Mobile Edge Vertical Computing over 5G Network Sliced Infrastructures: an Insight into Integration Approaches

Roberto Bruschi[*], Raffaele Bolla[*£], Franco Davoli[*£], Anastasios Zafeiropoulos[§], Panagiotis Gouvas[§]

[*] S3ITI National Lab., CNIT, Italy [£] DITEN, University of Genoa, Italy [§] UBITECH, Greece

Abstract–5G is designed to leverage on network softwarization technologies, like Network Functions Virtualization (NFV) and Mobile Edge Computing (MEC), to expose customized network instances and resources, at the edge of the infrastructure to vertical stakeholders. Most of the 5G success will depend on the ability to attract vertical stakeholders acting in the cloud, enabling them to smoothly port cloud applications to 5G, and to add performance and cognitive capabilities not supported in cloud environments. To this end, this paper provides an insight on the possible architectural approaches to fully integrate Vertical Applications (vApps) into the 5G infrastructure. The paper follows a top-down approach. First, it provides an outlook on the state-of-the-art in cloud application design, and on the MEC and NFV new capabilities. Then, on this basis, the analysis is devoted to identify integration issues not yet fully addressed in standard specifications. Two alternative architectural approaches are discussed.

I. INTRODUCTION

A key 5G objective resides in the enablement of a new class of vApps with heterogeneous and extremely challenging requirements [1-2]. To this end, the 5G community is embracing well-known technologies, like NFV and MEC. Both these frameworks are based on the unrestrainable "softwarization" process, which is going to transform network operators' infrastructures into distributed datacenters with advanced virtualization and software-driven capabilities.

MEC and NFV frameworks will have clear and well-separated objectives. As stated by the ETSI MEC working group (WG) in [3], "***MEC*** *uses a virtualisation platform* ***for running applications*** *at the mobile network edge.* ***NFV*** *provides a virtualisation platform to* ***network functions****."* As the infrastructure requirements of both approaches are quite similar, the use of a converged virtualization infrastructure would be beneficial.

These frameworks will be key enablers for flexible customization of mobile *network slices* to the needs of vApps [4-5] and their provision with full network-awareness and zero-perceived latency. Radically new applications [6-7] can be made viable through the joint adoption of these technologies. As defined by 3GPP and MGMN [2,4], a network slice is a logical end-to-end network providing specific 5G network services, offered as-a-Service by a Telecom Service Provider (TSP) to Over-The-Top (OTT) players, such as Vertical Industries. The TSP should support multiple network slices from different OTT players at the same time, and dynamically realize each of them through the composition of shared/isolated 5G functions' instances [8].

Notwithstanding the high complementarity between NFV and MEC, their integration and interplay in the 5G ecosystem is still largely unexplored. The objective of this paper is to identify the possible approaches, to highlight their main advantages and drawbacks, as well as to introduce relevant integration issues. To this end, the paper will follow a top-down approach: starting from current trends in cloud application design, we will identify the main evolution routes towards 5G, mapping them onto the emerging technological paradigms offered by the 5G infrastructures. From this mapping, two possible architectural approaches to integrate vApps with edge computing and NFV facilities are finally introduced and discussed, by outlining the roles of involved stakeholders, and the main induced benefits and drawbacks.

The remainder of this paper is organized as follows. Sect. II provides a short discussion on cloud vApps, and on how they have to evolve in the 5G scenario. Sect. III introduces the main building blocks composing the 5G architecture, and Sect. IV discusses possible integration approaches. Conclusions are drawn in Sect. V.

II. FROM CLOUD TO 5G-READY APPS

In cloud computing, state-of-the-art software engineering trends are based on the microservice concept. To achieve high scalability and agility levels, applications are decomposed into a *mesh of "cloud-native" microservices*, each one with specific and "small-scope" processing objectives, instantiated even multiple times, and packaged on independent virtual execution environments [9-10]. A central entity, named "*application orchestrator*," is in charge of realizing the application business logic, by managing the lifecycle and the mesh interconnection of such microservices over cloud resources.

