# Disaggregating a 5G Non-Public Network via On-demand Cloud-Native UPF Deployments

Jorge Baranda, Sergio Barrachina-Muñoz, Rasoul Nikbakht, Miquel Payaró, Josep Mangues-Bafalluy
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA)*, Barcelona, Spain
{jorge.baranda, sergio.barrachina, rasoul.nikbakht, miquel.payaro, josep.mangues}@cttc.cat

***Abstract*—Effective and real-time management of data planes becomes of paramount importance to support the ever-increasing challenging use cases and scenarios foreseen in 5G and beyond. In this demonstration, we present a comprehensive showcase of a cloud-native open-source 5G core deployment, realised as a disaggregated non-public network (NPN) spanning multiple locations. Ranging from cloud to edge points of presence, the demonstration establishes an end-to-end experimental platform with over-the-air transmission capabilities, specifically highlighting the on-demand creation and deletion of User-Plane Functions (UPFs). This dynamic deployment approach harnesses the advantages offered by edge locations, empowering the mobile network to adapt and scale as per its specific requirements. Furthermore, our demo elucidates how to use Open Source MANO (OSM) for easing the management operations of network administrators.**

***Index Terms*—UPF, 5G, Management and Orchestration, Cloud-Native, Disaggregation, Open5GS**

## I. Introduction

Recent trends in network management, like Network Function Virtualization (NFV) and Software Defined Networks (SDN), are allowing the transformation of mobile networks. Its enabling technologies, like virtualization/cloudification and the use of open interfaces, are essential to introduce the expected flexibility, dynamicity, and automation capabilities to manage and orchestrate B5G/6G networks. Such capabilities are allowing the introduction of new concepts in the architecture and deployment schemes of the mobile infrastructure, like the Service Based Architecture approach (SBA) proposed for the mobile core, the introduction of new stakeholders in the mobile network ecosystem, like the vertical industries or new service providers, and the appearance of new business models.

One of these new business models is the development of Non-Public Networks (NPNs), where a 5G/6G network infrastructure (i.e., core and radio access network) can be built for an individual enterprise [1]. These deployments provide many operational benefits and can replace former network deployments based on cable or Wi-Fi. In addition to this, such NPN deployments empower the development of novel vertical use cases with ultra-low latency, reliability, and high bandwidth requirements like the remote driving of automated guided vehicles or applications based on virtual, augmented, or mixed reality.

This work has received funding from EU Horizon 2020 grant agreements No 951867 (5G-ROUTES) and No 101017171 (MARSAL), from Spanish MINECO grants TSI-063000-2021-56/TSI-063000-2021-57 (6G-BLUR) and Grant PID2021-126431OB-I00 funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe" and Generalitat de Catalunya grant 2021 SGR 00770.

In this demonstration, we dig into the benefits and possibilities of the SBA approach, considering the disaggregation of mobile core Network Functions (NFs), to automate the distributed deployment of a cloud-native 5G mobile infrastructure to realise an NPN scenario. In the literature, there are several recent examples dealing with the deployment of cloud-native 5G infrastructure using open source software [2]–[5]. In [2], the 5G core network is deployed as a single container in a Kubernetes cluster using Open Air Interface software. In [3], [4], the 5G core network, based on Open5GS [6] software, is deployed as multiple containers in multiple nodes of the same Kubernetes cluster. However, the distribution of the different NFs of the mobile core is predefined and static for the given Kubernetes cluster, requiring adaptations for each environment. The same approach, a manual deployment with Docker of a Free5GC open source cloud-native mobile core, is considered in [5]. Although [3], [5] deal with mobile core network deployments considering multiple User-Plane Functions (UPFs), which is not the case of [4], they are also deployed as a whole with the rest of control-plane NFs of the mobile core.

