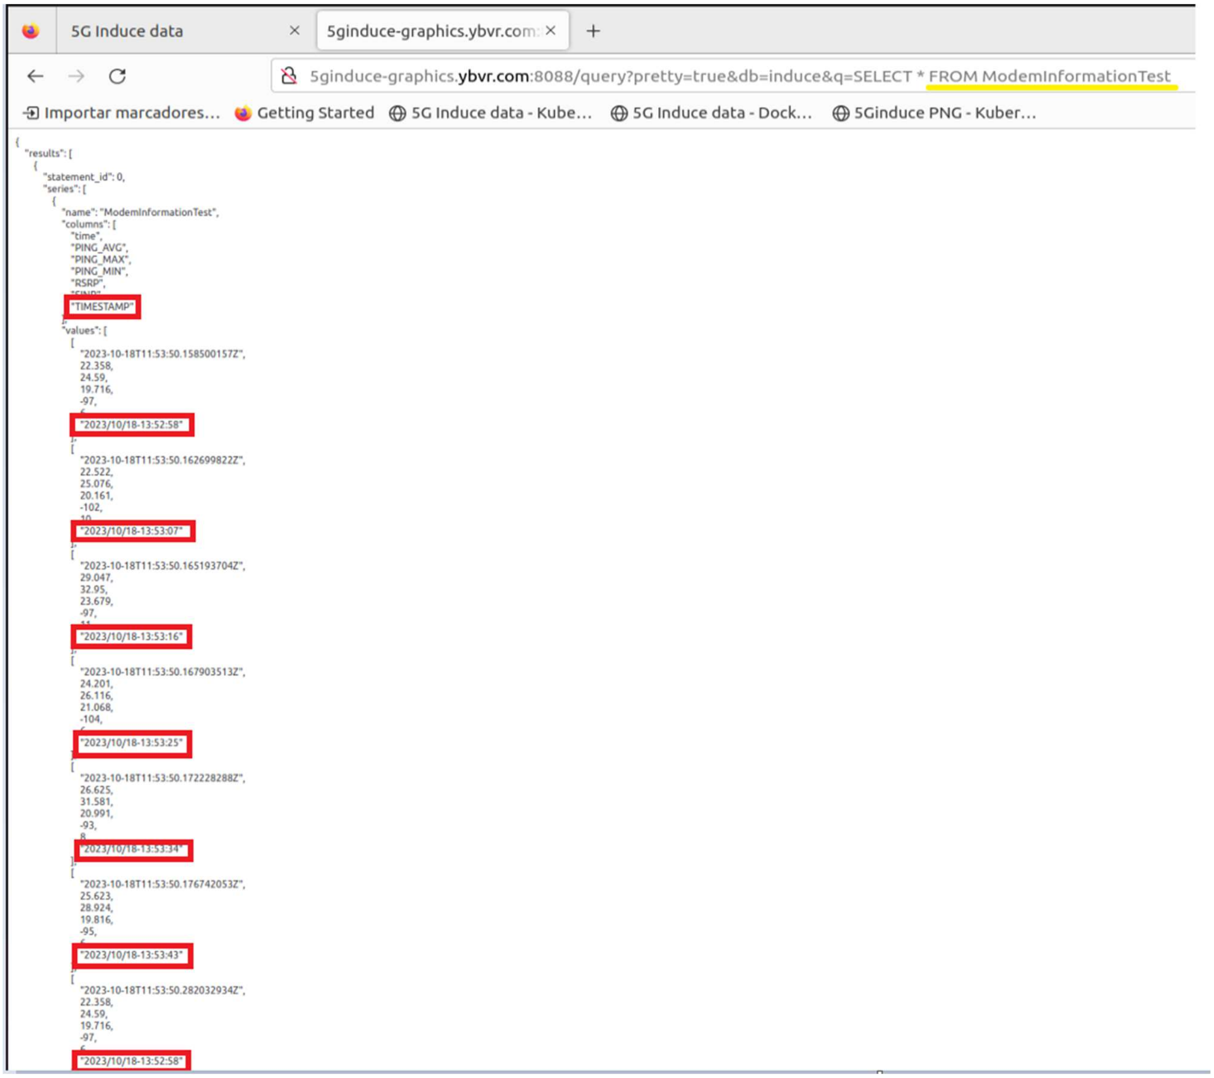

```
{
  "results": [
    {
      "statement_id": 0,
      "series": [
        {
          "name": "ModemInformationTest",
          "columns": [
            "time",
            "PING_AVG",
            "PING_MAX",
            "PING_MIN",
            "RSRP",
            "TIMESTAMP"
          ],
          "values": [
            [
              "2023-10-18T11:53:50.158500157Z",
              22.358,
              24.59,
              19.716,
              -97,
              "2023/10/18-13:52:58"
            ],
            [
              "2023-10-18T11:53:50.162699822Z",
              22.522,
              25.076,
              20.161,
              -102,
              "2023/10/18-13:53:07"
            ],
            [
              "2023-10-18T11:53:50.165193704Z",
              29.047,
              32.95,
              23.679,
              -97,
              "2023/10/18-13:53:16"
            ],
            [
              "2023-10-18T11:53:50.167903513Z",
              24.201,
              26.116,
              21.068,
              -104,
              "2023/10/18-13:53:25"
            ],
            [
              "2023-10-18T11:53:50.172228288Z",
              26.625,
              31.581,
              20.991,
              -93,
              "2023/10/18-13:53:34"
            ],
            [
              "2023-10-18T11:53:50.176742053Z",
              25.623,
              28.924,
              19.816,
              -95,
              "2023/10/18-13:53:43"
            ],
            [
              "2023-10-18T11:53:50.282032934Z",
              22.358,
              24.59,
              19.716,
              -97,
              "2023/10/18-13:52:58"
```

Figure 26 Modem database timestamp.

Comparing the TIMESTAMP field of the database inputs, it is inferred that the modem database is refreshed every 9 seconds.

In the case of the AGV database:

The AGV reports data every 5 seconds.

Table 18 SVM-10 UC3

KPIs-UC3	SVM-10
Data refresh	9 sec

This value is clearly over the target of 2 seconds, but this has a low impact on the functionality of the use case and it's considered easy to improve in future developments.

- SVM-11: Displayed Data Latency

To calculate this, the following times were measured:

The time to open the overlay HTML webpage: it is around 50 ms

The time to get data from the DB: it is around 100 ms

The time to generate the screenshot (download the HTML webpage and save it on a file): around 120 ms

The time to download the screenshot: 50 ms

Aggregating all the figures, a final measure of 320 ms is obtained, which is less than the KPI target (500ms).

Table 19 SVM-11 UC3

KPIs-UC3	SVM-11
Displayed data latency	320 msec

- SVM-12: CPU Load

In terms of CPU load our Network Application is not demanding because the Network Application just restreams the video coming from the AGV to the clients requesting the video. There aren't video re-encoding tasks. The next Figure 27 Figure 27shows the CPU usage of the container during part of the testing:

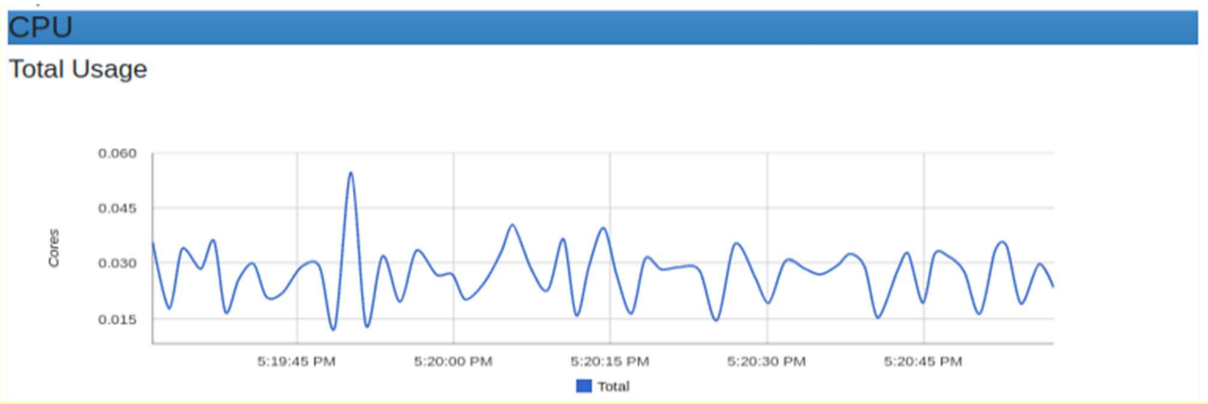

Figure 27 UC3 MEC CPU usage.

It checks out that the usage of the CPU is around 3.5% that fulfills by far the KPI of 80%.

Table 20 SVM-12 UC3

KPIs-UC3	SVM-12
CPU load	3.5%

- SVM-13: Video frame loss

Unfortunately, some technical limitations in the player analytics release, out of YBVR control, did not allow to measure this KPI.

Table 21 SVM-13 UC3

KPIs-UC3	SVM-013
Video frame loss	Not available

- SVM-14: Video buffering

YBVR’s player reports some usage metrics to the YBVR’s platform; among others, the buffer status (Figure 28):

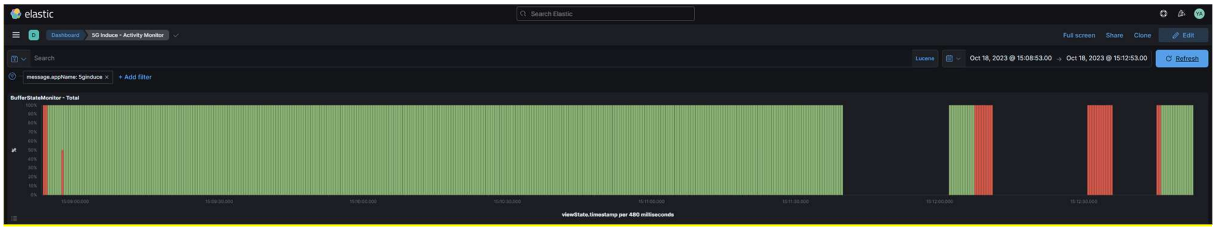
message.appName: 5ginduce

Oct 18, 2023 @ 15:08:53.00 → Oct 18, 2023 @ 15:12:53.00

Figure 28 Buffering status.

5 minutes time slot from FORD's ExFa tests were selected to calculate this KPI, obtaining that the player was buffering video for 38,000 ms and playing video for 275,000 ms (Table 22**Error! Reference source not found.**).

Table 22 Buffering stats

State	ms	Minutes	% time
Buffering	38,000	0.63	12%
Ready	275,250	4.59	88%

This value is clearly over the maximum expected measurement of 1% of buffering. This value is commented on in the next paragraph (Analysis of the results).

Table 23 SVM-14 UC3

KPIs-UC3	SVM-14
Video buffering	12%

- Other metrics

It is important to enlarge this analysis with other measures taken from the network side by the Ericsson probes.

Bandwidth measurement of all TCP traffic is gathered in Figure 29.

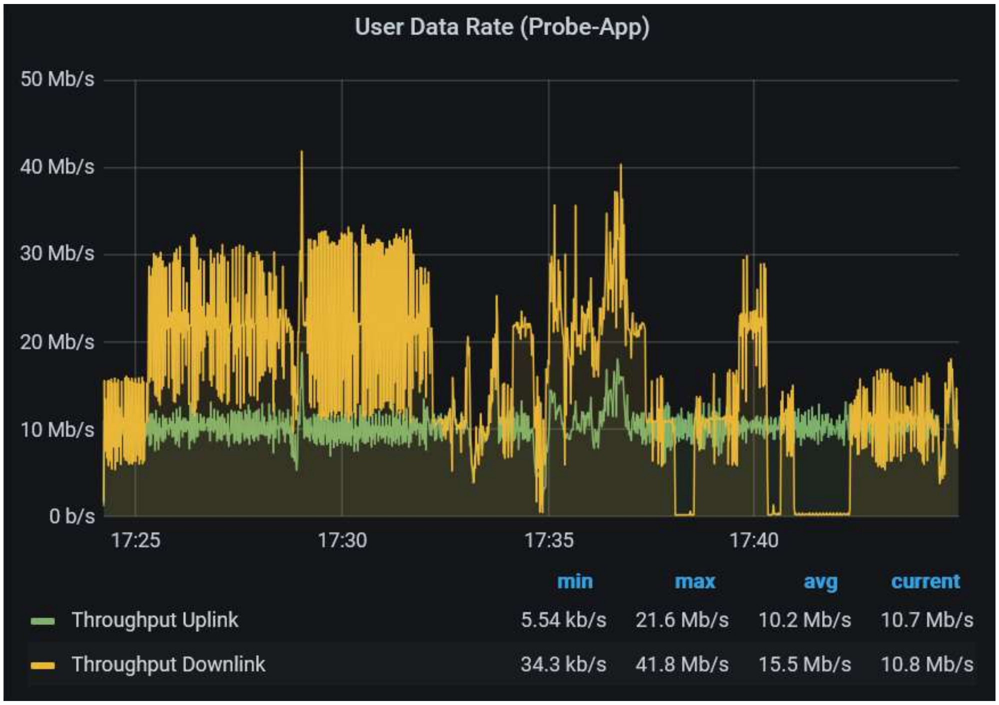

Figure 29 All TCP traffic aggregated.

TCP download traffic latency registered is represented in Figure 30.

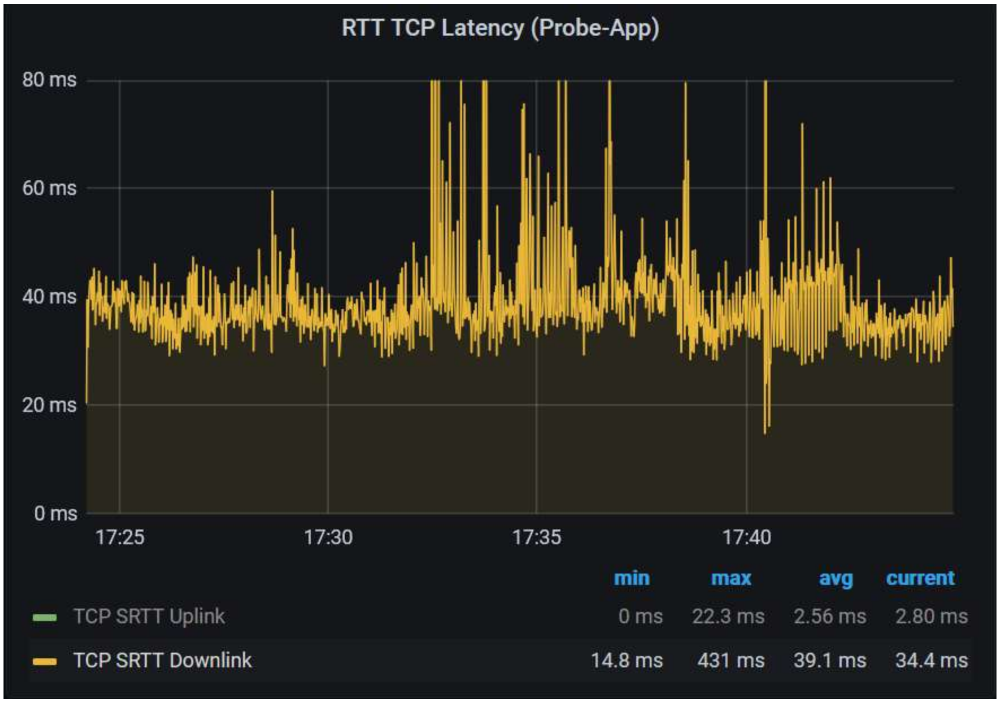

	min	max	avg	current
TCP SRTT Uplink	0 ms	22.3 ms	2.56 ms	2.80 ms
TCP SRTT Downlink	14.8 ms	431 ms	39.1 ms	34.4 ms

Figure 30 Download TCP Latency.

Download jitter is represented in Figure 31.

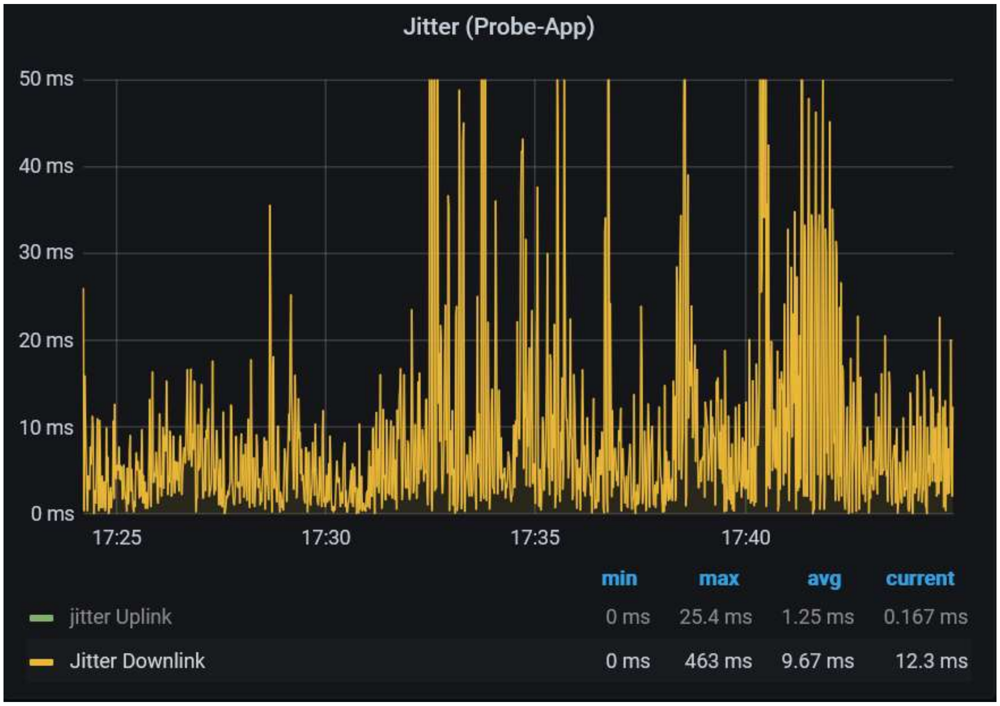

	min	max	avg	current
jitter Uplink	0 ms	25.4 ms	1.25 ms	0.167 ms
Jitter Downlink	0 ms	463 ms	9.67 ms	12.3 ms

Figure 31 Download jitter.

2.3.3 Analysis of the results

The results of this project are offering some questions to explore, opening new paths to research and directions for the evolution of video support in cellular networks.

TCP traffic aggregated graphics (Figure 29 shows very clearly the final setup of video streaming, with a bit rate of 10 Mbps in the contribution feed from the camera to the server (uplink), and the different measurements on the distribution feeds, depending on the number of devices playing the video simultaneously (x10 Mbps in downlink). This bitrate setup was defined after the confirmation of some problems (freezes and delay) with higher bitrates.

Then, the first conclusion has been that **uplink 5G bandwidth has not been a bottleneck in this use case**, contrarily to what was expected initially.

Buffering measurement, surprisingly high, gives a clue about the limitations of streaming in this environment. It alerts about the lack of frames available to feed the decoder. Why can it happen? Two causes could produce this effect: packet loses or high jitter. Jitter graph showed above (Figure 31) is offering downlink values of 9.67 ms on average, and 463 ms maximum. Notice that more than 30 ms is frequent in the graph.

Here, two configuration details must be considered: video streaming was configured as 30 frames per second, which means a new frame every 33 ms. Besides, this graph only gathers the downlink part of the path maybe a similar value for the uplink. The server is not doing any coding of the video; then, uplink and downlink jitter are contributing incrementally to the disordering of the frames. It could explain the high figure of buffering errors detected. The reduction of the bitrate was tested as a good factor to avoid video freeze and delays (the lower bitrate, the larger the jitter tolerance).

The second conclusion of the testing has been the need to **explore the impact of jitter and packet losses in radio links on 4K video streaming** and improvement actions to carry out.

Despite these difficulties, the user experience was very successful. 5G network has been tested as a valid transport channel for immersive video, allowing the remote control of moving vehicles with the bandwidth enough for a fully functional video quality. 5G evolution promises new and improved features to go beyond in immersive video experiences and applications The experience of feeling to be teleported to the AGV, looking around the vehicle, with the updated overlay of the AGV and modem data feeds, has been very positive for the final users. **Exploring new applications of this immersive monitoring**, mainly in unattended or dangerous environments, is something worth to be considered in future projects.

3 Trial results in ExFa-GR

3.1 UC4

3.1.1 ExFa setup

For UC4, the overall components of the AI Engine of the use case specific dashboards that compose eventually 2 Network Applications that interact with each other have been deployed in the ExFa. In total these include the Edge Data Collector, the Edge Analytics Engine and the On-Prem Kafka Platform, the On-Prem Data Collector, the On-Prem Analytics Engine, the Analytics Workflow Designer, the Visualisation Engine and the Data Collection Configurator. Also, a web server that is hosting the different dashboard has been set up, which ingest data from the AI Engine Network Application.