Depending on the nature of the application, the orchestrator can alter the application graph [10], by adding/removing: (*i*) new types of microservices to enable/upgrade specific application functions, (*ii*) instances of existing types of microservices to scale the overall application processing capacity, where needed.

The above operations are usually supported through a suitable mesh interconnection management. For instance, horizontal scaling operations are enabled by layer-7 traffic load-balancing/proxying inside the same microservice (i.e., in the so-called "microservice sidecar" [9]), or requested as-a-Service to the Virtual Infrastructure Managers (VIMs). DevOps upgrades can be realized similarly through the "*canary deployment*" procedure: when a new version of a microservice is deployed, only a small traffic share is redirected to it. If the new component version works properly, the traffic volume redirected to it is smoothly increased up to the entire workload.

The mesh of a cloud application is usually terminated by a single front-end point towards the public Internet, where only microservices interfacing end-users and connected things are placed. Under the perspective of cloud application providers, 5G represents an invaluable milestone to cut the edge of cloud technological limitations and to offer radically new services. This milestone could be achieved by reducing the end-to-end latency between applications and connected users/things to enable real-time procedures, as well as by making applications cognitively react to the networking context. Thus, part of the microservices should be moved from remote datacenters to the 5G edge, and the interconnectivity customized through 5G network slices, leading to key differences from today's applications, as summarized in the following.

II.A. Network Awareness

Network awareness enables the adaption of vApp operations according to network-level performance and events. To this end, starting from Release 13, 3GPP included the *Service Capability Exposure Function* (SCEF) to expose 4G network information to vApps through Application Programming Interfaces (APIs). SCEF has been designed to abstract services, like requests for certain Quality of Service (QoS) levels, monitoring of network resources, positioning of User Equipment (UE), etc. In the 5G core architecture, the *Network Exposure Function* (NEF) was introduced to evolve SCEF capabilities with particular attention to network slicing and to hide sensitive information on UEs and the network (e.g., topology) to third parties. If properly authorized, through the NEF, an external application might request specific traffic steering in the user-plane of a dedicated 5G network slice.

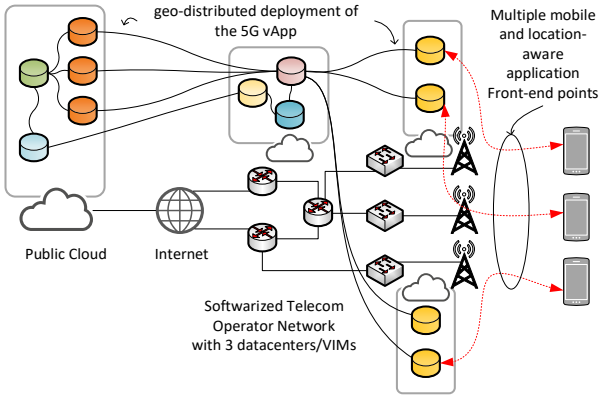

Figure 1. Example of application based on microservices as deployed in 5G facilities through edge computing technologies.

By consuming such APIs [8], the *Vertical Application Orchestrator* (VAO) can request specific network behavior/QoS levels, properly react to UE mobility events, and manage the lifecycle of the microservice mesh to enable/scale/modify application capabilities according to the number and the type of connected devices, their radio link status, etc.

II.B. Locality of Application components

To ensure zero-perceived latency times or to reduce the volume of traffic delivered to the Internet, components in the application mesh should be deployed close to the UEs. To this end, computing facilities might be applied at various aggregation levels of the mobile network [12]. The locality of application components happens at the price of a geographically limited serving coverage (Fig. 1): if deployed into an e/gNodeB, microservices can serve only locally attached UEs; if deployed at higher network aggregation levels, more UEs can be served.

While in common cloud scenarios the scaling of microservice instances is decided according to the overall incoming workload, in 5G the VAO has to independently manage this process in each geographical zone/datacenter where the application has to be made available. Moreover, the tighter the locality constraint is, the closer to the network edge the component will be allocated, reducing the need of horizontal scalability for performance adaptation.

