Use Case #2

Multi-agency and multi-deployment mission critical communications and dynamic service scaling

Overview and Objectives

In the framework of 5G-EPICENTRE, NEM's MCX solution is an innovation branch, which still must pass through internal quality processes. During the 5G-EPICENTRE project, its containerisation has been completed in order to take advantage of the Cloud-Native benefits of the new 5G context. Although the KPIs shown below have gone through an optimisation process, both from a 5G scenario and from vertical service point of view, they should be considered as a first approximation of NEM's future MCX containerised solution performance.

The experimentation objectives have followed two different lines: on the one hand to analyse the service performance (under those deliverable's definitions), and on the other hand to study interesting 5G functionalities for the PPDR sector, such as Slicing and QoS management, which are aligned with UMA's scenario objectives.

Use Case Description

MCXs provides a way for robust connection against errors, ensured predetermined quality of service and network priority over other communications. They are used from public safety and emergency (police, fire, health, etc.) to operators and industrial environments (Oil&Gas, mining, transport, etc.).

The most widespread public safety communication solutions such as TETRA or P25 are based on legacy private radio technology, with proprietary deployments. This is why both public administrations and private companies have to cover the costs of deploying these networks. In addition, there is a mix of providers and differences between networks that end up locking administrations into one provider, with the resulting cost overruns.

On the other hand, the current digital revolution demands new multimedia capabilities, high-speed data access and new functionalities. The emergency communications sector has not been excluded in this sense, and concepts such as remote video assistance, augmented reality, video emergency calls, etc. are beginning to gain relevance. However, the development of this type of services would not be possible over current TETRA or P25 networks, due to their technological limitations. Thus, the capabilities of 5G technology represent a key milestone in the evolution of the MCX sector, as it brings improved bandwidth, low latency with ultra-reliable service and massive machine-to-machine communication (eMBB, URLLC and mMTC) capabilities. In addition, 5G networks create the concept of "Network Slicing", which opens up the possibility of splitting the 5G mobile operator's network into something similar to a sub-network. These 5G "Slices" can be managed in a semi-independent way of each other, enabling a parallel MCX virtual sub-network, so even if the regular network is saturated by traffic, the MCX service can continue working normally.

This has therefore been the research focus of the last few years for NEMERGENT, to deepen the understanding of 5G technology and create a solution that can ultimately deliver enhanced emergency communications services. Therefore, NEMERGENT has taken advantage of the need for a technological leap in the emergency communications sector, along with the

This project has received funding from the European Union's Horizon 2020 Innovation Action programme under Grant Agreement No 101016521.

opportunity that new 5G improvements offered, to develop solutions that can be deployed as a vertical service in public, private or both hybrid 5G networks.