The roles of these components are depending on their placement in the overall deployment. In more detail the role of the Edge Data Collector, the On-Prem Kafka Platform, the On-Prem Data Collector has to do with the collection of data and their ingestion in the AI Engine; the Edge Data Collector is tasked to ingest data coming from the different machines that will be monitored to provide predictive maintenance suggestions, and these data after an initial preprocessing (via the Edge Analytics Engine) are forwarded to the AI Engine once collected by the On-Prem Data Collector that is operating over the On-Prem Kafka Platform that is used to transfer the data from the edge. The execution of the analytics happened over the On-Prem Analytics Engine that is used for detecting anomalies, predicting breakdowns or machine performance degradation and proposing slots for predictive maintenance, as well as sending a predictive maintenance signal back to the machine which can be then trigger various automations such as controlled machine and production shutdown or maintenance mode entry, kickstarting maintenance operations, etc (these actions depend on the use case and on how each operator executes their predictive maintenance procedures, and have not been implemented during the project). The overall configuration of the data collection and the AI pipeline is performed by the Data Collection Configurator and the Analytics Workflow Designer.

Finally, specific dashboards relevant to the use case have been set up, which ingest the data coming out of the execution of the AI pipelines happening in the AI Engine and can provide to machine operators a visual view on the status of the machine as well as the predictions provided by the AI Engine. These are part of the Visualisation Engine, but have been also ported to the Web Interface offered to factory personnel, that is served thro

The general architecture of the placement of the component is illustrated in the following Figure 32.

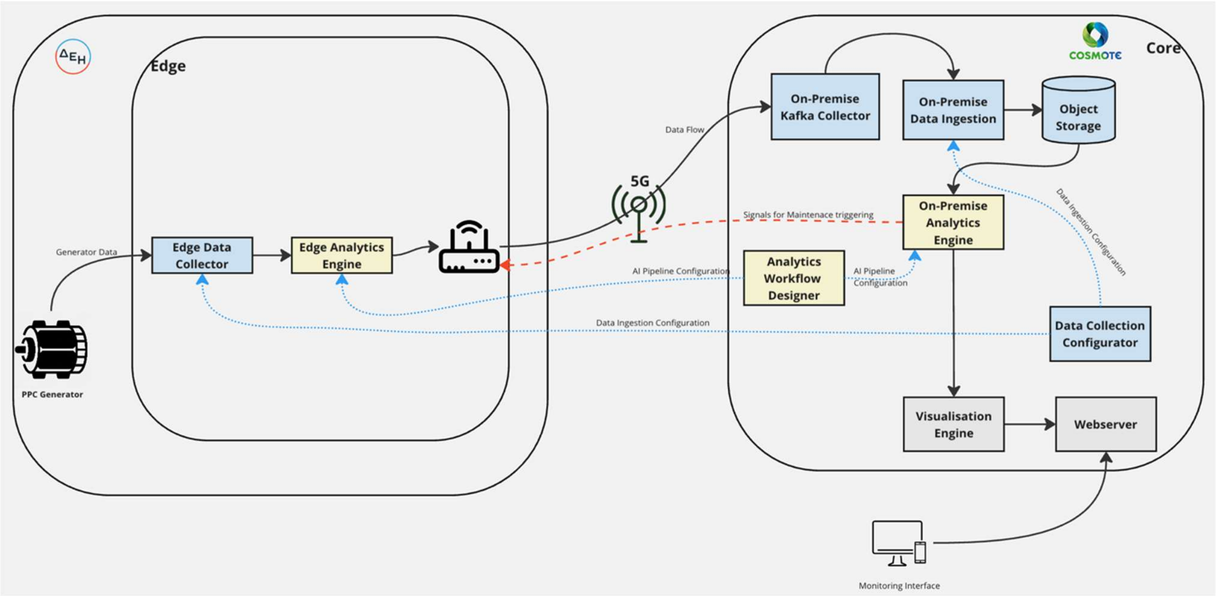

Figure 32 *General Architecture UC4.*

3.1.2 Service-level validation at ExFas and KPIs measurements

The experiment was carried out in PPC Innovation Hub where a scenario was set up ingesting data from a Steam Turbine Electric Generator that was the machine chosen for the predictive maintenance analysis. A 5G gNB was connected to the 5G core located in OTE Academy and allowed the data ingestion nodes placed at PPC to send the acquired data to the AI Engine deployed in the core, and the latter to send back a signal whenever a predictive maintenance action has been suggested. Before any test was performed, a remote onboarding and deployment of this network app had been carried out by Ubitech.

Figure 33 Power Generator at PPC premises used in UC4.

The experiment consisted of several phases which resulted in the development of the analytics engine that can provide predictive maintenance monitoring and failure(alarm) prediction in almost real time.

Preliminary Evaluation of Algorithms

In the initial phased of the project, and prior to the actual data ingestion phase, and since no dataset was available, in order to evaluate the performance of the different algorithms to be used, Eight Bells supported Suite5 by assessing multiple established predictive Machine Learning (ML) and Artificial Intelligence (AI) models to offer preliminary views on their effectiveness, using a synthetic dataset [1] modelled after an existing machine and consisting of 10 000 data points from a stored as rows with 14 features in column

- UID: unique identifier ranging from 1 to 10000
- product ID: consisting of a letter L, M, or H for low (50% of all products), medium (30%) and high (20%) as product quality variants and a variant-specific serial number
- type: just the product type L, M or H from column 2
- air temperature [K]: generated using a random walk process later normalized to a standard deviation of 2 K around 300 K
- process temperature [K]: generated using a random walk process normalized to a standard deviation of 1 K, added to the air temperature plus 10 K.

- rotational speed [rpm]: calculated from a power of 2860 W, overlaid with a normally distributed noise
- torque [Nm]: torque values are normally distributed around 40 Nm with a SD = 10 Nm and no negative values.
- tool wear [min]: The quality variants H/M/L add 5/3/2 minutes of tool wear to the used tool in the process.
- a 'machine failure' label that indicates, whether the machine has failed in this particular datapoint for any of the following failure modes are true.

The machine failure consists of five independent failure modes, some of which were similar to those of the generator. In more details, those were the following:

- tool wear failure (TWF): the tool will be replaced of fail at a randomly selected tool wear time between 200 - 240 mins (120 times in our dataset). At this point in time, the tool is replaced 69 times, and fails 51 times (randomly assigned).
- heat dissipation failure (HDF): heat dissipation causes a process failure, if the difference between air- and process temperature is below 8.6 K and the tools rotational speed is below 1380 rpm. This is the case for 115 data points.
- power failure (PWF): the product of torque and rotational speed (in rad/s) equals the power required for the process. If this power is below 3500 W or above 9000 W, the process fails, which is the case 95 times in our dataset.
- overstrain failure (OSF): if the product of tool wear and torque exceeds 11,000 minNm for the L product variant (12,000 M, 13,000 H), the process fails due to overstrain. This is true for 98 datapoints.
- random failures (RNF): each process has a chance of 0,1 % to fail regardless of its process parameters. This is the case for only 5 datapoints, less than could be expected for 10,000 datapoints in our dataset.

The algorithms chosen for evaluation comprise K-Nearest Neighbours (KNN), XGBoost (XGB), LightGBM (LXGB), AdaBoost (ADA), Decision Trees (DTC), Random Forest, Gradient Boosting Classifier (GBC), Support Vector Machine (SVM), Linear Regression (LR), and Multi-Layer Perceptrons (MLP).

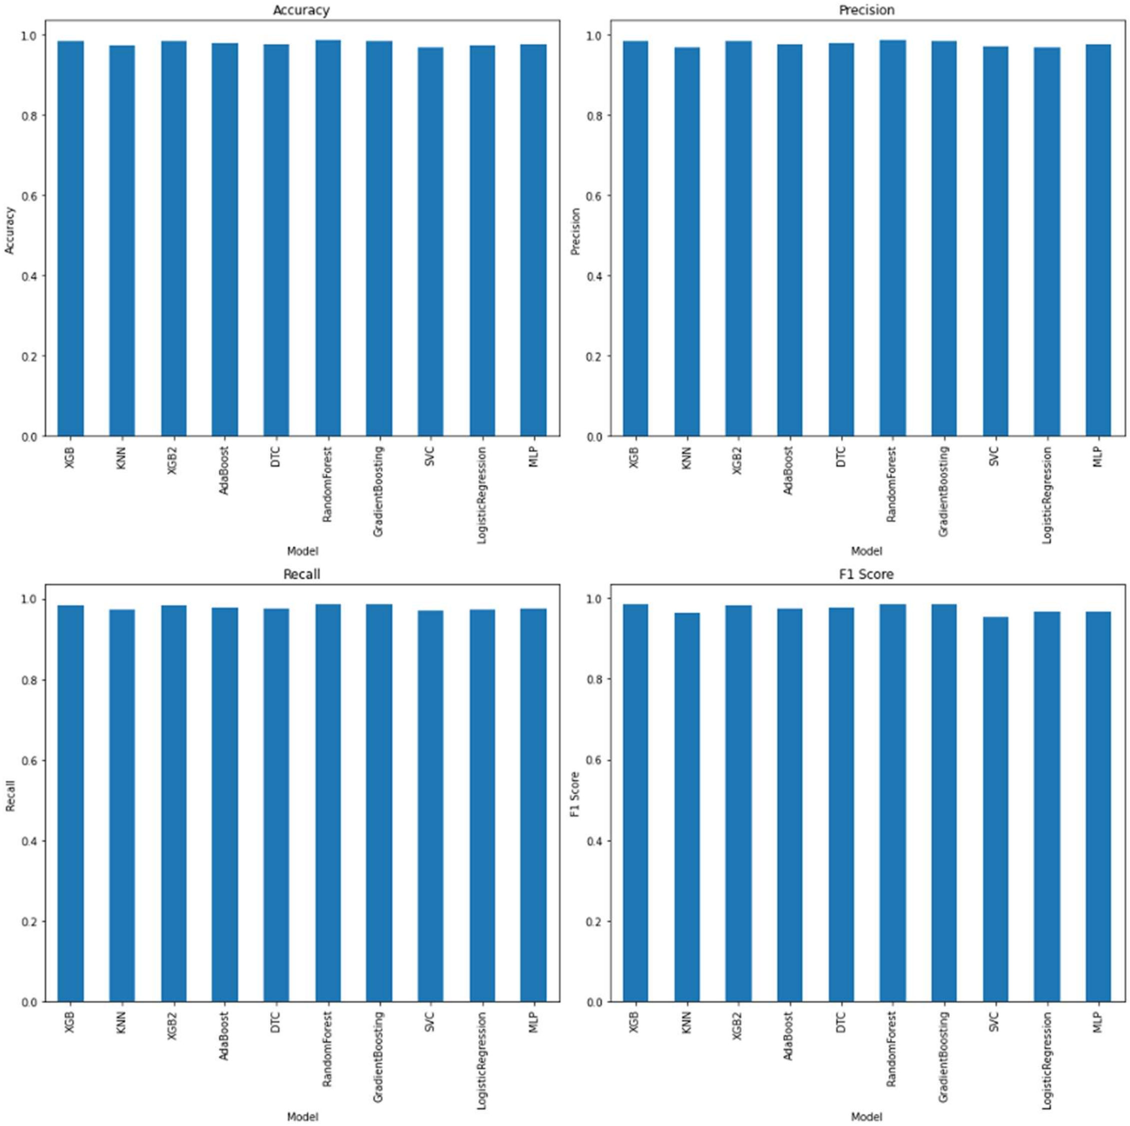

Figure 34 Graphical Representation of the Algorithms Comparison for UC4.

The evaluation was performed using known criteria such as Accuracy, Precision, Recall, and F1 score. The graphical outcomes of this comparison are illustrated in Figure 34, while the numerical data is displayed in Table 24.

All of the assessed algorithms exhibited satisfactory performance, with the Gradient Boosting family displaying the most favourable outcomes across all measures. It is worth mentioning that these algorithms demonstrated impressive performance even without any adjustments or refinements. The findings displayed are derived from the "fine-tuned" models; yet, the disparities observed were negligible. The main objective of this process was to evaluate the overall performance of these models and choose the one that would be used as part of the AI pipelines to be designed for the use case.

Table 24 Numeric Result of the Algorithms Comparison

Algorithm	Accuracy	Precision	Recall	F1 Score
XGB	0.984333	0.983684	0.984333	0.983947
KNN	0.972667	0.968867	0.972667	0.963401
lightXGB	0.984333	0.983111	0.984333	0.983259
Ada	0.978333	0.975464	0.978333	0.975712
DTC	0.976667	0.978453	0.976667	0.977455
RF	0.986667	0.985826	0.986667	0.985706
GB	0.985667	0.984658	0.985667	0.984684
SVC	0.969333	0.970274	0.969333	0.954565
LR	0.973667	0.968297	0.973667	0.967977
MLP	0.975333	0.974779	0.975333	0.967695

Data Collection Phase - Data Streams from Generator

The sensors on the steam driven electric generator produce and send data with a granularity of 1 recording per second. The analytics engine receives the data and creates the appropriate results in the same frequency since the input processing and inference is made on windowed data sets (except for the cold start period).

The first phase was dedicated to the collection, cleaning and analysis of historical data that were used to train our models. The data included sensor measurements and logs of maintenance events, for a full year (2023). Table 25 shows a list of the sensor measurements and their description.

Table 25 *Sensor Measurements from Steam Turbine Electric Generator*

Measurement	Description
Steam Flow (tn/h)	Measurement of the flow of steam which powers a turbine that drives the generator.
Steam Temp bef SV (C)	The temperature of steam which causes the shaft of the turbine to rotate.
Steam Pressure (bar)	The pressure of steam that pushes on the many angled blades of the turbine, causing the shaft to rotate.
Vertical Vibration Bearing No 1 (mm/s)	The measurement of the effect of vertical vibrational load on the bearings
Horizontal Vibration Bearing No 1 (mm/s)	The measurement of the effect of horizontal vibrational load on the bearings
Condenser Absolute Pressure (bar)	The pressure created when cooling water causes the steam to condense.
Rotor Speed (rpm)	The current running speed of the generator's rotor.
Generator Active Power (MW)	The useful power produced by the generator

The event log, on the other hand, is the recording of events with a “start” and “end” datetime, along with the type of event. The type can get one of the following values: emergency maintenance (EM), manual preventive maintenance (MP) or scheduled downtime (SD). In the failure prediction scenario, we focus on the EM type only, since it is the only one that is worth predicting.

Analytics Design and Execution Phase - Failure prediction

The measurement data were cleansed and resampled at 1 min frequency and a Feed Forward Neural Network (keras) was trained to predict a possible event 1 minute before happening, using the knowledge derived from the last 10 mins of measurements. Since the EM events in the logs were scarce, the actual training data were mostly selected to be around the time periods that those events occurred.

The final set of 18 features used to predict failure at time step *t+1* were: *Steam Flow, Steam Temp, Steam Pressure, Vertical Vibration Bearing, Horizontal Vibration Bearing, Condenser Absolute Pressure,* at time step *t*, and for each one of these features we computed the rolling average of the last 10 minutes, as well as the trend (increasing or decreasing) of the last minute.

The next phases included the validation, testing, fine-tuning and evaluation of the model using a part of the historical data that was not used for training.

The final model was set in the production environment and was able to ingest the data stream, accumulate information and produce the outcome in the form of an alarm for the next minute.

Analytics Design and Execution Phase - Sensor Monitoring and Forecasting

This scenario involves the sensor measurements in raw form and attempts to predict their behaviour in the seconds to follow, which can be very helpful for monitoring purposes. From an analytics point of view, It is a common time series forecasting task that employs lagging values from the current time step in order to predict the next value.

The measurement data were cleansed and resampled at 10 secs frequency and an XGBoost regressor was trained for each of the sensors involved, namely *Steam Flow, Steam Temp, Steam Pressure, Vertical Vibration Bearing, Horizontal Vibration Bearing and Condenser Absolute Pressure*. Using the lag features of 10, 20, 30, 40, 50 and 60 seconds ago, each model is able to accurately forecast the measurement value of the next step i.e. 10 seconds ahead.

The experiment KPIs' are the following:

- SVM-01: Service deployment. This KPI measures the time taken for service deployment at the ExFa site using the 5G-INDUCE platform. The target deployment time was set at less than 4 minutes, and the actual deployment time was measured at an average time of 220 seconds (3,5 minutes). It's important to note that this deployment time was measured after all container images had been fully downloaded, ensuring a fair assessment of the deployment process efficiency.

Table 26 SVM-01 UC4

KPIs-UC4	SVM-01
Service deployment time	220 seconds

- SVM-02: Functional tests. This KPI measures the ability to effectively transfer data to/from the nApp, focusing a) on the delivery of the complete set of data from the machine to the storage facility and b) on the correct export of the data from the AI engine to other sources.

Table 27 SVM-02 UC4

KPIs-UC4	SVM-02
Functional tests – a) Data Transferred to nApp	100% of generated data - PASS
Functional tests – b) Data Exported from nApp	100% of AI output data - PASS

- SVM-15: Data Transfer Time. This KPI compares data transfer time between the Shopfloor and the nApp for each specific measurement that is forwarded to the analytics engine, where in the first case it is part of the NAO and in the second it is hosted on the cloud. For our experimentation, apart from the NAO provided by 5G Induce, we experimented with 2 cloud providers, namely DO (DigitalOcean) and Google Cloud, by having the Analytics Engine deployed at those infrastructures as well. The timings we also equal when sent back to the source collection point, where we had installed a docker application that was capable to capture signals coming from the analytics engine that could trigger maintenance operations.

Table 28 SVM-15 UC4

KPIs-UC4	SVM-15
Average Data Transfer Time (NAO)	38.12 msec
Average Data Transfer Time (Cloud – DO)	157.26 msec
Average Data Transfer Time (Cloud – Google Cloud)	143.61 msec

Figure 35 presents a comparison of the transfer time measured over the different deployments/infrastructures.

Figure 35 Transfer time of sensor data to/from different infrastructures tested.

- SVM-16: AI/ML Pipeline Execution Duration. This KPI compares the time to execute the AI/ML pipelines designed between the NAO nApp and the cloud-based nApp. For these metrics, we deployed similar instances with similar resources for the analytics engine, with the only exception that the scaling options in the cloud providers were slightly higher (in terms of resources that could be added) than those of the NAO, as these were pre-defined by the hosting packages and not subject to change/configuration.

Table 29 SVM-16 UC4

KPIs-UC4	SVM-16
Average AI/ML Pipeline Execution Duration (NAO)	289 msec
Average AI/ML Pipeline Execution Duration (Cloud - DO)	252 msec
Average AI/ML Pipeline Execution Duration (Cloud - Google)	234 msec

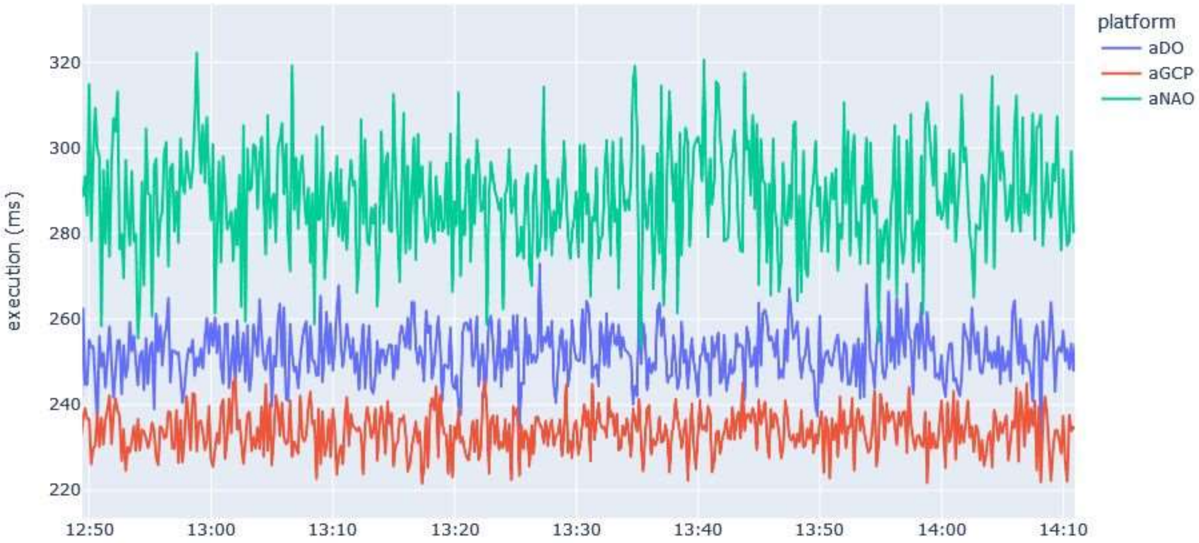

Figure 36 Execution Time recorded over the different infrastructures tested

3.1.3 Analysis of the results

The results of the experimentation with UC4 under this ExFa proved the claim that the placement of AI solutions over 5G infrastructure could prove useful for certain applications, such as the anomaly detection and consequently the predictive maintenance of machines used in industry settings which are capable of continuously streaming their operational condition parameters and environmental and production data that could be used to identify when certain maintenance operations should take place.