The application components' locality triggers a further relevant problem related to UE mobility [3]: when a UE performs a handover between cells covered by different computing facilities, its connectivity towards the microservices in the new local datacenter should be provided without interrupting/resetting the application session.

Therefore, specific network solutions have to be

supported to steer application traffic between UEs and the closest microservices, as well as to migrate the "state" of application sessions among the instances acting in the new and the old computing facilities. Migration operations might be triggered by the same application components deployed in the new facility, or by the VAO reacting to handover event notifications (see Sect. II.A).

II.C. User-Plane Integration

In the latest-generation cloud scenarios, microservices (and their sidecars) communicate among themselves using flexible interfaces like Representational State Transfer (REST), through virtual networking resources (e.g., virtual networks) usually provided by a single VIM. The external reachability of the application is usually guaranteed by legacy Domain Name System (DNS) services in the public Internet. Thus, while advanced network technologies (e.g., software-defined networking) are used by VIM providers as "private" means to optimize their infrastructure, cloud applications are currently designed to rely over classical network protocol stacks.

5G vApps are envisioned to go beyond this model, and to include explicit network QoS constraints (e.g., on latency) in the interconnection links, not only between components, but even among components and UEs. The positioning of the application components into VIMs available in the network should be performed to cope with such constraints (see Fig. 1). Differently from their cloud relatives, 5G vApps are going to heavily rely on multiple geographically-distributed VIMs, and explicitly attach their front-end components to mobile UEs, before the public Internet.

Regarding UE attachment, 3GPP 5G core specifications introduced the *User-Plane Function* (UPF) [8], which is in charge of realizing (through multiple instances potentially placed at various network aggregation levels) all the user-plane operations in a (shared/isolated) network slice. A further function, the *Session Management Function* (SMF), is meant to configure the UPF forwarding behavior. If enabled, a vApp can request the SMF (through the NEF – see Sect. II.A) to configure UPFs to steer predetermined traffic flows towards locally-attached external data-networks [8]. These data-networks represent the user-plane "*attach points*" [15] between the vApp and the mobile network domains. In 4G, there are no 3GPP-standard means to realize these attach points in edge facilities [11].

Regarding multi-VIM deployment, the VAO cannot play a direct role, since the selection of deployment VIMs and the setup of end-to-end network services require sensitive data on UEs (e.g., positioning) and on the TSP (e.g., network topology). 3GPP designed the NEF to hide this network-internal information to third-parties. As a result, the VAO should expose the deployment requirements of the application graph to the TSP, which uses them to setup the proper user-plane network resources. These services include 5G network slices and inter-VIM connectivity, as well as control-plane interfaces (e.g., NEF). In the example in Fig. 2, two NFV services are considered. The first service realizes a 5G network slice, and includes a couple of UPFs, a SMF, and a NEF, to which the VAO (acting as a 3GPP Application Function) is interfaced to monitor network operations and to request configurations. The user-plane interconnectivity between the 5G service and the vApp is realized through UPFs' external data networks. The second NFV service is not in the 3GPP specification domain, and it is used to realize inter-VIM connectivity.

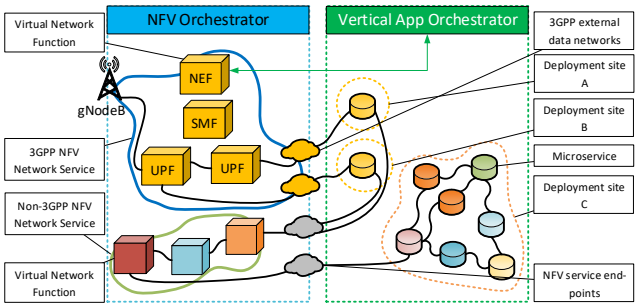

Figure 2. Example of a vApp deployed in three VIMs and attached to two NFV services.