Differently, this demonstration shows, to the best of our knowledge, the first truly practical disaggregated deployment of a cloud-native mobile core based on Open5GS. This deployed mobile core is interconnected with two radio access networks to build a multi-gNB end-to-end experimental platform with over-the-air transmission capabilities. More specifically, the control-plane functions of the cloud-native mobile core are strategically deployed in a centralized point of presence. In contrast, the user-plane functions can be dynamically deployed on-demand across the available and distributed points of presence. These control and user-plane functions are treated as individual Network Services (NSs), leveraging an ETSI NFV Management and Orchestration (MANO) stack, thereby enhancing the automation and flexibility capabilities of the deployment process. This flexibility enables the use of edge locations, hence exploiting their capabilities such as low latency and data privacy in NPN scenarios. Moreover, it facilitates organic growth, allowing for the on-demand creation and deletion of UPFs of the mobile core network in response to evolving requirements and potential new locations/sites of an NPN.

## II. System Architecture

Fig. 1 illustrates the experimental setup deployed in the CTTC 5G EXTREME testbed. On one hand, there are different

978-3-903176-59-1 ©2023 IFIP

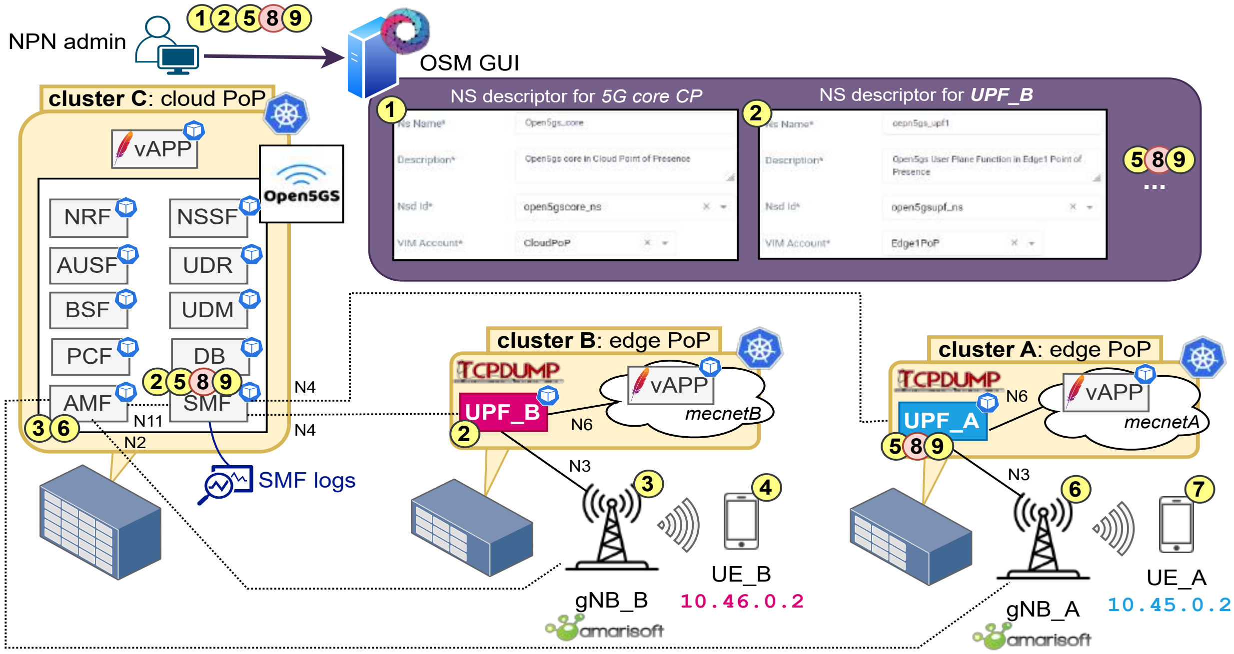


Fig. 1. Distributed 5G Mobile Core deployment under demonstration.

commercial-of-the-shelf servers hosting the different entities of the MANO framework: one node running an instance of ETSI Open Source MANO (OSM) Release 13 and three independent Points-of-Presence (PoPs), where all-in-one Kubernetes cluster in each PoP act as a Container Infrastructure Manager (CIM). This MANO framework performs the deployment of virtualised NSs at the three PoPs representing a cloud and two edge locations, respectively.

On the other hand, there is the equipment related to the provision of Radio Access Network (RAN) capabilities: one Amarisoft Callbox Ultimate and one Amarisoft Callbox Mini hardware act as stand-alone capable 5G gNodeBs (gNBs), while the Amarisoft UE Simbox equipment provides User Equipment (UE) emulation capabilities. A detail of this equipment is represented in Fig. 2.

