EBU

OPERATING EUROVISION AND EURORADIO

TR 079

BANDWIDTH REQUIREMENTS FOR BROADBAND DISTRIBUTION OF AUDIOVISUAL MEDIA SERVICES

Technical Report

Geneva
February 2024

Document History

EBU Committee	TC	
Drafting Group	Strategic Programme Distribution	
First published	Feb. 2024	
Revised		

Publication Keywords: Broadband Distribution, OTT, Internet, Network Topology, Bandwidth Requirement, Peak Demand.

Acknowledgement

EBU technical publications are the work of experts from EBU Members, Associate Members and third parties consisting of international standards bodies, industry partners, academic institutions and independent consultants.

Their contribution to EBU technical publications is a very generous act by the individuals concerned and by their employers. The EBU appreciates their efforts and thanks them most sincerely.

This report has been produced with the assistance of the following entities:

EBU Members: Roland Beutler (ARD/SWR), Alan Boyle, Andrew Murphy, Paul Tweedy (BBC), Michael Nugent (ERT), Gerjo Bruntink (NPO), Karl Petermichl (ORF), Vittoria Mignone (RAI).

EBU Project Manager: Darko Ratkaj.

Acronyms and Abbreviations

app	Abbreviation of application (on a smart device)
ATSC	Advanced Television Systems Committee (US DTV committee)
BEP	Building Entry Point
CDN	Content Delivery Network
DAB	Digital Audio Broadcast system
DOCSIS	Data Over Cable Service Interface Specification
DSL	Digital Subscriber Line
DSLAM	Digital Subscriber Line Access Multiplexer
DVB-T	Digital Video Broadcast - Terrestrial system
DVB-T2	Digital Video Broadcast - Terrestrial system, second generation
ISP	Internet Service Provider
IXP	Internet Exchange Point
HD	High Definition (television)
LAN	Local Area Network
OTT	Over-The-Top
PON	Passive Optical Networking
PSM	Public Service Media (organization)
QoS	Quality of Service
SAI	Service Area Interface
SCP	Subscriber Connection Point
UHD	Ultra-High Definition (television)
WAN	Wide Area Network

Contents

Bandwidth Requirements for Broadband Distribution of Audiovisual Media Services

1. Introduction

Public Media Service (PSM) content and services are increasingly being distributed over the Internet by means of over-the-top (OTT) models. Thereby, PSM organizations are following the general trend in the media industry.

In the past, PSM organizations have been offering predominantly linear audio and video services that could be distributed by dedicated broadcast networks such as terrestrial, satellite or cable. However, PSM organizations nowadays are responding to audience expectations for on-demand content and services in various ways such as creating their own online platforms and making extensive use of third-party platforms. Users can download corresponding apps to access video content, audio podcasts and engage with PSM organisations on their social media channels. Most PSM apps also allow access to live-streaming linear programmes.

For the time being, the major share of viewing and listening of PSM content is represented by linear TV and radio services delivered over conventional broadcast networks. Depending on national conditions the ratio between linear and on-demand service consumption of broadcast content currently ranges between 80—95% in favour of linear but with a tendency to increasing on-demand consumption.

PSM organizations are carefully monitoring this development in order to take timely strategic decisions with regards to future distribution of their content and services. Some expect to eventually abandon distribution over broadcast networks in the near future, while others see broadcast networks as a fundamental pillar of their distribution strategy even in the long run.

The debate about future distribution strategies often culminates in statements such as “the overwhelming trend is that media consumption is shifting to the Internet and IP networks; in the future broadcast networks may no longer be needed.” Whether or not such claims hold true is yet to be seen. However, it seems to be important to critically scrutinize whether the exclusive distribution of all PSM content and services, both linear and on-demand, over broadband networks is feasible.

To this end, feasibility has both technical and financial aspects. Firstly, it is necessary to analyse if broadband networks are technically capable of carrying all PSM content and services. Secondly, given their technical feasibility, PSM organizations need to understand the economic implications of purely broadband distribution.

This report deals primarily with the technical capabilities of broadband networks. Economic aspects will only briefly be touched upon as this topic deserves a separate in-depth analysis.

2. Scope

“What if all PSM linear and on-demand services were to be exclusively distributed over the Internet in terms of OTT?”

This report attempts to answer this question from the technical point of view. This means to give some insight into the technical implications regarding the capabilities of the broadband networks which are expected to carry all content and services of PSM organisations to users. Other aspects, such as economic, regulatory, prominence, etc. have been briefly touched upon in EBU TR062 [1] and BPN 127 [2] and are not covered in this document.

Figure 1 gives a high-level overview of the current distribution of PSM content and services.