III. ANATOMY OF THE 5G NETWORK AND vAPP ORCHESTRATION

The operational behavior of a 5G vApp is affected by multiple subsystems owned by different stakeholders, which might act in autonomous fashion, at different layers, and with diverse objectives. As outlined in [4] and in [2], the main stakeholders actively involved will be three: the *vertical industry* owning the application, the *TSP(s)* offering 5G services, and the *telecom infrastructure provider(s)* offering computing and communication facilities. We describe the main control and management blocks, acting in these stakeholders' domains. The *VAO* manages the lifecycle of the graph of microservices composing the vApp, and acquires network and computing resources as-a-Service. The *Business and Operational Support Systems* (BSS/OSS) provides resources (e.g., network slices) of the TSPs to

vertical industries as-a-Service. These modules expose the 5G network to verticals in terms of network slices [4]. At the southbound, the BSS/OSS is supposed to interface the NFV Orchestrator (NFVO) to request the (de)activation/modification of NFV services [14]. The *NFVO* manages the network services composing the network slices activated by the BSS/OSS. In detail, the NFVO is in charge of deploying and managing the lifecycle of both 3GPP services and functions (e.g., NEF, UPF, etc.) of any activated network slice, and non-3GPP network services (e.g., for the multi-VIM interconnectivity). The *Mobile Edge Orchestrator* (MEO) manages the embedding of mobile edge applications and their lifecycle. It is triggered by the BSS/OSS [3]. The VIMs expose the resources (especially computing and storage) of datacenters mainly to the NFVO and to the MEO/VAO. The *Wide-area Infrastructure Managers* (WIMs) realize the logical interconnectivity among sets of service/vApp components instantiated in different datacenters or towards UEs.

The VAO acts in the Vertical Industry domain, the BSS/OSS, the MEO, and the NFVO in the domains of the TSPs, and the VIMs and WIMs act in the Infrastructure Provider's domain (Fig. 3). Most of these control blocks should also support advanced "*multi-tenancy*" (i.e., hosting multiple overlaying systems) and "*multi-domain*" (i.e., exploiting the resources from multiple underlying systems) capabilities. WIMs and VIMs, especially the ones derived by cloud computing (e.g., OpenStack), already provide complete multi-tenancy capabilities, and provide VAO and NFVO with the possibility of managing different graph instances into separated and isolated tenant spaces. Notwithstanding its fragmented nature (see the use case in [7]), the overall ecosystem should target a fully automated control of all the resources/services at any layer to allow the 5G vApp lifecycle management: from their planning (Day-0) and their first deployment (Day-1), through their in-life operations (Day-2) – e.g., upgrade, scaling, etc. – to their termination.

In Day-0 operations, the VAO asks the BSS/OSS for 5G network slices and edge computing resources, indicating QoS/locality requirements of the vApp; and the BSS/OSS computes a suitable deployment plan (e.g., it selects the datacenters and the wide-area resources to be applied).

In Day-1 operations, the BSS/OSS requests wide-area interconnectivity from the WIM(s), and triggers, if needed, the NFVO (and the MEO) to setup and/or properly configure network services and edge computing resources. Upon the successful fulfilment of the previous operations, the VAO can start deploying the vApp.

In Day-2, all the services are running. The VNFs and the vApps' components are monitored by the NFVO and the VAO, respectively, which might independently scale them to cope with the incoming workload, perform self-healing procedures in case of problems, manage upgrade operations, etc. As previously stated, the VAO can also interact with the 3GPP NEF for acquiring data/events of UEs in its network slice, request changes in traffic steering rules, and suitably adapt the vApp graph to enable/optimize application capabilities. In Day-1 and Day-2, the involved orchestrators should interact with VIM(s) in order to manage the resources needed to execute VNFs'/vApps' components. While these interactions and the ones between the BSS/OSS and NFVO are well-specified [14], further aspects related to the integration of the 5G ecosystem and vApps are still open. Among the most relevant ones, we can cite: *(1) The VAO-MEC duality*: in the current specification, the ETSI MEC framework and, consequently, the MEO, only allow "monolithic" vApps, with no explicit support for multiple microservices and an external VAO. The MEO has been designed as a sort of simplified VAO, but residing in the TSP domain rather than in the Vertical Industry one. Lifecycle management capabilities are consequently more limited than in a traditional orchestrator, but they include APIs to expose mobile network information, which overlap with the ones provided by the same 3GPP 5G network through the NEF. For this reason, we consider the ETSI MEC framework (and the MEO) to be an interesting alternative for simple vApps with limited deployment requirements. More complex vApps