The specific case of the experimental generator in the premises of PPC serves as a use case where edge processing of information could facilitate predictive maintenance operations and in turn safeguard the operation of machinery used in critical infrastructures, as the overall health of such devices does not only depend on constantly monitoring the status of the machine, but also on timely detecting and predicting anomalies and be in a position to rapidly trigger the initiation of corrective actions (such as requesting less output or generator shutdown) that will help to prevent catastrophic results for the machine itself, and that may also impact the overall critical infrastructure.

As the analysis of the KPIs shows the main things that need to be considered aside from the deployment metrics, are those of the data roundtrip and the AI execution duration. The sum of those two factors provides the time interval between the collection of the data, the processing and the triggering of an event that can reach back the generator to initiate the maintenance process.

As conventional AI/ML applications run on cloud infrastructure, in our analysis we compared these timings as recorded over the 5G-INDUCE infrastructure with those of having the same AI/ML scenarios operated over public cloud infrastructures (experimented both with Google Cloud Platform and Digital Ocean platform).

As shown in the graphs in the previous section, data roundtrip over the NAO is faster than roundtrips recorded when the machinery (generator in this case) had to send the data to a public cloud platform and then receive the signals to initiate the maintenance automation procedures.

In parallel, the data processing and AI/ML preprocessing timings were also recorded and compared between the 3 different cases. As the graphs suggest, the performance of the AI/ML engine in the case of the NAO was slightly slower than those recorded over the different cloud infrastructures, as the latter have the ability to scale more than the resources available in the NAO, while the initial resources used were slightly more than those used over the NAO.

By combining these numbers, as shown in Figure 37, it is proved that for the specific case, the 5G deployment offers an advantage over the public cloud infrastructure, as the loss in execution time (for the scale of analytics developed for the specific solution) can be compensated by the data transfer times. This difference would be even bigger in case the analytics engine was used to ingest data for various machines at the same time over the same factory setting, assuming that in such a scenario the scaling configuration of the Ai Engine between the NAO and the cloud resources would be identical.

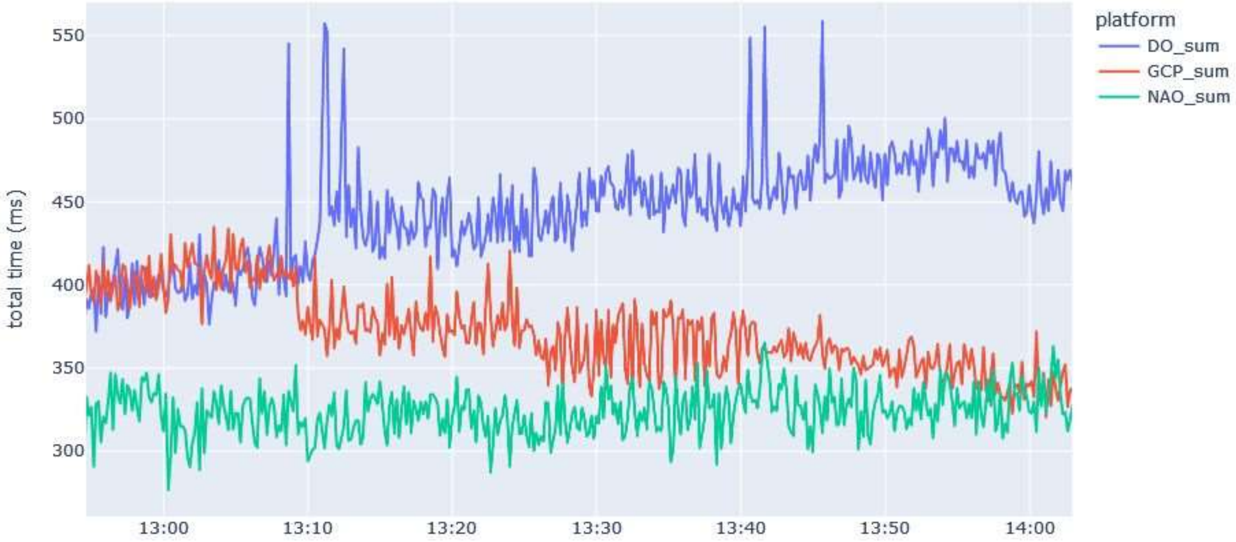

Figure 37 Combined Timings for Data Transfer and Analytics Execution over the different infrastructures tested.

Finally, it needs to be stated that the benefits of using a 5G infrastructure extend further than those recoded through the above mentioned KPIs. The ability to host applications over the NAO and the slicing of resources allows the proper use of already available computing resources and can be seen as one of the key benefits of using 5G for accommodating AI/ML applications that can run on the edge. Moreover, data confidentiality and privacy issues can be easily addressed, as in a 5G deployment scenario as the one used in this use case, no data needs to leave the setting of the factory (or the core network) and thus all processing happens over trusted resources.

3.2 UC5

3.2.1 ExFa setup

In Use Case 5, two Network Applications are developed and deployed, namely the Corrosion Detection application and the Intruder Detection application. Both applications leverage an unmanned aerial vehicle (UAV) system enhanced by cutting-edge AI technologies deployed at the edge and core of 5G networks.

The Corrosion Detection Network Application is designed to inspect tanks and pipelines, identifying corroded areas on their surfaces. The primary objective is to efficiently detect corrosion, leading to a significant reduction in corrosion maintenance costs. Moreover, the Intruder Detection Network Application conducts real-time area surveillance with minimal end-to-end latency. This capability enhances the protection of critical industrial infrastructures by promptly detecting intruders.

Figure 38 illustrates a general architecture overview of Use Case 5 in the format of a graph for both Network Applications including all the essential components, communication flows, and key players and stakeholders. During the operation, the pilot will transmit video footage from the UAV to the Network Applications, which will identify corroded areas and intruders. The pilot will receive this information and it will be shown on top of the smartphone screen along with the video footage. On the other side a Commander entity supervises the video streamed from the UAV and the corresponding detections.

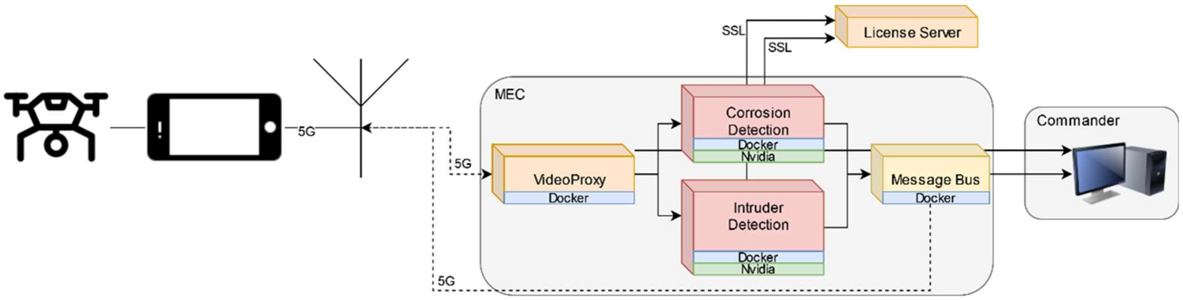

Figure 38 General Network App architecture overview for UC5.

This use case is composed of four different components to be deployed by the 5G-INDUCE platform. The description of each component and the technologies used to deploy and use specific hardware such as NVIDIA is described in deliverable D4.2. An external component has been added in order to provide security and licensing to the network application. A license key will be given to every user that will be able to use an Intruder/Corrosion Detection Network App. These Network Applications will ask the License Server component for permission to perform high computational expensive AI algorithms to detect corrosion or intruders.

Every network application component was dockerised and stored in the internal container library of the project in GitLab. Besides a YAML (docker-compose) file that specifies the correct deployment information is shared. The 5G-INDUCE platform will download and deploy each of the components dockerised in the corresponding machine. Figure 2 presents the final system’s architecture deployed with the protocols and interfaces mapping of each component when it was tested in the Greek Experimentation Facility. Standard interfaces for protocols such as AMQP (Port: 5672/tcp) or RTMP (Port: 1935/tcp) will be mapped to the ones available (decided by the 5G-INDUCE platform). Moreover, each component deployed in the system will have an accessible port exclusively reserved for the 5G-INDUCE platform. These ports function as health checks or heartbeats, enabling the platform to ascertain the operational status of the components—whether they are still active or have undergone termination.

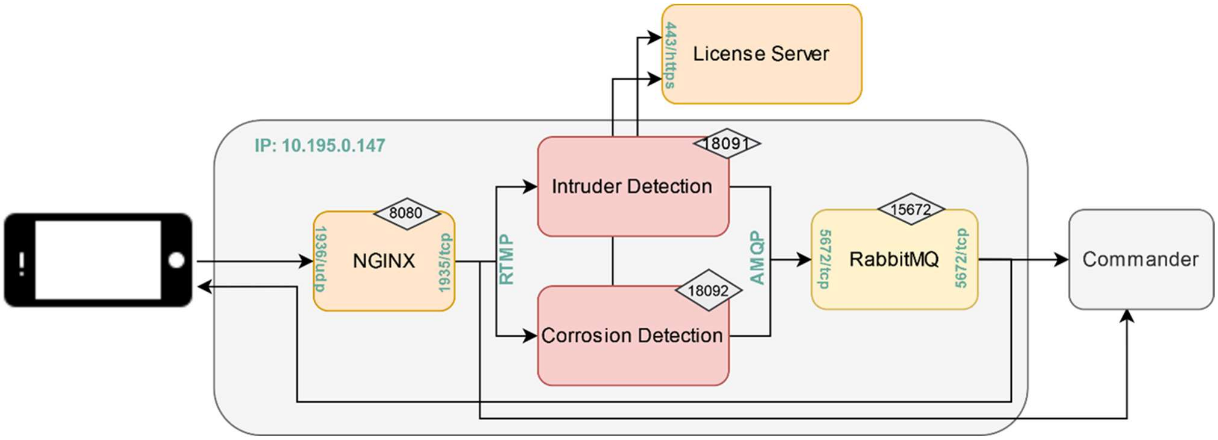

Figure 39 Greek ExFa UC5 deployment.

Once the platform has successfully deployed each of the components, the pilot can connect to the services. First, the UE will check if every component is deployed and if it is correctly working. Next, it will connect with the UAV and forward the video stream from the UE to the network application. Figure 39 presents the UC5 deployed in the trial.

3.2.2 Service-level validation at ExFas and KPIs measurements

The experiment was carried out in PPC Innovation Hub where a scenario was set up with different pipes with corrosion and multiple people acting as "intruders". A 5G gNB was connected to the 5G core located in OTE Academy and allowed the pilot's smartphone to connect to the network application deployed in the core. Before any test was performed, a remote onboarding and deployment of this network app had been carried out by Ubitech.

The experiment KPIs' are the following:

- SVM-01: Service deployment. This KPI measures the time taken for service deployment at the ExFa site using the 5G-INDUCE platform. It evaluates the duration from when a Network Application instantiation request is initiated by the end user through the NAO to when the confirmation of service instantiation appears on the NAO front-end's service management tab. The target deployment time was set at 60 seconds, but the achieved result was notably faster at 33 seconds. It's important to note that this deployment time was measured after all container images had been fully downloaded, ensuring a fair assessment of the deployment process efficiency (Table 30 SVM-01 UC5).

Table 30 SVM-01 UC5

KPIs-UC5	SVM-01
Service deployment time	33 s

- SVM-02: Functional tests – Connection establishment. This KPI the proper establishment and validation of connections among Network Application components. It ensures that video streams are shared, and detection results are accurately displayed on screen. The KPI's target is the successful validation of these connections (Pass/Fail). There was a successful validation. Each component of the Network Application performed as expected during the testing phase. Screenshots (Figure 40, Figure 41 and Figure 42) were captured as evidence of the successful validation of this functional test (Table 31).

Table 31 SVM-02 UC5

KPIs-UC5	SVM-02
Functional tests – Connection establishment	Pass

Figure 40 Corrosion detection of different pipes with areas non-corroded.

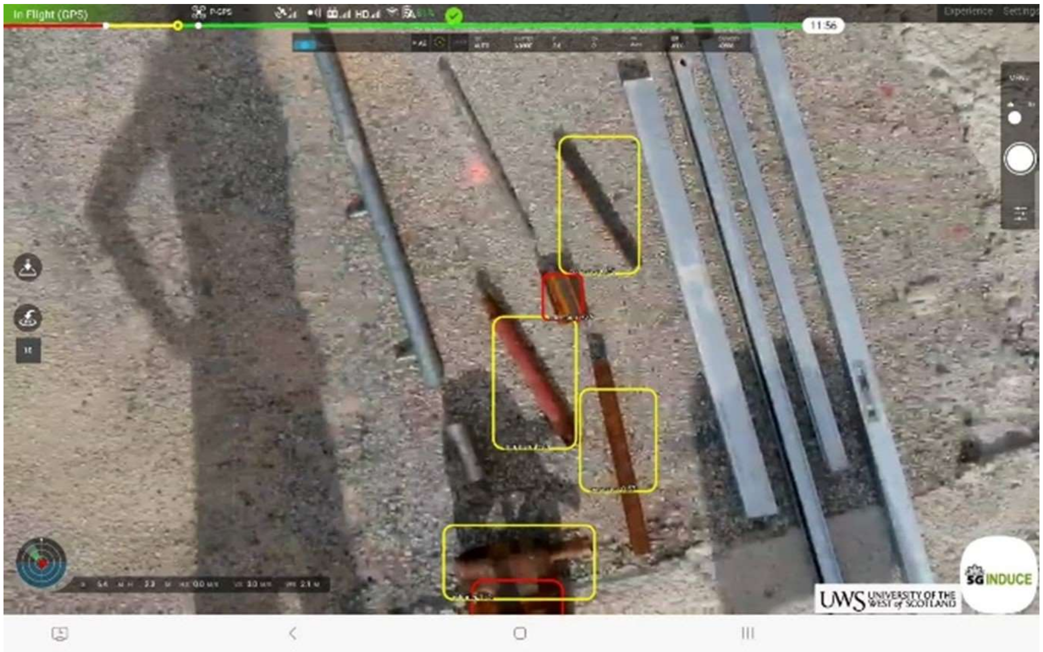

Figure 41 Corrosion detection of pipes of different sizes and materials.

Figure 42 Intruder detection of all the 5G-INDUCE partners involved in the testing

- SVM-03: Application-specific configuration time. Application-related configuration time after a certain monitoring parameter that triggers the configuration is detected. The KPI was set at 60 seconds, with actual results showing an improvement, achieving configuration in 48 seconds. The recorded time represents the average duration taken by pilots or users to configure the UE application to successfully establish connection with the Network Application (Table 32).

Table 32 SVM-03 UC5

KPIs-UC5	SVM-03
Application-specific configuration time	48 s

- SVM-06: HR video quality (real-time operation). This metric evaluates the quality of high-resolution real-time video streamed from UAVs to the Network Application. The KPI was set at 8 Mbit/s or HD video quality. The achieved result meets this standard, with HD video being transmitted during operations. Specifically, the video streamed has a resolution of 1280x720 pixels, exceeding the baseline HD resolution of 1080x720 pixels (Table 33).

Table 33 SVM-06 UC5

KPIs-UC5	SVM-06
HR video quality (real-time operation)	HD video is transmitted

- SVM-07: End user perceived latency. This metric assesses the perceived latency experienced by end users from the moment the UE begins streaming a video to when the pilot starts detecting events. During evaluation, the Key Performance Indicator (KPI) involves pilots assigning a score ranging from 0 to 5. The trial results indicate an average score of 4.6 for the time taken to start detecting events and an average score of 4.8 for video latency. It's noteworthy that feedback was collected from five participants during the trial (Table 34).

Table 34 SVM-07 UC5

KPIs-UC5	SVM-07
End user perceived latency – Start detecting	4.6
End user perceived latency - Latency	4.8

- SVM-16: AI/ML pipeline execution time. This metric measures the time required for the Network Application to execute its AI/ML pipelines. Specifically, it evaluates the processing time per frame for different AI models. The KPI was set at 50 milliseconds per frame. The results show that the Corrosion AI model executed in 23.56 milliseconds per frame, while the Intruder AI model took 40.53 milliseconds per frame (Table 35). These timings are illustrated in Figure 43.

Table 35 SVM-16 UC5

KPIs-UC5	SVM-16
AI/ML pipeline execution time – Corrosion AI	23.56 ms
AI/ML pipeline execution time – Intruder AI	40.53 ms

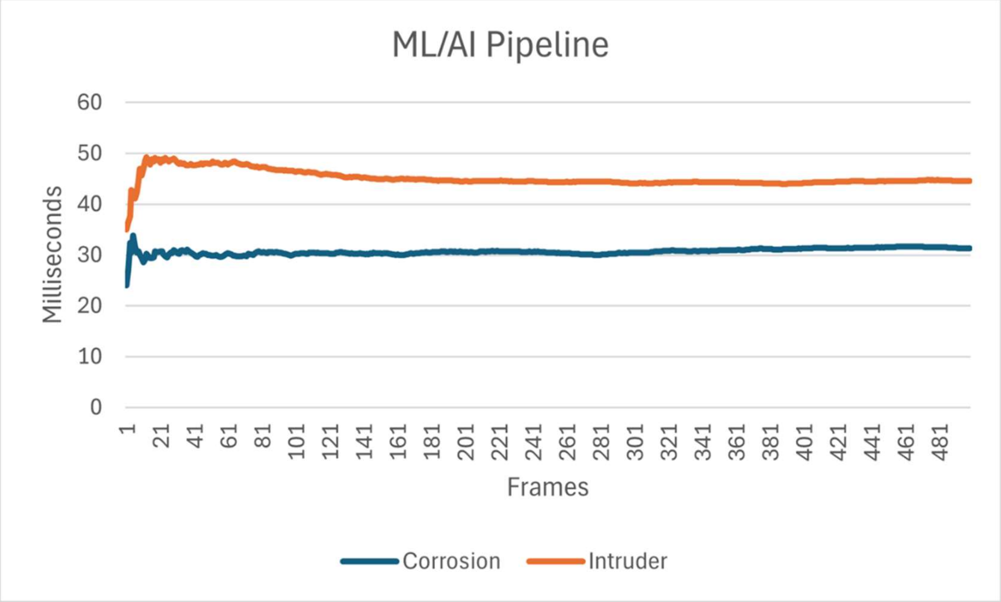

Figure 43 Execution speed of the frames processed over the Corrosion and Intruder detection components.

3.2.3 Analysis of the results

Several conclusions can be extracted after the validation and testing trial performed for the Use Case 5 in the 5G-INDUCE platform at the Greek ExFa. As can be appreciated, the fast deployment (SVM-01) is provided by the 5G-INDUCE platform allowing high scalability and orchestration when ensuring the optimal resource allocation whilst deploying the different services regarding its requirements. In computationally expensive applications such as UC5, where two instances of AI models require GPU capabilities, this resource efficiency by preventing over-provisioning can lead to cost savings.

In addition, one of the great advantages highlighted over the tests is the capability of choosing between edge computing and core computing. Using edge computing allows the users to process their video streaming close to the data source resulting in a reduction in video streaming latency. This became very useful when deploying this system in critical infrastructure such as in this validation (PPC) e.g. for intruder detection. Nevertheless, if the use case is less critical, the allocation of the AI models may be ported to the core, improving the overall system performance in terms of accuracy.

The attainment of functional test (SVM-02) results, high-resolution video streaming (SVM-06), and optimal end-user perceived latency (SVM-07) are intricately linked to the strategic implementation of the 5G network in combination of the 5G-INDUCE platform. Particularly through the utilization of network slicing, which plays a pivotal role in tailoring and allocating the network capabilities to meet the specific requirements of each application, enhancing the overall efficiency and performance of the network. This approach allows a more targeted and customized allocation of network resources, ensuring that the functional tests, high-resolution video streaming, and latency targets are met with precision and adaptability within the 5G framework.

Overall, the trial was a success with both Network Applications validated and all the target KPIs met.

4 Trial results in ExFa-IT

4.1 UC4

4.1.1 ExFa setup

The setup of UC4 - predictive maintenance was executed through the preparation and activation of the required data pipeline from Whirlpool's legacy environment.