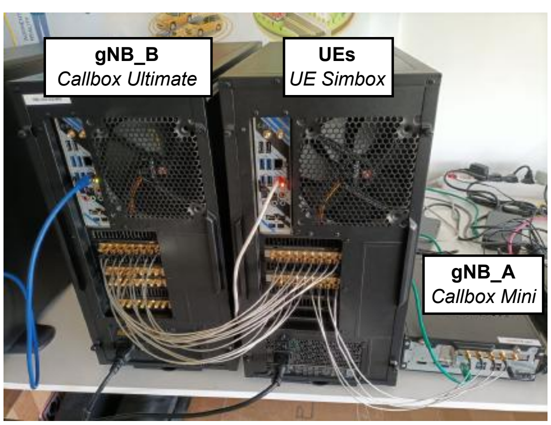


Fig. 2. Amarisoft equipment for realising the RAN components: two gNBs (Callbox) and one UE emulator (Simbox).

In this demonstration, we show the capabilities of this setup to perform a disaggregated deployment of the entities of the Open5GS [6] mobile core software, enabling the on-demand deployment of single UPF-based network services (NSs). Thanks to this, the network infrastructure can grow as required (e.g., a new factory can be added to an NPN deployment if a new location is opened or moved) and edge vertical applications can truly exploit the benefits of edge locations (e.g., low latency).

To achieve this, we have developed two different configuration blueprints (i.e., Helm charts) for the disaggregated deployment of the 5G mobile core (Open5GS v2.6.4) network functions[1]. One blueprint contains the Kubernetes' manifests to deploy the control-plane functions (AMF, AUSF, BSF, databases, NRF, NSSF, PFC, UDM, UDR, SMF). This blueprint acts as the centralised point of management of the whole mobile network, and it can be deployed in a cloud PoP. The other blueprint consists of the manifests belonging to the User-Plane Function (UPF). This blueprint provides flexibility to this mobile core setup, as different instances can be activated/deactivated at the considered edge PoPs when required to provide connection to nearby UEs to the internet or to vertical applications with edge-like requirements.

The flexibility and easy templating capabilities of Helm charts allow the parametrization of the different 5G mobile core network functions. For the inter-pod communications of the *control-plane* blueprint, most of NFs use Kubernetes *clusterIP* service type, so the associated pods are only accessible within the Kubernetes cluster (i.e., providing security to prevent non-required external access). However, this first blueprint required to provide external interface access to

[1]Blueprints for the disaggregated deployment of the Open5GS mobile core are publicly available at: https://zenodo.org/record/8333529.

connect with the RAN segment and UPF entity (i.e., AMF N2 interface and SMF N4 interface). For such external access to the AMF and SMF cluster pods, the *LoadBalancer* service type is used. This type of service is also used in the user-plane blueprint. This kind of Kubernetes service favours the distribution of mobile network entities across independent endpoints and allows the use of the standard ports (i.e., port 38412 for SCTP connection on N2 interface and port 2152 for GTP tunnelling connection on N3 interface), helping to the inter-operability with other equipment (e.g., Amarisoft RAN devices).

Finally, the corresponding OSM descriptors (VNF and NS packages) and OSM parameter configuration files exploiting Helm chart parametrization capabilities were developed to increase the degree of automation in the deployment process. This allowed the generic creation of multiple instances of the UPF in different edge PoPs using a single Helm chart definition and the selection of the deployment PoP.

