

“Anything as a Service” for 5G Mobile Systems

Tarik Taleb, Adlen Ksentini, and Riku Jäntti

ABSTRACT

5G network architecture and its functions are yet to be defined. However, it is generally agreed that cloud computing, network function virtualization (NFV), and software defined networking (SDN) will be key enabling technologies for 5G. Indeed, putting all these technologies together ensures several advantages in terms of network configuration flexibility, scalability, and elasticity, which are highly needed to fulfill the numerous requirements of 5G. Furthermore, 5G network management procedures should be as simple as possible, allowing network operators to orchestrate and manage the lifecycle of their virtual network infrastructures (VNIs) and the corresponding virtual network functions (VNFs), in a cognitive and programmable fashion. To this end, we introduce the concept of “Anything as a Service” (ANYaaS), which allows a network operator to create and orchestrate 5G services on demand and in a dynamic way. ANYaaS relies on the reference ETSI NFV architecture to orchestrate and manage important services such as mobile Content Delivery Network as a Service (CDNaaS), Traffic Offload as a Service (TOFaaS), and Machine Type Communications as a Service (MTCaaS). Ultimately, ANYaaS aims for enabling dynamic creation and management of mobile services through agile approaches that handle 5G network resources and services.

INTRODUCTION

5th generation mobile networks (5G), also referred to as beyond 2020 mobile communications systems, represent the next major phase of the mobile telecom industry, going beyond the current Long Term Evolution (LTE) and IMT-advanced systems. In addition to increased peak bit rates, higher spectrum spectral efficiency, better coverage, and the support of potential numbers of diverse connectable devices, including machine type communications (MTC) devices, 5G systems are required to be cost-efficient, flexibly deployable, elastic, and above all programmable. The need to lower mobile infrastructure costs and render their deployment flexible and elastic has become critical for the sustainability of mobile operators worldwide, mainly in light of the ever growing mobile data traffic on one hand and the stagnant (rather falling) average revenue per user (ARPU) on the other hand. With current mobile network designs, such required flexibility and elasticity are all but impossible to realize, particularly due to the traditional usage of specif-

ic-purpose networking equipment that can neither dynamically scale with mobile traffic nor be easily upgraded with new functions.

Along with recent and ongoing advances in cloud computing and their support of virtualized services, it has become promising to design flexible, scalable, and elastic 5G systems benefiting from advanced virtualization techniques of cloud computing and exploiting recent advances relevant to network function virtualization (NFV). Indeed, with regard to the latter, and thanks to the numerous advantages it offers in terms of network configuration flexibility, scalability, and elasticity, NFV has emerged as one of the key visions of the future 5G architecture. Researchers in both industry and academia have been designing new architectures for elastically composing and operating a virtual end-to-end network platform, on demand, on top of fragmented physical infrastructures provided by federated networked cloud [1]. Software defined networking (SDN) techniques have been seen as promising enablers for this vision of carrier cloud, which will likely play a crucial role in the design of 5G wireless networks.

Indeed, the 5G architectural design should be as flexible as possible by relying on NFV, SDN, and cloud computing. This flexibility consists in building, on-demand and within very short times [2], instances of virtual network infrastructures (VNIs) along with their composing virtual network functions (VNFs), e.g., packet data network gateway (PGW) and serving gateway (SGW), i.e., to scale up to sudden traffic growth. Meanwhile, the trend toward the decentralization/distribution of the cloud architecture, in the format of cloudlet or the so-called federated networked cloud (i.e., edge or fog computing), should allow the placement of VNFs nearby end-users. This will reduce the end-to-end communication path between mobile users and corresponding servers, which will definitively improve users’ quality of experience (QoE). The 5G network system should also allow mobile operators to create and orchestrate, in an easy and dynamic way, network services targeting a specific optimization of network resources, e.g., caching, traffic offload, or optimization of a specific application. In this vein, we devise in this paper the concept of ANYthing as a Service (ANYaaS), wherein on-demand creation and orchestration of 5G services are specified and enabled exploiting the benefits of both cloud computing and NFV. ANYaaS allows the instantiation of one or multiple mobile network-related VNIs, along with their integration, in order to build a specific 5G service targeting a specific network resource optimiza-

Tarik Taleb is with Sejong University and Aalto University.

Adlen Ksentini is with EURECOM.

Riku Jäntti is with Aalto University.

Digital Object Identifier: 10.1109/MNET.2016.1500244RP

tion. Among the service instances that ANYaaS may launch are the following.

Mobile CDN as a Service (CDNaaS): It creates and allows the management of a service instance of a virtual CDN whereby CDN VNF instances (e.g., IPTV, video on demand, video caches) are instantiated and strategically placed over the federated networked cloud nearby users and supported by smart caching strategies based on, among others, content popularity and viewers' geographical distribution and mobility patterns. The latter may also serve as a trigger for the instantiation of new CDN VNFs so that video content of interest is cached at nearby cache VNFs following the mobility of an individual or a group of viewers, similar in spirit to the concept of follow me cloud (FMC) [3].

Traffic Offload as a Service (TOFaaS): It creates and allows the management of a service instance of a VNI, consisting of local gateway (L-GW) VNF instances (e.g., alternatively VNF instances of small-scale PGW), over the underlying federated networked cloud. L-GW VNFs are assumed to have the ability to selectively offload IP traffic toward specific networks as determined by the policies of mobile operators. Exploiting the benefits of a distributed cloud architecture, L-GW VNFs may be instantiated at cloudlets nearby the radio access network (RAN) to achieve efficient utilization of network resources, similar in spirit to the mobile edge computing (MEC) concept [4].

The envisioned ANYaaS concept may create one single service instance operating individually, or multiple correlated service instances ensuring efficient integration between them. As an example of multiple services, CDNaaS and TOFaaS can be created, in an efficiently integrated fashion, to complement the role of each other. Further details on this use case will be given later.

This paper is organized as follows. We detail the ANYaaS concept and describe its orchestration system and how it is integrated within the ETSI NFV reference architecture [5]. We take the offering of CDNaaS and TOFaaS as a use case to describe in detail the working of ANYaaS. We discuss technological trends that allow the creation of a proof of concept (PoC) of ANYaaS. Using CDNaaS as an example, the section also investigates which virtualization approaches (virtual machine or container) is appropriate to CDN VNF. Finally, the paper concludes.