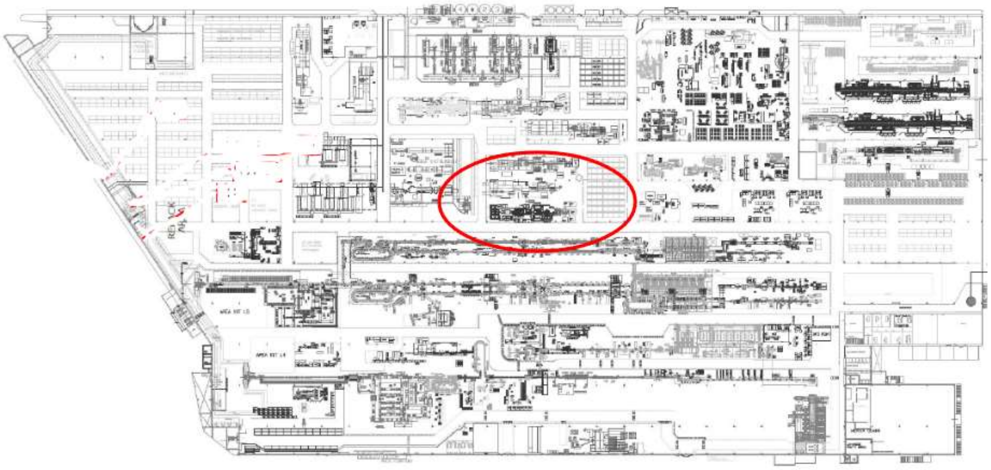

Figure 44 Whirlpool Factory Execution Setting of UC4.

In principle, as the this experiment is based on the same nApp that is used in the PPC case for UC4 (see section 4.1 of the document at hand), the same components have been deployed, and the architecture for this UC was almost identical with the one used in the PPC, as shown in Figure 45.

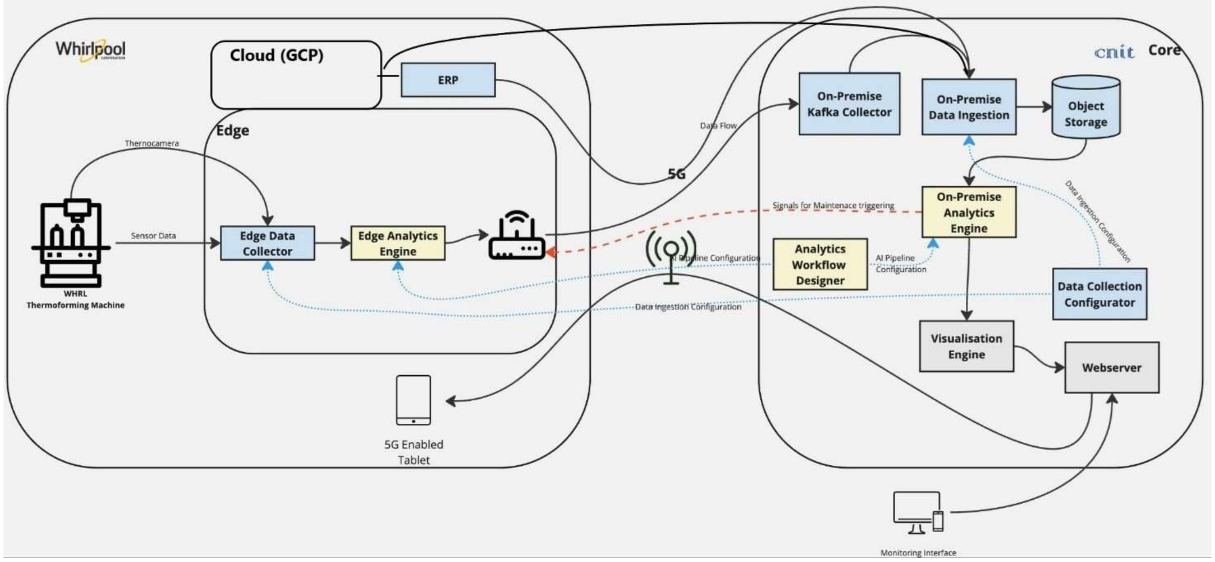

Figure 45 Whirlpool UC4 nApp Architecture.

As such, the components used in this case were the same as those of the PPC case, namely the Edge Data Collector, the Edge Analytics Engine and the On-Prem Kafka Platform, the On-Prem Data Collector, the On-Prem Analytics Engine, the Analytics Workflow Designer, the Visualisation Engine and the Data Collection Configurator. Also, a web server that is hosting the different dashboard has been set up, which ingest data from the AI Engine Network Application.

The main difference between the two cases have been the data collection methods, and the different analytics pipelines that have been developed, as each one tackled a problem with different parameters and constraints, while also different dashboards were created in each case.

Specifically, for this use case, in order to produce the predictive analytics results, the following data related to the thermoforming machine (COMI5) have been prepared to be ingested by Suite5 predictive maintenance platform:

- operational data (production data, sensor data, diagnostic data,...) extracted from machine PLC
- OEE data (efficiency data) collected by NEXT legacy system
- Machine ledger data (AM/PM planned maintenance data) collected by NEXT legacy system
- PM data (maintenance work orders data) collected by SAP PM legacy system
- Vision system data (thermal distribution images) collected by thermal camera vision system in the heating area of the machine.

These data, related to 2022 and 2023, have been collected, analysed, cleansed and used for algorithms training while the data 2024, ingested through pipeline, have been used for prediction generation.

Figure 46 Shop Floor Personnel inspecting the machine and consulting the analytics outputs.

Figure 47 UC4 nApp Analytics Dashboard displayed on a table on the shop floor.

The visualisation dashboard on 5G mobile device has been developed to provide evidence for the prediction and of the rough data to users near the machine in the thermoforming department (Figure 48, Figure 49).

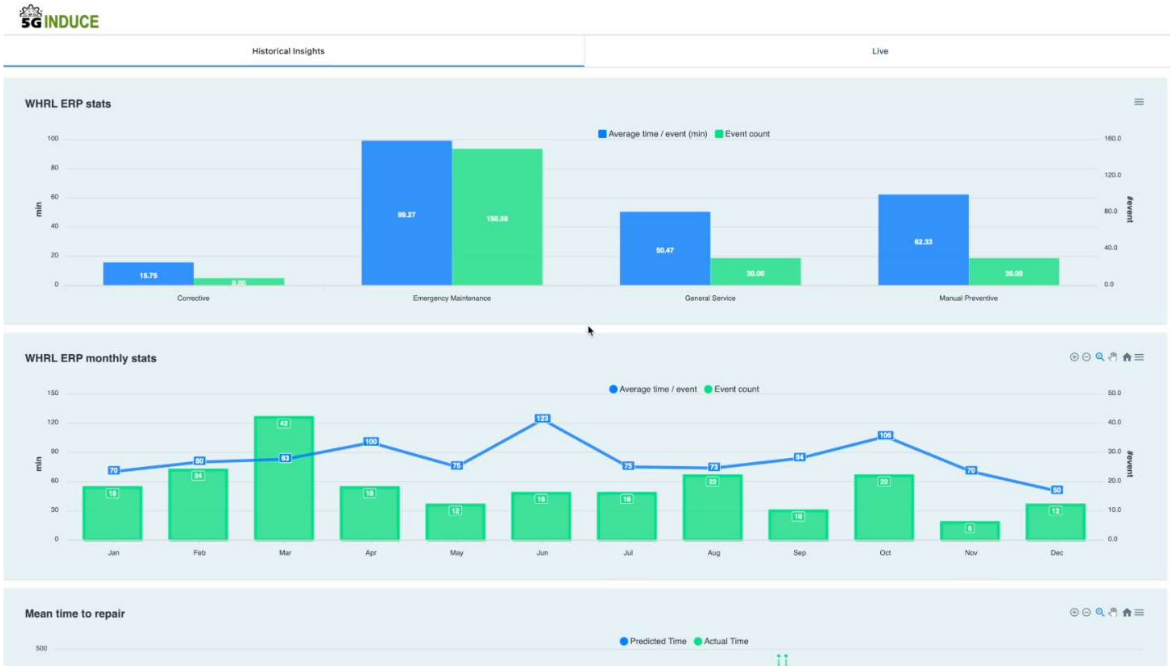

Figure 48 UC4 nApp Analytics Dashboard - ERP data.

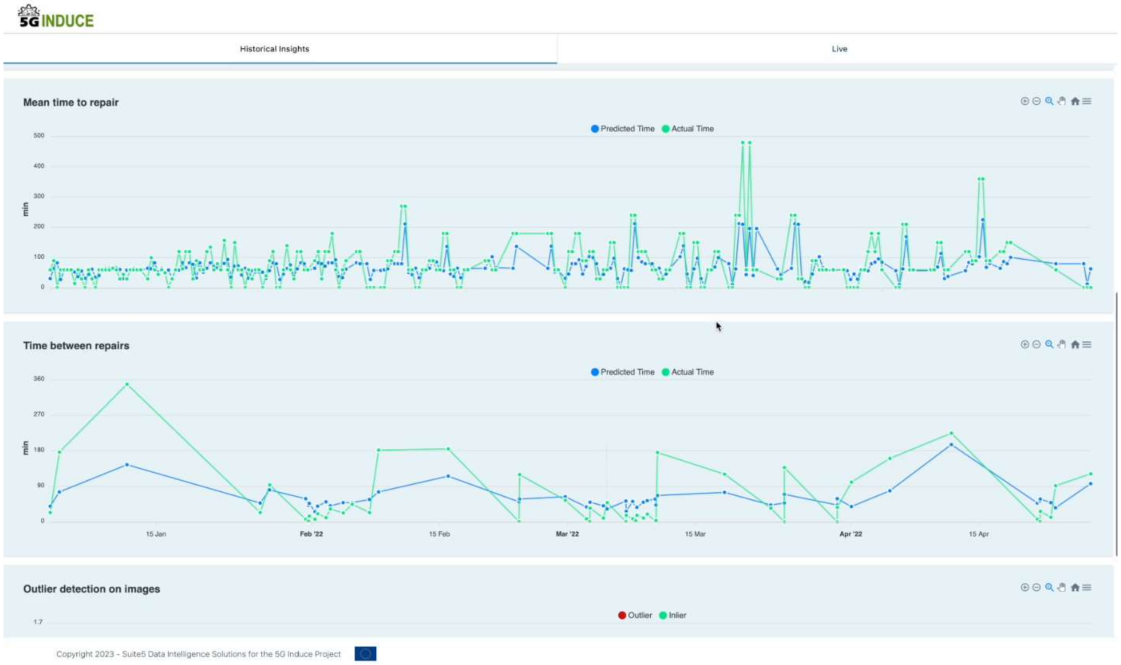

Figure 49 UC4 nApp Analytics Dashboard - Predictive Maintenance relevant Metrics.

In both options an Android device (Lenovo P12 tab pro) has been used and it was connected to 5G network through 5G SIMs: the test has been executed both with option1 using Wind3 SIM and with Option2 using CNIT dedicated SIM connected to the mobile testbed. For the option 2 demo the test could not be executed near the thermoforming machine but in the UC7 demo area to remain within mobile testbed coverage range.

4.1.2 Service-level validation at ExFas and KPIs measurements

This experiment achieved to take advantage of the various heterogeneous groups of data available and produce an event-based informative predictive maintenance application for the Shopfloor and the thermoforming machine (COMI5) in particular.

The past data of two full years (2022 and 2023) were utilised as training sets for two different types of ML models. The first was dedicated to predict the time between repairs (TBR), which indirectly predicts the time until the next failure, and the second responsible to estimate the mean amount of time required to repair the next failure (MTTR). The data employed were from the machine ledger and maintenance work orders data collections.

MTTR prediction

An XGBoost regression model was trained on the training set using the information of previous maintenance actions, as well as calendar features to predict the “mean time to repair” for the next event. The machine ledger data were cleansed and analysed, while a few new features were extracted such as the actual repair duration and the hour of the day each event occurred.

The next MTTR is predicted using a set of 7 features: the actual time required to repair the latest 3 incidents, as well as their original scheduled time, and the hour of the day of the last incident categorized in 2 shifts (day, night).

TBR prediction

Another XGBoost regressor was employed for this task. In addition to machine ledger data, the maintenance orders and the stock info was also added to the training set so that the type of maintenance is also considered for this prediction. The model was trained on the same period as the MTTR one, so that both models are aligned.

The next TBR is predicted using the following input features: the type of last maintenance action, the previous TBR, the hour of day the last action took place, as well as the types of components it required and how many remain in stock.

The validation, testing, fine-tuning and evaluation of both models was performed in the three last months of the historical data that was not used for training.

Similarly to the PPC case, the experimentation compared results between the NOA deployment, and the deployment of the same solution on the DO (Digital Ocean) and Google Cloud platform.

The experiment KPIs' are the following:

- SVM-01: Service deployment. This KPI measures the time taken for service deployment at the ExFa site using the 5G-INDUCE platform. The deployment time is the same as in the PPC case, as the same components are deployed.

Table 36 SVM-01 UC4

KPIs-UC4	SVM-01
Average Service deployment time	220 seconds

- SVM-02: Functional tests. This KPI measures the ability to effectively transfer data to/from the nApp, focusing a) on the delivery of the complete set of data from the machine to the storage facility and b) on the correct export of the data from the AI engine to other sources.

Table 37 SVM-02 UC4

KPIs-UC4	SVM-02
Functional tests – a) Data Transferred to nApp	100% of generated data - PASS
Functional tests – b) Data Exported from nApp	100% of AI output data - PASS

- SVM-15: Data Transfer Time. This KPI compares data transfer time between the Shopfloor and the nApp, where in the first case it is part of the NAO and in the second it is hosted on the cloud.

Table 38 SVM-15 UC4

KPIs-UC4	SVM-15
Average Data Transfer Time (NAO)	24.47 msec
Average Data Transfer Time (Cloud - DO)	228.08 msec
Average Data Transfer Time (Cloud -Google)	207.22 msec

The next diagram shows the timings as recorded over the different deployments in the same period.

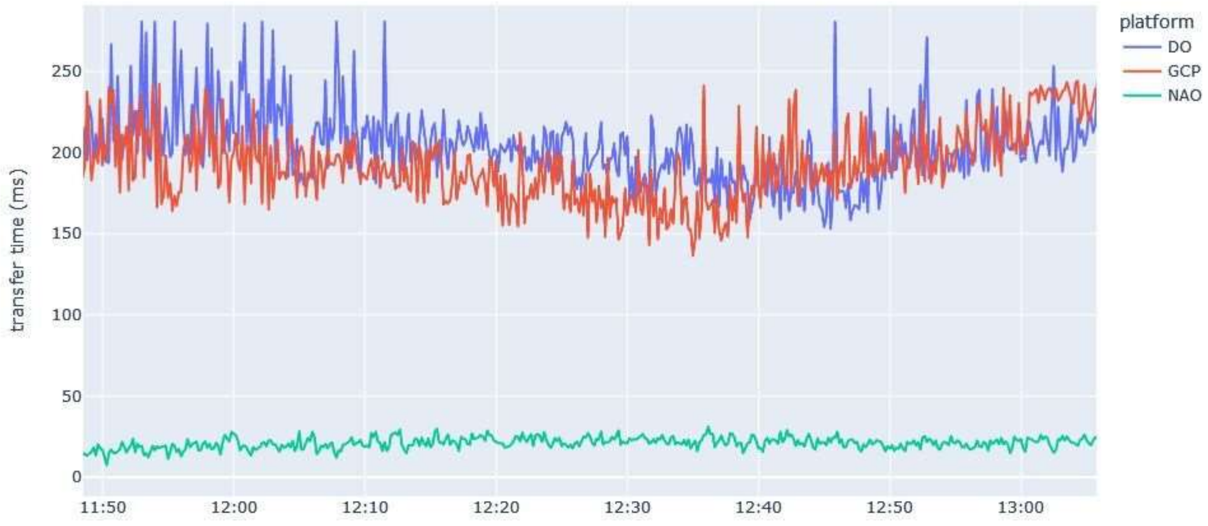

Figure 50 Transfer time of sensor data to/from different infrastructures tested.

- SVM-16: AI/ML Pipeline Execution Duration. This KPI compares the time to execute the AI/ML Pipeline designed between the NAO nApp and the cloud-based nApp.

Table 39 SVM-16 UC4

KPIs-UC4	SVM-16
Average AI/ML Pipeline Execution Duration (NAO)	246 msec
Average AI/ML Pipeline Execution Duration (Cloud - DO)	212 msec
Average AI/ML Pipeline Execution Duration (Cloud - Google)	198 msec

The next diagram shows the execution timings as recorded over the different deployments in the same period.

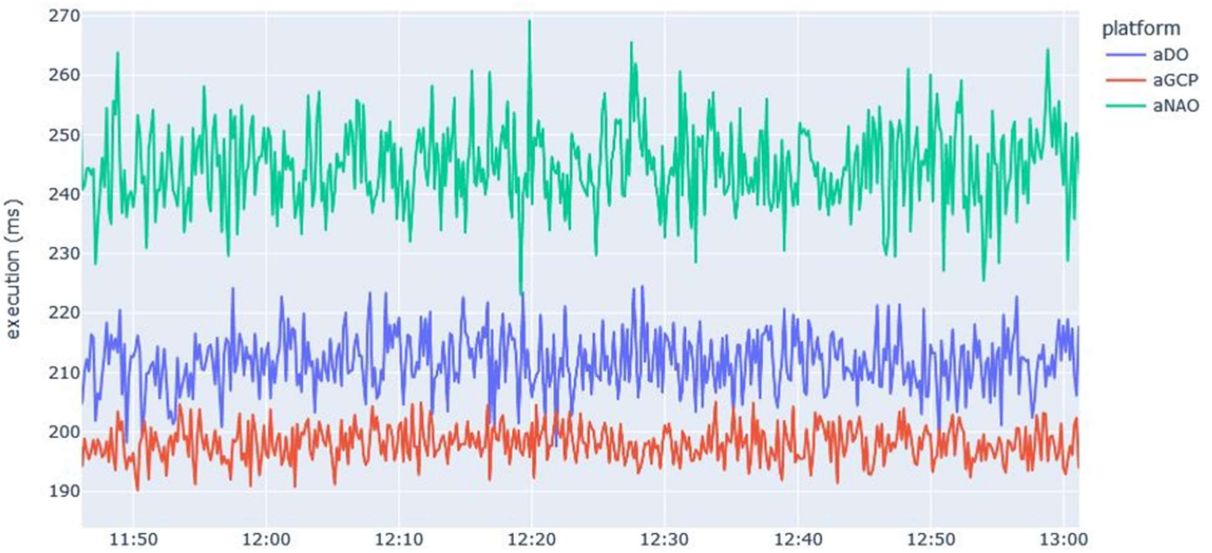

Figure 51 Execution Time recorded over the different infrastructures tested.

4.1.3 Analysis of the results

The results of the experimentation with UC4 under this ExFa proved the claim that the placement of AI solutions over 5G infrastructure could prove useful for certain applications, as also verified in the PPC ExFa as well.

As conventional AI/ML applications run on cloud infrastructure, in our analysis we compared these timings as recorded over the 5G-INDUCE infrastructure with those of having the same AI/ML scenarios operated over public cloud infrastructures (experimented both with Google Cloud Platform and Digital Ocean platform).

Data roundtrip over the NAO is faster than roundtrips recorded when the machinery (generator in this case) had to send the data to a public cloud platform and then receive the signals to initiate the maintenance automation procedures. In parallel, the data processing and AI/ML preprocessing timings were also recorded and compared between the 3 different cases. As the graphs suggest, the performance of the AI/ML engine in the case of the NAO was slightly slower than those recorded over the different cloud infrastructures, as the latter have the ability to scale more than the resources available in the NAO, while the initial resources used were slightly more than those used over the NAO.

By combining these numbers, as shown in the next figure, it is proved that for the specific case, the 5G deployment offers an advantage over the public cloud infrastructure, as the loss in execution time (for the scale of analytics developed for the specific solution) can be compensated by the data transfer times. This difference would be even bigger in case the analytics engine was used to ingest data for various machines at the same time over the same factory setting, assuming that in such a scenario the scaling configuration of the i Engine between the NAO and the cloud resources would be identical.

When considering that the nApp is capable of working on datastreams coming from multiple machines, the advantages of 5G are becoming evident as in such a case the bandwidth that needs to be available at the shop floor should be quite high, due to the huge amounts of data that would need to be transferred and analysed. Also the frequency of new data ingestion and prompt presentation of the analysis results is relevant in order to minimize the impact on defective production or on potential stoppages due to a late intervention due to the current quality control based only on statistical sampling.

Figure 52 Combined Timings for Data Transfer and Analytics Execution over the different infrastructures tested

Furthermore, as also identified in the PPC case, data privacy and sovereignty is a very crucial aspect that is also valid in the Whirlpool case, and as such the deployment of the nApp over a 5G infrastructure as the one promoted by 5G Induce is highly preferred to the deployment of the same application over a public cloud provider.