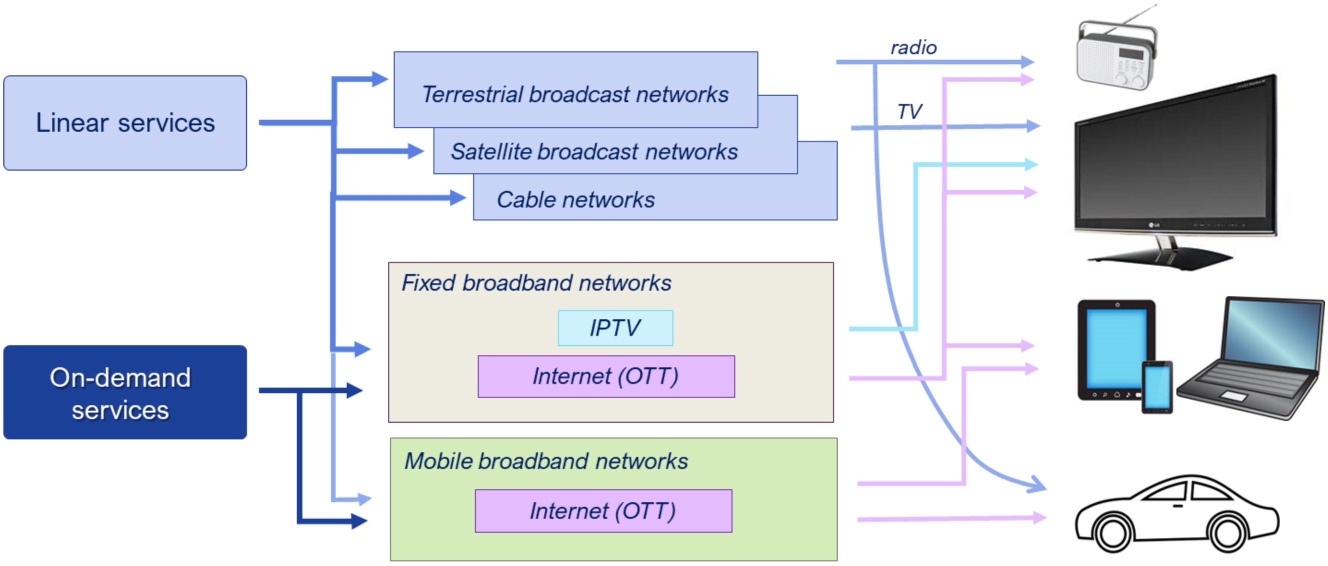

Figure 1: Current distribution paths of linear and on-demand services

Linear services are unidirectional by definition. Hence, only downlink capabilities are required to deliver them. Broadcast networks have been optimized exactly for this purpose. However, linear services can also be distributed over broadband networks as live streams.

On-demand services presume user interaction, which requires bidirectional communications. Broadband networks allow for communication in both directions, i.e., from service provider to user and vice versa. Therefore, broadband networks are natural candidates for delivering on-demand services.

There are instances where broadcast networks are used to push content to user devices, but they have never been commercially successful and therefore will be neglected in this report.

Figure 1 shows another important element. Broadcast networks target dedicated radio and/or TV receivers, e.g., a DAB+ signal can only be received and played back by a device that contains appropriate DAB+ circuitry. The same holds true for TV systems such as DVB-T/T2 and ATSC. In contrast, broadband networks target any device enabled to deal with IP-based signals. This includes multi-purpose devices such as computers, smartphones, and tablets but also many modern radio and TV sets. It is not the case the other way round, i.e., computers, smartphones and tablets do not inevitably have broadcast receivers.

If all PSM content and services were only to be delivered over broadband networks in the future, Figure 1 would condense to what is shown in Figure 2.

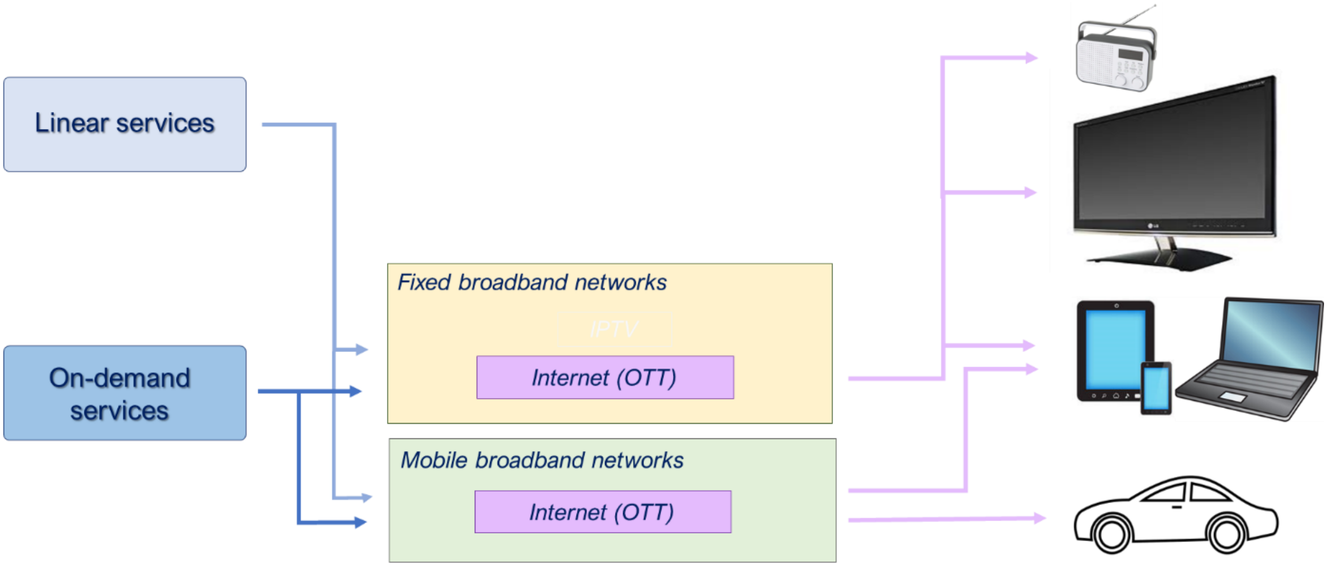

Figure 2: Distribution scenario in an online-only world

The analysis of a distribution scenario as shown in Figure 2 gives rise to several issues, such as:

- How to assess if the bandwidth available in broadband networks, in particular in access networks, would be sufficient to support universal live-streaming of linear TV services at peak times?
- What would be the expected increase of online data traffic resulting from moving linear services to online distribution?
- What would be the daily peak bandwidth requirement to support live TV viewing over broadband networks?
- How to assess exceptional / unforeseen peaks in demand?

The following sections will try to tackle these questions by first giving an overview of the basic Internet distribution mechanisms for audiovisual content and services. The subsequent analysis aims to assess the potential demand and provide some insight into the above questions.

3. Broadband Distribution Chain

Distribution of audiovisual content and services over the Internet is very complex. At a high level its structure can easily be appreciated, however, the more one dives into the detail the more complicated things appear.

On the highest level of abstraction, online distribution can be subdivided into three different domains:

- The Origin, which acts the PSM's online interface to the Internet.
- CDNs replicate content and services from the Origin on their cache and edge servers. Then, CDNs make content accessible to ISPs.
- The third element is ISPs' access networks. These are the networks to which users have a subscription. There are fixed access networks such as DSL, DOCSIS or fibre, but also mobile networks such as 4G and 5G.

Figure 3 illustrates this high-level structure.

