A QUIC-based proxy architecture for an efficient hybrid backhaul transport

M. Luglio, M. Quadrini, C. Roseti and F. Zampognaro
University of Rome "Tor Vergata",
dpt. of Electronics Engineering
Rome, Italy. E-mail: zampognaro@ing.uniroma2.it

S. P. Romano
University of Naples "Federico II",
dpt. of Electrical Engineering and ICT
Naples, Italy.

***Abstract*—Network Function Virtualization (NFV) allows a fast deployment of customized solutions when and where needed, supporting breakthrough operational models compliant to multi-tenancy and slicing paradigms. This approach enables challenging communication scenarios overcoming possible performance flaws as well as enabling new disruptive services within future 5G based networks. In this framework a virtual Performance Enhancing Proxy (vPEP) is proposed, to foster the efficient use of satellite links either as a supplementary or complementary backhauling. The vPEP objective is to efficiently manage both switch-over (for performance optimization) and failover (for service continuity) of end-to-end Web traffic within an hybrid terrestrial/satellite backhaul network. To this aim, vPEP leverages several features of the new QUIC protocol, such as the 0-RTT establishment of secure communications and the connection-oriented flow management over a UDP-based connectionless transport. In this paper we describe the vPEP implementation as well as a Proof-of-Concept (PoC) demonstrator used for the validation activities.**

***Index Terms*—NFV, QUIC, Satellite, vPEP.**

I. INTRODUCTION

5G fosters a seamless integration of the widest range of services and technologies ever witnessed in past wireless mobile networks. One of the keys for such revolution is the introduction of Network Function Virtualization (NFV), which includes the set of technologies capable to replace physical implementation of network devices and appliances (routers, firewall, ACL, etc.) with software implementation running on legacy multipurpose hardware. Accordingly, the satellite role can be reviewed by properly adopting such new technologies to overcome traditional performance flaws.

In this context, the VIBeS project (Implementation of VNFs for Broadband Satellite networks) [1], funded by the European Space Agency (ESA), addresses the design and the implementation of a virtualized Performance Enhancing Proxies (vPEPs) for an efficient management of a hybrid terrestrial-satellite network [2]. The reference network architecture is inspired from 5G specifications and envisage a mobile Radio Access Network (RAN) enhanced with virtual edge computing capability. The latter in turn allows the allocation of local virtual applications to breakout traffic or to provide advanced services, such as the target vPEP one.

The present paper specifically focuses on the set of virtual functions, included into vPEP instances, aimed to optimize Web-based traffic. In more details, we implemented several distributed virtual network functions (VNFs) based on some features extracted from the QUIC protocol [3], recently standardized within HTTP/3 [4] by the Internet Engineering Task Force (IETF). We exploited the QUIC connectionless transport to enable pre-established internal tunnels between vPEP agents with the aim to introduce flexibility in the use of multiple inner (backhaul) links, by supporting either dynamic switchover or failover procedures. We can establish a number of QUIC-based tunnels to multiplex web traffic, which can be opportunistically "moved" on the most suitable backhaul link on the basis of both the application requirements and the overall congestion status. Please note that alternative approaches and solutions exist, and are discussed for instance in [5].

As a preliminary validation, we prepared a virtualized Linux based Proof-of-Concept (PoC) Testbed to confirm the creation and management of such tunnels across different network segments (backhaul links), within different administrative domains.

II. REFERENCE SCENARIO

We assume that between the access network (client-side) and the public network (Internet) several backhaul links (L-1, L-2,..., L-*n*) are available. A backhaul link can be defined as an independent infrastructure operated by a network operator, with its own IP addressing and NAT/firewall strategies. In other terms, we assume a backhaul link as a layer 3 (IP) transport to the public network. In the context of future 5G networks, virtualization technologies allow to activate and/or remove such links on-demand as required to satisfy variable load conditions or Quality of Service (QoS) requirements. The different links can have different physical characteristics and can be enabled in an opportunistic way according to some pre-defined policies. In the present work, the management of such links is delegated completely to the vPEP: the motivation of this choice is described more in details in [6].