4.2 UC6

4.2.1 ExFa setup

The preparation of UC6 demo was conducted through the detailed identification of the Oculavis platform configuration tasks and of the data sets required to be loaded into the system to support user experience.

The machines identified for the experimentation were 2 twin cutting and bending machines, COSMA1 and COSMA2, in the Refrigeration factory stamping department.

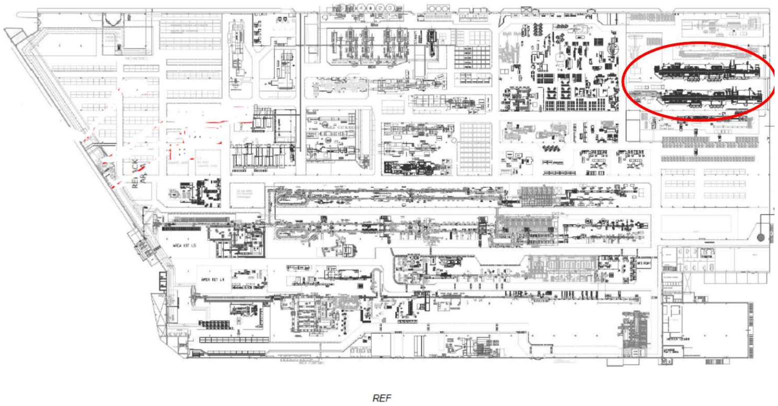

Figure 53 Whirlpool Factory Execution Setting of UC6.

The smart glasses have been configured to be connected to the mobile device, IOS tablet (iPad Pro 11) in this case. The test was executed only in option1 mode, that means connected to the Wind3 public network.

Figure 54 Equipment used in the DEMO.

Different business scenarios have been tested and validated by users:

- remote support from expert technician, connected via laptop, to maintenance operator near the machine to support complex diagnostic activity
- access to the Oculavis platform to visualise maintenance Procedures, maintenance plans, capture pictures and videos, input maintenance intervention data.

The users' feeling about response time was very positive with a system prompt to commands and fluid in video management.

The Network Application facilitates the transmission of high-quality video streams from shop floor technicians to remotely located maintenance experts. To enhance latency and bandwidth, the Media Server

VNF was migrated to the 5G infrastructure. The following Figure 55 illustrates the application architecture for UC6, deployed at the Italian Experimentation Facility (IT-ExFa).

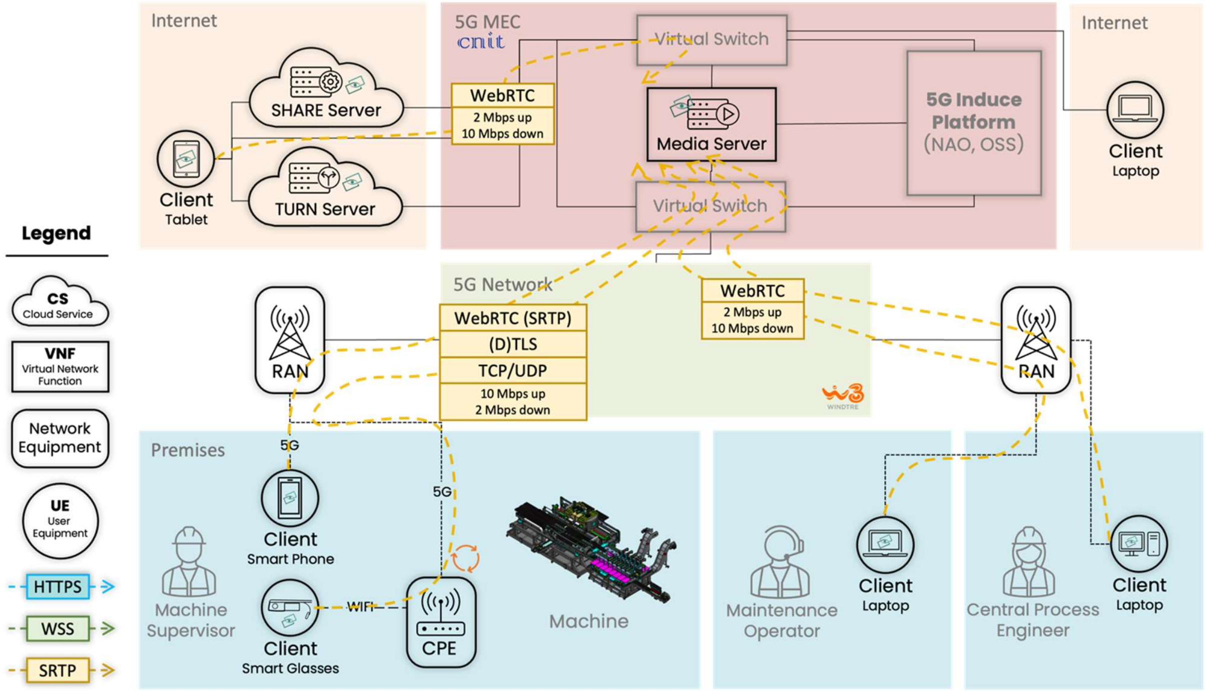

Figure 55. UC6 application graph.

The research project setup requires that the media server operates within an isolated network. Services running inside the isolated network are inaccessible from the Internet. However, the SHARE Call Service (SHACS) running on the Internet must still be able to access the media server VNF's signalling plane using the WebSocket Secure (WSS) protocol. Therefore, the media server VNF needed to be modified to allow a SSH reverse tunnel connection.

The goal of the reverse SSH tunnel is to establish an outgoing connection from the media server VNF (JANUS) running within the 5G network to the SHARE Call Service (SHACS). In this scenario, the 5G network firewall only requires configuration to allow outgoing connections from JANUS on port 2222 to SHACS's IP address.

The following Figure 56 illustrates the SSH reverse tunnel established by JANUS to SHACS.

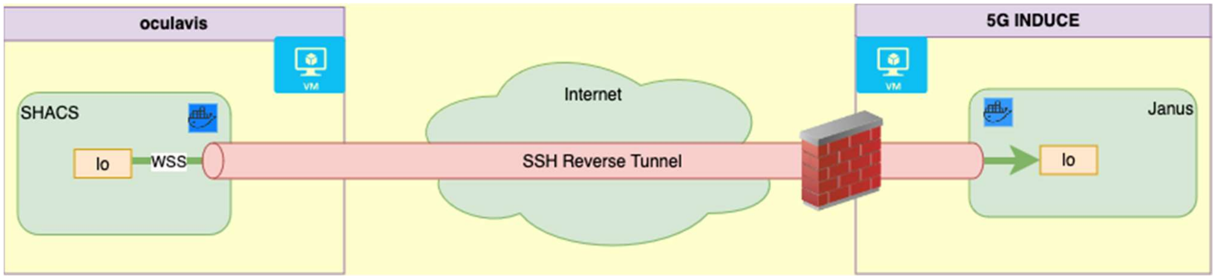

Figure 56 SSH Reverse Tunnel Janus to Call Service Connection.

To modify certain video quality parameters for the different test runs the machine operator has the possibility to change the video resolution, frame rate and sent bitrate via the iOS mobile application which is shown in the following Figure 57:

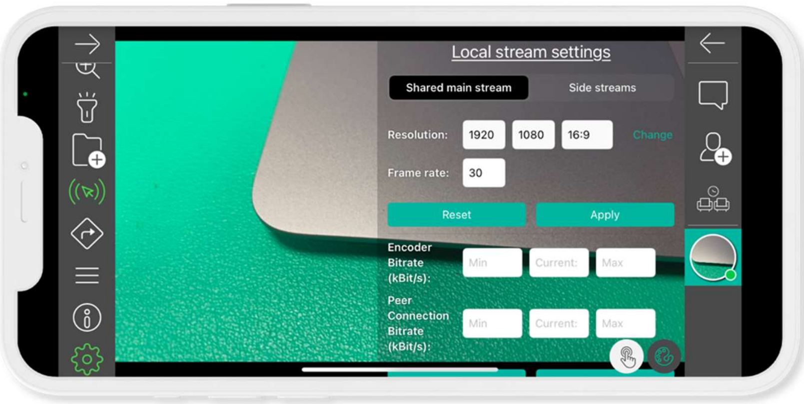

Resolution: 1920 1080 16:9
Frame rate: 30
Encoder Bitrate (kBit/s):
Peer Connection Bitrate (kBit/s):

Figure 57 iOS mobile application with advanced WebRTC video quality settings.

While the video quality statistics can be monitored and exported via a chrome browser (website: “chrome://webrtc-internals/”) the iOS application also supports a direct monitoring which is shown in Figure 58:

Own Stream Stats

Video Bitrate: 2097.0 Kbit/s
Available bitrate: 3573.0 Kbit/s
Limitation reason: bandwidth
Resolution: 960.0 x 540.0 px
Frame rate: 30.0 fps
Audio Bitrate: 13.0 Kbit/s

Device Stats

CPU Usage: 62%
Memory Usage: 107 MB

Figure 58 iOS mobile application with advanced WebRTC vide.

4.2.2 Service-level validation at ExFas and KPIs measurements

The service validation was conducted in two phases. The first phase, focusing on pre-validation of application settings, involved evaluating various resolutions, bandwidths, and frame rates to determine their optimal values. In the second phase, the deployment scenarios were validated through a realistic walking and working scenario at the Cosma Factory Area which are shown in Figure 61.

Pre-validation of applications settings

To evaluate the right application settings the resolutions were increased according to the in D5.1 explained testing procedure until a limit is reached. The tested resolutions can be found in the following Table 40.

Table 40 Resolutions and frame height and width in pixels

Resolution	Frame width	Frame height
HD	1280	720
FHD	1920	1080
QHD	2560	1440
UHD	3840	2160

Figure 59 illustrates the results of the initial test run, where the application operated in a standard non-5G environment with services deployed in MS Azure, Frankfurt, Germany (referenced in D5.1 as B1 deployment). These measurements serve as a baseline for comparison with those from the 5G-INDUCE deployment.

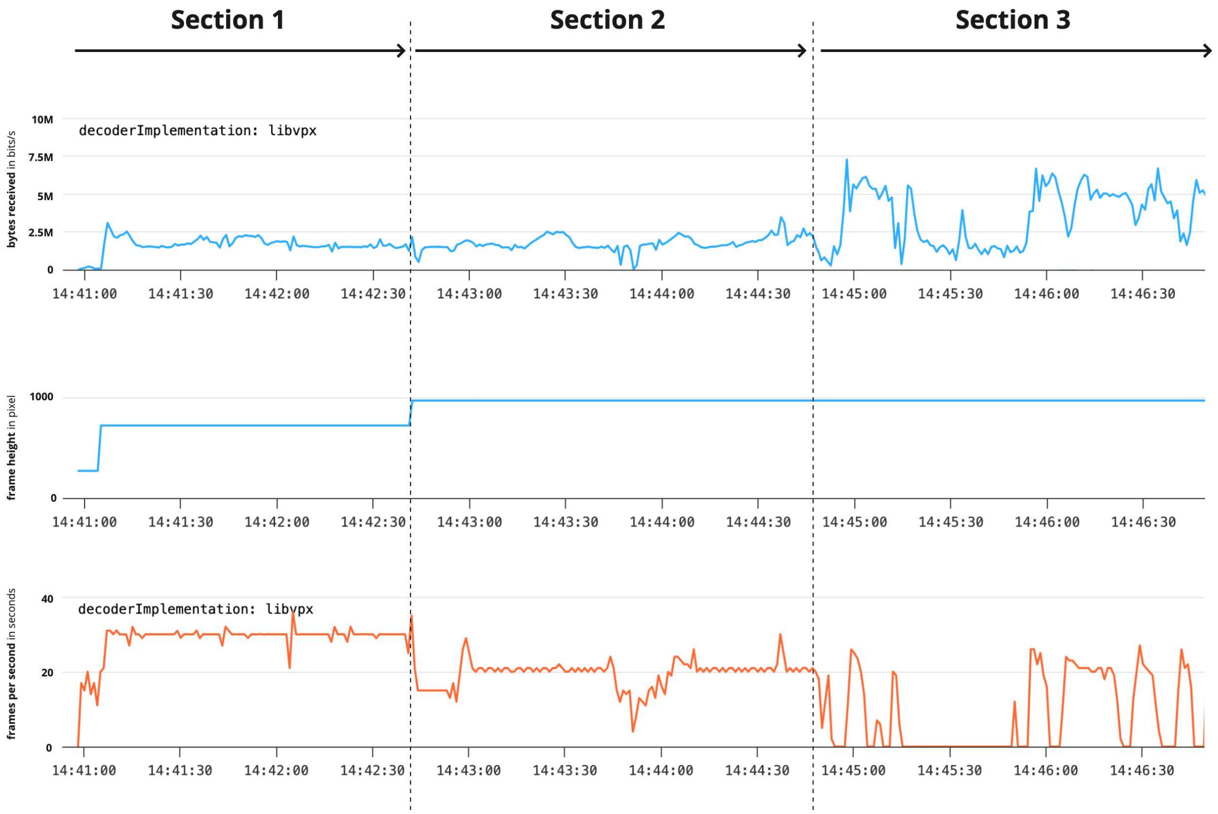

Figure 59 Measurements of bitrate, frame height and frames per second in standard MS Azure deployment (B1).

The measurements were divided into three sections, showing that the application maintained stable performance at an HD resolution with 30 frames per second, utilizing approximately 2 Mbit/s for video streaming. When the resolution was increased to FHD, stability persisted, but the frame rate automatically dropped to 20 frames per second. Attempts to further increase the resolution and bitrate settings failed due to insufficient available bandwidth. The encoder bandwidth was highly unstable and did not reach a consistent level, resulting in significant frame rate drops when trying to increase the resolution. Throughout the tests, the mobile operator observed an available bitrate on the device of approximately 4 Mbit/s by monitoring the statistics on the device.

The following Figure 60 presents the measurements obtained with the media server deployed in the 5G-INDUCE CNIT cluster (referenced in D5.1 as the B2 deployment):

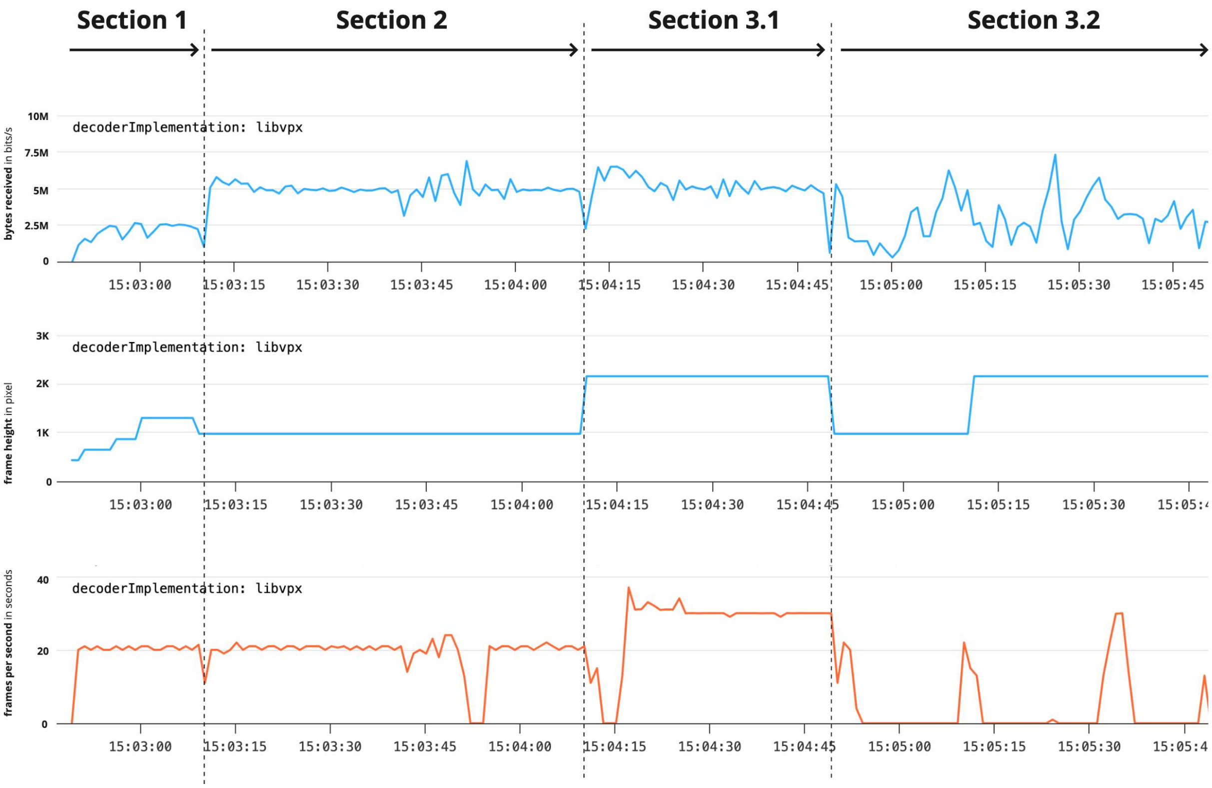

Figure 60 Measurements of bitrate, frame height and frames per second with 5G-INDUCE deployment (B2).

The measurements were again divided into three sections. With this deployment, we successfully increased the encoder bitrate to 5 Mbit/s, resulting in a more robust and stable performance at FHD resolution (see section 2) compared to the initial test without the 5G-INDUCE deployment. Additionally, this setup generally allowed us to reach a 4K (3840 x 2160 pixels) resolution while maintaining a stable bitrate of 5 Mbit/s (see section 3.1), although this stability lasted only for about 45 seconds. When the device was moved, the encoder bitrate became highly unstable, leading to the termination of the tests.

The measurements result in the following KPIs for the:

- SVM-07: High-resolution Real-time Video Quality

Table 41 SVM-06-UC6

KPIs-UC6	SVM-06	
2K resolution @30fps video with 10Mbit/s	Stable	FHD resolution @30fps video with 5Mbit/s
	Maximum	4K resolution @30fps video with 5Mbit/s

- SVM-05: Streaming bandwidth

During the measurements the mobile operator observed the in-app monitoring and detected a maximum available bandwidth of 10Mbit/s which result in the following KPI result:

Table 42 SVM-05-UC6

KPIs-UC6	**SVM-05**
20 Mbit/s	10Mbit/s

Validation of Deployment Scenarios over Cosma Factory Area

To assess perceived operational latency, five end users were asked to rate their experience on a scale from 1 (very poor) to 5 (excellent). This evaluation was conducted using an iOS device on the shop floor of the Bekos Cosma cutting line. During the test, users followed a predefined walking path, as illustrated in Figure 61.

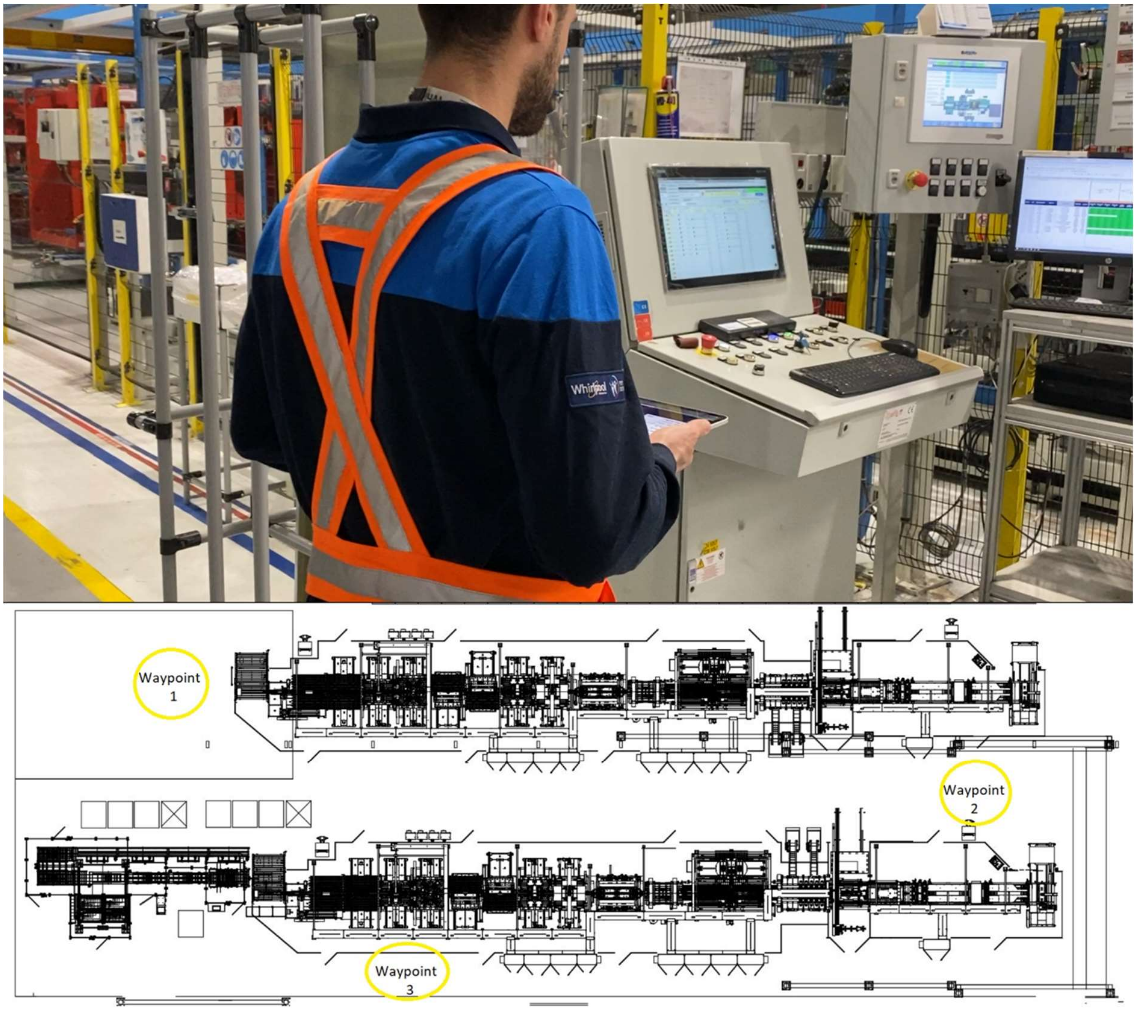

Figure 61 Mobile operator waypoints for validation over factory area.