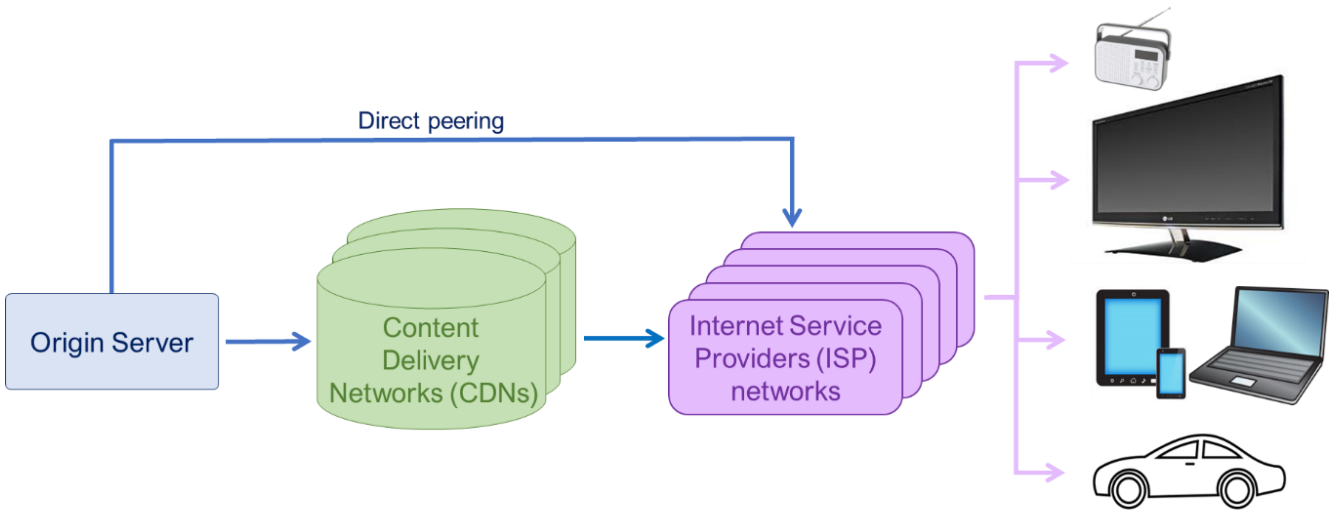

Figure 3: High-level visualization of Internet distribution

Figure 3 depicts several CDNs and a set of ISP networks. In practice, many PSMs employ more than one CDN to balance the traffic, obtain better throughput and potentially contain CDN costs. Also, most countries have more than one ISP that offers access to the Internet. When congestion occurs, it is likely that this will happen in access networks if they do not have sufficient capacity to support the concurrent user demand.

The first complication arises in that some PSMs have direct peering arrangements with ISPs themselves in order to bypass transport through CDNs.

When trying to understand if broadband networks are capable of carrying all PSM content and services it is necessary to delve deeper to unveil more details. Figure 4 gives a more granular overview. In this case the focus is on the access networks and to this end the issue of PSMs using both CDNs and direct peering to deliver their content and services to caches and edges of ISPs has been neglected.

The visualization in Figure 4 is based on the assumption that there is more than one ISP in the geographical area under consideration. Secondly, each ISP supports different types of access networks, i.e., DSL, fibre, and a mobile network. Clearly, this does not necessarily reflect the actual situation in a given country, but in principle, any combination of ISPs supporting different access networks is conceivable. An ISP may only be active in fixed access, i.e., offering DSL and fibre access, or only mobile networks, or both.

The underlying assumption in Figure 4 is that all types of access networks operate in a layered mode. This means that connections ramify from the interconnection points between CDNs and ISPs into some kind of hub from which households are served. The number and density of hubs varies from ISP to ISP and may reflect the number of subscribers in a given geographical area or the differences between cities and the countryside. Moreover, the number of households served by individual hubs varies across hubs.

The network topology shown in Figure 4 is specific for each ISP and in general is not publicly known. Some information may be retrieved from national regulators, for example the number of base stations, i.e., mobile cells. But it may be difficult to obtain information about the number of households served by individual hubs or the capacity available in a hub.

