Requirements to Support Cloud, Video and 5G Services on the Telecom Cloud

A. Asensio*, M. Ruiz, and L. Velasco

Optical Communications Group (GCO). Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
*aasensio@ac.upc.edu

***Abstract*—The growth of data traffic related to datacenter (DC) interconnection and the expected explosion of services requiring intensive use of telecom networks, such as video and mobile services, are contributing to the increment of the traffic that transport networks need to support. In fact, Cisco has recently forecast that in the next few years IP traffic related to video services will correspond up to the 80% of the total IP traffic and that mobile data traffic will reach about 367 EB per month by 2020. Aiming at reducing Capital Expenditure (CAPEX) and Operational Expenditure (OPEX) required to satisfy demand growth, service providers are exploring new approaches in contrast to traditional ones. From the transport network perspective, aiming at supporting such traffic growth, telecom network operators need to satisfy services requirements without incurring in network performance degradation and in a cost-effective manner. Therefore, cloud-ready transport networks are on the scope of different network providers, including intelligent control architectures and elastic data planes; i.e. based on novel optical technologies. Considering the telecom cloud, where telecom network resources can be offered as a service, in this paper we present three use cases related to DC interconnection, video distribution and 5G networks, and identify their basic connectivity needs from the transport networks according to realistic scenarios.**

***Keywords*—Cloud Services; Telecom Cloud; 5G Networks; Video Distribution.**

I. INTRODUCTION

Different cloud-based services are expected to co-exist in telecom cloud-based environments: from scenarios where services are hosted in datacenters (DCs) that are owned and operated independently from the network provider but require connectivity from the transport network, to scenarios where services can be hosted in the telecom network provider's facilities, taking advantage of the telecom network infrastructure globally.

Among the different cloud-based services requiring connectivity from the transport network, those contributing to the growth of the amount of data that transport networks need to convey are of particular interest. Two main components of traffic leaving from/arriving to DCs can be distinguished: traffic among DCs (DC2DC) and traffic between DCs and end-users (DC2U). Cisco's Global Cloud Index [1] forecasts DC2DC traffic and DC2U traffic to reach 905 exabytes (EB) and 1.9 zettabytes (ZB) per year in 2019, respectively. Regarding the sources, according to [2], 80% of the IP traffic will correspond to video traffic by 2019. In addition, in 2016 traffic from wireless and mobile devices is expected to surpass traffic from wired devices. In fact, forecast mobile data traffic will reach about 367 EBs per year by 2020.

From the network perspective, transport networks are currently configured with big static fat pipes based on capacity over-provisioning aiming at guaranteeing traffic demand and Quality of Service (QoS). The capacity of each optical connection is dimensioned in advance based on some foreseen traffic demand. However, demands from cloud require new mechanisms to provide reconfiguration and adaptability of the transport network to reduce the amount of over-provisioned bandwidth. *Cloud-ready transport network* [3] was introduced as an architecture to handle the dynamic cloud and network interaction. Moreover, to deal with applications requests for end-to-end (e2e) connectivity provisioning, the IETF has standardized the Application Based Network Operations (ABNO) [4]. ABNO's northbound interface can accept connection requests from the service-layer, facilitating thus dynamic connection requests. In addition, network resources virtualization, allows network operators to support connections' requests from services requiring certain service-specific parameters. Based on ABNO and aiming at facilitating network resources virtualization, the IETF is working on the Abstraction and Control of Transport Networks (ACTN) framework [5].

However, the evolution towards cloud-ready transport networks is based not only on intelligent control architectures (e.g. ABNO-based architectures) but also on elastic data planes that can satisfy cloud requirements efficiently for both network and cloud operators. Hence, elastic optical networks (EONs) play an important role to support cloud services.

In this paper, we focus on identifying a set of basic connectivity requirements from cloud-services to support them on the telecom cloud, independently from the underlying optical network technology. To that end, we study three realistic use cases where different needs may arise. These use cases are related to: *i*) DC interconnection, *ii*) video distribution and *iii*) future 5G networks.

The remainder of this paper is organized as follows: section II introduces our previous work related to the cloud-services that are considered in the use cases. Section III describes the use cases and the related background and in section IV connectivity requirements are devised for each use case from the results obtained in our previous works. Finally, section V concludes the paper.