The scenario was evaluated using both the standard MS Azure deployment (B1) and the 5G-INDUCE deployment (B2), in accordance with the test plan outlined in Deliverable D5.1.

The results of the five end users evaluating the perceived operation latency for both deployment options are given in the following Table 43.

Table 43. End user evaluated score for perceived operation latency from 1 (very poor) to 5 (very good)

Users	Score (B1)	Score (B2)
User 1	4	4
User 2	2	5
User 3	4	4
User 4	3	5
User 5	5	5
Average	**3.6**	**4.6**

- SVM-07: Perceived operation latency

The final KPI values for perceived operational latency with the 5G-INDUCE UC setup (B2) are presented below:

Table 44 SVM-07-UC6

KPIs-UC6	SVM-07
> 4	4.6

The measured values for the bitrate and frames per second of this test scenario are shown in the following Figure 62:

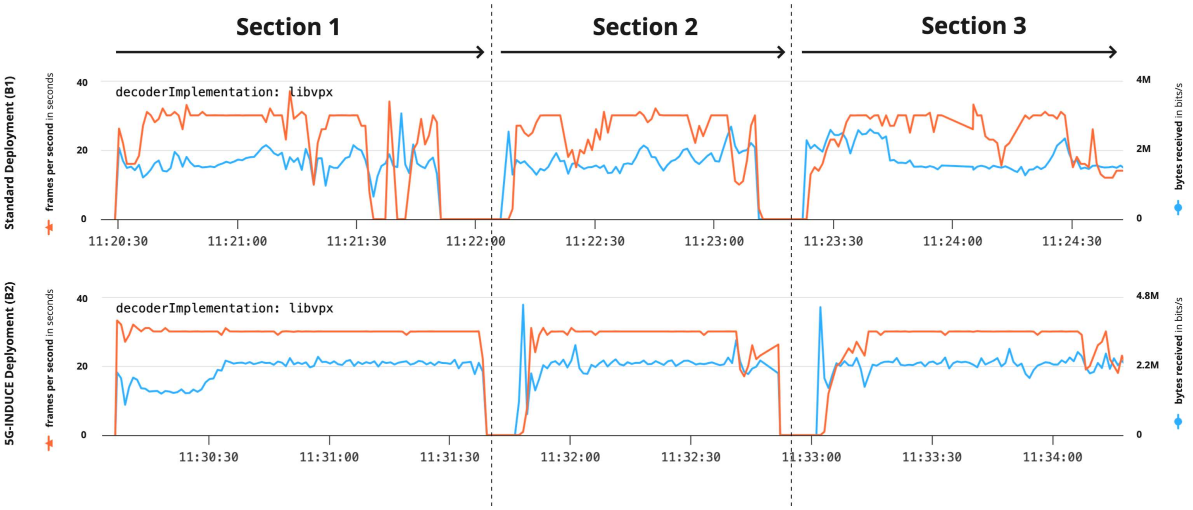

Figure 62 Measurements of bitrate and frames per second comparing deployment scenarios (B1, B2) over factory area location.

Since the mobile operator's walking speed varied between the two test scenarios, the camera was deactivated at each waypoint to ensure accurate and comparable measurements. Generally, the 5G-INDUCE deployment resulted in higher and more stable transmitted bandwidths, which contributed to a more consistent frame rate of 30 frames per second. An unstable frame rate and occasional frame loss in Section 1, from Waypoint 1 to Waypoint 2, led to lower perceived operational latency scores from some users for the non 5G deployment.

4.2.3 Analysis of the results

The testing at the Cosma Factory Area under the 5G-INDUCE platform revealed key benefits of 5G, particularly in enhancing video streaming stability and quality. Under a non-5G environment, the application could only maintain HD resolution at 30 frames per second with 2 Mbit/s, and increasing to FHD reduced the frame rate to 20 fps. However, with 5G, the application sustained a stable FHD resolution at a higher bitrate of 5 Mbit/s. Under certain test scenarios a supported 4K resolution was achieved, showcasing 5G's capabilities, although it was not possible to conduct the full test with these streaming capabilities due to instable network conditions and bad handling of those within the application.

The 5G network also significantly improved perceived operational latency and user experience in terms of consistent frame rates. In contrast to the non-5G setup, where users reported poor latency due to unstable frame rates, the 5G deployment consistently delivered a robust 30 fps, enhancing real-time performance. Even in dynamic environments during realistic scenarios, 5G provided more stable bandwidth, ensuring reliable application performance.

The results clearly demonstrate 5G’s benefits for video streaming applications, particularly in improving quality, stability, and latency. These enhancements make 5G a vital component for high-performance, low-latency applications, especially in demanding environments. The insights gained from the enhanced stability, scalability, and latency improvements observed with the 5G-INDUCE deployment provide a solid foundation for further research and development.

4.3 UC7

4.3.1 ExFa setup

UC7 demo has been held in a specific warehouse area in Refrigeration factory where the movement of forklifts and tuggers was particularly intensive with high risk of collision among mechanical vehicles and with moving people.

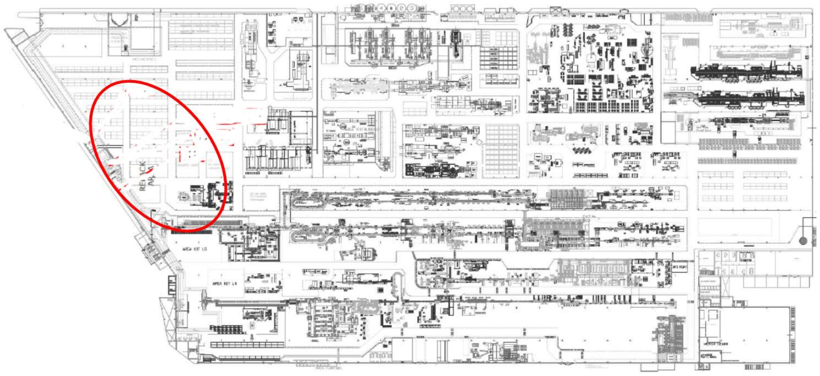

REF

Figure 63 Whirlpool Factory Execution Setting of UC7.

The test has been prepared since the end of 2022 with the installation of 6 anchors, delimiting testing area, 1 collector and 1 5G gateway, equipped with SIM.

The original infrastructure included an area with materials stocked on the floor and corridors like in the picture below. The 6 anchors positions were identified through a testing session to ensure the full coverage of the area (Figure 64).

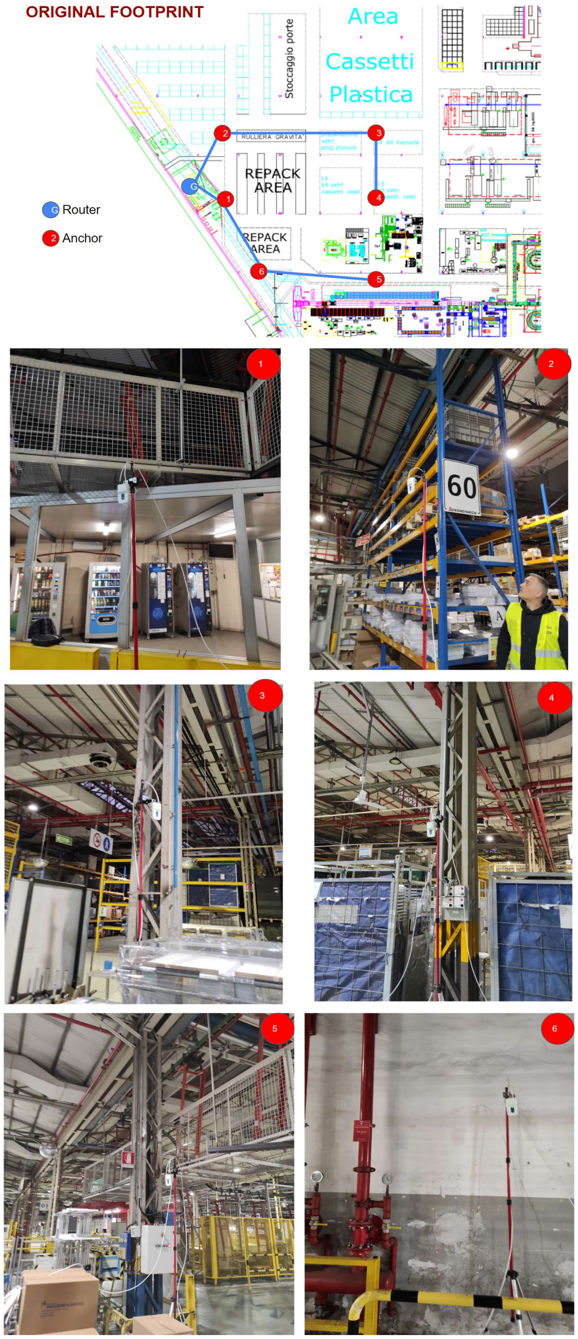

Figure 64 Anchors positions.

The whole area was strongly redesigned during 2023 with the installation of high shelves and metallic grids and the infrastructure had to be reviewed and adjusted to ensure the best performances.

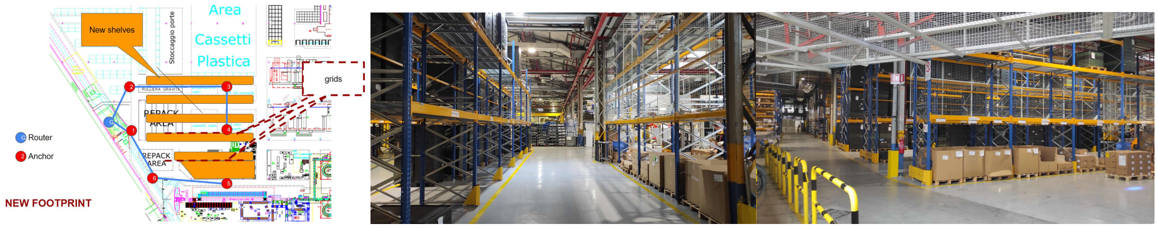

Figure 65 ExFa environment.

The demo session were held with both options, using both Wind3 public network 5G connectivity and CNIT mobile testbed 5G connectivity.

The devices used were the 6 anchors wired to the gateway, an android tablet (lenovo p12 tab pro) equipped with 5G SIM and UC7 nApp, used by forklift driver to visualise the traffic light, some wearable tags to track persons position to be used by forklift drivers and by people moving in the controlled area and lastly 5G smartphone, equipped with SIM and UC7 nApp too, to be used by human workers moving in the controlled area to manage the alerting for collision risk and visualize all moving assets in the area.

Figure 66 devices used in UC7.

Option 1: the ILINK gateway was equipped with wind3 SIM with public APN. In the original position, the Wind3 coverage was not stable and strong enough to ensure smooth execution of the demo. So, after a measurement session held by Wind3, the 5G devices of UC7 and UC8 have been moved out of the production rack, nearer to a window in order to leverage on outdoor coverage. The new positioning resolved the strength issues but the instability of the performances were still present, even if with no impact on the business validation of the use case

Figure 67 Original position of the gateway.

4.3.2 Service-level validation at ExFas and KPIs measurements

The validation of the UC7 nApp operation, at the laboratory included WiFi and 5G operational tests where critical delays have been monitored. For the needs of the Project, several tests have been conducted in iLINK's and testbed's premises. Wi-Fi and 5G (telecom provider) coverage have been provided at the laboratory and used as the network communication layer. Suitable network cabling ran across the lab control area, guaranteeing proper connectivity among the UWB Anchors, to the main switch, up to the 5G router. The 5G router has been initially connected to the local Wi-Fi, acting as a LAN router and then tested over 5G coverage as provided by the telecom provider.

The detailed steps taken towards the end-to-end deployment are presented as follows (based on the KPIs):

- Step1: Set up and run the nApp (a) in the laboratory infrastructure and (b) on the Cloud
- Step 2: Attach the UWB anchors to their lab position (open space) and provide 2 sets of UWB tags. 1 set of UWB tags for actual moving entities and a set of simulator units. Associate the mobile app on the cell phone with the respective UWB tags (through the admin UI)
- Step 3: Execute the location positioning routine using a nApp deployed at localhost (Wi-Fi connectivity). Repeat the location positioning routine using a nApp deployed at a remote Cloud Provider (Azure) - over 5G connectivity
- Step 4: Visualise the position of each entity on her cell phone application based on the data received at the Location VNF. Visualise the position of all entities on the administrators map (Map VNF).
- Step 5: Run collision detection scenarios, originating from entities (simulator or actual) moving at several directions and at varying velocities. Collect and analyse positioning/movement data at the Collision VNF before being inserted at the DB-Storage VNF. Identify the audio-visual triggers on the cell phone application as well as the proper traffic light indication (messages sent by the Message-Bus VNF at the Core). Visualize imminent crashes at the respective points of the lab control area on the administrators map (Map VNF). Collect and analyse historical collision data as produced by the Collision VNF.
- Step 6: Repeat Steps 3-5 for different infrastructure configurations at the lab, by prohibiting pure LOS between anchors and by gradually removing anchors from the configuration.

Targeted KPI	Measurement (iLINK cloud) ms	Measurement Option 1) ms	Measurement Option 2) ms (CNIT-VPN) ms
100% entities (min 3) visible on cell phone applications and administrator's map under 400 ms	96.03 and 41.54	177.84 and 54.26	
Time between UWB tag positioning data is received at the Location VNF till the response is depicted on the administrator map (Map VNF) <300ms.	6.26	9.34	
Time between location data is delivered to the Collision VNF and the collision algorithm estimates an imminent crash with a respective probability <150ms.	1.21	1.33	
Time between a collision is detected at the Collision VNF and the alarm is triggered at the mobile application, through the Message-Bus VNF <300ms.	62.42	136.16	
Depict colored areas on the administrator map based on the number/priority of imminent crashes detected (3 zones).	True	True	
Repeat all SVM02 calculations based on the different configuration of the height of the corridors (limited LOS) and the number of UWB anchors deployed at the configuration.			
Generic KPI, estimates the total delay between a location package (originating from a UWB tag) is delivered to the network, till the imminent crash is audio visually notified to the cell phone of the owner of the tag (from the NAO container) <500ms.	99.02	183.65	

Table 45 UC6 testing results

4.3.3 Analysis of the results

The UC7 of the 5G-INDUCE Project successfully helped in the validation of the significant benefits of 5G technology in the domain of Industry 4.0. More precisely, our use case aimed at the design and development of an advanced positioning and industrial safety system. The use case has been tested in multiple stages within iLINK's premises and the final tests have been conducted in the Whirlpool, IT premises.

The proposed safety monitoring system required low latency due to the industry's time-critical aspects. In this context, the moving entities, including forklifts and persons, had to provide their real-time location with the minimum delay possible. This has been absolutely crucial for the collision avoidance algorithm built on top of this component. Since the specific solution can raise immediate alerts directly to the entities facing the danger of an imminent crash, it is essential to have a seamless and fast data transmission. Additional qualitative visualisations such as heatmaps with the incident frequency, provided useful outputs for advanced operations management and decision-making within the facilities.

Due to some implications with the 5G SA network, there has been an additional round of tests with VPN connection to the Italian testbed utilising the NAO, defined as Option 1. Despite the increased latency due to VPN, the nApp managed to remain fully-functionable within the time limits set by the KPIs. This also highlights the importance of the 5G connectivity compared to 4G. Finally, Option 2 referred to the establishment of a base station within the Italian ExFas, provided by CNIT. This option did not require a VPN bridge to access NAO but presented several issues during the deployment.

4.4 UC8

4.4.1 ExFa setup

Validation of the use case requires deploying UC8 application components to support measuring IP performance and radio KPIs collection on two types of devices:

- ININ's 5G industrial gateway for continuous measurements and
- Samsung S22 5G with ININ's measurement application for on-demand, mobile/walk/drive and drone-based measurements.

The overall UC8 architecture is presented in Figure 68.

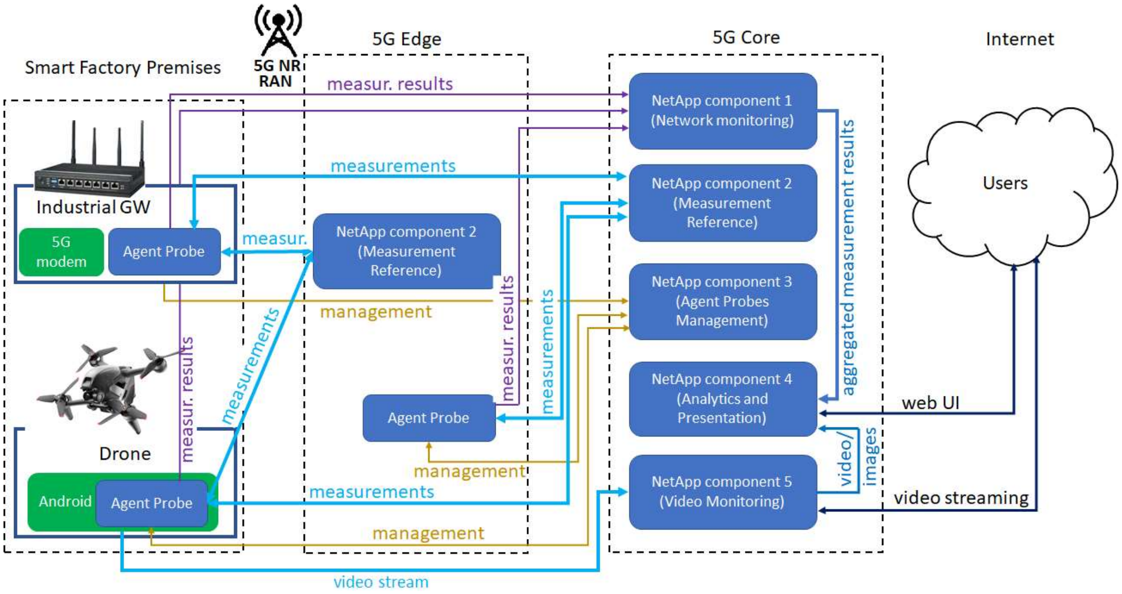

Figure 68 UC8 - Overall Use Case Architecture.

The validation itself was performed in three phases:

- Initial testing (2023) – using public 5G-NSA provide by WINDTREE with application components being deployed via VAO @CNIT and exposed on the public IP addresses.
- Final testing – Option 1 (2024) - using public 5G-NSA provided by WINDTREE with application components being deployed via VAO @CNIT and exposed at private IP addresses and accessible via VPN on 5G UE devices.
- Final testing – Option 2 (2024) – using private 5G SA provided by CNIT's mobile testbed where also application components were deployed and reachable from 5G UE directly on the private IP addresses.

Additionally, all phases also included a reference server deployment used in ININ's testbed to measure services over the public internet. The overview of the final testing architecture for both options is presented in Figure 69 and Figure 70.