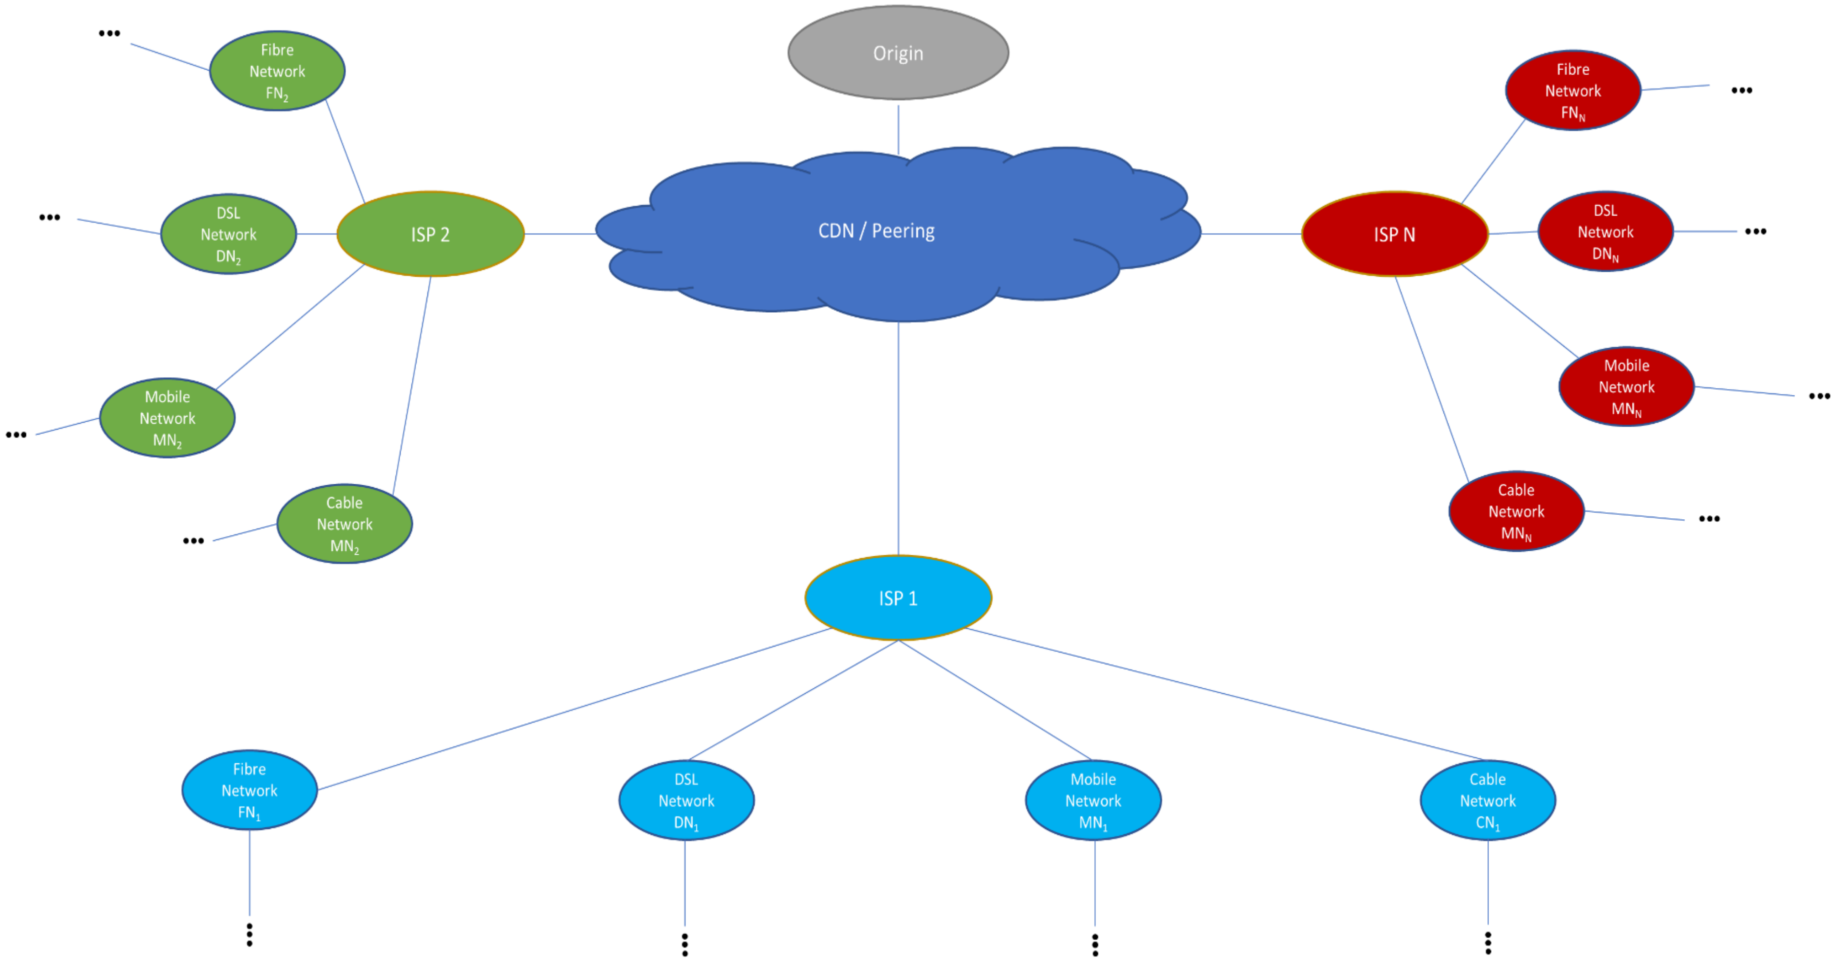

Figure 4: More detailed overview about distribution of audiovisual service over the Internet
The dots indicate the connection to households.