The research leading to these results has received funding from the Spanish MINECO SYNERGY project (TEC2014-59995-R) and from the Catalan Institution for Research and Advanced Studies (ICREA).

978-1-5090-0296-2/16/$31.00 ©2016 IEEE

II. RELATED WORK

In [6] and [7], we studied services in DC federations that require connectivity from the transport network. Specifically, in [6] a *follow-the-work* approach was considered, where a certain volume of data needs to be transported among DCs aiming at placing workloads close to the users. In addition, since energy expenditure is a predominant part of the total Operational Expenditure (OPEX) for DC operators, in [7] we studied scheduling algorithms trying to minimize energy and communication costs while guaranteeing good Quality of Experience (QoE) to end-users.

In addition, IP-based services could also take advantage of the telecom cloud not only requiring connectivity from the network, but also from resources hosted in the telecom network provider's facilities, e.g. traditional central offices. Although common connectivity requirements may be found for those services, Service Level Agreement (SLA) parameters need to be considered for certain of them. Among the different IP-based services, in our previous works in [8] and [9] we studied live-TV distribution and Cloud-Radio Access Network (C-RAN) [10] to support 5G networks, respectively as services that can be deployed in the telecom cloud and that are contributing to the growth of IP traffic that transport networks need to convey.

From the results obtained in our previous studies, in this paper we identify relevant connectivity requirements to support DC interconnection and services on the telecom cloud.

III. USE CASES

In this section, three use cases representing different scenarios where DCs and cloud-services require connectivity from the transport network are presented.

A. Data Transferences in DC Federations

Placing DCs in geographically diverse locations allows cloud operators to move workloads among their DCs to optimize some utility function while ensuring certain QoE to end-users accessing to the services that they host.

Big Internet companies, e.g. Google, own their infrastructures consisting of worldwide distributed large DCs interconnected through a wide area network and can take advantage from placing workloads in the most appropriate DC [11]. However, there are a large number of smaller independently operated infrastructures, which cannot perform such worldwide workload migration operations. Those medium-sized DCs can cooperate by creating DC federations [12]. DC federations allow them to reduce Capital Expenditure (CAPEX) needed to build large infrastructures and to reduce OPEX, e.g. thanks to moving workloads aiming at reducing energy expenditures. In addition, DC operators can take advantage of the wide area coverage that DC federations provide and bring cloud-based services closer to the users aiming at ensuring a good QoE. In this scenario, connectivity from the transport network is required to manage federated DCs interconnection.

Part of the DC2DC traffic between federated DCs is as a result of elastic operations in the cloud; workloads, encapsulated in virtual machines (VM), can be moved from one DC to another. In addition, another source of DC2DC traffic is for synchronization among databases hosted in the different federated DCs.

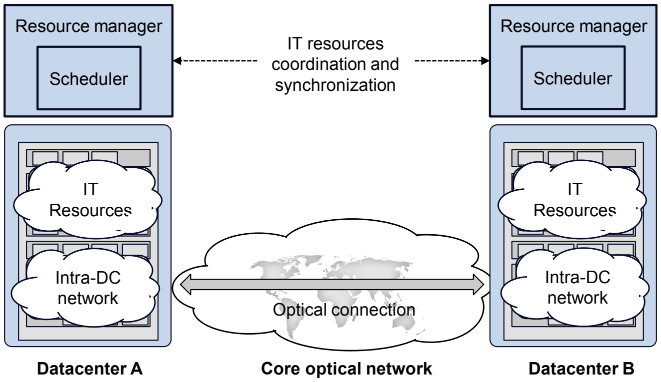

Fig. 1. Scenario with two federated DCs requiring connectivity from the transport network.

