Demo: Deploying Transparent Applications at the Network Edges with *LightEdge*

Estefanía Coronado*, Zarrar Yousaf‡, and Roberto Riggio*
*Smart Networks and Services, Fondazione Bruno Kessler, Italy; Email: {e.coronado, rriggio}@fbk.eu
‡NEC Labs, Germany; Email: zarrar.yousaf@neclab.eu

***Abstract*—The deployment of applications and services at the edge of the mobile network under the Multi-access Edge Computing (MEC) paradigm is an excellent option when strict latency requirements must be satisfied. In this work we introduce *LightEdge*, a lightweight, ETSI-compliant transparent MEC solution for 4G and 5G networks. *LightEdge* requires zero modifications to the operator's environment During the demo we will demonstrate how a typical web application can be offloaded to a MEC-enabled 5G network without impairing the effectiveness of the application.**

***Index Terms*—NFV, MEC, 5G, Edge Computing, Offloading**

I. Introduction

The ultra-low latency and the increased bandwidth of 5G systems are opening the door to multiple innovative applications. In line with this trend, Multi-access Edge Computing (MEC) allows offloading computing-intensive and latency-sensitive tasks to the edges of the network. However, despite the benefits brought by 5G, it will take some time until the existing 4G system transitions to a full 5G system. Commercial 5G networks are initially being deployed in Non-StandAlone mode, meaning that the 5G Radio Access Network (RAN) still interfaces with the 4G Evolved Packet Core (EPC). While this approach makes available higher bitrates, it impedes the deployment of new applications and services at the network edges due to the lack of proper support in the standard EPC, which is instead a native part of the 5G Core. The challenge is thus to develop a transition mechanism allowing the deployment of MEC systems in the 4G architecture to immediately enable their capabilities and features to the mobile users.

In this work we demonstrate *LightEdge*, a lightweight, ETSI-compliant MEC solution for 4G and 5G networks. The main goal of *LightEdge* is to provide Mobile Network Operators (MNOs) with a MEC platform that can immediately bring the advantages of edge computing to the 4G end-users, while enabling a seamless transition over the evolutionary path from 4G toward a full 5G architecture. The key strength of *LightEdge*'s design is its transparency to the existing components of a 4G network, therefore requiring zero modifications to the MNO's environment. During the demo we show how a typical web application can be offloaded to a MEC-enabled 5G network without impairing its effectiveness while reducing the experienced latency. All software is made available under a permissive open APACHE 2.0 license for non-commercial use. On line resources available at: http://lightedge.io/.

978-1-7281-4973-8/20/$31.00 ©2020 IEEE

II. *LightEdge* System Architecture: An Overview

The *LightEdge* system architecture, depicted in Fig. 1, is designed to allow UEs to consume applications and services located at the network's edges. *LightEdge* follows the bump in the wire architecture proposed by ETSI [1], thereby placing the MEC host between the RAN and the EPC of the 4G system. This enables the interception and interpretation of UE requests and the utilisation of applications and services at the edges, rather than routing them to remote locations.

This can be accomplished through various procedures. One way consists in intercepting DNS requests from users before reaching the intended DNS server and resolving them to the virtual IP address of a local MEC application. Notice that this approach will not work if secure alternatives to DNS, such as DNSSEC, are employed. Once a UE requests the virtual IP address of a local MEC application, *LightEdge* takes over the communication and the UE traffic is steered towards the MEC host. Standard stateful L4 or L7 load balancers (not depicted in the figure to improve readability) can be used to distribute the load among multiple MEC application instances. UEs thus get served from the local MEC host at the edge.

A. MEC Platform

The design of *LightEdge* extends the ETSI reference architecture [2] and encompasses the functionalities required to run MEC applications at the edges of the network.

1) Service Registry: The Service Registry is an ETSI MEC functional component that contains the catalog of services and applications that are available on the MEC platform.

2) RNIS: The Radio Network Information (RNI) Service [3] provides MEC applications with real-time information about the RAN including channel quality measurements (e.g. RSRP and RSRQ) and notifications of various RAN events (e.g., UE attach/detach, RAB setup/tear down).

3) Message Queuing Telemetry Transport (MQTT) Service: The MQTT service forms the communication nexus between the various sub-components of *LightEdge*. It works following a publish-subscribe paradigm where services and applications can publish/subscribe to a certain topic.

4) Traffic Rule Manager: The Traffic Rule Manager is in charge of (re)configuring the L3 switch to route the traffic among applications/services and the 3GPP network.

5) DNS Resolver: The DNS Resolver allows mapping requests originated from UEs to local IP addresses that are routable inside the MEC domain.

