POLITECNICO DI TORINO
Repository ISTITUZIONALE

Automated deployment and scaling of automotive safety services in 5G-Transformer

Original
Automated deployment and scaling of automotive safety services in 5G-Transformer / Baranda, J., Avino, G., Mangues-Bafalluy, J., Vettori Ricardo Martnez, L., Chiasserini, C.F., Casetti, C., Bande, P., Giordanino, M., Zanzola, M.. - STAMPA. - (2019). (IEEE NFV-SDN'19 Dallas (USA) 12-14 November 2019) [10.1109/NFV-SDN47374.2019.9039990].

Availability:
This version is available at: 11583/2752835 since: 2019-09-19T09:53:46Z

Publisher:
IEEE

Published
DOI:10.1109/NFV-SDN47374.2019.9039990

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

19 November 2026

Automated deployment and scaling of automotive safety services in 5G-Transformer

Jorge Baranda*, Giuseppe Avino⋆, Josep Mangues-Bafalluy*, Luca Vettori*, Ricardo Martínez*,
Carla F. Chiasserini⋆, Claudio Casetti⋆, Paolo Bande†, Marina Giordanino†, Marco Zanzola†
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain;
⋆Politecnico di Torino, Italy; † Centro Ricerche Fiat (CRF-FCA), Torino, Italy

***Abstract*—There is a growing interest of verticals (in this case, the automotive industry) to reap the benefits of 5G networks. At the same time, there is a clear trend of the telco industry to understand their needs. These are also some of the main goals of the EU 5G-TRANSFORMER (5GT) project. This demo focuses on the need of verticals to dynamically deploy services at the edge and to adapt the vertical service to network operational conditions. In particular, it is presented the extended virtual sensing (EVS) service, which deployed on demand at the distributed computing infrastructure (i.e. in the network), complements sensing and processing functions running in the car to detect the risk of collisions and take appropriate action, even if there is no direct communication between cars. The stringent latency constraints imposed by the EVS network service leave a limited processing budget at the vertical service level. Since such processing time is correlated with the CPU consumption of a virtual machine running a VNF of the EVS network service, in this demo we also show how the vertical service exploits the automated scaling capabilities offered by the 5GT service orchestrator to deploy a new instance of the EVS VNF upon reception of a CPU consumption alert generated by the available 5GT monitoring platform.**

***Index Terms*—5G, E2E service orchestration, scaling, NFV-SDN, edge computing, vertical automotive use case**

I. Introduction

5G networks are adapting to serve the needs of multiple vertical industries over the same shared infrastructure by applying concepts such as network slicing. The 5G-TRANSFORMER (5GT) platform, released as open source code at [1], is designed to map multiple vertical service deployment requests into a shared underlying infrastructure featuring multiple heterogeneous technologies for computing and networking. This is done through four main building blocks. First, the Vertical Slicer (5GT-VS) is the 5GT entry point for vertical industries to support the creation and management of their transport slices. Second, the Service Orchestrator (5GT-SO) provides federation of transport networking and computing resources from multiple domains and allocation to slices. This is the focus of this demo. Third, the Mobile Transport and Computing Platform for Verticals (5GT-MTP) is the underlying unified controller of transport stratum for integrated fronthaul and backhaul networks, also including distributed computing resources. Finally, the monitoring platform (5GT-MON) provides input data to adapt to network conditions to comply with service level agreements (SLAs).

This demo focuses on the 5GT-SO capabilities to deploy an extended virtual sensing (EVS) network service (NS) related with safety at the edge (near the cars). After that, an increasing simulated vehicle density results in more messages and more potential collisions to be detected in the EVS NS. This increases the CPU consumption of the virtual machine (VM) implementing the EVS VNF, which risks not complying with the vertical service SLA. Therefore, an alert is generated by the 5GT-MON platform and the 5GT-SO automated scaling functionality is triggered so that a new instance of the EVS VNF is created to share the processing load among both. In brief, the value proposition of this demo is the validation of the 5GT architecture to effectively deploy a NS designed by an automotive vertical (CRF), being able to satisfy the stringent requirements in terms of quality of service (QoS) expressed in the network service descriptor (NSD) received from the 5GT-VS whilst adapting to operational conditions.