The authors in [13] presented the details of VM migration aiming at providing dynamic resource management. Among others, a periodic approach was discussed considering that a given DC, placed in a certain region of the world, may be highly loaded during daytime hours, whereas during the night hours its load may be considerably low. In line with this idea, a basic scenario requiring inter-DC connectivity can be easily identified considering a follow-the-work strategy in a DC federation. In this case, VMs are moved to DCs closer to the users when they need to be accessed (i.e. working hours), reducing the user-to-service latency. Scheduling algorithms in the cloud run periodically taking VM migration and DB synchronization decisions. Once the set of VMs to be migrated and DB synchronization needs are determined, cloud management performs those operations in collaboration with inter- and intra- DC networks. However, tight coordination between the cloud and the network is required since those data migrations need to be performed within the scheduled period to avoid QoE degradation. Fig. 1 represents an example scenario with two federated DCs and their corresponding resource managers, requiring an optical connection for DC interconnection.

In addition, in [13], consolidation was described as a technique that may result in energy savings. DC operators can use green energy and replace either partially or totally energy coming from brown sources. However, green energy is not always available (depending on the hour of the day, weather and season, among others). In contrast, brown energy can be drawn from the grid at any time, although its cost might vary along the day. A first optimization to reduce energy expenditures is to perform consolidation, placing VMs so as to load servers as much as possible and switching off those servers that become unused. Interestingly, DC federations can perform elastic operations, migrating VMs among DCs aiming at minimizing operational costs by taking advantage from available green energy in some DCs and off-peak cheap brown energy in others DCs while ensuring the desired QoE level to the users.

B. *Live-TV Broadcasting*

Video signal distribution is one of the more stringent and popular service that telecom networks need to support. From the broadcasting industry side, video technologies' evolution from standard definition (SD) and high definition (HD) quality formats towards Ultra-High Definition (UHD) formats and the expected growth of end-users' devices consuming video services have led to new paradigms that need to be solved. Among others, those paradigms are related to cost-effective scalable solutions and to the required capacities to transport video signals. Specifically, in live-TV distribution, uncompressed video stream formats are used before the video is produced. Once the video has been produced, distribution to end-users is based on compressed video streams, which quality is adapted to the one that fits better the user's device; i.e., compressed SD, HD or UHD formats. Fig. 2 illustrates an example representing the flow of uncompressed and compressed video streams. In the example, an uncompressed UHD video stream is sent from a source capturing the video to the facilities hosting the processing resources required. Individual compressed video streams are distributed to end-users according to their devices' quality requirements.

However, different drawbacks can be identified when considering traditional Serial Digital Interfaces (SDI) -based transmission environments, among others: *i*) the expensive cost of the specialized hardware required to process video signals and *ii*) the increasing capacities required by novel video formats, since the SDI capacity is limited to HD formats. The first can be palliated by the use of commodity hardware and virtualization techniques; the latter, requires from technologies different than SDI to support the required capacities, e.g. IP-based technologies.

From the network side, the authors in [14] reported 4K UHD TV video streaming over an IP network thus enabling the migration from traditional Serial Digital Interfaces (SDI) -based transmission to all-IP environments. Although they demonstrated that current IP network infrastructure can transport uncompressed 4K UHD flows, the authors in [15] showed that the current IP network infrastructure cannot scale to satisfy expected traffic-growth to support that video quality, being the huge energy consumption one of the main drawbacks; the authors proposed using optical circuit switching to reduce such energy consumption.

Telecom cloud can result beneficial for live-TV services by facilitating both the computing resources hosted in the telecom network facilities for video processing and the network resources required to convey video streams [16]; however, in this paper we focus on the latter to identify connectivity requirements from the transport network.

C. *5G Networks*

5G networks are expected to support a wide variety of mobile services; e.g. related to sensor networks and UHD video streaming in high-speed trains. In addition, the forecast mobile traffic growth has lead mobile operators to re-think the traditional Radio Access Network (RAN). C-RAN [10] was proposed to support next generation mobile networks in a cost-effective manner.

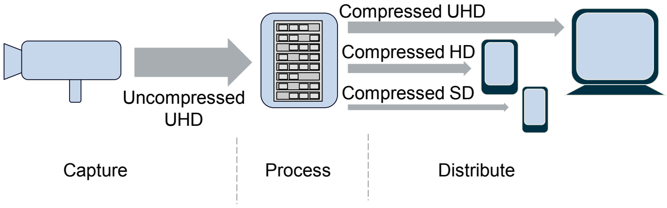

Fig. 2. Example of video streams considered for live-TV distribution.

