Arttu Purmonen, Vice President

CHALLENGES OF VIDEO DELIVERY

Shift to distributed access, Part 1

TELESTE

TERMINOLOGY

Please note that some terms have been simplified in this article. The article deals with video transmission, while being of the opinion that ignoring the irrelevant nuances of terms improves the reading experience. All of these limitations and their possible implications are discussed in the last section, 'Limitations'.

Auxiliary Core A Core that is used to deliver services (for instance: DOCSIS or broadcast video) and manage RPD resources assigned by a Principal Core. Several Auxiliary Cores can co-exist.

CAS Conditional Access System

CCAP Converged Cable Access Platform

CCAP Core We do not use this term in the article, as it may confuse readers. Instead, we use the term 'DOCSIS Core'.

Core This article uses the term Core together with other terms, for example Video Core and DOCSIS Core. In this article, the Core is a component that has a data plane (DEPI/UEPI) and control plane (GCP).

DA Distributed Access refers to any network solution in which the DOCSIS Core does not include Physical Layer functionality specified by Cablelabs. The functionality is in the Remote PHY devices. The white paper does not cover the Remote MACPHY alternative.

DEPI Downstream External PHY Interface

DOCSIS Core In this white paper, a DOCSIS Core is a CMTS platform that has data plane and control plane components. Although a DOCSIS Core is not always a Principal Core, this article assumes that the DOCSIS Core works as well in the role of Principal Core, unless otherwise expressed.

Engine Headend components without GCP capability. However, Engines can support a data plane (DEPI/UEPI).

GCP Generic Control Plane

L2TP Layer 2 Transport Protocol

NDF Narrowband Digital Forward

NDR Narrowband Digital Return

NIT Network Information Table

OOB Out-Of-Band; Certain legacy applications (e.g. FM radio, telemetry) delivered over distributed access networks.

Principal Core Responsible for the initial provisioning and configuration of the common parameters of RPDs; it may manage the resources of Auxiliary Cores and Engines. Several Principal Cores cannot co-exist except for redundancy purposes.

R-DEPI Remote Downstream External PHY Interface

R-UEPI Remote Upstream External PHY Interface

RPD Remote PHY device

Video HE Video Headend, a platform that is used to manage delivery of linear television services. It handles, e.g. multiplexing, encryption and PSI/SI table insertion.

VOD Core A Core that is used to manage and distribute Video-On-Demand services. It includes data plane and control plane components.

INTRODUCTION

Cable operators are enhancing their networks as consumers seek higher broadband speeds due to the attractive services available upon the IP layer. The R-PHY-based distributed architecture might be a tempting alternative for operators eager to take a quantum leap and build networks having substantially higher capacity. But how can cable television operators shift to R-PHY-based distributed access architectures cost effectively while still needing to run old legacy systems during and perhaps even after the transformation?

The question is complex and providing a comprehensive answer would require an entire book. Therefore, we will provide an answer via several white papers; part 1 here covers video transmissions when the goal is a Remote PHY-based network. Several paths lead to the goal and each entails certain pros, cons and options. We aim to paint a holistic picture and give ideas about how these paths differ and how operators can marshal the available options. Instead of claiming that one option is superior to another, our objective is to offer a framework for managing the shift. The shift illustrated in Figure 1 incorporates pitfalls raised by the framework.

The white paper targets cable television executives, network architects, DOCSIS specialists and people who keep linear television services running in the network. These seasoned and savvy players will make difficult decisions in a rapidly changing technology space, one where software is never ready and customers are hard to please. The article focuses on European challenges, but most of the solutions are valid in North America as well. However, we do not discuss Switched Digital Video, as its importance in Europe is negligible.

We hope that part 1 will provide new ideas for every reader. Please, let us know what particular questions part 2 should answer. We need to pave the way in many interesting areas and your feedback will influence the order of the articles.