II. System Architecture

Figure 1 presents the system under demonstration. The setup is split in two sites, the conference venue and the CTTC E2E 5G lab [2]. The SUMO simulator is used to simulate the mobility of cars and increase the vehicle density during the demonstration. The simulated car messages are sent through an openairinterface-based LTE network with a single UE emulating multiple cars. The EVS NS is deployed at the edge server which is under the control of an Openstack deployment running in the conference venue. The whole 5GT platform is deployed at CTTC, including a 5GT-VS receiving the requests from the vertical, translating it into an ETSI NFV IFA014 NSD, which is then instantiated by the 5GT-SO in close interaction with the underlying infrastructure, under the control of the CTTC-MTP, which, in turn, has visibility of the underlying VIMs, such as the Openstack deployed at the conference venue. Both sites are connected through VPN. Finally, vertical services can be requested from the GUIs of the various 5GT building blocks (5GT-VS, 5GT-SO, CTTC-MTP, 5GT-MON) and the deployment and operation of the service can be shown real-time. The 5GT-SO [3] is capable of integrating multiple MANO platforms through wrappers. In this case, OSM [4] is used to handle the instantiation of VNFs in the underlying points of presence (PoPs), including the edge server. The 5GT-MON is based on Prometheus to collect information and Graphana to display the monitored information in custom and dynamic created dashboards. As for the design of the EVS NS, it is composed of four VNFs, as shown in Figure 1. Next, its operation is explained. Cars periodically send Cooperative Awareness Messages (CAM)