Different from the traditional distributed RAN architecture, where radio frequency (RF) and baseband processing hardware are co-located in the cell's site, in C-RAN, RF processing hardware, referred to as Remote Radio Head (RRH), is kept in the radio cell, whereas baseband processing hardware, referred to as Base Band Unit (BBU), is centralized in BBU pools. In addition, in C-RAN, BBUs can be shared among different sites along the time and even virtualized [17]. However, C-RAN entails new paradigms related to the fixed access and metro networks supporting the mobile network.

Considering the Long Term Evolution -Advanced (LTE-A) technology, to interconnect RRHs and BBUs in remote BBU pools, connections requiring stringent delay (RTT latency lower than 400 μs is required) and supporting capacities ranging from Gb/s to hundreds of Gb/s need to be considered when the widely adopted radio interface protocol Common Public Radio Interface (CPRI) is used. Therefore, mobile fronthaul based on optical networks have been proposed in the literature (e.g. [18] and [19]). In addition, BBU pools may need to connect to remote BBU pools (through X2 interfaces) for co-ordination purposes and need to connect to the corresponding Mobility Management Entity (MME) and the Serving Gateway (S-GW), through S1 interfaces with control and users' data. To that end, the backhaul network needs to be considered. The required connection's capacity to support the X2 and S1 interfaces is lower than the required in the fronthaul links and their delay constraints are less restrictive, in the order of few tens of ms.

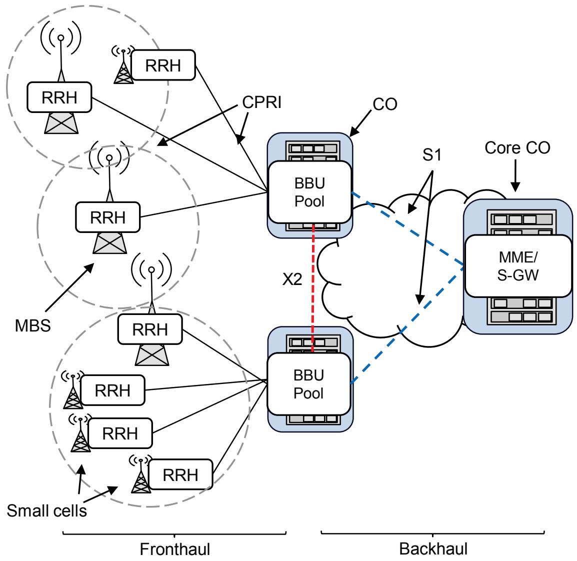

Fig. 3. Example of a C-RAN-based scenario.

According to [20], centralized RAN architectures can be implemented in different variants, including BBU cloud. Considering the cloud variant, virtualized BBU pools are hosted in different central locations and can be flexibly configured. Moreover, RRHs can be served from different virtualized BBU pools each time. In fact, C-RAN pooling gains have been demonstrated under certain cell's traffic assumptions. Specifically, the highest BBU pooling gains have been observed when cells supporting different traffic profiles that vary along the time are considered [21].

Fig. 3 depicts an example of C-RAN, where a set of RRHs corresponding to macro base stations (MBSs) and small cells cover different areas and are served from centralized BBU pools. It is worth noting that, due to the strict RTT latency required in CPRI links, placement of BBU pools is limited to distances lower than 40 km from any RRHs that they need to serve; e.g. BBU pools could be virtualized in central offices (COs) in the network metro segment. However, MME and S-GW could be virtualized in core COs [17]. Connections supporting S1 interfaces may need to be considered in the transport network.

IV. CONNECTIVITY REQUIREMENTS

In this section, we identify a set of connectivity requirements for each use case presented in the previous section that need to be taken into account from the network side. Requirements are devised from realistic scenarios and results in our previous works.

A. Federated DCs Interconnection

As introduced in the previous section, once in operation, scheduling algorithms inside cloud management run periodically trying to optimize some cost function, such as energy costs, and organize data transferences. To ensure that data among DCs is conveyed in the required time, DC operators can contract static connections among their DCs. However, the capacity of such connections needs to be over-dimensioned to guarantee that data transferences finish within the desired time.