ANYTHING AS A SERVICE: CONCEPT DESCRIPTION

Carrier cloud reveals a new way to ease the management of mobile networks, adding to them flexibility, elasticity, and programmability features, provisioned as potential 5G services. Carrier cloud is about enabling the on-demand edification of a carrier-grade mobile network on the cloud in an elastic way. In other words, the aim is to enlarge the service domain of cloud computing from the traditional way, whereby storage and computing are provided as a service at data centers, to the provision of mobile connectivity (as well) as a service. In this vein, we introduce the concept of ANYaaS, wherein solutions to optimize 5G resources are presented to mobile operators through a unique framework, allowing the on-demand instantiation and orchestration of services such as dynamic caching of video data

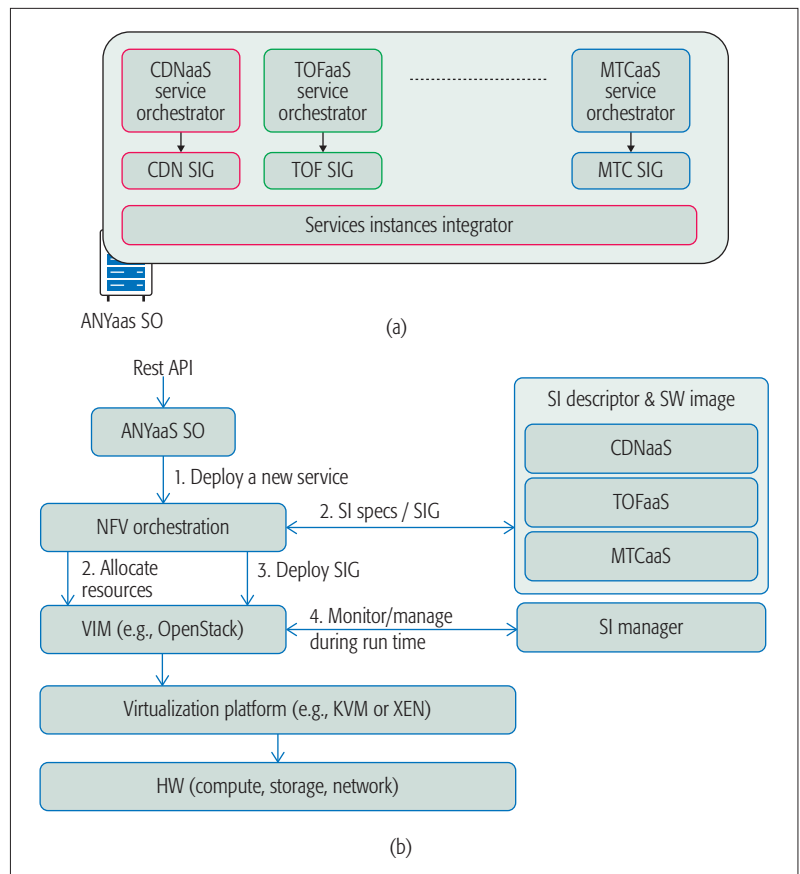


FIGURE 1. ANYaaS concept: a) abstract level architecture of the ANYaaS service orchestrator; and b) management lifecycle of Service Instances.

(CDNaaS), offloading specific traffic (TOFaaS), or building lightweight EPC to handle the traffic of specific applications (e.g., MTC) [2]. An abstract level architecture of the ANYaaS service orchestrator (SO) is depicted in Fig. 1a. The ANYaaS SO comprises a number of specific-purpose SOs that create graphs for their offered service instances (SI), namely service instance graph (SIG). A SI is a single instance of a certain service type (e.g., CDN, MTC, and TOF). It may consist of multiple service instance components (SIC), e.g., cache VNF in case of CDN and L-GW VNF in case of TOF. The SO of each SI gives as output the corresponding SIG, which is a network map of nodes (and links between them) hosting the SICs of a SI, over the underlying federated networked cloud. SI Graphs are formed based on different VNF placement algorithms [6] and taking different metrics relevant to the SI (e.g., caching strategies in case of CDN, L-GW VNF placement in case of TOF considering the relative geographical distances between L-GWs and end-users on one hand, and L-GWs and caches of associated CDN on another hand). For instance, the TOFaaS SO provides a TOFaaS SIG that indicates where to place L-GW VNFs over the underlying federated cloud along with appropriate specifications of the VMs that will be running the L-GW VNFs in order to efficiently offload traffic of a specific application. ANYaaS SO receives monitoring information about the performance of each instantiated VNF to decide on the corresponding SIG. For example, it may instantiate more lightweight vEPC VMs to cope with a sudden increase in MTC traffic. It

may also release or migrate VMs when demand for a specific service changes, to reduce the number of cache VNFs of a CDN SI when the cached content is no longer popular or users interested in the cached content have moved to another location. The ANYaaS SO may be controlled by a specific high-level web interface or by providing a REST northbound API. The provided REST

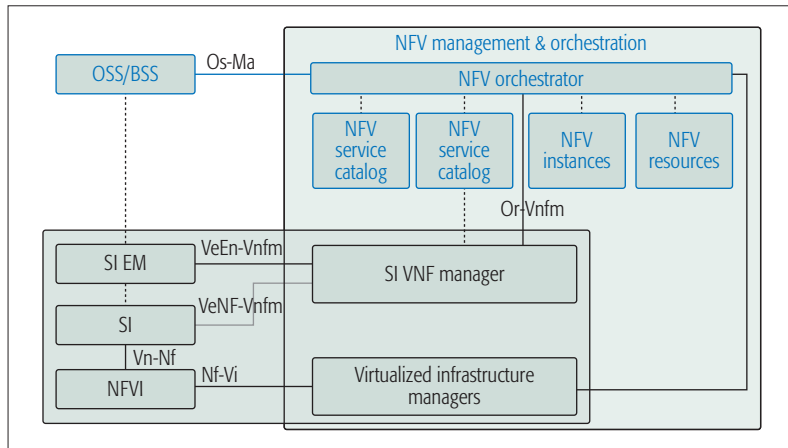


FIGURE 2. Possible integration of the proposed ANYaaS concept within the reference ETSI NFV architecture.