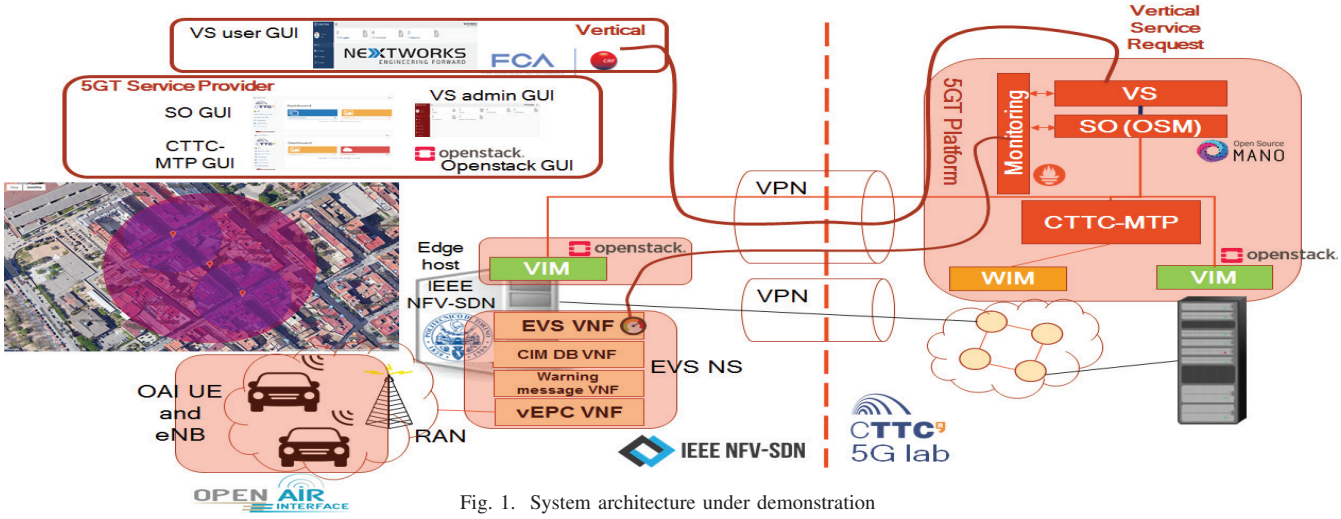


Fig. 1. System architecture under demonstration

with their position, speed and acceleration to the vehicle message database, also referred to as Cooperative Information Manager (CIM) DB. The EVS VNF continuously requests such information to the database and runs the algorithm for detecting potential collisions. If a risk of collision between two vehicles is detected, corresponding DENM (Decentralized Environmental Notification Message) messages are sent to both cars by the Warning message VNF. The interaction between the car and the VNFs in the EVS NS is done through an LTE network. Therefore, a virtual EPC is also deployed at the edge as a VNF part of the service. All these building blocks are deployed as VMs in the edge host to satisfy the required latency constraints of the EVS NS.

III. Demonstration

We will demonstrate the automated deployment of the automotive safety EVS NS using the 5GT platform and the infrastructure of Figure 1. In addition to this, the demonstration also shows the automated scaling of the NS when CPU consumption increases so that the required service response time is at risk. This scaling action instantiates a new EVS VNF instance, which, once operational, will handle the cars in half of the original geographic region of the scenario whilst the original EVS VNF will handle the remaining ones in the other half of the region. The main steps of the demo are:

1) A NS request arrives to the 5GT-SO asking for the deployment of the automotive safety NS described previously.

2) The 5GT-SO processes the request as explained in [3], i.e., a) it calculates the appropriate placement (*edge host*) of the different VNFs to honor the requirement embedded in the request, b) communicates with the MANO platform to launch the different VMs at the underlying VIM and c) checks and manages the need for setting up potential links for VNFs placed in different PoPs. d) In addition to this, the 5GT-SO, due to the existence of monitoring parameters and autoscaling rules in the NSD, interacts with the 5GT-MON to create the appropriate exporters and alerts to react in front of a violation of the conditions expressed in the NSD.

3) Once the NS is ready, the communication between the eNB and the deployed vEPC VNF is activated, so the UE (emulating multiple cars) connects to the mobile network.

4) After that, we launch the Vehicle Emulator, a C-based code that reads the output of the SUMO simulation (consisting of a scenario with two intersections and an increasing number of vehicles) to periodically generate CAMs coming from several vehicles populating the area. These messages reach the CIM VNF through the vEPC VNF. The EVS VNF queries the CIM to evaluate the risk of collisions between all cars of the scenario. In parallel, the CPU consumption of the EVS VNF is continuously monitored by the 5GT-MON platform to guarantee the low latency requirement of the NS, that is, verifying that the processing time of the EVS NS is below a threshold, i.e 5ms. Since there is a correlation with CPU consumption of EVS VNF and EVS NS response time, this is the value we need to monitor to set the scaling threshold.

5) As the emulation progresses, the number of CAM messages increases. Hence, the CPU load at the EVS VNF starts increasing. The 5GT-MON notices it and produces an alert to the 5GT-SO to trigger the autoscaling action, which consists of a *scale out* operation. Then, the 5GT-SO launches a new instance of the EVS VNF. With the new EVS VNF instance, the CPU load of both EVS VNF instances is below the required threshold, hence guaranteeing the low latency requirement of the vertical service.

Acknowledgments

This work has been partially funded by the EC H2020 5G-TRANSFORMER Project (grant no. 761536), by MINECO grant TEC2017-88373-R (5G-REFINE) and Generalitat de Catalunya grant 2017 SGR 1195.

References

[1] 5G-TRANSFORMER platform source code. Available at: http://5g-transformer.eu/index.php/dissemination/open-source/

[2] J. Baranda et al., "Orchestration of end-to-end network services in the 5G-crosshaul multi-domain multi-technology transport network," IEEE Communications Magazine 56(7), July 2018.

[3] J. Mangues et al, "5G-TRANSFORMER Service Orchestrator: Design Implementation and Evaluation", in Proc. of EUCNC'19, June 2019.

[4] ETSI. Open Source MANO project. Available at: https://osm.etsi.org/