As reference, a high-speed fiber link is defined as the default one, while traffic may be migrated to a secondary one (e.g., satellite, or another terrestrial link) in case of failures or in case of temporary congestion. The (default) terrestrial link speed is assumed 20 Mbit/s, with a secondary satellite link of 10 Mbit/s with a Round Trip Time (RTT) of about 500 ms. This set up is compliant with the framework envisaged by the VIBeS project,

978-1-7281-5127-4/20/$31.00 ©2020 IEEE

considering common access capacities with different backhaul links types.

Within this communication framework, we consider traffic generated by Web applications classified into two QoS categories: "high-priority" and "low-priority". We decided to use host IP addresses for the classifications; any other IP/application parameter can be used as well in real-life scenarios, as well as using more classes. The proposed vPEP virtual functions will handle such traffic categories separately through two QUIC tunnels and have full control on the use of the available backhaul links.

III. The QUIC-based proxy solution

The vPEP is a virtual network element composed of several virtual functions. It splits end-to-end TCP connections with the end-systems, namely Web clients and Web servers, with the aim of enabling the inner communication through QUIC-based tunnels. Such tunnels act as virtual paths for different categories of traffic and represent the chosen granularity level for the "migration" operations. The QUIC protocol is considered as the key enabling technology for our vPEP, as it is capable to map different web objects to several *Stream IDs* within the same *Connection ID*, which is the unique identifier of the tunnel. The number of QUIC tunnels is a parameter set by the service operator on the basis of the target QoS management policies, and it is set to 2 within this paper. The higher the number of tunnels, the greater the flexibility in the management of heterogeneous traffic. On the other hand, a higher multiplexing on a single tunnel guarantees consistent traffic dynamics on which operator can take decisions.

In case of congestion, low priority traffic can be quickly "moved" to the supplementary (satellite) backhaul link in order to reduce the latency experienced by the traffic in the high priority one. This operation is allowed because the vPEP can be configured to change the backhaul link in use by a QUIC-tunnel, while keeping unchanged the *Connection ID*; this operation is therefore robust to possible changes in the source IP address and transparent to the connections already established. In general, tunnel "migration" can be controlled and managed by a controller entity at the application level, meeting the scope of future software-based network functions in the frame of 5G [7], implementing a customized QoS logic.

IV. PoC Testbed description

To validate the QUIC-based proxy architecture described above, we set up a Linux KVM based Testbed, divided in three main parts: client-side (Access), proxy-side (Backhaul) and server-side (Internet). The client side and server side host legacy and unmodified HTTP applications based on TCP transport and leveraging standard HTTP/2 client and server software. The backhaul segment is suitable to host all the defined vPEP functions, which includes proxy functionalities and enabling of QUIC transport. In particular, we included a custom implementation of QUIC based on quic-go[1]. The overall Testbed configuration is depicted in Figure 1.

[1] L. Clemente, https://github.com/lucas-clemente/quic-go

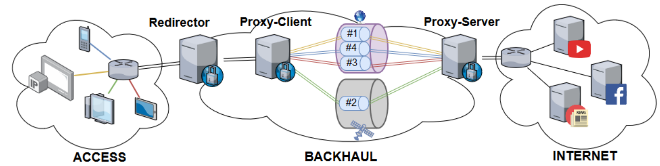

Fig. 1. QUIC-based testbed architecture.

A. vPEP detailed deployment

The chain of virtual functions for the vPEP is spread along the two edges of the hybrid backhaul link to provide a seamless and efficient backhaul service to the end-systems. These are:

– Redirector (RD) function
– Proxy-Client (PC) function
– Proxy-Server (PS) function

The RD function connects the web-client access network to the PC element. The RD is based on a open-source software (i.e., Squid) and its role is to make the QUIC proxy-chain and operations transparent to the Web clients, avoiding changes in their default configuration. Basically, RD intercepts ongoing connections, terminates them locally (splitting) and redirects all the traffic to the corresponding (as defined in the initial configuration) QUIC-based tunnel. Each tunnel is managed by a separate PC, which controls three network interfaces: one towards RD and two towards the available backhaul links. In the proposed configurations, two tunnels are defined.