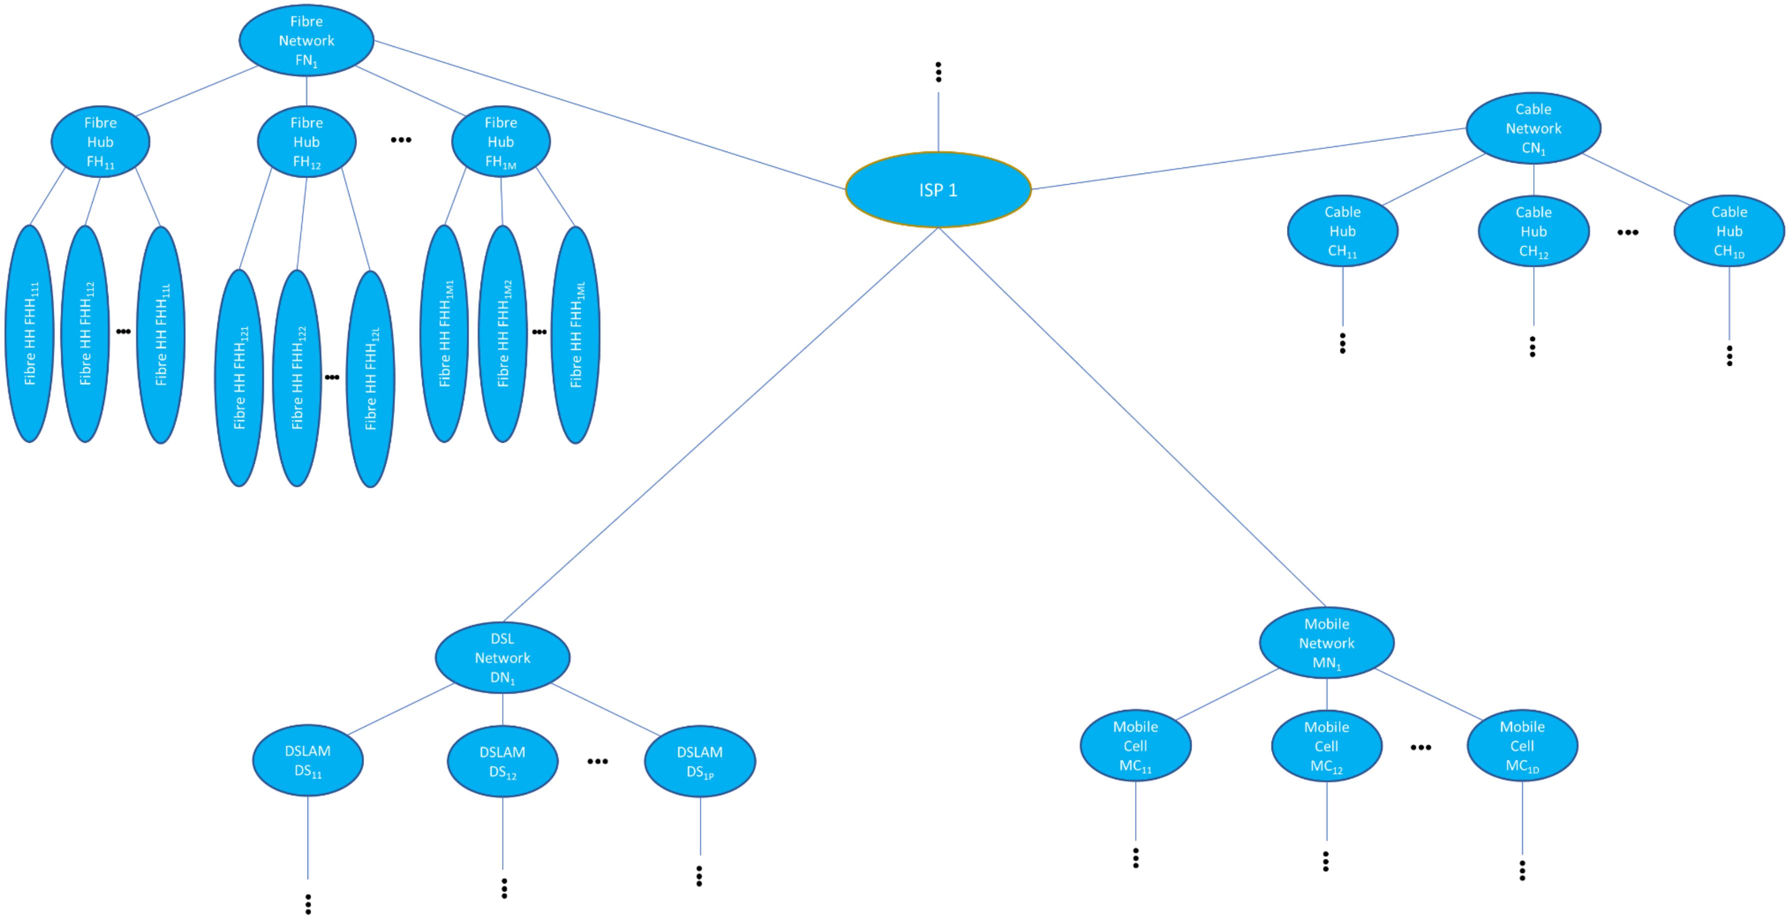

Figure 5: Simplified network topology of the access networks of a single ISP

For the sake of simplicity, the household layer is shown for one fibre network part, only, but needs to be replicated for each hub, DSLAM and mobile cell shown.

Another complication has also been omitted in Figure 4. As shown, it appears as if each ISP operates its own technical infrastructure independent from other ISPs. However, there is a trend that ISPs share their infrastructures to bring down investment and operation costs. This is usually subject to national regulation. Some countries prohibit infrastructure sharing, while others encourage it to boost broadband rollout (in particular) in rural areas.

Figure 5 depicts a simplified topology of Internet distribution under the discussed assumptions. However, it does not say anything about how data is carried from source to device.

Broadband networks are based on IP communication. As such they offer different modes of communication. The default mode is *unicast*, meaning each device is served by a dedicated bidirectional communication session. However, there are also *multicast* and *broadcast* modes that can be configured.

Multicast refers to a situation where the same data stream is delivered to a set of devices that can ‘subscribe’ to it. These devices are all known to the sender. At the same time, a parallel *unicast* connection can be established between individual devices and the source to provide feedback about the quality of transmission.

Broadcast mode is equivalent to what is known from dedicated broadcast networks. A single data stream is sent from one source to all devices which may be in position to receive the signal. They do not need to be known to the transmitting entity. Broadcast traffic is almost always limited to a logical LAN and is never routed over a WAN.

The default way of dealing with an on-demand request in a broadband network is by means of end-to-end unicast communication; this means from user device through all intermediate entities to the source of content and back. Again, this is a simplified description as in reality the content may not be delivered straight from the origin but may be retrieved from some cache server closer to the user. Unicast has traditionally been the only way to reliably deliver content across the Internet.

If many users request the same content or service at the same time, for example a live stream of a football match, a dedicated unicast connection is established for each of them. The total network traffic is the sum of the data transmitted over all unicast streams, even if different streams carry the same content. In case of a popular live event such as a football match this may result in huge traffic. Such traffic peaks could be dramatically reduced if multicast or even broadcast modes be used. While such modes can be deployed, they rely on network-level support that cannot always be guaranteed across involved entities. This is the reason why previous attempts at delivering multicast services over the open Internet to mass audiences across multiple operators have not been successful at scale.

Furthermore, current ISPs' and CDNs' business models are based on unicast. Hence, they may not have the commercial incentive to implement multicast or broadcast.

The business relations in the peering market are not transparent to users such as PSMs. The situation is similar with regards to ISP networks where there is no evidence that multicast / broadcast modes are employed with OTT delivery.

Beyond the strictly technical view, a number of other topics are relevant for broadband media distribution, such as the EU Digital Media Act (regulating gatekeepers), the two-sided-market initiative by telecom operators demanding network use fees to be paid by large content providers, net neutrality, and sustainability. Although these issues are not in the scope of this report, they are highly relevant to the discussion about 'broadcasting versus streaming” and are considered elsewhere in the EBU.

4. Last Mile Network Topology

The situation shown in Figure 5 is already very complex, and yet it does not fully capture the essential characteristics of real-life network configurations. What is called a hub, DSLAM or a cell is a very rough description of the network topology of the so-called “last mile” connecting households. While many users are still waiting for their homes to be connected to fibre networks, it is probably safe to assume that in the core networks and deeper in the access networks, most components are connected by fibre links. Moreover, there seems to be a trend that fibre connections are coming closer to households in Europe, even though it may take many years yet before every household is directly connected to a fibre network. Some may never be, because they are in rural areas where other broadband technologies, such as mobile, satellite or fixed wireless access might be commercially more viable to bridge the last mile.

Generally speaking, the last mile (wired) infrastructure for classical telephone lines could be illustrated as shown in Figure 6 [3].

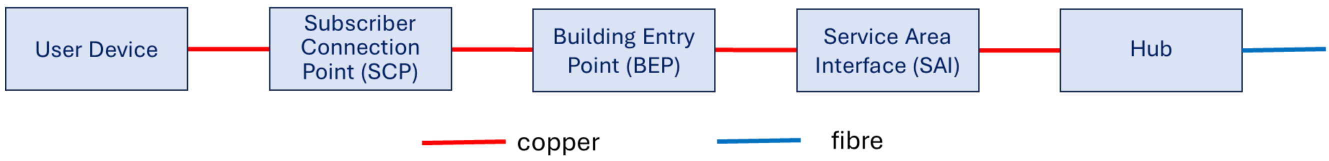

Figure 6: Overview of classical last mile infrastructure elements

Even though the structure depicted in Figure 6 refers to a classical set-up it can be used to visualize the situation for other cases as well. On the right-hand side of Figure 6 there is the hub which could also be a DSLAM, or mobile cell, as shown in Figure 4. This hub serves a number of service area interfaces (SAI), which correspond to the metal boxes at the side of the roads in city quarters. These are connected to the buildings by the so-called building entry point (BEP). From there the line goes to the subscriber connection point (SCP) in the house or flat. This corresponds to the socket where user devices such as routers can be connected. In single-family houses there may not be a difference between the BEP and SCP. However, in apartment buildings they are different. Finally, user devices are reached by means of in-house network infrastructure.