Name	Type	Function
NfV orchestrator	Functional block	Orchestrates NFVI resources across multiple VIMs and does lifecycle management of network services
NfV service catalog	Functional block	Represents the repository of all VNF packages
NfV instance repository	Functional block	Holds information about available/reserved/allocated NFVIs
VNF manager	Functional block	Responsible for the lifecycle management of VNF instances
Virtualized infrastructure manager	Functional block	Control and manage the NFVI compute, storage, and network resources
Os-MA	Interface	Used for data exchange between OSS/BSS and NfV orchestrator (e.g., SI lifecycle management)
Or-Vnfm	Interface	Used for data exchange between NfV orchestrator and VNF manager (e.g., VNF instantiation and NFVI resource allocation for a VNF)
Or-Vi	Interface	Used for data exchange between the NfV orchestrator and VIM (e.g., NFVI resource reservation/release, VNF software image addition/deletion/update)
VeEn-VFm	Interface	Used for data exchange between element management (EM) and VNF manager (e.g., VNF instantiation and VNF instance query)
VeNF-VFm	Interface	Used for data exchange between VNF and VNF manager (e.g., VNF instantiation and VNF instance query)
Vn-Nf	Interface	Used for data exchange between VNF SI and the NFV infrastructure (e.g., southbound API)
Nf-Vi	Interface	Used for data exchange between VIM and NFV Infrastructure (e.g., allocation of VMs with specific compute/storage resources and VM migration)

TABLE 1. Functional blocks and interfaces used in Fig. 2 as per the MANO definition.

API should allow an abstraction of the configuration of the ANYaaS service. Through the service instances integrator, the ANYaaS SO orchestrates several SIs, considering synergies between them and updating, when necessary, their corresponding SIGs. While Fig. 1a shows the case of specific service orchestration (i.e., CDNaas, TOFaas, and MTCaaS), ANYaaS would integrate several other services that aim to optimize network and cloud resources.

The management lifecycle of SIs is performed by ANYaaS SO through interaction with a NfV orchestration system, as schematically depicted in Fig. 1b. A user (i.e., an individual as well as an enterprise user) first communicates with an ANYaaS SO through a web-based interface using the provided REST API. Upon request for the creation of a service, the ANYaaS SO forwards the request including the needed configuration parameters to the NfV orchestrator to create that service instance. Based on the service profile available at the ANYaaS SI portfolio, which includes the configuration and descriptor of each available SI, the NfV SO asks the virtual infrastructure manager (VIM) to allocate resources and deploy the service on the underlying federated cloud. It should be noted that the descriptor and configuration files are service specific. The deployment will be based on the SIG formed by the corresponding SI SO using adequate VNF placement algorithms. The service descriptor also includes the image/software repository of VNFs to be launched. The VIM represents a cloud controller platform, such as OpenStack. The VIM is in charge of the provisioning and deployment of cloud resources, together with the management and monitoring of the deployed SI and its SICs. The SI manager is responsible for the VNF lifecycle, in terms of software update, and scaling and monitoring VNFs throughout the service lifecycle. Monitoring information is made available to the ANYaaS SO and NfV orchestrator through a standardized interface (e.g., an Or-VNFM interface, as depicted in Fig. 2) or through a specific API interface to the ANYaaS SO. Based on this monitoring information, the ANYaaS SO may take decisions on whether to instantiate new SI/SICs or to release resources. It is worth noting that the envisioned architecture of the ANYaaS SO is fully compliant with the ETSI NFV reference architecture as defined by the Management and Orchestration (MANO) working group [5]. Figure 2 shows how the ANYaaS SO can be integrated within the ETSI NFV reference architecture. In the envisioned architecture, the SI manager includes both the SI VNF manager and the SI element manager (EM, i.e. monitoring). Unlike the MANO recommendation, in the proposed architecture, the SI VNF manager is not responsible for scaling a service up or down. It is rather the role of the ANYaaS SO to take this decision and communicate the determined action to implement (i.e., scale up or down) to the VNF manager. The communication between the ANYaaS SO and SI VNF manager may be carried out through a proprietary interface. The NfV service catalog is similar to the descriptor of the ANYaaS SIs. One SI VNF manager is created per each ANYaaS SI. The associated SIG is communicated to the VIM through the Nf-Vi interface. Note that all the reference interfaces are the same as those defined

in the reference architecture. Table 1 summarizes the functional blocks and interfaces used in Fig. 2 as per the MANO definition.

ANYaaS: TOFaaS AND CDNaAS USE CASE

To describe the functionalities of ANYaaS SO and the lifecycle management of its offered service instances, we consider the use case whereby a mobile operator needs to cope with increasing amounts of video traffic. For this purpose, the mobile operator intends to create a cost-efficient mobile CDN whereby popular videos are strategically cached (i.e., CDNaAS). The mobile operator also wants to have requests to popular videos handled by the created mobile CDN, and the corresponding video traffic to be offloaded at traffic offload points nearby its RAN (TOFaaS) to avoid congestion of the scarce resources of its core network. For video caching, several techniques have been proposed to place popular video content at several locations in the mobile network, i.e. RAN (eNB) or P-GW [7, 8]. Depending on their strategies for video placement in the network, caching techniques could considerably reduce the load on the mobile core network. More details on improving mobile network performance through content caching could be found in [8]. Throughout the remainder of this section, we rather focus on how the envisioned CDN and TOF service instances would be orchestrated and managed using the ANYaaS SO.