PC therefore represents an end-point of the QUIC tunnel and then it encapsulates all the coming traffic from the client-side in a QUIC tunnel. Finally, PS acts as the remote end-point of the QUIC tunnel, receiving traffic from PC across any of the available backhaul links and forwarding it to the target web servers over legacy HTTP/TCP connections.

For the sake of validation of the proposed solution, we configured two links with very different physical characteristics, whose specific parameters are described in section II.

V. Validation and performance results

We present a preliminary test, with the aim to demonstrate the feasibility and the potentiality of the proposed vPEP solution based on QUIC tunnels. In particular, we drill down through the dynamic management the QUIC-based "migration" and how it can be helpful to improve the overall service performance. Basically, to reproduce the communication scenario introduced in section I, we generated an overall traffic pattern varying over time, first below the terrestrial backhaul capacity (20 Mbit/s), then at 300 s overloading it (in average). In the benchmark case, only the terrestrial backhaul is available (vPEP off). Figure 2 collects simulation results in terms of:

– total rate measured on the terrestrial backhaul;
– RTT of a probe TCP connection crossing the terrestrial backhaul.

While total rate is below the bottleneck capacity (simulation time < 300 s), the RTT is kept constant to its minimum value, then preserving optimal performance conditions for all the active web application flows. We remark that most of the Web traffic is composed of quite a high number of objects downloaded with correlations with dependencies managed at the session and application layers. The size and the scheduling of web objects download makes the resulting data rate oscillating over time and most likely below the available bottleneck capacity. Then, the main performance indicator is the object download time, which in turn will be proportional to the end-to-end latency, herein expressed as RTT of a probe TCP connection. After 300 s, we overloaded the bottleneck terrestrial link by increasing the number of web page downloads. The download rate is bounded to the available capacity so the RTT values drastically increase. This means that traffic competition inflates the latency experienced by each and every flow, impairing the overall web application performance.

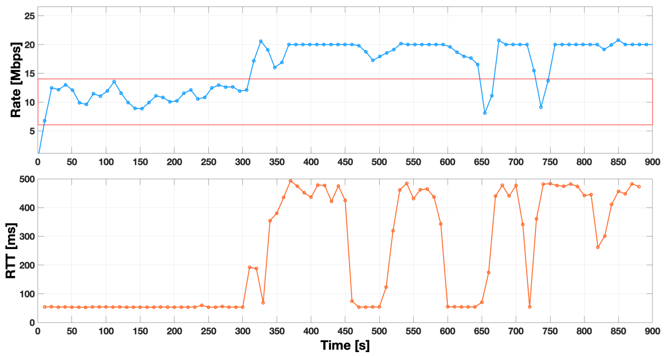

Fig. 2. Benchmark performance with a single backhaul link (terrestrial)

By enabling the proposed vPEP instead, there is an improvement on the service by efficiently leveraging the supplementary satellite capacity when congestion is detected. We configured the vPEP to establish and maintain two QUIC-based tunnels to map "low-priority" and "high-priority" traffic, respectively. Such a classification has been done based on IP addresses by using a custom access list in RD, so that the traffic is equally split in 2 tunnels. By default, both tunnels go through the default terrestrial backhaul. Then, we set two thresholds to trigger automatic tunnel "migration":

- *Threshold 1* equal to 70% of the terrestrial backhaul capacity. When the overall traffic load exceeds this threshold, the "low-priority" tunnel is moved over satellite.
- *Threshold 2* equal to 30% of the terrestrial backhaul capacity. When the overall traffic is below this threshold and the "low-priority" tunnel is running over satellite, this latter is moved back over the terrestrial backhaul.