Figure 6 can be used to visualize the different levels of fibre deployments. They are differentiated by the point up to which fibre is available. Figure 7 shows the different configurations [3].

In addition to the complexity of the last mile network structure there is another aspect which needs to be taken into account when trying to explore the capabilities of existing or future access networks. This issue is usually discussed under the label “shared medium”.

Generally speaking, shared medium means that several households have to share the bandwidth of a given access technology. This does not refer to sharing the bandwidth within a single household, e.g., the parents watch a video stream on the large TV in the living room while the kids use their own devices to play some online games in their respective rooms. In this case, obviously the family has to share the bandwidth offered by their Internet contract.

Properly, shared medium refers to multiple households sharing the available network bandwidth. Even though this is a known issue in cable networks, it actually applies to any kind of access network. The only difference is the amount of available bandwidth that has to be shared.

Referring to Figure 6, the crucial element seems to be the SAI in the last mile network chain.

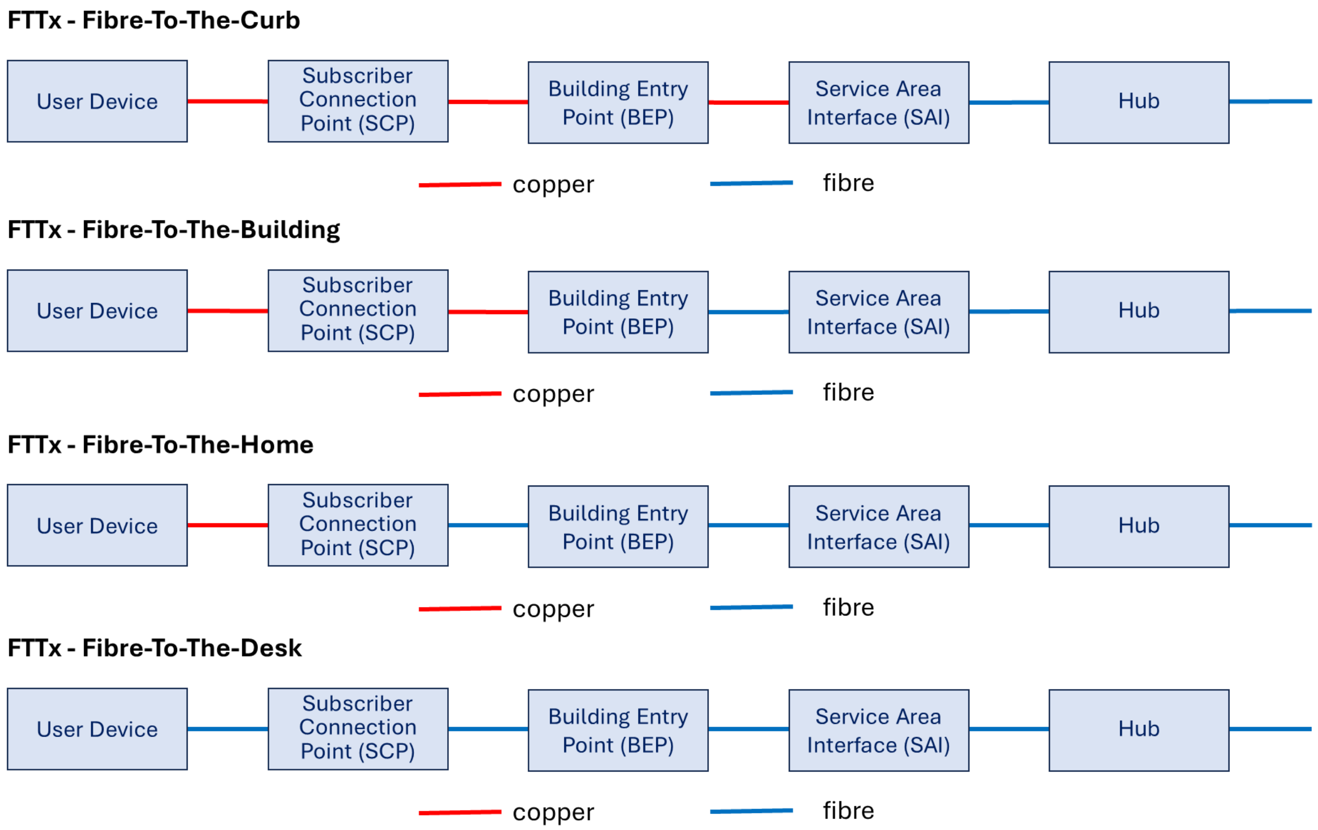

Figure 7: Different levels of fibre deployment

Typically, SAIs offer a total data rate of the order of 1.7 Gbit/s for cable networks or 2.5 Gbit/s for fibre [4]. In cable networks a single SAI currently serves between 100 and 1000 households. If all subscribers use the Internet at the same time the available data rate to each household may drop to values which are insufficient for e.g., streaming HD video content without degradations. Fibre is better in this respect, not only because the total available data rate to be shared is larger, but also because typically only 32 or 64 households are connected to a single SAI.

Connecting more households to an SAI than the accumulated bandwidth that the corresponding contracts would require is called oversubscription. This is justified by the statistical evidence that not all households require the maximum data rate concurrently, and hence there should always be some headroom to accommodate all bandwidth requests. However, as everyday practice shows, this is not always true, in particular in times of popular events such as sport or breaking news. Actually, this issue will become more pronounced and critical in the light of the analysis carried out in this report which tries to illuminate issues arising from moving all PSM content and service exclusively to broadband distribution.

5 Estimation of bandwidth requirements

When trying to decide whether broadband networks are capable of carrying all linear and on-demand PSM services it is necessary to understand real user behaviour and the resulting bandwidth demand.

This demand is not constant over the period of a day. Moreover, there are seasonal variation of usage. During summertime media consumption is usually lower than in the winter. Nevertheless, the daily usage pattern appears to be stable.