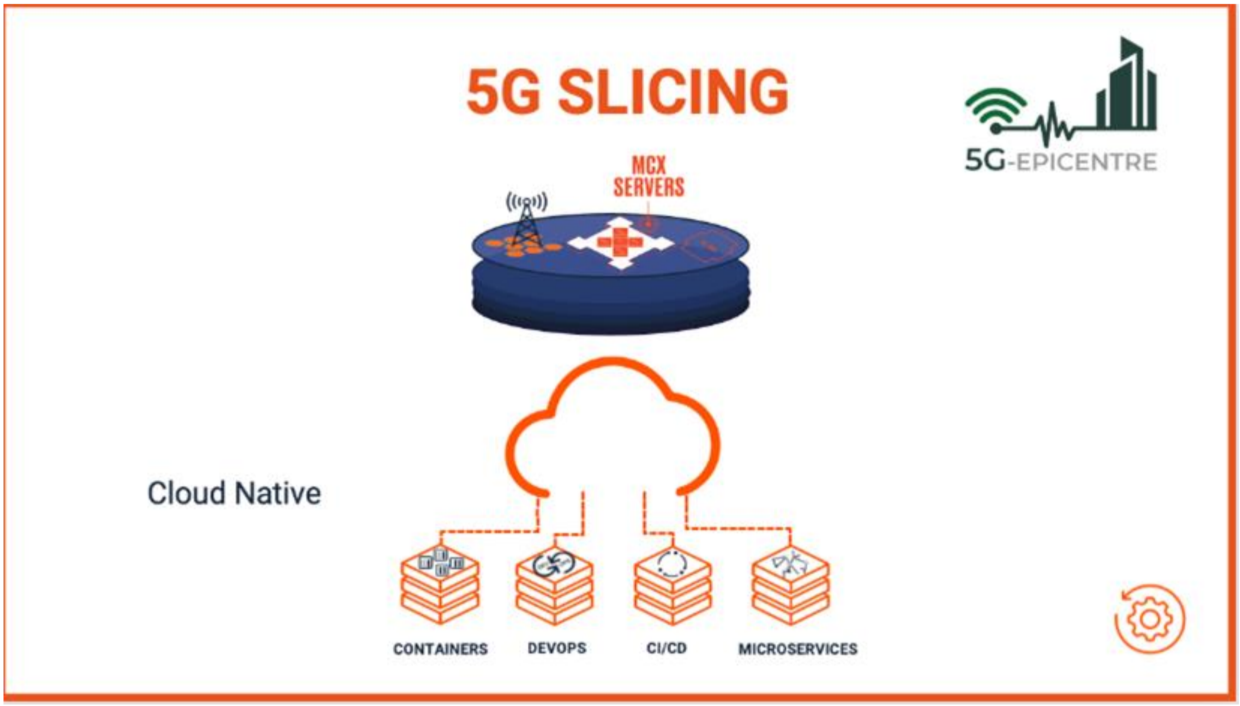

Figure 1: UC2 Architecture

In this framework (Figure 1), a platform such as the one proposed in the 5G-EPICENTRE project acquires significant relevance. NEMERGENT intends to use the infrastructure and services created in the project to continue offering technological advances in its MCX solutions, and to experiment with new functionalities. The Use Case (UC) aims to address the need for coordination between Public Protection and Disaster Relief (PPDR) agencies by implementing a dynamic solution based on location or network conditions. It also seeks to explore Quality of Service (QoS) management and slicing for MCX communications. To this end, a fully microservices-based MCX solution has been developed, fitting with the cloud-native nature of the 5G-EPICENTRE platform. This technology, used in both the MCX solution and the 5G-EPICENTRE platform, will increase the deployment agility, enhancing the experimental needs raised in this UC.

Experiment Setup/Methodology/Deployment

In order to prepare for the experimentation, a Kubernetes-based environment managed by UMA hosted the deployed MCX services. Additionally, a minimum of two PTT UE devices were used, where MCX clients were installed to perform E2E MCX interactions. The measures were collected by one of these clients during the different iterations, and then uploaded to the MCX server-side services. Then, these measurements were MQTT-queued, as service-associated results, and added to the 5G-EPICENTRE platform's analysis processes.

All KPIs have managed to reach optimal values, thus improving the forecasted performance.

As far as the deployment is concerned, NEM's MCX solution is entirely based on Kubernetes. The service can be deployed using Helm, a package manager for Kubernetes. It helps to deploy applications reading the templates, deploying the services accordingly. If the cluster environment does not support Helm, it can be used from an external machine pointing out to the cluster. Currently, the MCX deployment has the following dependencies:

- **Network:** Container Network Interface (CNI) (any should work);
- **Storage:** Container Storage Interface (CSI) (any should work);
- **Seamless monitoring:** Kubernetes Prometheus operator is required. External access to services, which can be configured to use LoadBalancer or NodePort. The system is by default configured to use LoadBalancer:
  - If LoadBalancer is to be used, the cluster where the services are deployed must have a LoadBalancer instance. So far, metalLB has been chosen.
  - If NodePort is to be used, it does not have any requirement.
- **RTPengine** requires to be able to deploy it in hostnetwork mode; and
- **Deploying tool:** Helm is required. It is not necessary to use Helm on the host machine. Helm can be used in a remote machine pointing out to the host machine.

Experiment Execution and Results

Table 1 provides the performance experimentation and the obtained KPIs. As a result, practically all KPIs have managed to reach optimal values, thus improving the forecasted performance. Additionally, the "Advanced SC3 scenario experimentation" provides further detail on the experimentation from a functional point of view, analysing the impact of QoS management and Slicing. The following sequence of statistical data has been calculated under the following conditions.