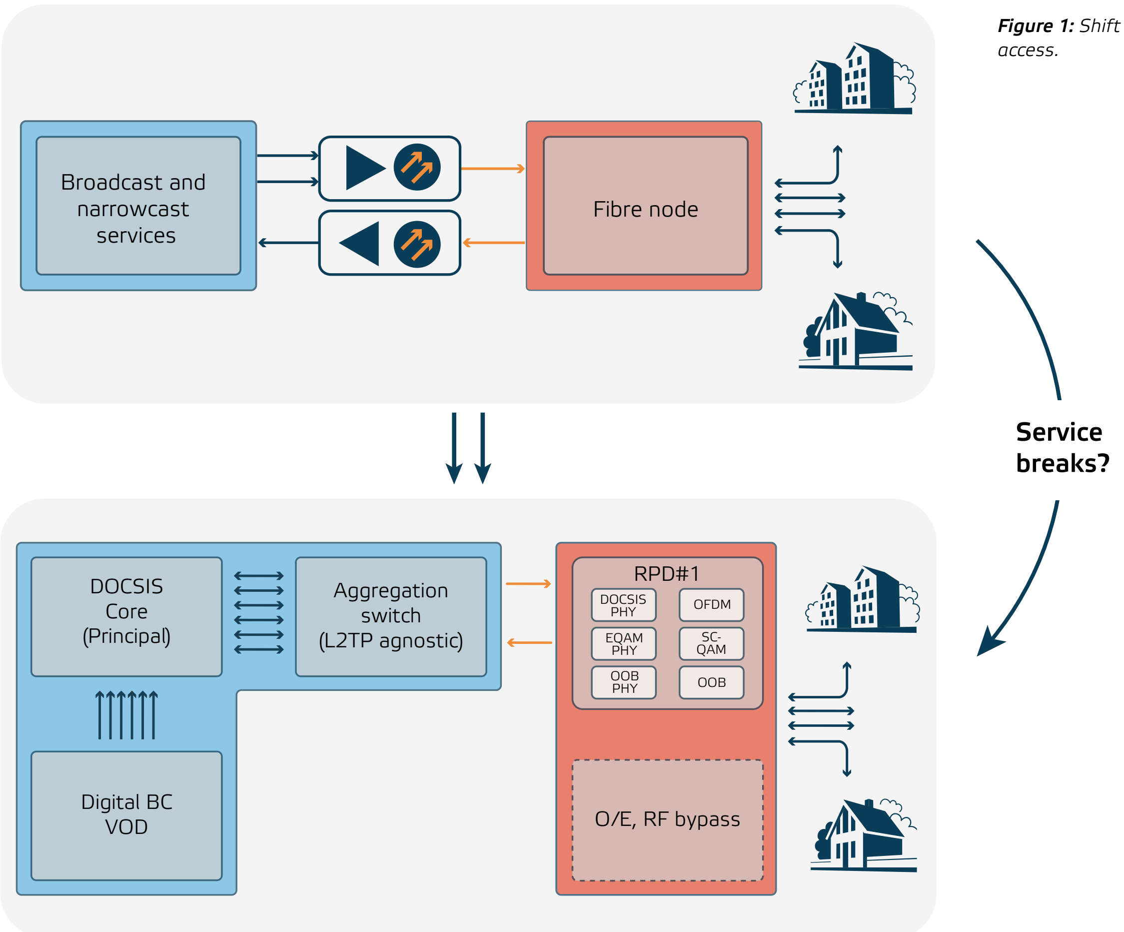

Figure 1: *Shift to distributed access.*

> **Can the distributed access architecture be switched on at the same time that the old infrastructure is still in use?**

CHALLENGES OF VIDEO DELIVERY

It is difficult to find a scheduled service break time slot. Networks are in use 24/7, thus our goal is to answer the question: Can the distributed access architecture be switched on at the same time that the old infrastructure is still in use? The implementation of Remote PHY technology-based access networks will have an impact on the existing Video Headends, existing HFC networks, backbones, network management tools and dozens of operational processes. But that is not all. Now several teams, previously operating their 'own' devices, need to understand each other as new platforms are being shared. If a DOCSIS Core is used for linear television transmission, the frequent software upgrades might surprise linear TV experts who dislike nonlinear behavior during the upgrades. In turn, DOCSIS specialists may find it strange to load Ethernet links with content without knowing if consumers are even watching it.

We dig deeper into such questions as:

- What does offering linear TV services over the Remote PHY-based networks require?
- What kinds of architectural options exist and what are their pros and cons?
- How does IP multicast work over the distributed access networks?
- Can existing broadcast Video Headends feed the distributed access networks, or are new solutions required?
- Is the analogue shutdown mandatory or optional?