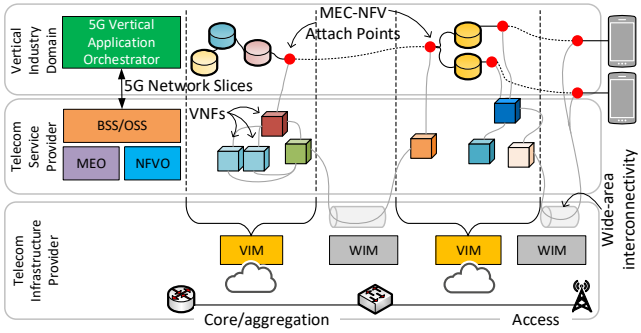

Figure 3. Example deployment of a vApp into a 5G infrastructure, main involved stakeholders and related architectural key building blocks.

(like the ones already running in the cloud) might not fit this framework. Therefore, we decided to omit the MEO presence in the following, by including implicitly its functionalities into the VAO.

(2) User-Plane integration and deployment isolation: as shown in [8] and anticipated in Sect. II.B, vApp components running in different datacenters should be connected to UEs and among themselves. This interconnection should be provided by the NFVO, which sets up and manages NFV services with 3GPP functions towards the UEs, while inter-datacenter connectivity might rely on generic non-3GPP NFV services. As identified also by the ETSI MEC WG [15], the main open issue is to attach the user-planes of vApps and NFV services, which might be hosted in different isolated tenant spaces of VIMs. On one hand, this isolation permits VAOs and NFVOs to work on VIM-level partitions without competition on shared network and computing resources; on the other, also user-plane traffic cannot be exchanged easily between two isolated tenant spaces, making the realization of "*attach points*" between vApps and NFV services [15] non-trivial. In detail, such attach points correspond to virtual networks created in the VIM, which hosts vApp components and VNFs to be interconnected. As shown in [8], 5G specifications allow UPF to locally connect these 3GPP-external networks.

(3) TSP-level information hiding: as outlined also in the 3GPP NEF definition, TSPs might consider topology-related and NFV data (e.g., which VIMs and VNFs are used) as sensitive information not to be disclosed, but exposed in anonymized fashion only. Specific solutions should be undertaken to prevent the VAOs from directly accessing the infrastructure-level control blocks (VIMs and WIMs), and to hide the identities of these blocks. Moreover, the overall architecture should also prevent the VAO from accessing information on third-party execution containers (e.g., virtual machines, Linux containers, etc.) running in the TSP or other vertical industry domains. Even in this respect, a suitable isolation among the stakeholders is crucial.

IV. INTEGRATION ALTERNATIVES

This section introduces two main viable approaches that can be undertaken to integrate 5G vApps into the 5G ecosystem. As far as network awareness is concerned, the VAO can be interconnected with the 5G NEF in both approaches without any particular problems.

IV.A. Vertical applications embedded as VNFs

This architectural scenario corresponds to the one proposed in [15]. In in this scenario, as shown in Fig. 4, the entire edge computing platform is embedded as a VNF in an NFV service (e.g., realizing the 5G Core) in order to easily integrate the vApp and the NFV user-planes. This "edge computing" VNF is composed by multiple virtual machines, each one hosting vApp components (also from different vertical industries). The lifecycle and the interconnection of these virtual machines are managed by the *VNF Manager* (VNFM), which, according to NFV specifications, acts as a sort of "driver" between the high-level commands of the NFVO and specific commands to VNF-internal operations. To perform such internal operations, the VNFM has access to the VIM, and specifically to the same tenant space as the NFVO. To enable the VAO to manage the lifecycle of vApp components (e.g., start a microservice) [15], the edge computing VNF should expose suitable APIs through its *Element Manager* (EM) [14], which, in its turn, can reflect VAO requests to the VNFM, and then to the VIM. This EM can be made reachable by the VAO in a direct fashion (e.g., using a public IP address), or through the OSS/BSS. Both the edge computing VNFM and the EM are meant to be unique for multiple VAOs/vApps. Further details are in [15]. The main advantages of this approach are related to the user-plane integration and to the hiding of the TSP topology/infrastructure information.