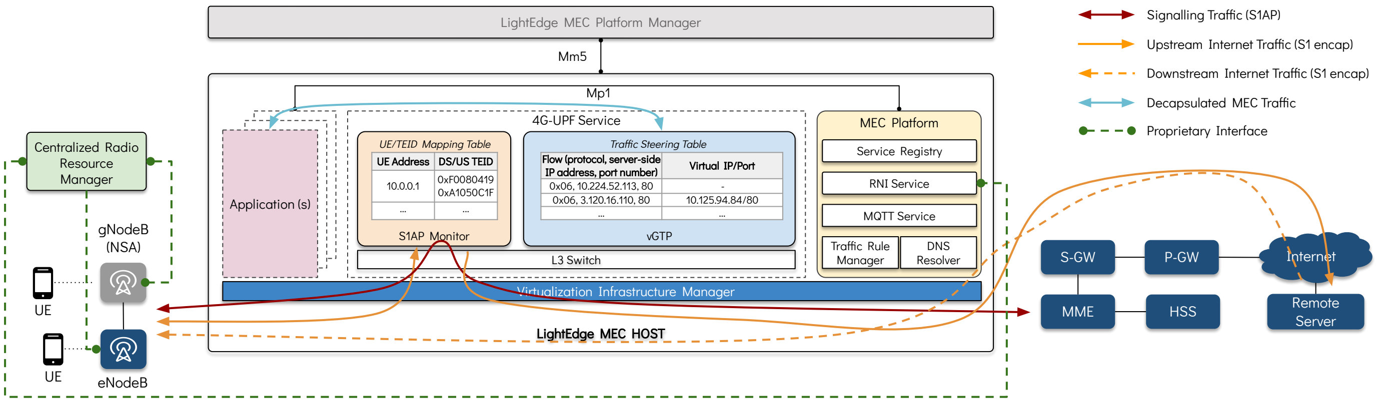

Fig. 1: The *LightEdge* reference system architecture.

B. 4G-UPF Service

This is the core service that is complemented and leveraged by the *LightEdge* service extensions in the MEC platform. A detailed view of this component is included in Fig. 1. This service comprises the functional elements described below.

The *L3 Switch* is in charge of steering the traffic between eNB/EPC and the MEC services under the control of the *Traffic Rule Manager*. The S1AP channel is steered to the *S1AP Monitor* module by matching the IP protocol type (SCTP is $0x84$) while the GTP-u stream is steered to the *vGTP* module by matching the UDP port (2152 for GTP-u).

The *S1AP Monitor* receives the SCTP-encapsulated S1AP traffic and tracks the upstream and the downstream TEIDs. This is done by monitoring the *InitialContextSetupRequest* and the *InitialContextSetupResponse* messages. The former assigns the upstream TEID while the latter assigns the downstream TEID. Once both messages are collected for a UE, a new rule is added to the *UE/TEID Mapping Table*.

The *vGTP* manages the stateful GTP encapsulation/decapsulation between eNB and SGW. The inner IP flow of the upstream GTP-u traffic is matched against the *Traffic Steering Table*. If a match is found on the 3-tuple (protocol, server-side IP address, and port number), the GTP protocol stack headers (GTP, UDP, and IP) are removed.

III. Implementation Details

LightEdge has been designed with cloud-native principles, where each component is deployable using container technologies and the platform itself is natively compatible with Kubernetes. We have developed a prototype of *LightEdge* and deployed it on an LTE testbed. The RAN part comprises a 3GPP-compliant LTE stack provided by srsLTE while as EPC we use nextEPC. It must be noted that *LightEdge* is vendor-agnostic and can be used with any eNodeB/EPC components (including commercial ones). The *4G-UPF Service* is implemented as a native application while the MEC platform is a simple Python agent. The *Traffic Rule Manager* and the DNS Resolver are implemented using, respectively, *ipables* and *dnsmasq*. The 5G-EmPOWER Software-Defined RAN controller [4] provides the RNIS functionality. 5G-EmPOWER can interface with commercial and open-source eNBs and extract relevant channel quality information such as UE RSRP/RSRQ measurements. Finally, the *MQTT Service* is implemented using RabbitMQ.

IV. Demo

In this demonstration we illustrate the traffic offloading capabilities enabled by *LightEdge*. The demo is divided into two scenarios to showcase the latency reduction provided by *LightEdge* while maintaining the performance level. In the first scenario a user accesses a remote web service running on a Microsoft Azure instance using a standard smartphone without making use of the traffic offloading features offered by our platform. A web dashboard allows the user to graphically see the latency experienced by accessing these services from a remote site. By contrast, in the second scenario a new entry is added to the *Traffic Rule Manager* in order to allow the user to consume another instance of the same web service that is running instead on *LightEdge*. Again using the web dashboard, the user can verify that the latency has been significantly reduced in this case.

Acknowledgements

This work has been performed in the framework of the European Union's Horizon 2020 project 5G-CARMEN co-funded by the EU under grant agreement No 825012. The views expressed are those of the authors and do not necessarily represent the project. The Commission is not liable for any use that may be made of any of the information contained therein.

References

[1] European Telecommunications Standards Institute (ETSI), *MEC Deployments in 4G and Evolution Towards 5G*, Std. ETSI White Paper 24, February 2018.

[2] European Telecommunications Standards Institute (ETSI), *Multi-access Edge Computing (MEC); Framework and Reference Architecture*, Std. ETSI GS MEC 003, March 2019.

[3] European Telecommunications Standards Institute (ETSI), *Mobile Edge Computing (MEC); Radio Network Information API*, Std. ETSI GS MEC 012, July 2017.

[4] E. Coronado, S. N. Khan, and R. Riggio, "5G-EmPOWER: A Software-Defined Networking Platform for 5G Radio Access Networks," *IEEE Transactions on Network and Service Management*, vol. 16, no. 2, pp. 715–728, 2019.