Table 1: UC2 KPIs

KPIs	Results expected	Experimentation results
UC 2.1	Network RTT	12,65 ms - Optimal
UC 2.2	MCPTT Access Time	37,7 ms - Optimal
UC 2.3	MCPTT E2E Access Time	224,3 ms - Optimal
UC 2.4	Service Availability	99,95% - Optimal
UC 2.5	Re-instantiation time	76,56 s - Acceptable
UC 2.6	Merged talk groups creation time	1 s - Optimal

Table 2: UC2 experiment considerations

Experiment	Name	Samples	Comments
E1	Network RTT	25 iterations, 100 samples per iteration	Performed under normal network conditions, 5G reference value.
E2	MCPTT Access Time	25 iterations, 100 samples per iteration	Time elapsed since user requests to speak in a call, and when user gets signal to start speaking.
E3	MCPTT E2E Access Time	25 iterations, 100 samples per iteration	Time elapsed since user requests to speak and when user gets signal to start speaking, including call establishment and ack from receiver user(s).
E4	Service Availability	25 iterations, 30 samples per iteration	Availability of the service as a percentage measured in a period of time.
E5	Re-instantiation time	10 iterations, 10 samples per iteration	Time elapsed since start of instantiation of the MCX service until the service is fully instantiated.
E6	Merged talk groups creation time	25 iterations, 100 samples per iteration	Time the MCX service takes to create a merged group and users are notified.

The statistical results are shown in Figure 2, framing them with a reference measure based exclusively on 5G performance, such as UMA's RTT calculation.

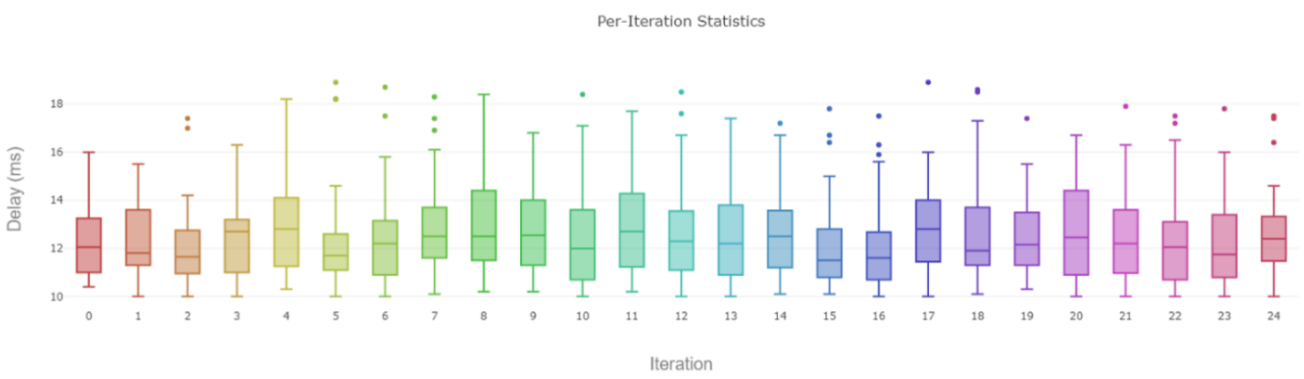

Figure 2: RTT measurement per iteration

Subsequently, service measures are provided in Figure 2- Figure 5, where MCX standard-based KPIs are analysed, providing the basic performance which is further elaborated on the Advanced SC3 scenario experimentation, which is described after the conclusions of the present factsheet.