Figure 3 illustrates how the ANYaaS SO instantiates/deploys a CDN and a traffic offload network as services. While Fig. 3a shows the step of monitoring the resources of the underlying cloud infrastructure and the end-user's geographical distribution and behavior in terms of video viewership, Fig. 3b illustrates the determination of the two CDN and TOF SIGs and their enforcement/orchestration through adequate cloud management platforms (e.g., OpenStack). The ANYaaS SO is assumed to receive feedback from monitoring tools that gather data information about the performance and resource availability of the underlying cloud infrastructure along with the behavior of mobile users. Information on available resources and users' behavior patterns are used to create profiles on users and cloud performance. If a specific traffic, consuming large network resources, is detected, the ANYaaS SO may request from the NFV orchestrator the instantiation of a TOFaaS service instance. This request may also be a follow-up to an explicit request from the mobile operator through the REST API or based on policies pre-configured in the ANYaaS SO. The trigger may also include the traffic pattern, the users' profiles, and their locations over the network. The TOFaaS configuration profile may include, among many parameters, the description of the virtual container or the VM description (i.e. the L-GW VNF), the number of TOFaaS instances to deploy, the specifications of VMs to run the different SICs forming the TOFaaS instances, their placement over the underlying federated cloud, the allocated IP ranges, and the S1-MME interface to enable users to connect to the MME. It is worth noting that:

- The S1-MME interface allows connection to the MME to handle user mobility.

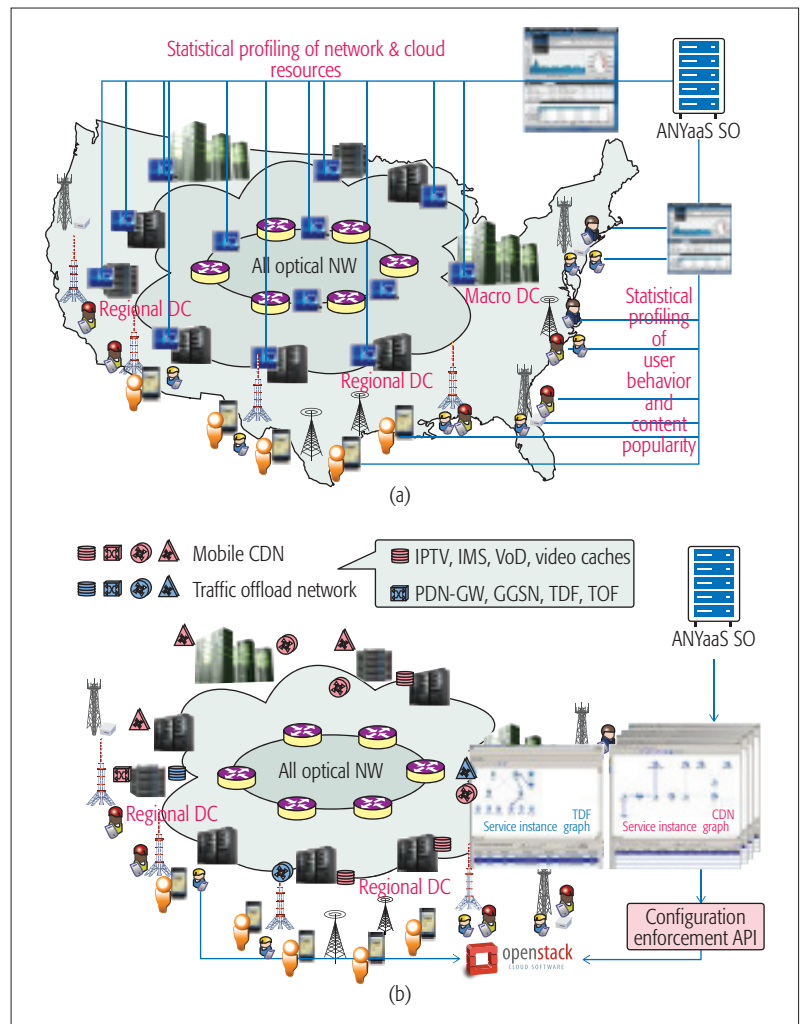


FIGURE 3. Monitoring-based CDN and TOF SIG creation and enforcement in the envisioned ANYaaS SO: a) the monitoring step; and b) the SIG enforcement and deployment step.

- The VNF placement in the network is carried out using specific algorithms such as those introduced in [8, 9].

Finally, the VIM takes over and deploys the VMs in the selected DCs according to the SIG provided in the previous step. The VMs are connected to the network through SDN by creating appropriate routing rules through the southbound API (e.g., OpenFlow). A TOFaaS manager is also created, which is in charge of the lifecycle and the monitoring of the deployed VNF. The TOFaaS informs the ANYaaS SO of any changes in the VNF state in order to scale up or down. Moreover, the monitoring information should help detect the case of VNF failure. The TOFaaS manager is connected to a VIM monitoring tool such as Ceilometer, available in OpenStack. The ANYaaS SO can also trigger the migration of a L-GW VNF from one DC to another, in case traffic has moved to another location, e.g. users are moving from one location to another while connecting to the same service.

The deployment of a mobile CDN and its caches depend on the network operator policies and its business relationship with the content providers. Regarding of the latter, a content provider may install a CDN server in the mobile core net-

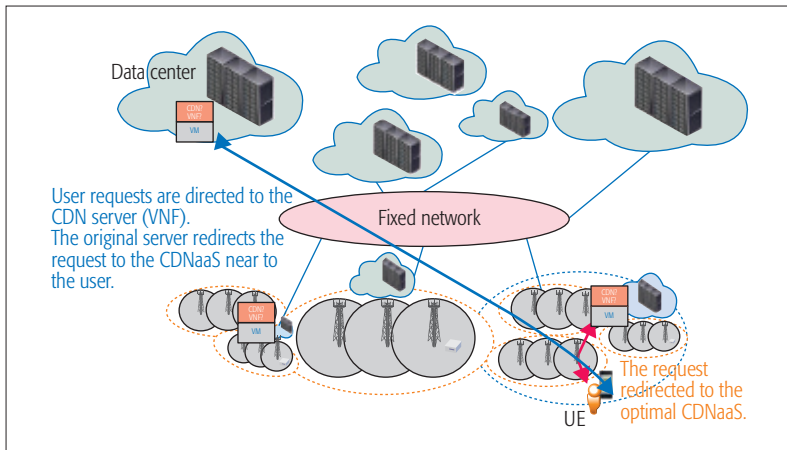


FIGURE 4. CDNaaS in action.

work, e.g. a standalone server or a virtual instance hosted in the cloud (as depicted in Fig. 3). Through DNS (domain name system) redirection, all users' requests are redirected to the local CDN server hosted in the mobile core network. If the server has the requested content, it will subsequently forward the content to the user. Otherwise, it will send the request to the original server and cache the response for future requests. To create a CDN SI, ANYaaS OS follows the same procedure as for the creation of a TOF service instance, making some important considerations. For example, the storage size is more important in CDNaaS than in TOFaaS since it has an impact on the amount of video content that a CDNaaS SIC can cache. Figure 4 shows an example CDN service instance, whereby users' requests are redirected to the closest CDNaaS SIC.