Figure 8 shows an example of user behaviour in Germany.

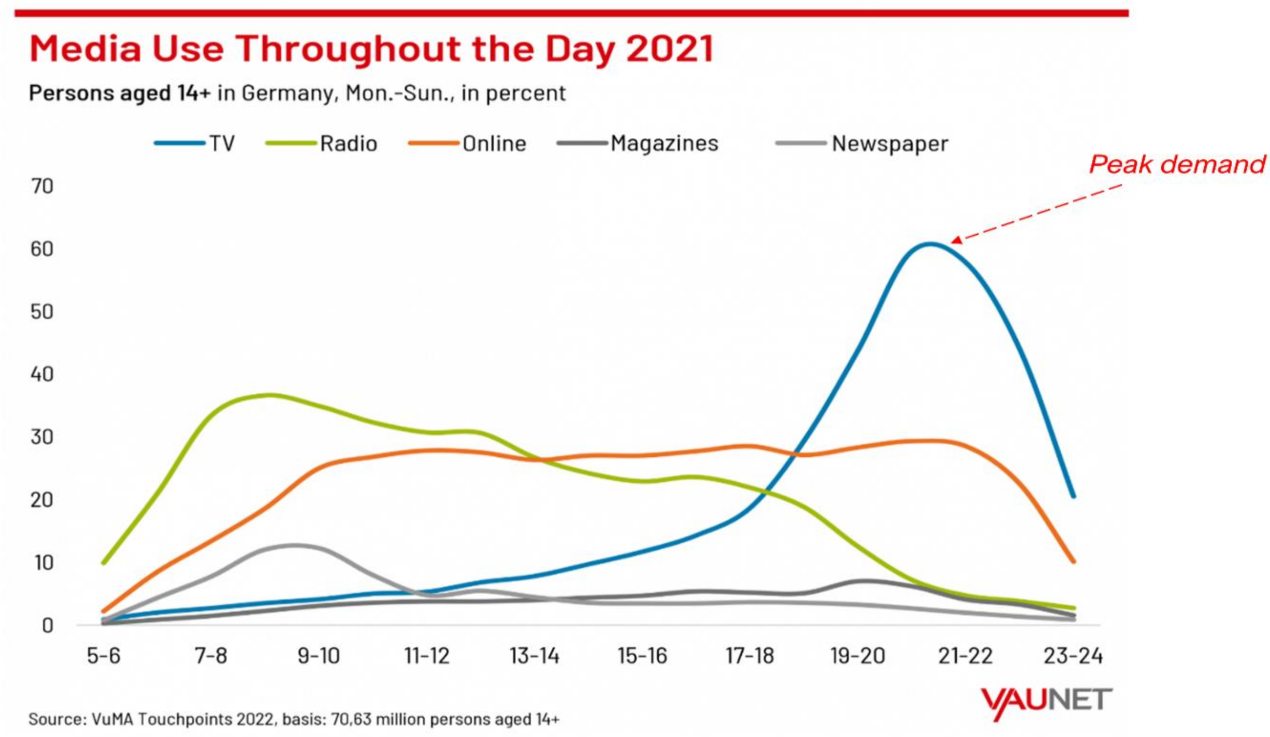

https://www.digitaltvnews.net/?p=37775

Figure 8: Example of different media usage over a day in Germany in 2021 for people aged ≥14

The shape of these curves is quite similar in other countries even though the details may vary. For the analysis presented here the interesting curves are the blue one for TV, the green one for radio and the red one for online usage. The latter can be interpreted as the percentage of people above 14 years of age who are using some kind of online service. This will include all social media usage, access to websites and audiovisual services and PSM audiovisual content and services. It can be assumed that live-streaming of TV and radio programmes is also part of the online consumption.

Radio services shown in Figure 8 obviously correspond to linear radio programmes that are not carried over broadband networks but rather by means of conventional radio broadcast networks such as FM, DAB+, cable or satellite. The same applies to the TV curve in Figure 8 which refers to broadcast networks, in this case DVB-T2, cable and satellite.

As PSM on-demand content and services already make part of the online curve, one could argue that in order to find an answer to the 'what-if' question in § 2 it would be sufficient to look at linear TV consumption taking place today over conventional broadcast networks. This is because the main increase in broadband network traffic would come from linear TV viewing. To a first approximation the impact of linear radio programmes can be neglected as they require substantially lower data rates than TV programmes.

To provide a satisfying user experience for linear TV services, a minimum data rate to each user device is required. Starting from the information contained in Figure 8 it is possible to estimate the total data rate needed to deliver linear TV services to all user devices. As all users in this scenario are served by unicast connections, the maximum total data rate will be reached during peak time of TV consumption which according to Figure 8 is around 21:00.

To estimate the total data rate during peak time the following calculation can be made:

- The German population ≥14 years is 70.9 million [5].
- 62% of them watch TV concurrently at peak time (see Figure 8).

- This is a peak audience of 43.96 million.
- On average there are 2.02 people per household (derived from [6]).
- This means TV is being watched in at least 21.76 million households at the same time.
- If only one device is used per household this means that 21.76 million devices need to be served.
- A typical data rate for HD TV streaming is 3.5 Mbit/s.

This results in a total data rate to serve all these devices concurrently of 7.62 x 10^{13} bit/s, or 76.2 Tbit/s at peak time. This is a huge number.

Actually, it is almost five times the maximum data rate ever reached at the world's largest Internet exchange in Frankfurt, which was 14.4 Tbit/s in 2022 [7]. Furthermore, this figure is additive to the other IP traffic in ISP networks.

Additionally, over the course of a year there are typically large, mostly sporting, events that attract huge audiences and that increase the demand for live transmissions and associated on-demand content. The largest audiences may be substantially higher than the daily peak analysed above.

Broadband networks must cope with these varying requests for content. While the average demand can probably be more or less readily taken into consideration when designing network capacities, a far bigger challenge is to cope with demand peaks arising from large events. However, from a PSM perspective it is important that even under these conditions, reliable delivery of content and services be guaranteed.

The bandwidth estimation model used above can be applied to other countries or regions, provided that the input values are known. It also allows sensitivity analysis with regard to service mix (e.g. HD vs UHD viewing) or QoS requirements.

However, the results of this kind of analysis need to be interpreted carefully. It is important to understand how the traffic is distributed within a given ISP network and across different networks. As has been shown in Figures 4 & 5, no single network or webserver has to serve all users. The question then is how to refine the analysis to reflect the capacity requirements more accurately. A way forward may be formed by examining Figures 4 & 5.