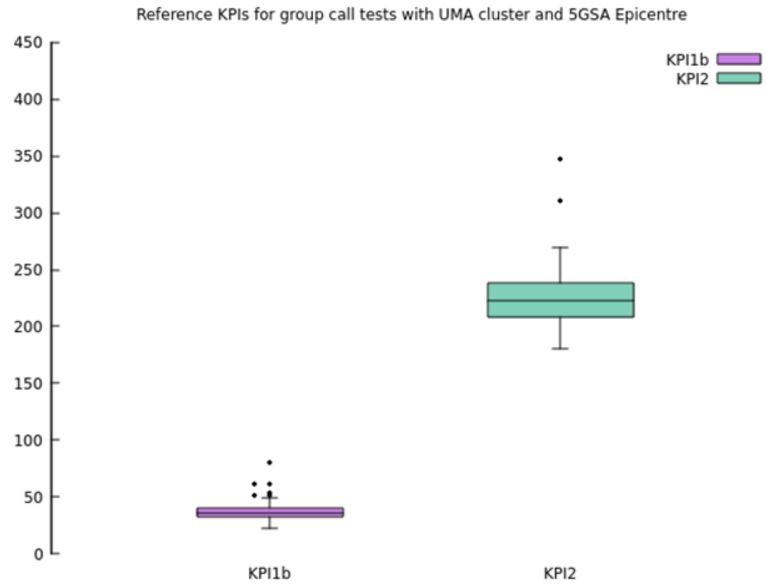

Figure 3: UC 2.2 Access time (KPI1b) and UC 2.3 E2E Access time (KPI2) at UMA

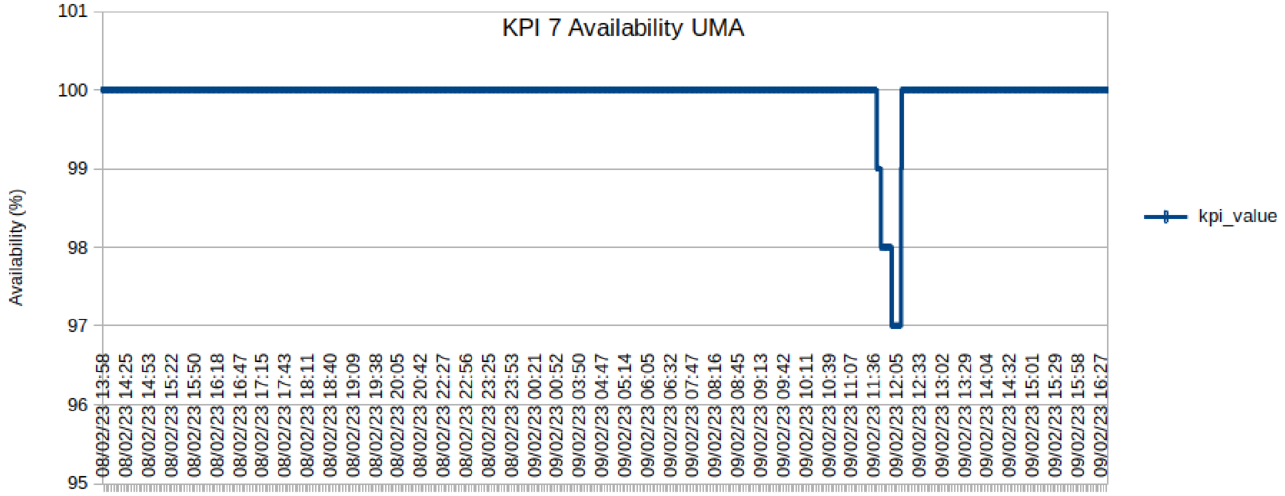

Figure 4: UC 2.4 Service Availability at UMA

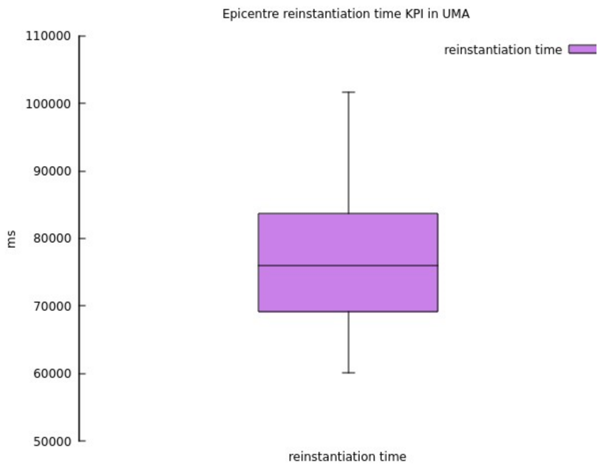

Figure 5: UC 2.5 Re-instantiation time at UMA

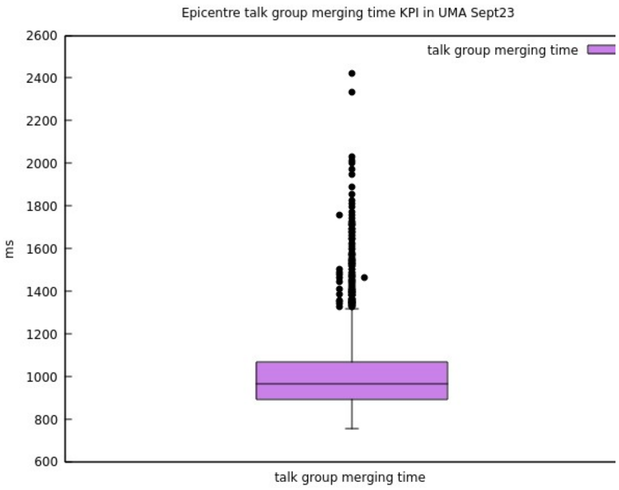

Figure 6: UC 2.6 Merged talk groups creation time at UMA

Conclusions

Almost all UC2 KPIs have reached the optimal values, having undergone optimisation processes by both the vertical and the underlying network managers. More KPIs than expected have reached optimal values, which is a major milestone for an innovation branch like the one tested. Furthermore, those KPIs that have obtained acceptable values, such as Re-instantiation time, have not been the focus of this scenario, so they are probably susceptible to improvement. These have been further elaborated in the context of SC4.

The proposed UC has also been intensively analysed from a functional point of view, exploring both slicing and QoS Management for service-network interaction, needed to guarantee MCX communication conditions in PPDR environments. Both of these concepts have a direct impact on how reliable 5G networks might be perceived by PPDR agencies', and these demonstrations encourage the adoption of 5G solutions among PPDR developers.

All these initiatives have had a corresponding demonstration, to show live how the MCX service works in the context under analysis. Therefore, the UC2 experimentation process can be considered satisfactory, since from a performance, conceptual analysis of advanced 5G functionalities and demonstration point of view, the experimentation has been a success.

Advanced SC3 scenario experimentation