Based on content popularity and its geographic distribution, the ANYaaS OS may specify the locations where CDNaaS SICs have to be deployed. The NFV orchestrator then builds the CDNaaS SI using the configuration descriptor. Based on the number of locations concerned by the CDNaaS SI to deploy, the CDNaaS descriptor specifies the number of VNFs to deploy, their configuration, and their placement over the federated cloud. The descriptor also indicates the CDN software to be used for CDNaaS VNFs. Notable examples of such CDN software are Squid¹ and NGINX.² The VIM then instantiates the VMs and creates the CDNaaS manager to handle the different CDNaaS SICs. The manager agent monitors the CDNaaS SICs and makes this information available to the ANYaaS SO so the latter can react (scale up or down) to any change in the network traffic or to the failure of a VM. Finally, the VMs are interconnected together through a suitable SDN technology (e.g., OpenFlow) building a hierarchical CDN, i.e. if one CDNaaS SIC does not have the content, it can ask another CDNaaS SIC before asking the central CDN server. Once the CDNaaS SICs are active, the CDN server is updated in order to take the new CDNaaS SICs when serving future user requests, as illustrated in Fig. 4.

So far, we have shown the situation where ANYaaS SIs are separately built, independent of each other. However, there are cases where two or more SIs may be correlated, and their respective performances may improve if their respective SIGs are jointly formed rather than being separate-

ly formed independent of each other. Throughout the remainder of this section, we show how the service instances integrator of the envisioned ANYaaS SO can improve the performance of the network when integrating SIGs of both CDNaaS and TOFaaS SIs. For this purpose, we consider the case whereby a mobile operator receives a request to stream a live event over a specific region targeting a specific population of mobile users. Such a use-case could represent a local event, e.g. a music concert not broadcast on TV. This use case is different from the CDNaaS use-case, presented in the previous section, as there is no agreement being established with a content provider, and accordingly no central CDN server is needed. Furthermore, caches should be carefully placed over the federated cloud. As stated in [10], the cost in terms of the communication path in the mobile network is not inversely proportional to the distance of the cache location to the eNodeB. The cost is rather a convex function of the relative distance of the locations of caches to eNodeBs. For instance, according to the presented results, placing a cache at two hops distance from an eNB and associating it to a L-GW represents the best solution. Accordingly, this use-case requires forming CDNaaS and TOFaaS SIGs together to optimize the cache placement. One of the main challenges in this use-case relates to the orchestration of both service instances, i.e., which service instance should be built first or whether the orchestration should create both service instances in an atomic way. In the following, we illustrate how the ANYaaS SO addresses this challenge.

Upon receiving a request from a content provider, the ANYaaS SO sends a notification to the NFV orchestrator to build the respective CDN and TOF SIs. The ANYaaS SO indicates that multiple SIs are needed and integration between these SIs is required. The SI integration is deemed by the ANYaaS according to some logic or explicitly indicated by the content provider or mobile operator. The NFV orchestrator then uses the service descriptor to indicate the creation of integrated services, CDNaaS and TOFaaS. The most critical step is done at the service descriptor and configuration where SIC placement in the federated cloud is decided. The SIC placement algorithm begins by determining where to place the TOFaaS SIC that maximizes the coverage of users who would be interested in the content. Furthermore, L-GWs should be, in general, at a distance from RAN nodes that optimizes the network, as stated in [10]. Upon the instantiation of the TOFaaS SI, the CDNaaS SICs are placed near to the L-GW SICs. After that, the NFV orchestrator builds the SIG map and sends a request to the VIM in order to deploy the SI according to this map. The VIM deploys the VMs to host the different SICs of the two TOFaaS, and CDNaaS SIs create the SI managers in charge of the lifecycle of the two SIs. The SI manager sends the monitoring information to the ANYaaS SO, which can ask to change the SI requirement if the SI placement is not efficient, for instance, if users' QoE is lower due to the fact that SIG is not optimal, i.e. SIs are not placed at the optimal positions or the users have changed location. Figure 5 summarizes the pseudo-algorithm used to decide whether to jointly or separately form SIGs.