To provide realistic capacity values required in such static connections, based the scenario described in our work in [6], Fig. 4a illustrates an example of the static connection's capacity required between two remote DCs for VM migration and DB synchronization targeting at different average time-to-transfer. In the example, scheduling algorithms running each hour are considered. To finish data transferences in less than 10 minutes on average, connection's capacity higher than 400 Gb/s would be required. However, aiming at reducing static connection's capacity over-dimensioning, considering 150 and 200 Gb/s connections results in 30 minutes on average time-to-transfer and would be enough to finish transferences without overlapping with the next scheduling period (i.e. in less than one hour). Taking into account such connection's capacities, Fig. 4b, represents the time-to-transfer for VM migration and DB synchronization. Clearly, data transferences finish in less than one hour at any time, being the peak time-to-transfer values above 50 minutes.

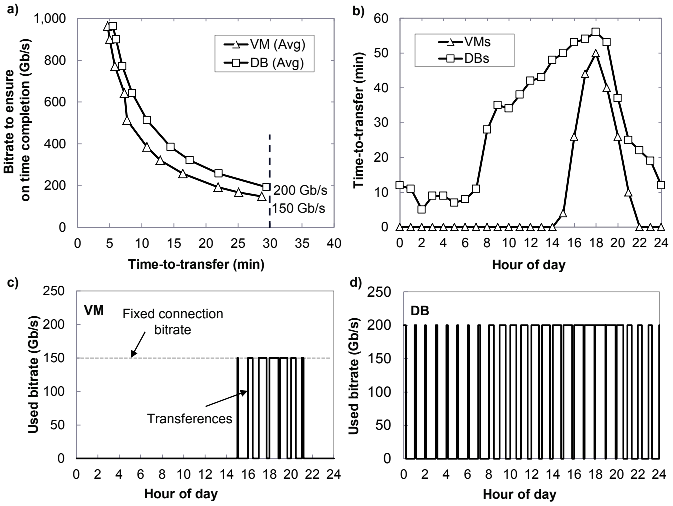

Fig. 4. Bitrate required in static connections against average time-to-transfer (a). Time-to-transfer required along the day (b). Used bitrate in static connections (c and d).

From the connection's capacity values showed above, the first connectivity requirement is devised: *huge capacity* (in the order of hundreds of Gb/s). However, even dimensioning connections to reduce the over-provisioned bitrate, such requirement derives in connection's underutilization when static connections are considered.

To illustrate the underutilization of static connections, Fig. 4c and Fig. 4d show the use of 150 Gb/s and 200 Gb/s static connections, for VM migration and DB synchronization respectively, along the day. It is clear now that such connections are underutilized during low traffic periods. Therefore, trying to minimize the high costs that static connections entail due to their over-dimensioning and underutilization, instead of using static connections dynamic connectivity can be considered. In [6], we found bitrate savings close to 60% in follow-the-work –based scenarios. In addition to energy costs, we showed in [7] that communication costs need to be taken into account when performing elastic operations in DC federations. Impact of the communications costs contributed to motivate dynamic connectivity, which is currently supported by the ABNO architecture.

Dynamic connectivity can then be identified with the following connectivity requirement from the network perspective: *bandwidth-on-demand*. Notwithstanding, some drawbacks arise when considering bandwidth-on-demand instead of static connections. One of the most critical is network resources availability at requesting time.

B. Live-TV Services on the Telecom Cloud

As introduced in section III, live-TV distribution in SDI-based environments is currently limited by HD capacities. In fact, uncompressed video streams in the 4K UHD TV format range from 6 to 48 Gb/s, according to ST 2036-1 [22], which is much more than the capacities required by HD formats (e.g. from 1.5 to 3 Gb/s). In addition, 4K UHD digital cinema has been standardized and commercialized in the movie industry, and 8K UHD quality is in the roadmap of some operators [15]; uncompressed real time 8K UHD transmission needs 72 Gb/s connections.