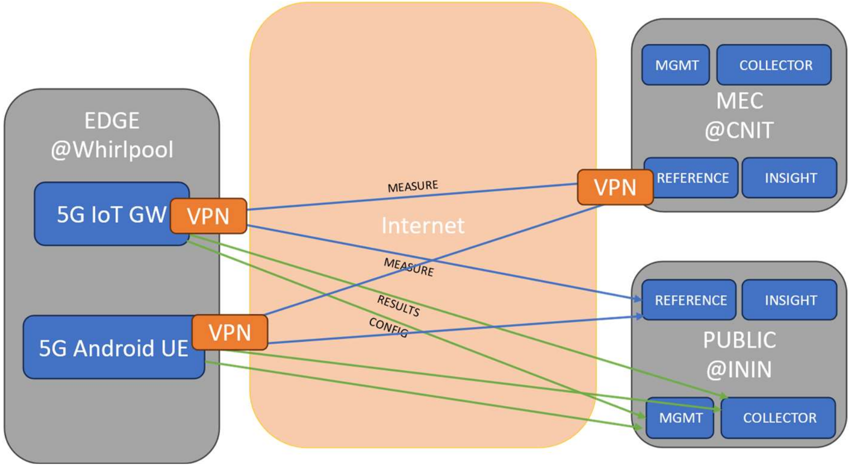

Figure 69: UC8 Final Testing – Option 1.

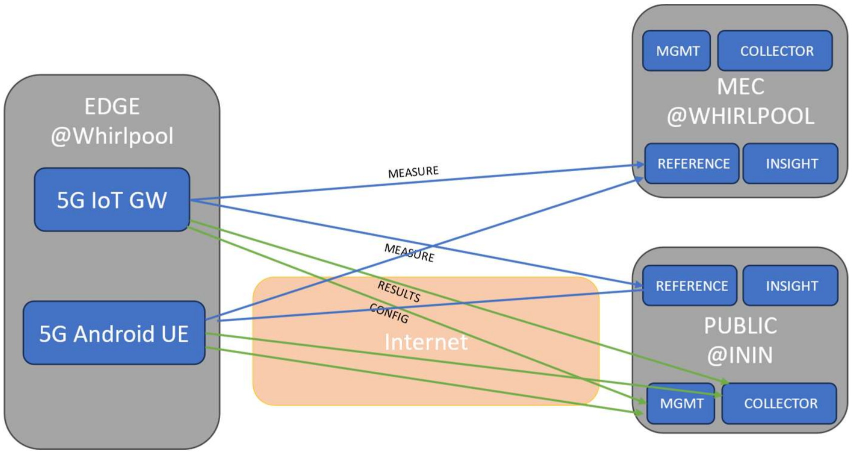

Figure 70 UC8 Final Testing – Option 2.

Due to regulatory constraints and the inability to secure a flight permit for drone operations within the Whirlpool factory premises, it was necessary to identify an alternative location for the drone performance monitoring evaluation. After careful consideration, the Luka Koper/Port of Koper in Slovenia was selected as the main reference industrial infrastructure for the following reasons:

1. Regulatory Compliance: The Port of Koper has established protocols and permissions for drone operations, ensuring compliance with aviation regulations and minimizing legal risks.

2. Infrastructure Similarity: The port's industrial infrastructure offers comparable characteristics (metal obstacles etc.) to the Whirlpool factory outdoor environment, such as operational complexity, layout, and scale, providing a relevant environment for testing.

3. Safety Considerations: The open and controlled environment of the port reduces the risk of interference with sensitive equipment (e.g., GNSS and gyroscope interference) and allows for safer flight paths.

4. 5G Validity: The Port of Koper's active industrial operations and dedicated 5G NSA network provide a dynamic setting that is conducive to robust data collection, closely emulating the conditions expected at the Whirlpool outdoor premises and other industrial environments.

5. Logistical Feasibility: Proximity to the project team and ease of access to the port facilities ensured efficient deployment and operation of the drone without significant delays.

In conclusion, the decision to conduct the drone performance monitoring evaluation at the Port was a strategic measure to ensure the integrity of the testing process while adhering to regulatory requirements and safety standards. This approach has allowed for the collection of valuable 5G NSA performance metrix in complex industrial environment that is applicable to the Whirlpool factory outdoor context.

Figure 71 UC8 – 5G UE with qMON Agent Preparation (figure left and middle) for Drone Flight (right figure) in Luka Koper[1].

Radio Nr Rsrp Dbm: -101,00 – -64,00
Radio Operator Code: (All)
Radio Nr Band: (All)
Radio Access Technolo..: 5G-NSA
Alias Hash: 1f249
Timestamp: 20/09/2023 0 20/09/2023 2
Radio Nr Rsrp Dbm: -103,00 – -64,00

Figure 72: UC8 – 5G Drone Flight (left figure) with KPI analytics (right figure) operating in Luka Koper/Port of Koper.

---

[1] Special thanks to Luka Koper/Port of Koper and the Hellenic Drones team for assuring the drone flight.

To verify the prepared 5G User Equipment (UE) with the integrated qMON agent, we conducted comprehensive testing using two distinct types of drones. The first type was a custom-built drone, provided by Hellenic Drones (Figure 71 and Figure 72), which was specifically designed to meet the extra load requirements. The second type was a commercial off-the-shelf drone from DJI (Figure 73), for which we developed and attached necessary extensions to enable it to carry an additional payload – 5G UE.

The custom-built drone from Hellenic Drones featured advanced capabilities, including enhanced flight stability, extended battery life. The commercial drone from DJI was selected for its widespread availability and proven reliability. ININ team engineered custom extensions that allowed the drone to carry the additional weight of the 5G UE without compromising its flight performance. This adaptation was crucial in demonstrating the versatility and adaptability of the qMON agent across different drone platforms. Prepared commercial drones underwent a series of tests that included:

- **Flight Endurance:** Measuring the maximum flight duration while carrying the 5G UE.
- **Signal Integrity:** Evaluating the strength and consistency of the 5G signal in different environments.
- **Data Transmission:** Assessing the speed and reliability of data transmission between the drone and the ground control station.

The prepared setups and the results of these tests are presented in the figures that follows (Figure 73 to Figure 76), providing a clear visual representation of the drones' capabilities and the successful integration of the 5G UE with the qMON agent.

Figure 73: UC8 – Commercial Drone with Mounted 5G UE and Integrated qMON Agent.

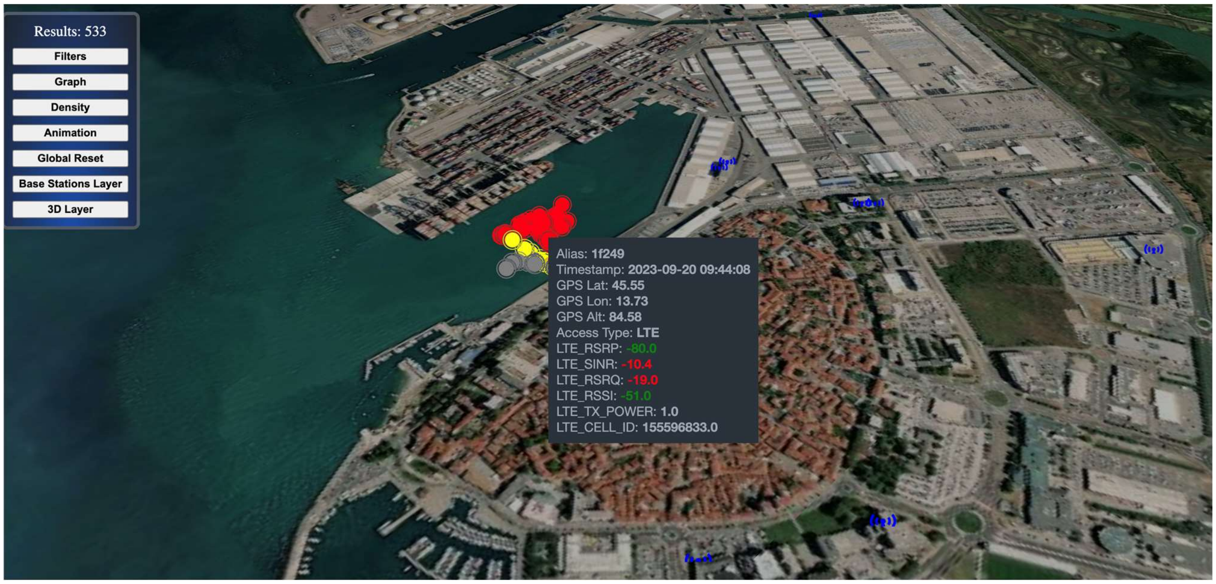

Alias: **1f249**
Timestamp: **2023-09-20 09:44:08**
GPS Lat: **45.55**
GPS Lon: **13.73**
GPS Alt: **84.58**
Access Type: **LTE**
LTE_RSRP: **-80.0**
LTE_SINR: **-10.4**
LTE_RSRQ: **-19.0**
LTE_RSSI: **-51.0**
LTE_TX_POWER: **1.0**
LTE_CELL_ID: **155596833.0**

Figure 74: UC8 – Drone-based 5G network performance monitoring in Luka Koper/Port of Koper.

Alias: **928dd**
Radio operator code: **29340**
Timestamp: **2024-06-01 16:27:56**
GPS Lat: **45.89**
GPS Lon: **13.84**
GPS Alt: **190.82999999999998**
Access Type: **5G-NSA**
LTE_BAND: **B3**
LTE_BAND_BW: **20**
LTE_PCI: **488**
LTE_CA_STATE: **ADDED**
LTE_CA_NUM: **1**
LTE_RSRP: **-55.0**
LTE_SINR: **16.0**
LTE_RSRQ: **-9.0**
LTE_RSSI: **-51**
LTE_TX_POWER: **-28.0**
LTE_CELL_ID: **405511**
NR_BAND: **n78**
NR_BAND_BW: **100**
NR_RSRP: **-102.0**
NR_SINR: **13.5**
NR_RSRQ: **-13.0**
NR_RSSI: **nan**
NR_TX_POWER: **22.0**
NR_PCI: **534**

Figure 75 UC8 – Drone-based 5G NR Propagation for Operational Base Station Verification.

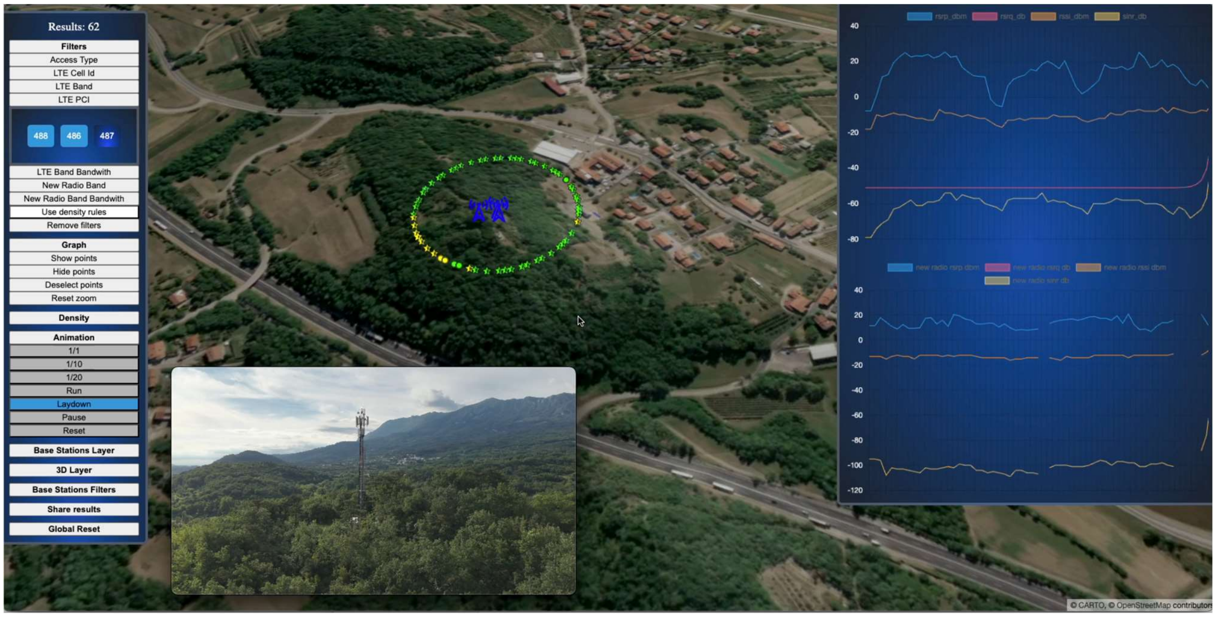

Figure 76 UC8 – Drone-based 5G Base Station Detailed Monitoring with Video Stream.

In the following link the video presentation of Drone-based performance monitoring can be viewed:

Drone-based network performance monitoring.

4.4.2 Service-level validation at ExFas and KPIs measurements

As presented in the previous chapter, two deployment scenarios (i.e. Option 1 and Option 2) were used validate the functionality of deployed application in both mode of measurement operations: continuous (gateway) and on-demand (android phone).

The collected KPIs are the following:

- SVM-01-Service deployment. This is the overall time from the instantiation of the application through the VAO until application components are deployed and services are available.

Table 46: SVM-01-UC8-Option1

KPIs-UC8 – Option 1	SVM-01
Total time (no slice creation)	3 min 15 sec

Table 47: SVM-01-UC8-Option2

KPIs-UC8 – Option 2	SVM-01
Total time	3 min 15 sec
Slice creation	1 min 33 sec
Application components deployment time	1 min 40 sec

- SVM-02-nApp specific functional test. Test is performed manually by checking if the operating services are behaving as expected, including UE connectivity, UE measurement agent registration to the management server, data collection validation and finally if the analytics dashboards are properly populated.

Table 48 SVM-02-UC8-Option1

KPIs-UC8 – Option 1	SVM-02
UE Connectivity	OK
UE measurement agent registration	OK
Data collection validation	OK
Analytics dashboard population	OK

Table 49 SVM-02-UC8-Option2

KPIs-UC8 – Option 2	SVM-02
UE Connectivity **	OK
UE measurement agent registration	OK
Data collection validation	OK
Analytics dashboard population	OK

** ININ's Samsung Android S22 phone does not support 5G SA mode.

- SVM-07-Perceived application operation related latency. This is end-to-end round-trip time from 5G UE to the Edge where the Reference server component VNF is installed.

Table 50 SVM-07-UC8-Option1

KPIs-UC8 – Option 1	SVM-07
Latency (mean) – on-demand	58 ms
Latency (mean) – continuous	88 ms

Figure 77: UC8 – SVM-07-Option 1 – on-demand.

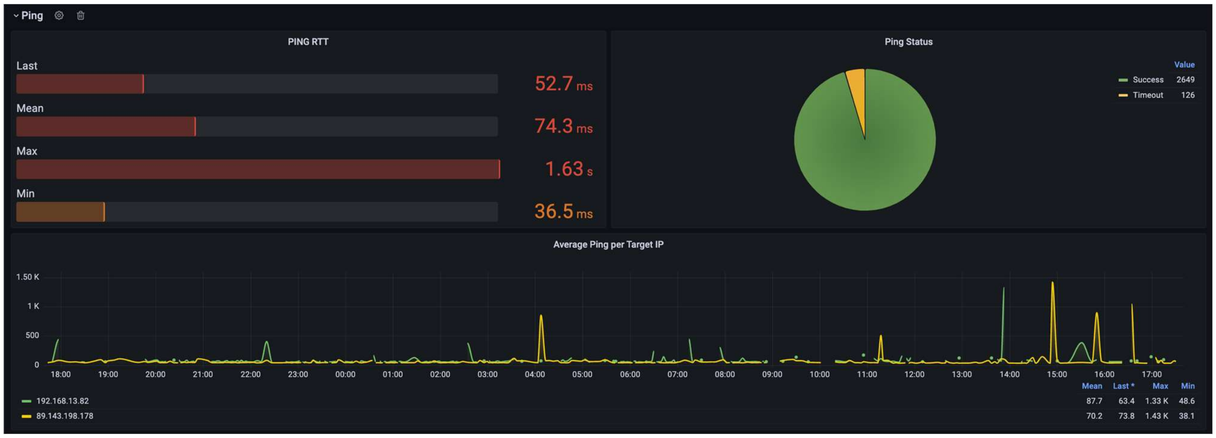

Figure 78UC8 – SVM-07-Option 1 – continuous.

Table 51 SVM-07-UC8-Option2

KPIs-UC8 – Option 2	SVM-07
Latency (mean) – on-demand	N/A
Latency (mean) – continuous	18 ms

Figure 79UC8 – SVM-07-Option 2 – continuous

.SVM-05 – Streaming bandwidth (download). It represents end-to-end IP download speed from the Edge where the Reference server component VNF is installed to the 5G UE.

Table 52 SVM05-DL-UC8-Option1

KPIs-UC8 – Option 1	SVM-05
Streaming bandwidth (download, mean)- on-demand	38 Mbps
Streaming bandwidth (download, mean)- continuous	23 Mbps

Table 53 SVM05-DL-UC8-Option2

KPIs-UC8 – Option 2	SVM-05
Streaming bandwidth (download, mean)- on-demand	N/A
Streaming bandwidth (download, mean)- continuous	64 Mbps

- SVM-05 – Streaming bandwidth (upload). It represents end-to-end IP upload speed from the Edge where the Reference server component VNF is installed to the 5G UE.

Table 54 SVM05-UL-UC8-Option1

KPIs-UC8 – Option 1	SVM-05
Streaming bandwidth (upload, mean)- on-demand	9.5 Mbps
Streaming bandwidth (upload, mean)- continuous	8 Mbps

Table 55 SVM05-UL-UC8-Option2

KPIs-UC8 – Option 2	SVM-05
Streaming bandwidth (upload, mean)- on-demand	NA
Streaming bandwidth (upload, mean)- continuous	36 Mbps

Figure 80 UC8 – SVM-05-Option 1 – on-demand.

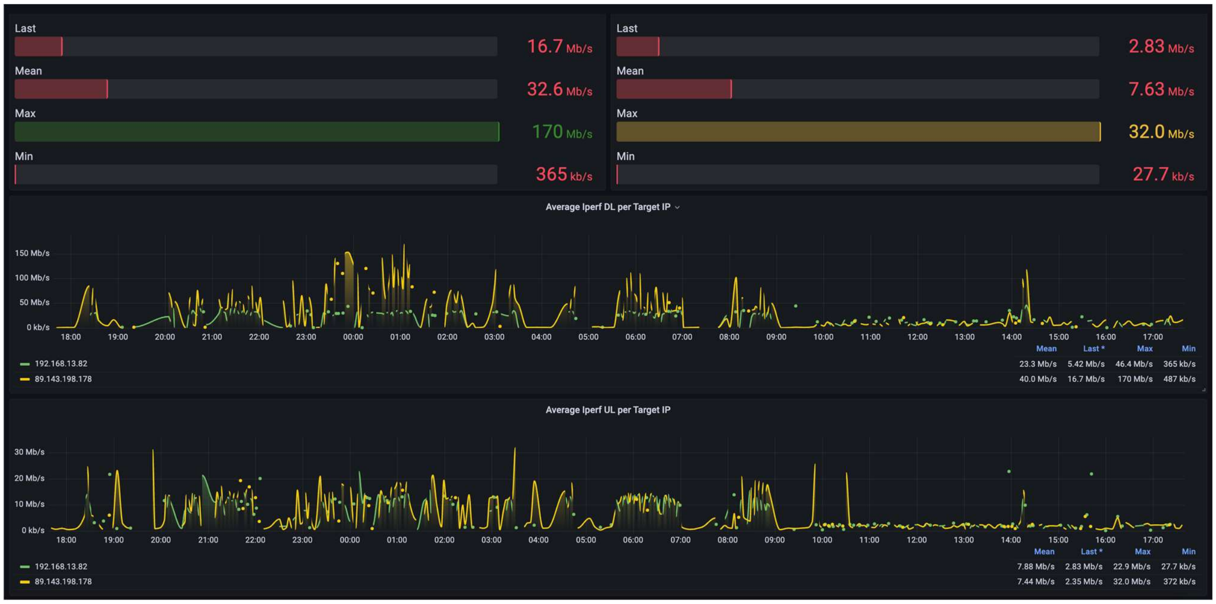

Figure 81 UC8 – SVM-05-Option 1 – continuous.

Figure 82 UC8 – SVM-05-Option 2 – continuous.

4.4.3 Analysis of the results

Collected radio KPIs (i.e. signal levels of RSRP, RSRQ, RSSI, SINR, NR RSRP, NR RSRQ, NR RSSI, SINR) clearly show how much the physical location of the 5G UE inside the industrial environment can affect the performance.

The first figure is showing collected radio KPIs on the Android phone (on-demand mode) at the UC7 location. One can spot that 5G-NSA was available at the beginning but later the phone modem consolidated on using LTE-only which can happen for various reasons (e.g. NR signal levels are too low, the battery consumption is higher when using NR) that affect decision logic on the device itself.