¹ www.squid-cache.org

² <http://wiki.nginx.org/NGinx>

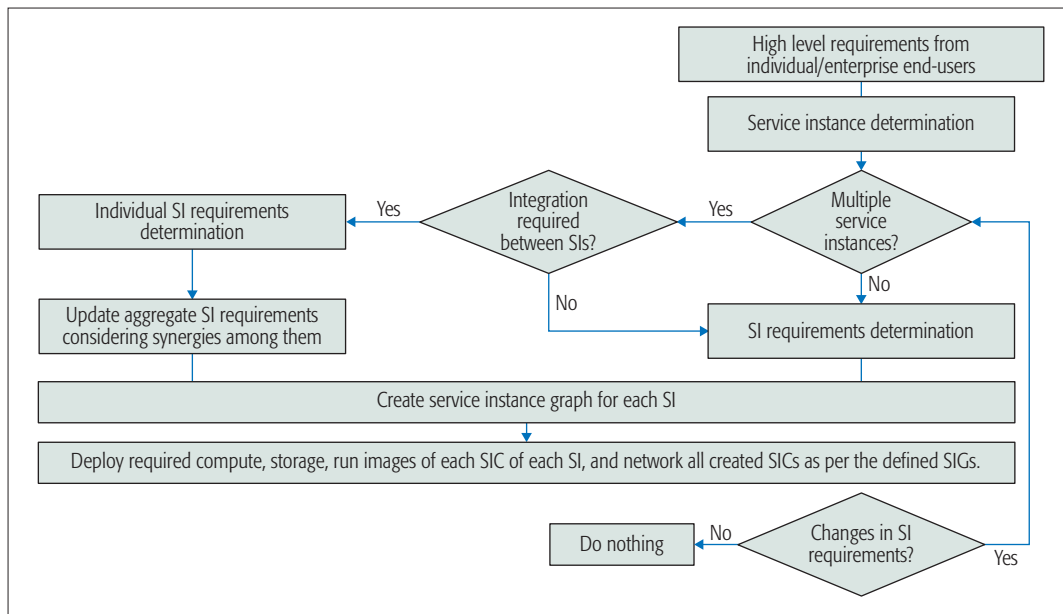


FIGURE 5. ANYaaS SO pseudo-algorithm to decide whether to jointly or separately form SIGs.

ANYaaS: VNF PERFORMANCE CHALLENGES AND SOLUTIONS

IMPLEMENTATION ISSUES

As illustrated in Fig. 3, ANYaaS SO's functional blocks are compliant with the ETSI NFV reference architecture. The envisioned ANYaaS SO uses mainly the same standardized interfaces to communicate with the referenced architectural blocks, e.g., the NFV orchestrator and the VNF manager. The Open NFV³ (or OP-NFV) initiative (under the Linux Foundation) is seeking to create integration projects to have a reference implementation of the NFV as per the ETSI requirement and architecture. The OP-NFV initiative combines upstream codes from projects such as Openstack (i.e., VIM), Opendaylight (i.e., SDN controller), and the Linux kernel, while carrying out extensive testing, custom configuration, and possibly upstream code patches. However, the main challenge remains with the performance of VNFs and their comparison to the performance of purpose-built network nodes. Although ETSI has specified the design patterns of VNFs ([11]), the VNF performance improvement is still an open issue. Indeed, the VNF implementation has to mainly deal with *network input/output* performance in the virtualized environment. This introduces non-trivial challenges when hosting the VNF on top of a hypervisor, as many low-level details, such as memory access patterns, cache locality, task allocation across different CPU cores, and synchronization primitives, may have a dramatic impact on overall performance. An alternative to VMs for running VNFs is the container-based virtualization technology. Containers have emerged as a new technology to isolate and run applications on top of an operating system. In this approach, the operating system's kernel runs on the hardware node with different isolated guest VMs installed on top of it, wherein the isolated guests are called containers. With container-based virtualization, no overhead is associated with having each guest running on a completely installed operating system. While

To create a CDN SI, ANYaaS OS follows the same procedure as for the creation of TOF service instance, making some important considerations. For example, the storage size is more important in CDNaaS than in TOFaaS since it has an impact on the number of video content that a CDNaaS SIC can cache.

this approach could improve performance, since just one operating system takes care of hardware calls, one disadvantage is that each guest should use the same operating system used by the host.

One recent initiative, namely ClickOS [12], aims at building a technology enabler for NFV based on open source tools. ClickOS is a minimal OS (operating system) based on the XEN software platform optimized for middlebox processing, where middlebox refers to all hardware-based network appliances used to run a specific network function, e.g., firewall, intrusion detection system (IDS), and network address translation (NAT). ClickOS includes Click, the software modular router, in order to process packets and act as a router or firewall. As one of the challenges of NFV is the ability to process packets as fast as hardware-based solutions, ClickOS leverages the XEN I/O subsystem by changing the back-end switch, virtual net devices, and back/front end-drivers. The results presented in [5] show that ClickOS is able to forward packets at around 30Gbps, proving that NFV could achieve the same performance as hardware-based solutions. Furthermore, ClickOS is able to boot in only a few seconds. Another emerging and well-established technology that facilitates the development of network I/O intensive applications and hence facilitates the VNF deployment is the Intel Data Plane Development Kit (DPDK). DPDK is a software framework that proposes a set of primitives that ease the creation of efficient VNFs on x86 platforms, in particular high-speed data plane applications. DPDK assumes that processes operate in pull mode in order to become more efficient and decrease the time spent by a packet traveling the server. Indeed, this requires that

³ <https://www.opnfv.org>

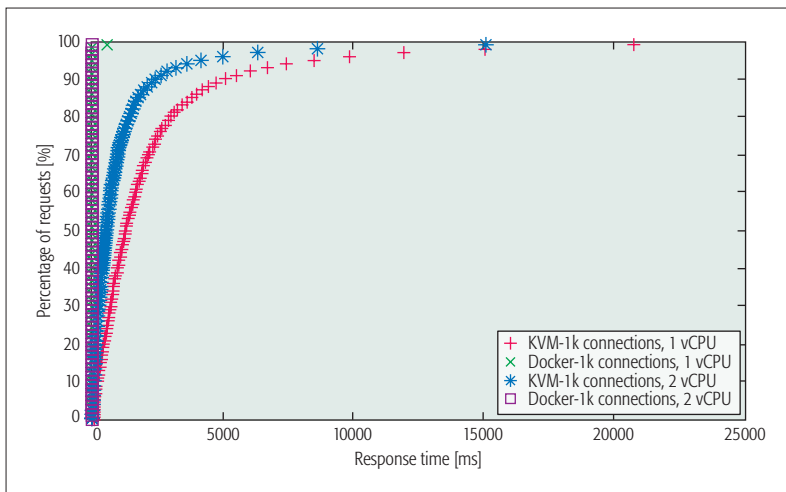


FIGURE 6. Empirical CDF of the response time.

each process needs to occupy one full CPU core. In other words, DPDK processes are attached to a specific CPU core for optimization reasons.

Therefore, putting all these features together, ClickOS and DPDK represent a highly relevant platform to implement ANYaaS services. Indeed, all mentioned functionalities required by ANYaaS services, such as assigning IP addresses, instantiating cache VNFs, and establishing connections to remote MTC servers, can be easily implemented and run on top of ClickOS and DPDK thanks to their ability to manage and forward a high number of packets. Moreover, in order to scale up, ANYaaS services have to be instantiated instantaneously, which could be achieved by ClickOS VMs as they boot in only few seconds.