The restults with such vPEP configuration is shown in Figure 3. The RTT over the terrestrial backhaul is equal to the minimum value for most of the time, also when traffic load increases (after 300 s). As a consequence, the "high-priority" traffic encounter ideal condition overall the simulation. The few RTT remaining spikes are due to the not-optimal "migration" algorithm and the corresponding threshold values, which can result not sufficiently responding up on sudden rate changes. The optimization of the algorithm is postponed to the future work.

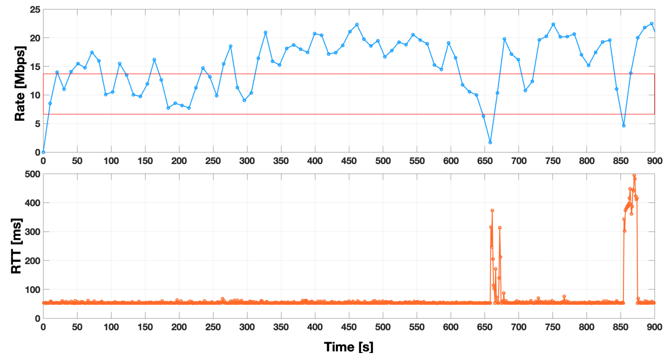

Fig. 3. vPEP managing two QUIC-based tunnels.

VI. Conclusion and Future Work

We presented a novel vPEP solution to exploit hybrid backhaul links, based on the deployment of virtual proxy functions. In particular, we implemented a pair of virtual proxy functions inheriting some features of the QUIC protocol to manage tunnels that can be dynamically moved over the available backhaul links. A PoC testbed has been built up to perform several tests aimed to demonstrate the feasibility of the proposed solution and assess the potential benefits on performance. As future work, we plan to design a better classification of traffic categories and advanced algorithms for QUIC-based tunnels migration.

Acknowledgment

The results presented herein are part of the outcomes of the ESA VIBeS project, Contract n.: 4000122991/18/UK/ND [1]. Responsibility of the content resides with the authors.

References

[1] ESA ARTES AT Project, "VIBES - implementation of virtualised network functions (VNFs) for broadband satellite networks," *ESA*, 2018. [Online]. Available: https://artes.esa.int/projects/vibes

[2] A. Abdelsalam, M. Luglio, M. Quadrini, C. Roseti, and F. Zampognaro, "QUIC-proxy based architecture for satellite communication to enhance a 5G scenario," in *2019 International Symposium on Networks, Computers and Communications (ISNCC)*, June 2019, pp. 1–6.

[3] J. Roskind, "QUIC: Quick UDP Internet connections - Design Document and specification rationale," Google, Tech. Rep., 2012.

[4] M. Bishop, "Hypertext transfer protocol version 3 (HTTP/3) - draft-ietf-quic-http-24," IETF Internet Draft, Tech. Rep., 2019.

[5] A. Abdelsalam, M. Luglio, C. Roseti, and F. Zampognaro, "Analysis of bandwidth aggregation techniques for combined use of satellite and xdsl broadband links," *International Journal of Satellite Communications and Networking*, vol. 37, no. 2, pp. 76–90, 2019.

[6] A. Abdelsalam, A. Bujari, M. Luglio, D. Munaretto, C. E. Palazzi, M. Quadrini, S. P. Romano, C. Roseti, and F. Zampognaro, "Implementation of Virtualised Network Functions (VNFs) for Broadband Satellite Networks," in *European Conference on Networks and Communications, EuCNC*, 2019, pp. 182–186.

[7] M. Luglio, S. Romano, C. Roseti, and F. Zampognaro, "Service delivery models for converged satellite-terrestrial 5G network deployment: A satellite-assisted CDN use-case," *IEEE Network*, vol. 33, no. 1, pp. 142–150, 2019.