This initial experimentation has been taken further, exploiting the advantages offered by this QoS and Slicing scenario. The impact of slicing on critical communications when faced with different underlying traffic profiles has therefore been explored.

An experimental context has been created in which, based on the 5G scenario above, the network has been stressed by inserting different traffic profiles. A real PPDR scenario has been enabled, facing the issues that may apply in real emergencies, when everyone tries to access the network at the same time (it would also adequately emulate the communications of PPDR agencies managing mass events, such as concerts or sports matches).

For this purpose, the experiment was split into two phases. First, a single slice was used as a control measure, showing the performance of a PPDR service against different traffic profiles. The results obtained show the behaviour of NEM's MCX service sharing network resources with emulated users, as it would be managed from legacy networks, such as 4G. The second phase shows how applying slicing policies enables PPDR traffic differentiation from other generic user's traffic, isolating PPDR communications from saturation.

Figure 7 shows a schematic of the experimentation carried out. Two different UEs (MCS1 and MCS2) have been used to manage the automated interactions during the experiment, addressing MCX group calls and collecting the KPIs. In addition, two different UEs (S1 and S2) were saturated with different traffic profiles, through the different Slices. Using the same radio unit (provided by UMA), and the same User Plane Function (UPF) (provided by HPE), emulated generic traffic, and MCX traffic were separated, avoiding the impact of saturation on PPDR services.