VIRTUAL CDN BENCHMARK RESULTS

In this section we shed light on results that compare the VNF implementation of a virtual CDN, running a NGINX HTTP server, on top of a container and a VM. Here, the choice of NGINX is motivated by its wide adoption in the CDN industry. For instance, the Netflix platform is based on NGINX software to deliver their video content based on HTTP streaming. In contrast, the container technology is based on Docker⁴ (with Ubuntu 14.04 image), while the VM OS is Ubuntu 14.04, running on a KVM-Qemu hypervisor. The host OS is using Ubuntu OS, running on top of an Intel Xeon (4 CPU core, 2.8 Ghz) with 8 Gbytes of RAM. The conducted experiments consist in emulating a CDN system based on HTTP streaming, whereby several requests are sent to the NGINX HTTP server that hosts the content. Although this experiment does not reflect the whole ANYaaS system, it enabled us to obtain a clear indication on the performance of the VNF hosting the virtual CDN instance. We emulated two scenarios: the NGINX server is run on top of Docker containers; and the NGINX server is run on top of a VM. The shown results focus on two important performance metrics for CDN services: response time and the number of requests handled per second. While the former gives an indication of the quality perceived by a user, as high values indicate high latency to access the service, the latter indicates for the service operator the number of services one VNF can handle.

⁴ www.docker.com

⁵ With Dynamic Adaptive Streaming over HTTP (DASH) technologies, multiple quality/bitrate representations of the same video are stored on a plain HTTP server. A video is segmented in chunks, and a Media Presentation Description (MPD) file describes chunk information. The client receives the MPD file and proceeds by retrieving the video chunk-by-chunk, potentially switching among available qualities.

Figure 6 shows the empirical cumulative distribution function (CDF) of response time; each point represents the percentage of requests that were served in a time less than or equal to a specific value on the x-axis. The number of concurrent connections to the server (VNF) was set to 1000, and the results are obtained when allocating one vCPU and two vCPUs to the VNF; one vCPU represents one single CPU core. The HTTP NGINX server hosts a low-quality video stream, wherein the chunk size is equal to 70 KBytes.⁵ We observe that Docker outperforms KVM, and achieves high responsiveness regardless of the number of vCPUs used. Docker maintains short response times, merely 10 ms and 20 ms, for the case of using two vCPUs and one vCPU, respectively. However, in the case of KVM, 40 percent and 70 percent of the connection requests experienced a response time longer than 1s when using two vCPUs and one vCPU, respectively. It reaches for some connections 15 s, which may be unacceptable for a CDN service. These results are attributable to the fact that Docker incurs less overhead for interacting with the operating system and also for network I/O (i.e., especially given our host-mode configuration, which gives it native access to the host's networking stack).

Figure 7 illustrates the performance of KVM and Docker in terms of the number of user requests handled per second, for different numbers of concurrent connections; high numbers of concurrent connections mean a highly loaded server. Besides serving a low-quality video, we considered in this test also serving a high-quality video (i.e., chunk size equal to 450 Kbytes). Moreover, only one single CPU (1 vCPU) is dedicated to the VNF. As for the response time, Docker outperforms KVM as it handles higher numbers of user requests per second for both video qualities. We observe that Docker accommodates more than 3000 requests per second in comparison to KVM for both qualities. Furthermore, it maintains approximately the same number of handled requests regardless of the number of concurrent connections.

Clearly, from these preliminary results it is highly recommended to use Docker as a virtualization approach to implement VNF for CDNaaS. We explain this by two factors:

- Docker is more lightweight and introduces less overhead when communicating with the host OS.
- Docker is highly suitable to application-based services (e.g., CDN), where there is no need for all low-level components of an OS.

We believe that we might find different results if the tested service requires low-level system components such as routing, which is, for instance, the case with TOFaaS.

CONCLUSION

In this paper we introduced the ANYaaS concept for the upcoming 5G mobile systems. ANYaaS relies on cloud computing and NFV to ease the management of mobile services that aim for optimizing the usage of network resources. ANYaaS allows a mobile operator, in conformance with the ETSI NFV reference architecture, to orchestrate and manage one or multiple service instances on-demand and in a dynamic way. As

a potential use case, we demonstrated how the ANYaaS orchestration system makes possible the creation and deployment of two important correlating services: CDNaaS and TOFaaS. Finally, we discussed key technologies that would enable the development of the ANYaaS proof of concept, and provided experiment results focusing on the performance of CDNaaS. In particular, we found that it is highly recommended to use the Docker virtualization approach to build the CDNaaS VNFs.

ACKNOWLEDGMENT

This work was partially supported by the TAKE 5 project funded by the Finnish Funding Agency for Technology and Innovation (TEKES), and in part by the Finnish Ministry of Employment and the Economy. It is also partially supported by the European Union's Horizon 2020 research and innovation programme under the 5G!Pagoda project with grant agreement No. 723172.

REFERENCES

- [1] T. Taleb, "Towards Carrier Cloud: Potential, Challenges, & Solutions," *IEEE Wireless Commun.*, vol. 21, no. 3, June 2014, pp. 80–91.
- [2] T. Taleb, A. Ksentini, and A. Kobbane, "Lightweight Mobile Core Networks for Machine Type Communications," *IEEE Access*, vol. 2, special section on 5G Wireless Technologies, 2014.
- [3] T. Taleb and A. Ksentini, "Follow Me Cloud: Interworking Distributed Clouds & Distributed Mobile Networks," *IEEE Network Mag.*, Sept. 2013.
- [4] Mobile-Edge Computing — Introductory Technical White Paper, available at https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf, last access 26/11/2015.
- [5] ETSI GS NFV-MAN 001, Network Function Virtualization (NFV) Management and Orchestration (MANO), V. 1.1.1, 12-2014.
- [6] M. Dabbagh *et al.*, "Online Assignment and Placement of Cloud Task Requests with Heterogeneous Requirements," *Proc. IEEE Int'l. Global Commun. Conf. (Globecom)*, Dec 2015.
- [7] X. Wang *et al.*, "Cache in the Air: Exploiting Content Caching and Delivery Techniques for 5G Systems," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 131–39.
- [8] M. Bagaa, T. Taleb, and A. Ksentini, "Service-Aware Network Function Placement for Efficient Traffic Handling in Carrier Cloud," *Proc. IEEE WCNC 2014*, Istanbul, 2014.
- [9] T. Taleb, M. Bagaa, and A. Ksentini, "User Mobility-Aware Virtual Network Function Placement for Virtual Mobile Network Infrastructure," accepted in *IEEE ICC 2015*, London, UK.
- [10] S. Spagna *et al.*, "Design Principles of an Operator-Owned Highly Distributed Content Delivery Network," *IEEE Commun. Mag.*, Apr. 2013.
- [11] ETSI GS NFV-SWA 001, "Network Function Virtualization (NFV) Virtual Network Function Architecture," V. 1.1.1, 12-2014.
- [12] J. Martins *et al.*, "ClickOS and the Art of Network Function Virtualization," *Proc. USENIX NSDI '14*, Apr. 2014, Seattle, Washington, USA.