Regarding user-plane aspects, the edge computing system being embedded as a VNF, the realization of attach points among vApps and VNFs coincides with the NFV-standard attachment procedure between VNFs. The hiding of instrastructure information is intrinsically achieved, since the resources are not directly exposed by the VIM(s), but through the EM of a VNF contextualized into the abstracted topology of an NFV service.

The disavantages might be potentially numerous. Firstly, this VNF is not part of the 3GPP 5G ecosystem; so, in rigorous terms, it shall not be part of a 3GPP NFV service. Secondly, there is no VIM-level isolation among the vApps' components and the TSP's VNFs. This means that, depending on the NFV service, also the vApp components deployed in the shared tenant space might have user-plane network access to third-party virtual machines (hosting other vApps or VNFs). Also quotas of VIM resources become

shared among the TSP and all the hosted vertical industries. Similarly, if not correctly handled by the EM or the BSS/OSS, a VAO can access the list of all the virtual machines running in the shared tenant space, including the ones hosting TSP VNFs or third-party vApps. Therefore, this solution might not hide TSP-level information. Moreover, to the best of the authors' knowledge, no EM currently provide similar capabilities.

Further key drawbacks concern the south-bound interfaces of the VAO. State-of-the-art VAOs have direct interfaces to VIM APIs, to run and configure microservices. To attract these vertical industries, this integration solution should provide the same capabilities. Therefore, EMs or BSS/OSS should expose VIM-like interfaces to the VAO(s), through handling the aforementioned isolation issues. Moreover, to provide access to microservices, the NFV service should provide "external networks" where the VAO(s) can connect. This can also make the entire NFV service more vulnerable to external attacks. As a final drawback, this approach permits to attach vApp components only to that single NFV service that embeds the edge computing VNF. Therefore, any interconnectivity service requested by the vApp should be part of the same NFV service.

IV.B. VIM-level integration

We devised a novel integration approach, adopted as reference design in the MATILDA project. The idea is to separate the tenant spaces of each vApp and NFV service in each datacenter, so that each orchestrator has its own isolated resources, quotas, external networks, etc. As depicted in Fig. 5, the user-plane integration is performed by using VIM-level virtual networks, and sharing them only between the NFVO tenant space and the one of the specific vApp to be attached. This sharing capability is available in almost all the enterprise-level VIMs and in OpenStack (through rule-based access control policies). The creation of these networks should be performed by the OSS during Day-1 operations: upon the initial selection of the datacenters, the OSS creates the tenant spaces for the vApp, the virtual networks, and shares them towards the created tenant spaces. Then, the OSS requests the NFVO to activate the needed network services. These requests are accomplished by using the shared networks as NFV service "end-points" [14]. If the service is a 5G core, these shared networks are the UPF-attached "external networks" (Sect. II.C), with no need of adding third-party VNFs. Therefore, this approach is fully