## III. Demonstration

This demonstration showcases the practical disaggregated and distributed deployment of the different NFs of a cloud-native 5G network mobile core based on Open5GS software [6] and packaged in different NSs. This disaggregation enables the on-demand activation and deactivation of edge sites thanks to the dynamic deployment of UPF entities. Initially, there are several instances of a cloud-native Apache service running in the Kubernetes cluster present at each PoP. These Apache instances represent different instances of vertical applications. After onboarding the associated descriptor packages containing the disaggregated 5G mobile core in the OSM MANO stack, the main steps of the demonstration are:

1) OSM receives the request to deploy the Open5GS control-plane NS in the Cloud PoP. In the configuration of the SMF, there are two data networks defined (e.g., *mecnetA* and *mecnetB*). Once the deployment of the 5G core components is up and running in the selected Kubernetes cluster, two UEs, namely UE_A and UE_B, are registered in the UDR database of Open5GS, so they can authenticate in the network.
2) Now, we proceed to create the first user-plane NS in an edge PoP. OSM receives the request to deploy the user-plane NS in cluster B, a different cluster than the rest of the mobile core NFs. Let's call it UPF_B. This UPF is configured to reach the already deployed SMF in the cloud PoP. The logs of the SMF show how UPF_B registers in the SMF.
3) The Amarisoft gNB_B is started, and it registers in the AMF NF located in the cloud PoP.
4) UE_B is powered on, registers to the AMF, and an IP address of the range of *mecnetB* is assigned. We check that UE_B can reach the internet as well as the Apache service running in the same Kubernetes cluster where UPF_B is running.
5) We proceed to create the second user-plane NS in another edge PoP. OSM receives the request to deploy the user-plane NS in cluster A. Let's call it UPF_A. Like the previous UPF, UPF_A is configured to register in the available SMF.
6) The Amarisoft gNB_A is started, and it registers in the available AMF entity.
7) UE_A is powered on, registers to the AMF, and an IP address of the range of *mecnetA* is assigned. As previously, we check that UE_A can reach the internet as well as the Apache service running in cluster A edge PoP.
8) Now, to further show the on-demand capabilities of the deployment, we proceed to delete the UPF_A. This edge location is deactivated and, as expected, UE_A loses connectivity but UE_B is still operative.
9) We recreate the UPF_A instance again by issuing the corresponding request to OSM. Once the UPF_A is up and running, UE_A is powered off and on to recover connectivity to the internet and the simulated edge applications running in cluster A edge PoP.

During the demonstration, all steps are shown through the graphical user interface of OSM and the logs provided by the different containerised NFs, where the procedures of the 5G mobile core can be observed. We will also show the connectivity of the UEs with the external entities through *tcpdump* traces running on the different UPF instances.

## IV. Conclusions

This demonstration shows the disaggregated and on-demand deployment of the elements of an open source cloud-native 5G mobile core. In the context of NPN use cases, the demo focuses on the sequential creation/termination of different user-plane NSs (UPFs) to enable distributed edge locations according to the needs of the mobile network. This is possible thanks to the design and implementation of the corresponding descriptors and the exploitation of the capabilities of open source tools like OSM and Kubernetes. Overall, this work serves as a starting point for the (multi-cluster) disaggregation of 5G/6G core network functions, paving the way for the development of more sophisticated scenarios incorporating multiple slices and exploiting data-driven approaches for proactive orchestration.

## References


[1] 5GPPP Technology Board, "Non-Public-Networks – State of the art and way forward", Available at: http://doi.org/10.5281/zenodo.7230191, November 2022.

[2] O. Arouk and N. Nikaein, "5G Cloud-Native: Network Management & Automation", in IEEE/IFIP Net. Operations and Man. Symposium (NOMS), Budapest, Hungary, April 20-24, 2020. IEEE, 2020, pp. 1–2.

[3] S. Barrachina-Muñoz, M. Payaró, and J. Mangues-Bafalluy, "Cloud-native 5G experimental platform with over-the-air transmissions and end-to-end monitoring", in CSNDSP 2022, Porto. IEEE, 2022, pp. 22–26.

[4] N. Apostolakis, M. Gramaglia, and P. Serrano, "Design and Validation of an Open Source Cloud Native Mobile Network", in IEEE Communications Magazine, vol. 60, no. 11, pp. 66–72, 2022.

[5] J. Wu, Y. Gao, L. Wang, J. Zhang, and D. O. Wu, "How to Allocate Resources in Cloud Native Networks Towards 6G", in IEEE Network, pp. 1–7, 2023.

[6] Open5GS 5G Core, Available at: https://open5gs.org/open5gs/docs/guide/01-quickstart/ [Accessed: July 2023].