MESON: Facilitating Cross-Slice Communications for Enhanced Service Delivery at the Edge

Konstantinos Katsaros*, Vasileios Glykantzis*, Panagiotis Papadimitriou†,
Georgios Papathanail†, Dimitrios Dechouniotis‡ and Symeon Papavassiliou‡

*Intracom Telecom, Greece
{konkat, vasgl}@intracom-telecom.com
†University of Macedonia, Greece
{papadimitriou, papathanail}@uom.edu.gr
‡National Technical University of Athens, Greece
{ddechou@netmode, papavas@mail}.ntua.gr

Abstract—**In this paper, we advocate the careful orchestration of cross-slice communications (CSC). We argue that multi-tenancy and service co-location presents unique opportunities for B2B interactions, inter-service communications and service composition, especially in the case of edge computing and location-based services. However, network slice isolation in its prevailing form raises significant concerns related to performance and resource utilization. In this context, we present *optiMized Edge Slice OrchestratioN* (MESON), a MANO framework aiming at facilitating cross-service/cross-slice communications and further identify and discuss key challenges towards the support of secure and optimized cross-slice communications at the network edge.**

I. Concept

Over the last years, the advent of virtualization technologies and network programmability, typically considered in the context of Network Function Virtualization (NFV) and Software-Defined Networking (SDN) paradigms, have substantially contributed to the emergence of *network slicing* capabilities, allowing (5G) mobile network operators to lease service-tailored bundles of network, compute and storage resources to their customers. This presents a unique opportunity for innovative services to get deeply integrated into the (edge of the) network infrastructure. Taking a step further, and in an analogy to the well-established cloud realm, we expect these technological advances to facilitate the emergence of a broader service ecosystem, allowing cross-service interactions. Such interactions may significantly vary, from existing operations, such as single sign-on systems, to currently emerging Function-as-a-Service (FaaS) deployments, to future service interactions in the context of location-based (edge) services, *e.g.,* touristic guide applications integrating advertisement or social networking information.

However, network slice provisioning commonly targets strong (traffic) isolation between the various vertical service providers, essentially preventing any *cross-slice communication* (CSC) and data access between service components (*i.e.,* virtualized network functions – VNFs) even if they are placed on the same (edge) datacenter (DC) or server. Obviously, though adhering to justified security and performance guarantees, such setups do not take advantage of multi-tenancy, potentially leading to significant latency benefits, as well as traffic overheads both in the backhaul/transport segments of the network and further to/from the Internet, *i.e.,* inter-domain traffic. While technical solutions are feasible and have been engineered in the past [1], [2], they present a limited view of the problem space. Indeed, CSC in the context of the emerging (edge) service ecosystem, goes beyond the traffic traversal of network slice boundaries, pointing to a wider set of challenges related to service placement decisions and forwarding/performance optimizations, resource utilization and B2B SLA-related policy enforcement; as well as automation, targeting a zero-touch management (ZTM) model for the support of CSC.

In this respect, *optiMized Edge Slice OrchestratioN* (MESON) aims at the design and prototype implementation of an enhanced management and orchestration (MANO) framework able to foster the establishment of cross-slice/tenant interactions and to facilitate the secure *and* optimized communication between next-generation services, across network slice borders. This corresponds to the establishment of the ZTM mechanisms required for the realization of cross-service interactions including service discovery, placement and communication establishment, subject to OSS/BSS-level procedures expressing the intent of service providers to establish a synergy.

The MESON framework focuses on the edge of the network infrastructure. First, the optimization of cross-slice communication, becomes particularly important when considering the focus of edge computing on the support of low-latency applications. At the same time, the support of service placement decisions and resource management policies that facilitate cross-tenant interactions, becomes particularly important when it comes to an inherently distributed environment, where the solution space grows substantially.

II. Use Cases

In this section, we briefly discuss three illustrative use cases for the envisioned framework. A first simple scenario is considered for the case of Virtual Network Operators (VNOs),

This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH - CREATE - INNOVATE (project code: T1EDK-02947).

which provide transparent caching solutions throughout their network so as to reduce traffic and improve content delivery. A synergy can be established to mutually benefit from their co-location at several areas of the network infrastructure. A symmetric cache peering relationship is realized, *i.e.,* when a video streaming request reaches a cache node of the first VNO and results into a cache miss, a (co-located) cache node of the second VNO is queried first. If there is a peering cache hit, the content is fetched from the peering VNO in order to reduce traffic load on the backbone and transport segments, otherwise the request is forwarded to the origin server.