BIOGRAPHIES

TARIK TALEB [S'04, M'05, SM'10] (tarik.taleb@aalto.fi) received the B.E. degree (with distinction) in information engineering and the M.Sc. and Ph.D. degrees in information science from Tohoku University, Sendai, Japan, in 2001, 2003, and 2005, respectively. Prof. Taleb is a professor at the School of Electrical Engineering, Aalto University, Finland. He was a senior researcher and 3GPP standardization expert with NEC Europe Ltd. He was then leading the NEC Europe Labs Team, working on research and development projects on carrier cloud platforms. Prior to his work at NEC, he worked as an assistant professor at the Graduate School of Information Sciences, Tohoku University. His current research interests include architectural

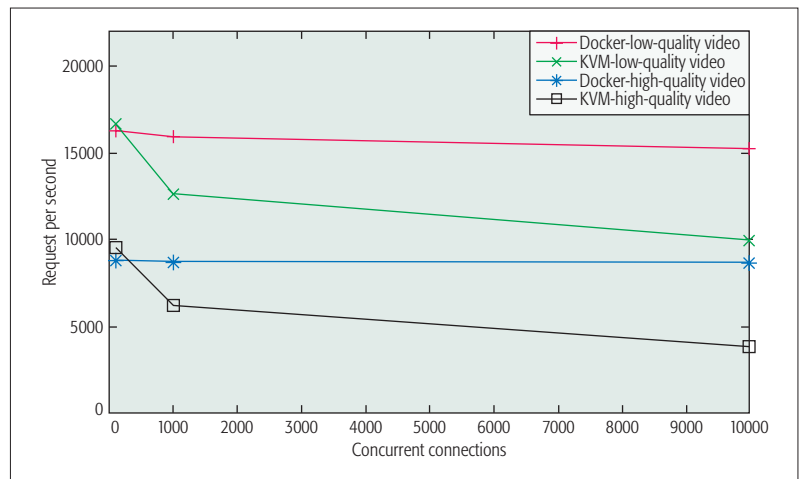


FIGURE 7. Number of requests handled per second vs concurrent connections.

enhancements to mobile core networks, mobile cloud networking, mobile multimedia streaming, and social media networking. He has also been directly engaged in the development and standardization of the evolved packet system as a member of 3GPP's System Architecture Working Group. He is an IEEE Communications Society (ComSoc) Distinguished Lecturer. He is a board member of the IEEE ComSoc Standardization Program Development Board. He is serving as the Chair of the Wireless Communications Technical Committee, the largest in IEEE ComSoc. He founded and has been the General Chair of the IEEE Workshop on Telecommunications Standards: From Research to Standards, which is a successful event that received the "Best Workshop Award" from IEEE ComSoc. He is/was on the editorial board of *IEEE Transactions on Wireless Communications*, *IEEE Wireless Communications Magazine*, *IEEE Transactions on Vehicular Technology*, *IEEE Communications Surveys and Tutorials*, and a number of Wiley journals. He has received many awards, including the IEEE ComSoc Asia Pacific Best Young Researcher award in June 2009. Some of his research work has also received Best Paper Awards at prestigious conferences.

ADLEN KSENTINI [SM'14] received the M.Sc. degree in telecommunication and multimedia networking from the University of Versailles Saint-Quentin-en-Yvelines, and the Ph.D. degree in computer science from the University of Cergy-Pontoise, in 2005, with a dissertation on QoS provisioning in IEEE 802.11-based networks. From 2006 to 2015 he worked at the University of Rennes 1, as an associate professor. During this period, he was a member of the Dionysos Team with INRIA, Rennes. Recently, he joined the Mobile and Wireless Networking Department at the EURECOM Institute as an associate professor. He has been involved in several national and European projects on QoS and QoE support in future wireless, network virtualization, cloud networking, and mobile networks. He has co-authored more than 100 technical journal and international conference papers. He received the best paper award from the IEEE ICC 2012 and ACM MSWiM 2005. He has been acting as TPC symposium chair for IEEE ICC 2016 and 2017. He was a guest editor for *IEEE Wireless Communications Magazine*, *IEEE Communications Magazine*, and two ComSoc MMTC letters. He has been on the Technical Program Committee of major IEEE ComSoc, ICC/Globecom, ICME, WCNC, and PIMRC conferences.

RIKU JÄNTTI [M'02, SM'07] (riku.jantti@aalto.fi) is an associate professor (tenured) in communications engineering and the head of the Department of Communications and Networking at Aalto University's School of Electrical Engineering, Finland. He received his M.Sc. (with distinction) in electrical engineering in 1997, and the D.Sc. (with distinction) in automation and systems technology in 2001, both from Helsinki University of Technology (TKK). Prior to joining Aalto (formerly known as TKK) in August 2006, he was professor pro tem in the Department of Computer Science, University of Vaasa. He is an associate editor of *IEEE Transactions on Vehicular Technology*. His research interests include radio resource control and optimization for machine type communications, cloud based radio access networks, spectrum and coexistence management, and RF inference.