978-1-5090-0296-2/16/$31.00 ©2016 IEEE

To study the scalability of UHD live-TV distribution in the telecom cloud, in [8] we compared a centralized approach against a distributed one. Fig. 5 illustrates the two approaches. In the centralized approach (Fig. 5a), an uncompressed video stream is processed in a single DC and flows with aggregated compressed video streams are conveyed through the core network to metro areas with the corresponding switches supporting a number of users, each consuming live-TV services with its own required quality. Differently, in the distributed approach (Fig. 5b), uncompressed video streams need to be conveyed through the core network from a primary DC to a set of secondary DCs, placed close to the user. Uncompressed video streams are processed in those secondary DCs and the resulting aggregated flows with compressed video streams are transported to the metro switches directly through the metro network when possible or using the core network when it is required.

Although the distributed approach resulted in remarkable network CAPEX savings, common connectivity requirements can be found for both approaches. Flows transporting uncompressed and aggregated compressed video streams require *huge capacities*; e.g. in the realistic scenarios we tackled, flows of 100 Gb/s were considered to transport uncompressed video streams corresponding to 8 TV channels requiring 12 Gb/s each or aggregated compressed video streams according to different qualities adoption scenarios. However, different values can be considered, ranging from tens to hundreds of Gb/s depending on several factors such as the proportion of users consuming different video qualities, the aggregation level in flows conveying compressed video streams and the number of TV channels transported, among others. Moreover, although connections supporting uncompressed video streams can be based on static connections, the capacity of the connections supporting aggregated compressed video streams depends on the time, since live-TV consumption strongly depends on the hour of the day, presenting peaks during prime time hours (usually at evening and night hours) and off-peak periods during office and night hours. Therefore, *bandwidth-on-demand* may be required in such connections.

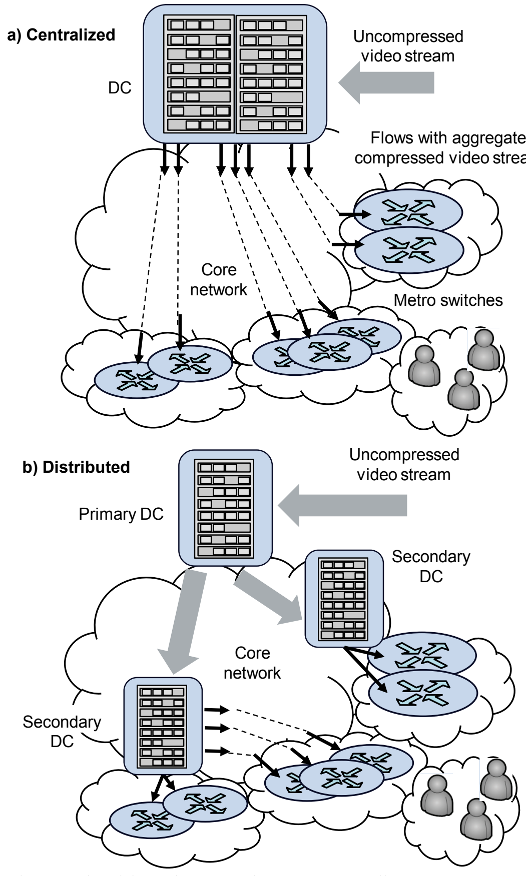

Fig. 5. Telecom cloud-based approaches to support live-TV.

Although the above mentioned connectivity requirements are common with the previous use case, differently from it, live-TV services are delay sensitive (e.g. jitter highly impacts on the service). Therefore, strict QoS requirements, in terms of *delay*, need to be considered for live-TV distribution in the telecom cloud.

C. C-RAN on the Telecom Cloud

To illustrate the C-RAN use case presented in the previous section and devise connectivity requirements from the transport network, we consider a scenario based on our previous work in [9]. RRHs are geographically distributed and cover areas with two differentiated traffic profiles varying along the day: business and residential. Moreover, RRHs corresponding to small cells can be activated or deactivated for off-loading purposes; i.e. to satisfy spikes in the demand at certain hours. In addition, a number of COs can be equipped to support BBU pools.