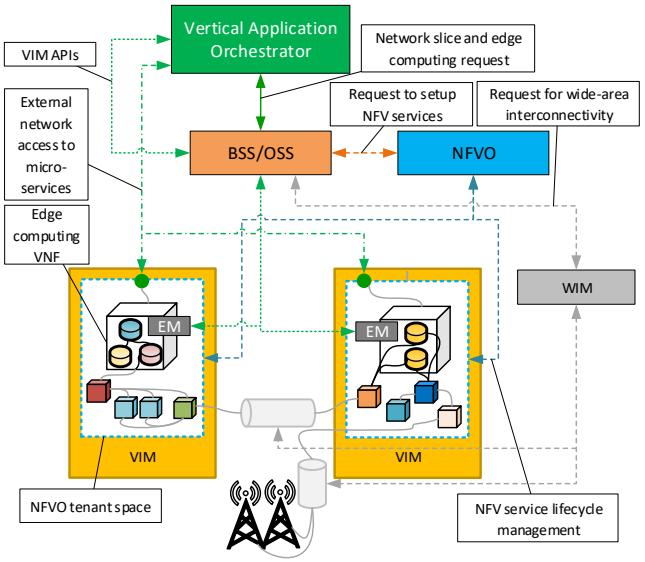

Figure 5. "MEC applications embedded as VNFs" architectural solution.

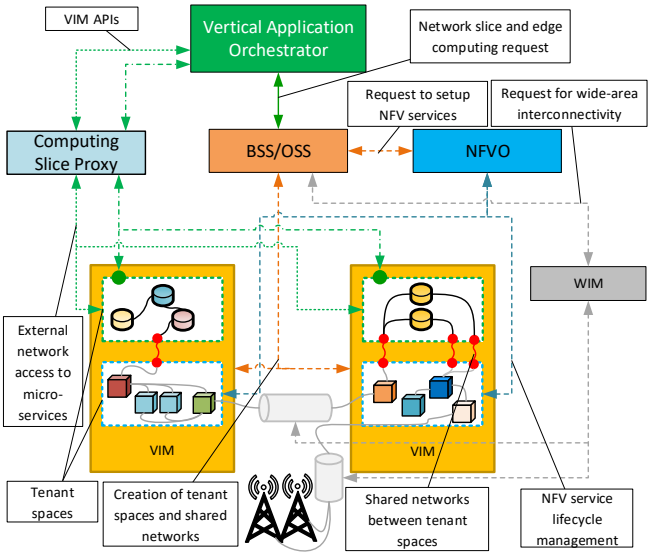

Figure 4. "VIM-level integration" solution.

compliant also with 3GPP 5G specifications.

On the vApp side, there is no main difference from today's cloud scenarios, since the VAO entirely owns its tenant space. Only the shared virtual networks are pre-populated by the OSS. As Day-1 operations complete, the VAO has only to attach front-end vApp components to the shared network in order to have interconnectivity towards UEs and further vApp components in other VIMs. The tenant spaces being completely separated, the "external networks" to enable the VAO controlling the microservices are less vulnerable than in the previous scenario. This because each external network provides access to a single vApp, isolated from the others at VIM level.

Regarding the TSP-level information hiding, this is partially accomplished by using separated tenant spaces, since vApp components can only access TSP's VNF instances through the shared networks, which can be protected with suitable firewall/security rules. To hide also information related to the infrastructure topology, we added a further architectural element in the TSP domain, named "*Computing Slice Proxy*" (CSP). This element has the role of transparently proxying the (REST) APIs between the selected deployment VIMs and the VAOs, anonymizing all the information that could reveal the identities and the owners of the edge datacenters. Therefore, the same southbound interfaces to VIMs used by VAOs in today's cloud scenarios can be maintained without any further changes. A potential drawback of this solution may arise if the TSPs act also as vertical application providers. In this case, they would incur an overhead owing to the presence of multiple tenant spaces.

V. CONCLUSIONS

This paper provided some insights on the integration of vApps within 5G infrastructures. Two possible architectural alternatives to cope with the above problems have been discussed along with their main benefits and disadvantages, which are summarized in Table 1. The analysis suggested that the VIM-level integration approach can better support a smooth porting of cloud vApps into the 5G ecosystem. This could be of paramount importance to make the 5G ecosystem more attractive to those vertical stakeholders running their applications in today's cloud.

ACKNOWLEDGMENTS

This work has been supported by the Horizon 2020 5G-PPP Innovation Action MATILDA (Grant Agreement no. 761898).

Table 1. Advantages and drawbacks of the analyzed approaches.