In the second scenario, an augmented reality (AR) service provider of a touristic/city guide AR service with stringent latency requirements allows users to upload live feeds from their phone camera and get back overlaid metadata for their current location. The service provider leases a network slice at the edge of the network and integrates information retrieved from co-located slices of third-party service providers, such as personalized recommendation/advertisement services or social network media services, in order to reduce latency.

Finally, inspired by Industry 4.0 challenges, a human-robot interaction assisted by AR/VR technologies is considered. Two co-located slices for human AR/VR guidance and robotic manipulation, which belong to different operators or have different Key Performance Indicators (KPIs), establish CSC to exchange information between humans and robots and also share common functionalities (*i.e.*, path planning, simultaneous localization and mapping, object recognition) in order to coordinate their activities and fulfill the ultra-low latency requirements of their interaction.

III. Challenges

The optimization of CSC in MESON seeks to take advantage of network slice and service co-location in the envisioned edge computing ecosystem. Our objective is to avoid unnecessary traversals of the network infrastructure when possible. The potential gains of such optimizations depend on the network topology and the location of the peering service components. One approach is to rely on shared network instances within NFV Points-of-Presence (PoPs), as realized by corresponding OpenStack Compute Hosts. Leveraging on this feature, a shared network instance can be created for one of the tenants/service providers. Subsequently, the MESON orchestration framework can mediate, so that the service provider is able to grant access privileges for the shared network to its peering service provider. The latter is then allowed to attach its peering VNFs to the shared network. The MESON MANO framework mediates in this process by enabling the exchange of VNF location and configuration information between peering service providers.

Policy enforcement is crucial in addressing a series of concerns related to security, resource and performance isolation in the context of CSC. Traffic isolation should ensure that only legitimate traffic is allowed to access resources across network slice borders. To this end, several techniques have been employed, such as the creation of shared logical networks and the enforcement of strict access control on top (*e.g.,* with OpenStack's Role Base Access Control). In addition, policy enforcement goes through the careful orchestration of each individual peering service, targeting the controlled sharing of compute and storage resources allocated to VNFs. In essence, the objective is to ensure that the performance of each separate service, as delivered to each end-user, does not deteriorate due to peering. This may require the instantiation of separate VNFs dedicated to the support of CSC, essentially confining peering resource consumption within VNF borders. In turn, this gives room to the firm control of resource utilization as exposed by the corresponding auto-scaling strategies for these VNFs. Traffic policing also plays an important role in allowing the control of incoming traffic across a CSC path. Such functionality may either be provided by the Slice Provider or the Slice Tenants. Across all the dimensions discussed above, monitoring has a central role in enforcing peering policies. This corresponds to the monitoring of (i) the resources allocated on each side of a CSC, with the purpose of enabling precise auto-scaling, policing, and supporting KPI inspection, (ii) the resources of the CSC path, with the purpose of avoiding performance bottlenecks between network slice borders and potentially enforcing traffic policing.

CSC further raises additional challenges in terms of slice placement. In particular, when a peering request is submitted before the slice deployment, existing placement methods, which generate mappings based only on resource and topology constraints, may lead to sub-optimal mappings from the perspective of CSC *i.e.,* the slice placement may not facilitate the peering with other slices. Therefore, slice placement should be coordinated with peering service discovery in order to enforce the co-location of slices with cross-service interactions.

IV. Conclusions

In this paper, we raised the need for cross-slice communication, when different service components are deployed across co-located but strictly isolated network slices. We presented the need for a new MANO framework that aims at optimized CSC and, thereby, opening up new (business) opportunities for cross-slice interactions. The potential gains from CSC are discussed in three use cases, whereas we also highlight the most critical challenges for our envisioned MANO framework.

References

[1] K. V. Katsaros, V. Glykantzis, and G. Petropoulos, "Cache peering in multi-tenant 5G networks," in *Second IFIP/IEEE International Workshop on Management of 5G Networks (5GMan 2016)*, 2017. [Online]. Available: https://doi.org/10.23919/INM.2017.7987446

[2] K. V. Katsaros and V. Glykantzis, "Experimenting with cache peering in multi-tenant 5G networks," in *21st Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, 2018. [Online]. Available: http://doi.org/10.1109/ICIN.2018.8401623

European Union
European Regional Development Fund

HELLENIC REPUBLIC
MINISTRY OF ECONOMY & DEVELOPMENT
SPECIAL SECRETARY FOR ERDF & CF
MANAGING AUTHORITY OF EPAnEK

EPAnEK 2014-2020
OPERATIONAL PROGRAMME
COMPETITIVENESS
ENTREPRENEURSHIP
INNOVATION

ΕΣΠΑ 2014-2020
ανάπτυξη - εργασία - αλληλεγγύη
Partnership Agreement
2014 - 2020

Co-financed by Greece and the European Union