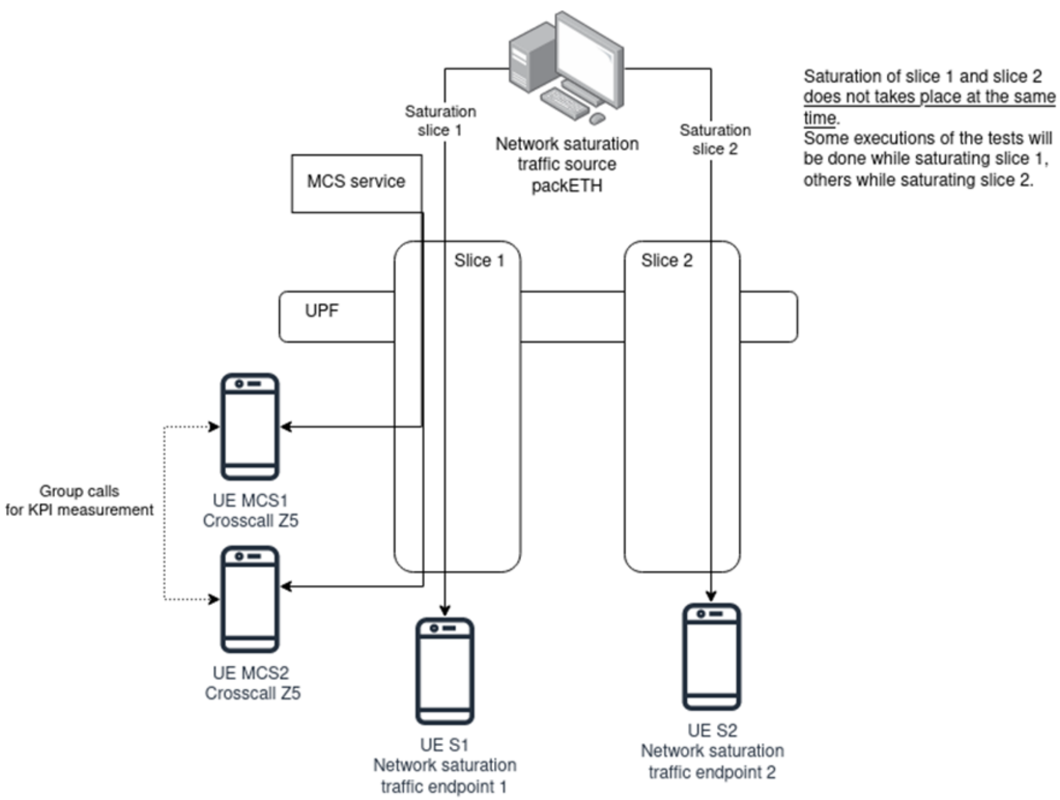

Figure 7: Advanced SC3 scenario architecture

Under these conditions, the KPI2 MCPTT E2E Access Time has been evaluated, as an extension to Experiment UC 2.3 (E3) KPI previously shown in Table 1, since it is considered the most indicative of those defined for UC2. Figure 8 shows the performance of the MCX service with different underlying traffic, with no slicing policies (left) and with slicing policies (right).