First, similarly as in [9], let us evaluate a scenario where two COs are equipped, with minimum CAPEX, aiming at serving RRHs along the day. Sites configuration is based on LTE-A 4x4 MIMO. Two cases are distinguished according to the bandwidth considered: 40 MHz and 100 MHz (see [9] for a detailed description). Focusing on the S1 interface (since it requires connectivity from the transport network) Fig. 6 represents the required capacity between each of the two BBU pools and the corresponding location hosting the packet core functions (e.g. MME/S-GW) against the hour of the day. In the example, S1 capacity per each site is about 0.6 Gb/s and 1.5 Gb/s corresponding to the 40 MHz and 100 MHz bandwidth cases described. Results in Fig. 6, show that *huge capacity* is required in both cases, being especially remarkable when 1.5 Gb/s per S1 interface are considered.

In addition, in [9] results showed that, although in the 40 MHz bandwidth scenario, equipping two COs resulted in the minimum CAPEX, in the 100 MHz bandwidth scenario, equipping more COs (each serving less RRHs) resulted not only in CAPEX savings but also in OPEX savings, since COs can serve active RRHs dynamically, allowing to reduce power consumption in those COs not serving any RRH during certain periods of time. Due to the variability on the connectivity and the bitrate required along the time in a given CO, *bandwidth-on-demand* should be considered to support C-RAN from the transport network side.

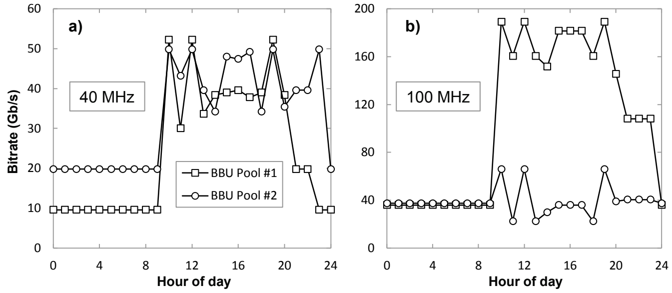

Fig. 6. Peak bitrate required to support S1 interfaces in two BBU pools.

However, service-specific requirements to support 5G mobile networks may arise since 5G needs to support services targeting at their specific performance goals. Therefore, although S1 delay constraints are less restrictive than for CPRI, services in the mobile network may lead to *strict delay* constraints in the connections supporting users' data. In addition, an eventual failure in a link supporting a connection transporting S1 traffic from a BBU pool would impact several sites simultaneously. Considering *bitrate guarantees* by means of diversity in those connections would avoid mobile service interruption in a cost-effective manner from the network provider side.

V. CONCLUSIONS

To devise connectivity requirements aiming at supporting DC interconnection and cloud-services on the telecom cloud, first we have presented three realistic use cases: *i*) DC interconnection in DC federations, *ii*) live-TV distribution and *iii*) C-RAN to support future 5G networks. Then, from the results obtained in our previous works, we have identified, per each use case, the basic requirements needed from the transport network.

Specifically, the following common requirements have been identified for the three use cases: huge capacity and bandwidth-on-demand. However, service-specific requirements have been identified for the live-TV distribution and C-RAN use cases. Live-TV distribution requires strict delay constraints as part of the QoS, whereas C-RAN requires not only strict delay constraints but also certain bitrate guarantees to avoid service interruption simultaneously in several sites, which can impact negatively to a wide range of users and services in next generation mobile networks. Finally, TABLE I summarizes the use cases and their connectivity requirements.

TABLE I. BASIC CONNECTIVITY REQUIREMENTS

	Use case		
Connectivity requirement	***DC interconnection***	***Live-TV***	***C-RAN***
Huge capacity	✓	✓	✓
Bandwidth-on-demand	✓	✓	✓
QoS (in terms of delay)	-	✓	✓
Bitrate guarantees	-	-	✓

REFERENCES

[1] CISCO Global Cloud Index (GCI)

[2] CISCO Visual Networking Index (VNI)

[3] L. M. Contreras, V. López, O. González, A. Tovar, F. Muñoz, A. Azañón, J.P. Fernández-Palacios, and J. Folgueira, "Towards Cloud-Ready Transport Networks," IEEE Communications Magazine, vol. 50, pp. 48-55, 2012.

[4] D. King and A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations," IETF RFC7491, 2015.