	Advantages/Drawbacks					
Integration Approach	3GPP 5G User-Plane Integration	3GPP 5G NEF integration	Full isolation among multiple third-party vApps	Need of API extensions in cloud VAOs	full isolation between NFV Services and VApps	TSP physical topology hiding
vApps embedded as VNFs	✓	✓	to be supported at EM/OSS	✓		✓
VIM-level integration	✓	✓	✓		✓	✓

REFERENCES

[1] X. Foukas, *et al.*, "Network Slicing in 5G: Survey and Challenges," IEEE Commun. Mag., vol. 55, no. 5, pp.94-100, May 2017.

[2] R. El Hattachi, J. Erfanian, "NGMN 5G White Paper," Feb. 2015, Online (Last Accessed on 30th Oct. 2018): https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf.

[3] "Mobile Edge Computing (MEC); Framework and Reference Architecture," ETSI GS MEC 003, v. 1.1.1, March 2016.

[4] The 3GPP Association, "Study on Management and Orchestration of Network Slicing for Next Generation Network," Tech. Report 28.801, v. 15.0.0, Sept. 2017.

[5] X. Zhou, *et al.*, "Network Slicing as a Service: Enabling Enterprises' Own Software-Defined Cellular Networks," IEEE Commun. Mag., vol. 54, no. 7, pp. 146-153, July 2016.

[6] M. Maier, *et al.*, "The Tactile Internet: Vision, Recent Progress, and Open Challenges," IEEE Commun. Mag., vol. 54, no. 5, pp. 138-145, May 2016.

[7] B. Rusti, *et al.*, "Smart City as a 5G Ready Application," Proc. 2018 Internat. Conf. Comm. (COMM), Oct. 2018, Bucharest, Romania, pp. 207-212.

[8] The 3GPP Association, "System Architecture for the 5G System," 3GPP Tech. Specification 23.501, Release 15, v. 15.2.0, Jun. 2018.

[9] P. Jamshidi, *et al.*, "Microservices: The Journey So Far and Challenges Ahead," IEEE Software, vol. 35, no. 3, pp. 24–35, May/Jun. 2018.

[10] A. Balalaie, *et al.*, "Microservices Architecture Enables DevOps: Migration to a Cloud-Native Architecture," IEEE Software, vol. 33, no. 3, pp. 42-52, May/Jun. 2016.

[11] F. Giust, *et. al.*, "MEC Deployments in 4G and Evolution Towards 5G," ETSI Whitepaper no. 24, First Edition, Feb. 2018. ISBN No. 979-10-92620-18-4.

[12] "Mobile Edge Computing (MEC); Market Acceleration; MEC Metrics Best Practice and Guidelines," ETSI Group Spec. MEC-IEG 006, v. 1.1.1, Jan. 2017.

[13] "Description of Network Slicing Concept," NGMN 5G P1, Requirements & Architecture, Work Stream End-to-End Architecture, v. 1.0.8, Sept. 2016.

[14] "Network Functions Virtualization, Architectural Framework," ETSI Group Spec. NFV 002, v. 1.2.1, Dec. 2014.

[15] "Deployment of Mobile Edge Computing in an NFV Environment," ETSI Group Report (GR) MEC 017, v. 1.1.1, Feb. 2018.

Roberto Bruschi serves as researcher at the CNIT National Lab. of Smart, Sustainable and Secure Internet Technologies and Infrastructures (S3ITI). His research interests include 5G, MEC, and NFV. Roberto co-authored over 100 scientific publications.

Raffaele Bolla is full professor of telecommunication networks at the University of Genoa. He focuses on green IT and 5G research. He has co-authored over 200 scientific publications.

Franco Davoli is full professor of telecommunication networks at the University of Genoa. His research interests are in 5G and green networking. He has co-authored over 350 scientific publications.

Panagiotis Gouvas is co-founder and R&D director of Ubitech, Ltd. His research interests are in 5G networks, cloud computing, and security. He co-authored over 35 scientific publications.

Anastasios Zafeiropoulos is currently a senior R&D architect with Ubitech, Ltd. His research interests include 5G, green IT, and SDN. He coauthored over 40 scientific publications.