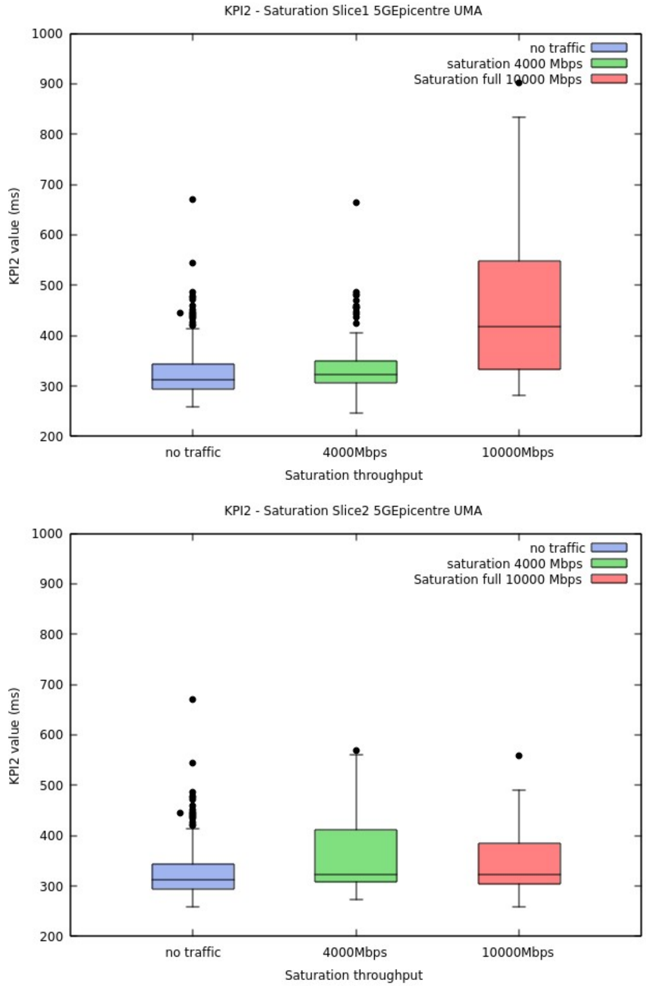

Figure 8: MCX KPI2, Benchmarking slicing's impact

In phase 1 (left), performance is affected when the introduced traffic starts to be significant, obtaining a degraded KPI2 average value compared to phase 2 (right), where slicing policies were applied. Figure 9 shows the impact of 10 Gbps channel saturation on the KPI2 sample dispersion, comparing phase 1 (i.e., with no slicing) with phase 2, where traffic was split into two separate slices.

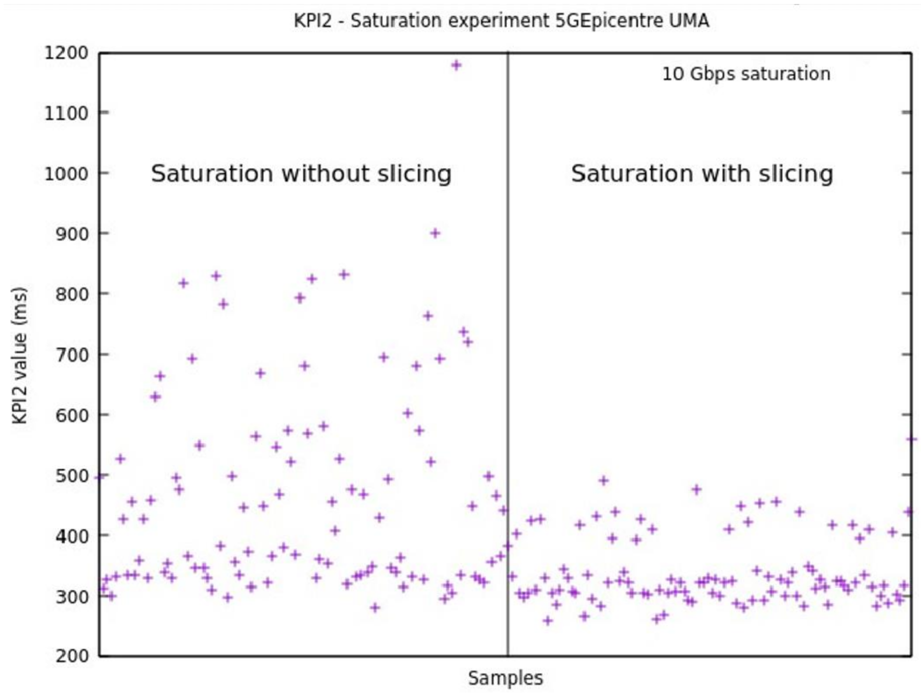

Figure 9: Impact of 10 Gbps channel saturation on the KPI2 sample dispersion

Figure 9 shows how a few samples in the worst-case scenario even overcome the limit value defined in the standard for this KPI2, exceeding 1000 ms. However, thanks to the slicing feature, the MCX service is not affected by the same traffic, thus demonstrating that 5G networks can guarantee the needs of the first responders in the most demanding conditions, enabling them to act swiftly even in complex emergencies.

In addition to the slicing processes, other forms of prioritisation of PPDR traffic over general user traffic have been studied. For this purpose, QoS management has been used through the N5 interface of the PCF. This interface enables the request of specific 5QI for trusted services. The trial has shown how the MCX service's Application Function (AF) communicates with the HPE's 5G Core Network (5GC) PCF through the N5 interface. The AF identifies itself as a service provider, and requests a specific QoS for media flows. The 5GC manages the request, checks its origin and legitimacy, allocates the new QoS to the MCX service, and propagates it to the rest of the 5G network elements. The trial demonstrated that the request and allocation of resources is successfully achieved, in accordance with the MCX requirements. The process has been as follows, performing the full end-to-end request and assignment in less than 73ms:

- The AF has requested for a new 5G 5QI level. From default 5QI level to specific 5QI level assigned to the MCX flows.
- The PCF has processed the request, and has sent a message accepting it.
- 5GC has requested the creation of new dedicated flow from gNB and received confirmation.

Both the N5 and N2 interfaces have been monitored during the results gathering process.