When access networks do not provide the QoS they are supposed to, this may be due to lack of coverage of a given geographical area, oversubscription at the aggregation point such as DSLAMs in DSL networks, or hubs in fibre networks, or too many users in a mobile cell. Hence, the last hop from a given hub to the user devices, as illustrated in Figure 5, is without doubt one of the most crucial bottlenecks.

Looking at the issue from this perspective suggests a need to concentrate on this part of the distribution chain in first instance. As the broadband distribution path from origin to user device ramifies into many concurrent connections (see Figures 4 & 5), it is clear that each of them only has to cope with a very small fraction of the total data rate (derived above) at peak time. However, trying to evaluate the capabilities of access networks this way presumes that information about the number of hubs, DSLAMs or mobile cells and the respective number of served devices is known.

To give an example, if it is assumed that in Figure 5 a given fibre hub is meant to serve 10000 households concurrently, a total data rate of 10000 x 3.5 Mbit/s = 35 Gbit/s has to be provided by the hub. This seems to be much more realistic and feasible even though hubs that actively support 35 Gbit/s are not yet commonly used today.

Another critical issue is the interconnection between CDNs and ISPs or, in case of direct peering, between the media service providers and ISPs. The total capacity at the interconnection points needs to be sufficient to support the peak traffic demand. The available data rates in public Internet exchange points (IXPs) appear to be substantially smaller than the traffic equivalent to the daily peak for linear TV. Most traffic currently goes either through private peering, e.g., between CDNs and ISPs, or from within ISP-embedded CDN caches, for which the information is normally not publicly available.

6. Discussion

In the scenario where all viewing and listening of PSM content and services would occur only on the Internet, the biggest contribution to the increase of the required bandwidth would come from linear TV. This is because the data rates required for radio services are substantially lower than for TV services. The on-demand viewing is already included in the current Internet traffic. Furthermore, as most linear TV viewing occurs at home, the traffic increase will mostly affect the fixed broadband networks rather than mobile networks.

Furthermore, the calculation of the total data rate entailed by distributing all linear TV online as a criterion to decide if broadband networks would be able to cope with the expected traffic increase, seems to be misleading. Deriving global figures for data rates does not help to understand the issues. Rather, the analysis carried out in this report calls for carefully trying to identify the bottlenecks along the entire broadband distribution chain. To this end, the details of the network topology are a primary factor. As discussed in § 5, the capabilities of the hubs of the last mile in the access networks serving households can be seen as one of those bottlenecks. Even though their capabilities still have to be enhanced, the necessary performance boost is not beyond reach in a reasonable time frame from a purely technical point of view.

Moreover, the increase of bandwidth provision over time in response to future demand will be determined by a combination of factors, for example, the increase of encoding efficiency against the take-up of higher-quality services.

Also, there is a concern that contemporary broadband networks would not be able to carry the equivalent of all DTT and satellite broadcast traffic due to capacity issues in the core of the Internet infrastructure.

However, the changes in user behaviour occur over a long period of time and in the past the Internet – both core and access - has grown organically to cope with steadily increasing demand [8]. It is likely that this behaviour will continue in the future and the entities within the content delivery chain, such as ISPs and CDNs, will adapt in due time to growth. In many countries there are massive investments being made in the rollout of fibre access networks which replace old copper networks. This will result in a big increase in the available bandwidth in access networks, in some cases far beyond the need for media services. Furthermore, while video has been a major driver of broadband traffic growth, other types of services and applications may be more significant in the future, such as gaming or the metaverse.

The last 25 years of organic Internet capacity growth has shown that the overall provision of broadband network capacity by the market will match the demand placed upon it by its subscribers. This demand is driven largely by the availability of content that the audience and therefore the subscribers want.

If, however, PSMs would like to transition to Internet distribution faster than organic growth this would entail severe issues with regards to costs, inadequate regulatory framework(s) and access to the audience.

The primary concern relevant to PSMs is having minimum required bandwidth available to all users throughout the entire country, which is not the case in most EBU Member countries as of today.

In a purely commercial infrastructure build-and-operate model, there will always be a percentage of the PSM audience that cannot be reached with high-bandwidth access technologies in an affordable manner, unless regulatory interventions such as direct subsidies or incentive schemes are made available.

Another concern is that the open interconnection/peering points (by which the entities within the OTT value chain exchange traffic) are not necessarily scaling to meet the demand placed on them by increasing consumption of TV services, as core network and peering upgrades are capital-intensive. The risk to PSMs is that this limits the choices available to reach its audiences and retain control over quality.

The open market model for connectivity and content delivery is driven by commercial interests, hence it only prioritizes efficiency improvements that are commercially viable. For example, improving efficiency decreases the temporal pressure to provide capacity growth. On the other hand, PSMs strive to improve efficiency across the distribution chain to better serve their audiences. This can mean, for example, the deployment of caches deeper into the access networks or optimizing the content itself through better distribution encoding.

Another risk to PSMs is that the continued reduction of content delivery costs per Gbyte in the face of increased traffic demand will eventually slow down. There are signs in the market this is already happening, as is provider consolidation, also. Therefore, the market cannot be relied upon to continue delivering the same competitive choice as in the past.

Access to content and services on the Internet is growing independent of PSM engagement. However, in the light of the analysis in this report, PSMs should continue to engage in the development and roll-out of broadband networks to make sure that their requirements are adequately met, in particular with respect to network coverage, costs and access regulation.

7. References

[1] EBU TR062: Over-the-top distribution of PSM content and services

[2] EBU BPN127: Requirements for distribution of PSM content and services

[3] https://www.elektronik-kompendium.de/sites/kom/1403191.htm

[4] https://www.inside-digital.de/ratgeber/shared-medium-kabel-internet-lte-vorteile-nachteile

[5] https://service.destatis.de/bevoelkerungspyramide/#!a=14,100&l=en&g

[6] https://www.destatis.de/EN/Themes/Society-Environment/Population/Households-Families/Tables/households.html

[7] https://www.de-cix.net/en/about-de-cix/media/press-releases/new-data-record-at-europes-largest-internet-exchange-de-cix-frankfurt-breaks-the-sound-barrier-at-14-terabits-per-second

[8] https://www.itu.int/itu-d/reports/statistics/2022/11/24/ff22-international-bandwidth-usage/