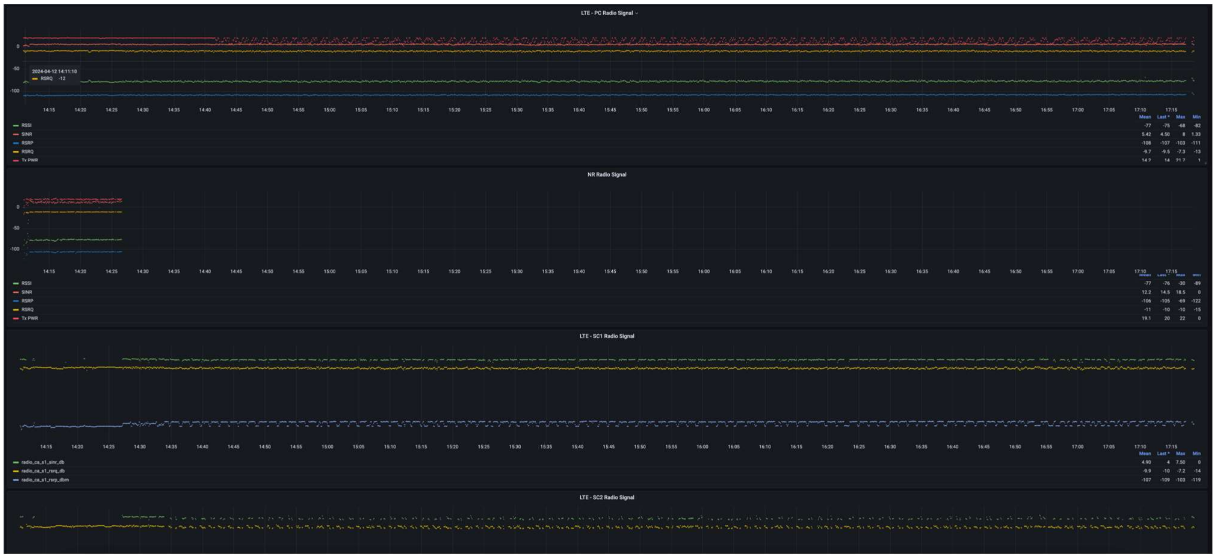

Figure 83 UC8 – Radio KPI analysis – on-demand.

On the other hand, the gateway (continuous mode) was placed in the office room at the shop floor where although the 5G-NSA was available all the time the performance and stability of the mobile network was significantly worse comparing to the UC7 location.

Figure 84 UC8 – Radio KPI analysis – on-demand.

Further analysis showed that RSRP/NR RSRP levels were at the “cell edge” threshold (i.e. < 105 dBm) which can trigger the UE device to do a cell handover. This is nicely shown in the following figure which shows IP throughput performance per cell station ID (being marked by a different color).

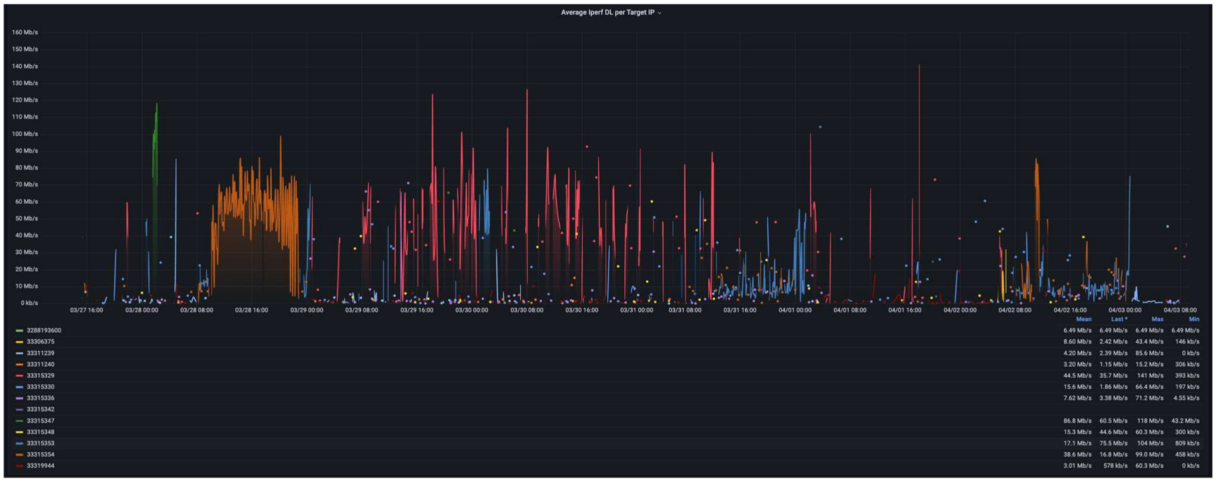

	Mean	Last *	Max	Min
3288193600	6.49 Mb/s	6.49 Mb/s	6.49 Mb/s	6.49 Mb/s
33306375	8.60 Mb/s	2.42 Mb/s	43.4 Mb/s	146 kb/s
33311239	4.20 Mb/s	2.39 Mb/s	85.6 Mb/s	0 kb/s
33311240	3.20 Mb/s	1.15 Mb/s	15.2 Mb/s	306 kb/s
33315329	44.5 Mb/s	35.7 Mb/s	141 Mb/s	393 kb/s
33315330	15.6 Mb/s	1.86 Mb/s	66.4 Mb/s	197 kb/s
33315336	7.62 Mb/s	3.38 Mb/s	71.2 Mb/s	4.55 kb/s
33315342				
33315347	86.8 Mb/s	60.5 Mb/s	118 Mb/s	43.2 Mb/s
33315348	15.3 Mb/s	44.6 Mb/s	60.3 Mb/s	300 kb/s
33315353	17.1 Mb/s	75.5 Mb/s	104 Mb/s	809 kb/s
33315354	38.6 Mb/s	16.8 Mb/s	99.0 Mb/s	458 kb/s
33319944	3.01 Mb/s	578 kb/s	60.3 Mb/s	0 kb/s

Figure 85 UC8 – Radio KPIs – cell handovers at static location in continuous mode.

This also corresponds to the actual measured IP performance. On the above figure, there is an orange segment which is showing no cell handovers for that specific time frame. If the timelines of IP throughput measurements are matched to the above timeline one can spot more stable and better performance on IP throughput download and upload.

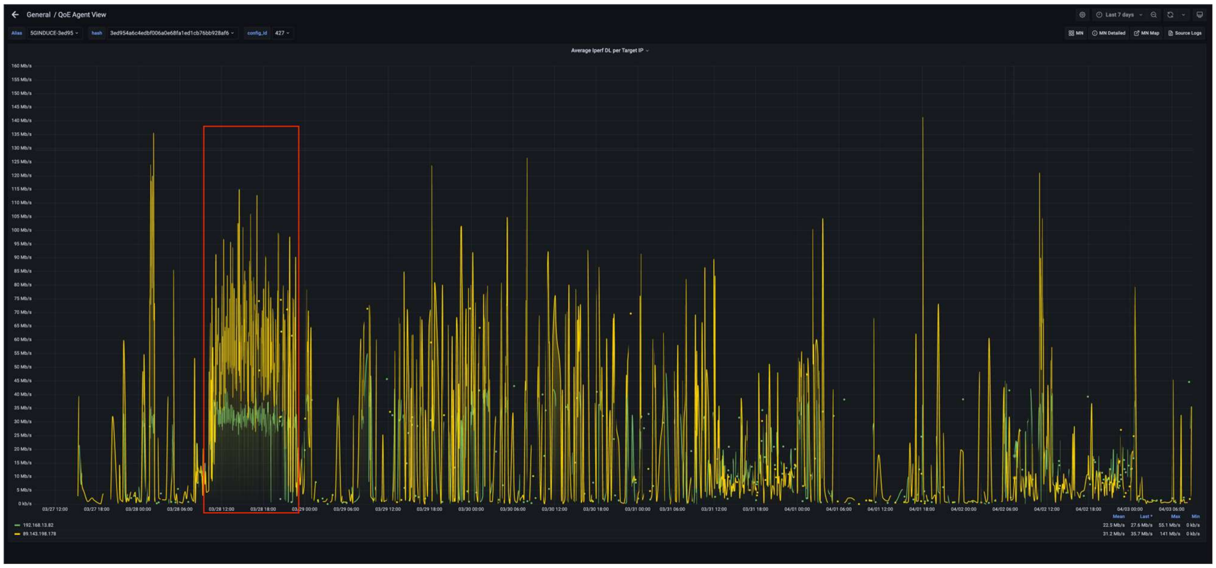

Alias: 5GINDUCE-3ed95
hash: 3ed954a6c4edbf006a0e68fa1ed1cb76bb928af6
config_id: 427

	Mean	Last *	Max	Min
192.168.13.82	22.5 Mb/s	27.6 Mb/s	55.1 Mb/s	0 kb/s
89.143.198.178	31.2 Mb/s	35.7 Mb/s	141 Mb/s	0 kb/s

Figure 86 UC8 –IP download throughput analysis.

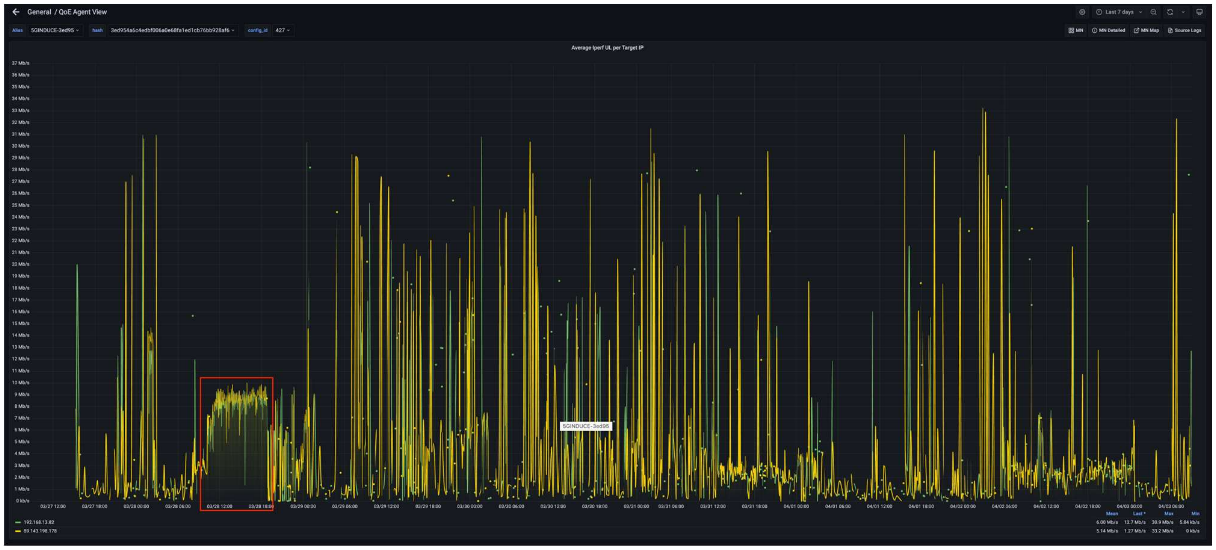

Alias: 5GINDUCE-3ed95
hash: 3ed954a6c4edbf006a0e68fa1ed1cb76bb928af6
config_id: 427

	Mean	Last *	Max	Min
192.168.13.82	6.00 Mb/s	12.7 Mb/s	30.9 Mb/s	5.84 kb/s
89.143.198.178	5.14 Mb/s	1.27 Mb/s	33.2 Mb/s	0 kb/s

Figure 87 UC8 –IP download throughput analysis.

5 Conclusion

The final tests conducted under the 5G-INDUCE project were presented in D6.2. Using the 5G-INDUCE platform and leveraging the nApps developed within the project, solutions targeted in Industry 4.0 were tested in 8 different use cases, implemented in 3 countries (Spain, Italy, Greece). The tests were carried out in industrial facilities, resulting in complex challenges, as expected in fully operational, dynamic environments.

The results of the tests led to useful conclusions about the possibilities of improving the operation of the industrial sector when modern advanced systems are integrated. More specifically:

- The utilization of 5G networks made possible the integration of services with increased requirements, such as low latency and high throughput. Especially for video streaming applications, 5G offered a significant benefit in improving quality, stability, and latency. These enhancements make 5G a vital component for high-performance, low-latency applications, especially in demanding environments.
- The efficient and immediate utilization of the data collected by the cameras and sensors is vital for the proper performance of the services provided. For example, when using a drone, a large amount of data is collected. This data must be used quickly as the drone moves from one area to another transmitting new data.
- The implementation of edge infrastructure offered a choice between edge and core computing. The use of edge computing allows users to process their data close to the data source, which has proven useful in video streaming cases where streaming latency has been reduced. On the other hand, if the use case is not critical, the distribution of AI models can be moved to the core, improving the overall performance and accuracy.

Overall, the test results of the 8 use cases demonstrate the significant potential and challenges of applying 5G in the Industry 4.0 domain, especially in scenarios involving image transmission and remote operations.

References

[1] Explainable Artificial Intelligence for Predictive Maintenance Applications," 2020 Third International Conference on Artificial Intelligence for Industries (AI4I), 2020, pp. 69-74, doi: 10.1109/AI4I49448.2020.0002

ANNEX

KPIs measurements for all ExFa-SP UCs

During the experimentation sessions at Ford, a demo was done with all Spanish UCs running simultaneously on November 17, 2023. This allowed the observation of KPIs for the 3 UCs at the same time.

Note that during the time of this demo the control flows such as the AGV control over UDP remain constant, as the AGVs will send and receive messages as long as they are switched on and connected to the network. On the contrary, some flows depend on the demands by the user. This is the case, for example, of the video traffic TCP flows, which is high when the user switches on the camera but drops when switched off. The flows of the UCs have been explained in previous sections. The figures below represent the results that could be measured from the network and validate that the network can fit perfectly all 3 UCs at the same time.

In Figure 88 all TCP aggregated traffic can be observed . Peaks of over 30 Mbps are reached in the downlink. The uplink demand is quite constant at 10 Mbps, with an occasional peak up to 20 Mbps.

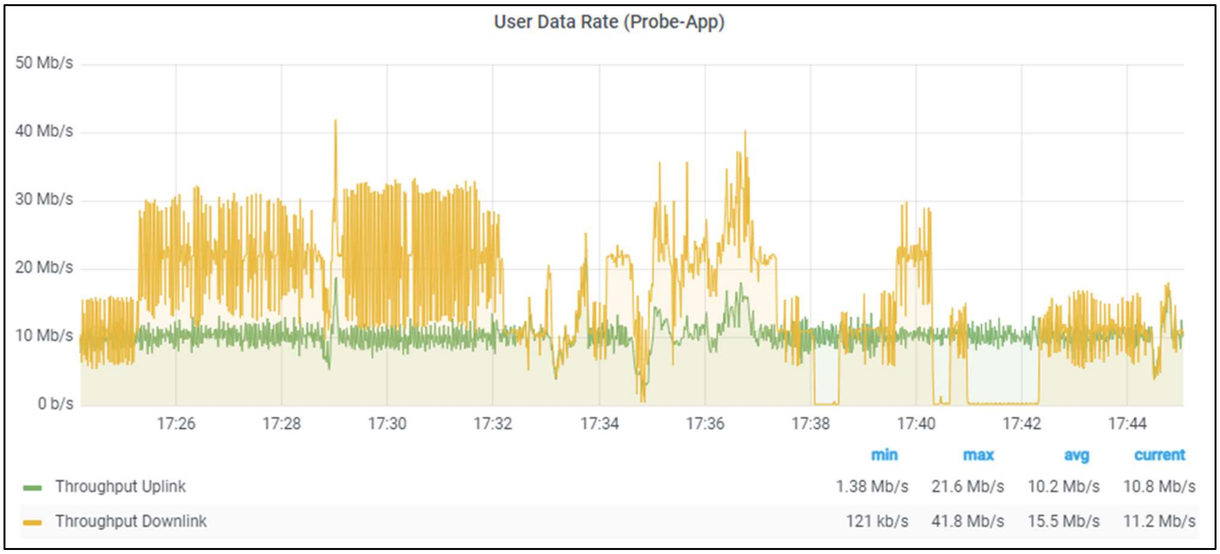

Figure 88 TCP aggregated traffic.

The aggregated UDP traffic was measured as well, shown in Figure 89. In this case, the uplink demands are higher, mainly ranging between 6 and 8 Mbps. The downlink observed is around 4 Mbps.

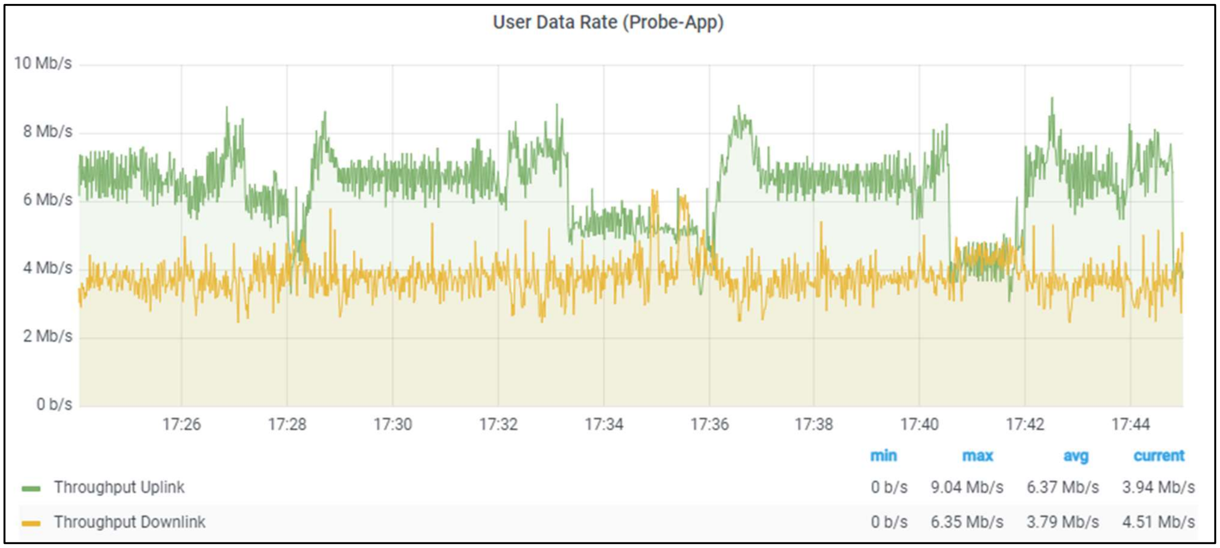

	min	max	avg	current
🟩 Throughput Uplink	0 b/s	9.04 Mb/s	6.37 Mb/s	3.94 Mb/s
🟨 Throughput Downlink	0 b/s	6.35 Mb/s	3.79 Mb/s	4.51 Mb/s

Figure 89 UDP aggregated traffic.

The downlink RTT latency from the application server to the device was also measured, as seen in Figure 90. The latency averages 39.2ms while the 3 UCs are running.

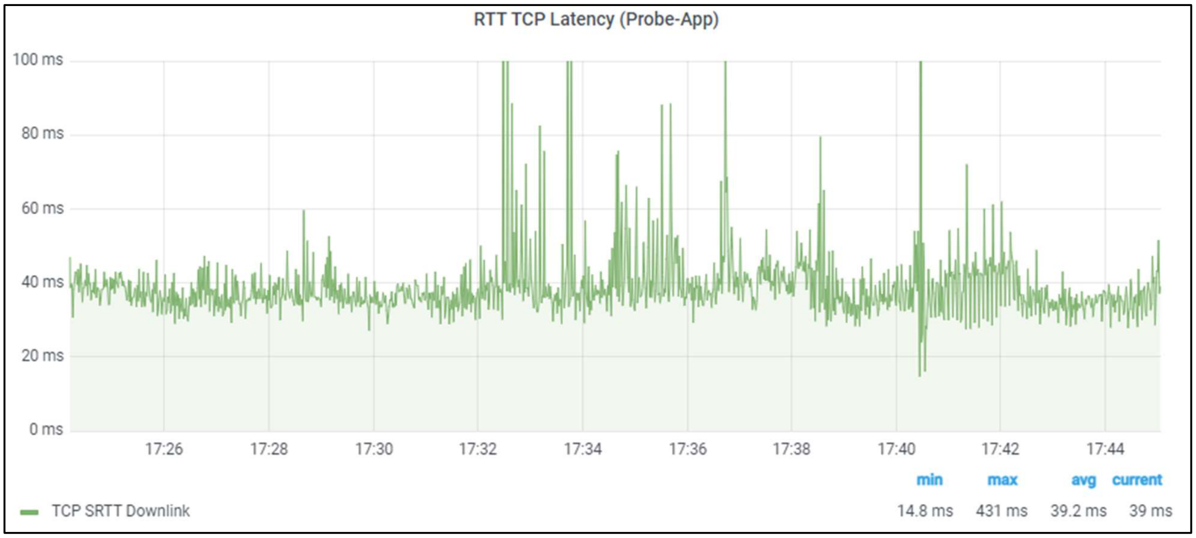

	min	max	avg	current
🟩 TCP SRTT Downlink	14.8 ms	431 ms	39.2 ms	39 ms

Figure 90 Downlink RTT latency.