[5] D. Ceccarelli and Y. Lee, "Framework for Abstraction and Control of Transport Networks," IETF draft, 2015.

[6] L. Velasco, A. Asensio, J.Ll. Berral, V. López, D. Carrera, A. Castro, and J.P. Fernández-Palacios, "Cross-Stratum Orchestration and Flexgrid Optical Networks for Datacenter Federations," IEEE Network Magazine, vol. 27, pp. 23-30, 2013.

[7] L. Velasco, A. Asensio, J. Ll. Berral, E. Bonetto, F. Musumeci, V. López, "Elastic Operations in Federated Datacenters for Performance and Cost Optimization," Elsevier Computer Communications, vol. 50, pp. 142-151, 2014.

[8] A. Asensio, L. M. Contreras, M. Ruiz, V. Lopez, and L. Velasco, "Scalability of Telecom cloud Architectures for Live-TV Distribution," in Proc. IEEE/OSA Optical Fiber Communication Conference (OFC), 2015.

[9] A. Asensio, P. Saengudomlert, M. Ruiz, and L. Velasco, "Study of the Centralization Level of Optical Network-Supported Cloud RAN," accepted in IEEE International Conference on Optical Network Design and Modeling (ONDM), 2016.

[10] "C-RAN the road towards green RAN," China Mobile Research, 2011.

[11] X. Zhao, V. Vusirikala, B. Koley, V. Kamalov, and T. Hofmeister, "The Prospect of Inter-Data-Center Optical Networks," IEEE Communications Magazine, vol. 51, pp. 32-38, 2013.

[12] I. Goiri, J. Guitart, and J. Torres, "Characterizing Cloud Federation for Enhancing Providers' Profit," in Proc. IEEE International Conference on Cloud Computing (CLOUD), 2010.

[13] M. Mishra, A. Das, P. Kulkarni, and A. Sahoo, "Dynamic Resource Management Using Virtual Machine Migrations," IEEE Communications Magazine, vol. 50, pp. 34-40, 2012.

[14] F. Poulin, T. Kernen, and A. Kouadio, "Ultra High Definition TV Over IP Networks," EBU Technical Review, 2014.

[15] S. Namiki, T. Kurosu, K. Tanizawa, J. Kurumida, T. Hasama, H. Ishikawa, T. Nakatogawa, M. Nakamura, and K. Oyamada, "Ultrahigh-Definition Video Transmission and Extremely Green Optical Networks for Future", IEEE Journal of Selected Topics in Quantum Electronics, vol.17, no.2, pp.446,457, 2011.

[16] M. Ruiz, M. Germán, L. M. Contreras, and L. Velasco, "Big Data-backed Video Distribution in the Telecom Cloud," accepted in Elsevier Computer Communications, DOI: 10.1016/j.comcom.2016.03.026, 2016.

[17] L. Velasco, L.M. Contreras, G. Ferraris, A. Stavdas, F. Cugini, M. Wiegand, and J. P. Fernández-Palacios, "A Service-Oriented Hybrid Access Network and Cloud Architecture," IEEE Communications Magazine, vol. 53, pp. 159-165, 2015.

[18] F. Ponzini, L. Giorgi, A. Bianchi, and R. Sabella, "Centralized radio access networks over wavelength-division multiplexing: a plug-and-play implementation," IEEE Communications Magazine, vol.51, no.9, pp.94-99, 2013.

[19] N. Carapellese, M. Tornatore and A. Pattavina, "Energy-Efficient Baseband Unit Placement in a Fixed/Mobile Converged WDM Aggregation Network," IEEE Journal on Selected Areas in Communications, vol. 32, no. 8, pp. 1542-1551, 2014.

[20] T. Pfeiffer, "Next generation mobile fronthaul and midhaul architectures," IEEE/OSA Journal of Optical Communications and Networking, vol.7, pp. B38-B45, 2015.

[21] A. Checko, H. Holm, and H. Christiansen, "Optimizing small cell deployment by the use of C-RANs," in Proceedings of European Wireless 2014, 20th European Wireless Conference, vol., no., pp.1-6, 2014.

[22] ST 2036-1 "Ultra High Definition Television —Image Parameter Values for Program Production," SMPTE, 2009.