Figure 2: *A regional headend.*

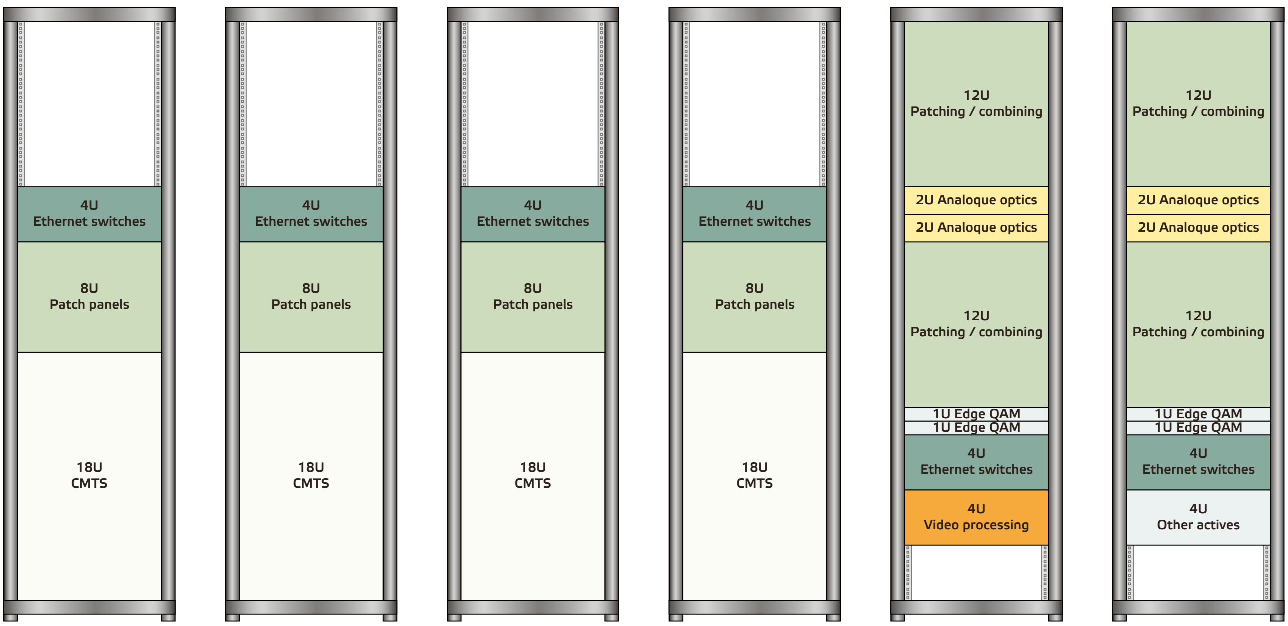

> **The traditional broadcasting transmission is expected to work and be simple.**

STARTING POINT

Linear television services, whether digital or analogue, have been robust workhorses and superior methods for distributing the same content to everyone. Traditional broadcasting transmission is expected to work and be simple. Figure 2 shows how a 'typical' regional headend might be constructed today. Saying 'typical' is misleading though, since, for instance, the presence of VOD services impacts construction a great deal. However, in many European countries comprehensive VOD services and shared edge QAMs (CMTS + VOD) are rare in comparison to North America.

As Figure 2 illustrates, the patch panels, RF combiners/splitters and CMTSs still require a great deal of headend space. However, traditional video processing, including local channel reception, does not require more than a couple of rack units and often only one rack unit is enough [1]. At least in Europe, it is still common to combine narrowcast and broadcast signals on the RF level, as Figure 3 shows, and full CCAP solutions are rare. The last combiners, which combine broadcast and narrowcast services, might be external or inbuilt to the optical receivers, as Figure 3 shows.

However, the inbuilt alternative offers higher performance for reasons described in the document '4K QAM downstream' [2].

The distributed architecture offers the option of using single carrier QAM (SC-QAM) modulators inbuilt into the Remote PHY devices instead of the Edge QAM modulators at the headend. Please see Figure 4 for further information. In such a case, linear TV broadcasting is transmitted over pseudowires using the Layer 2 Tunneling Protocol (L2TP). Unfortunately, only a few existing Video Headend platforms are able to stream IP video in L2TP frames. To the best of our knowledge, only one product successfully does it at the time this white paper has been published [3]. Should then broadcast video be streamed through the DOCSIS Core that encapsulates IP video packets, as Figure 4 illustrates? It is only an option. Next, we will discuss the pros and cons of five different options for managing video delivery over R-PHY-based distributed access networks.

Figure 3: *Example of a partial CCAP system today.*

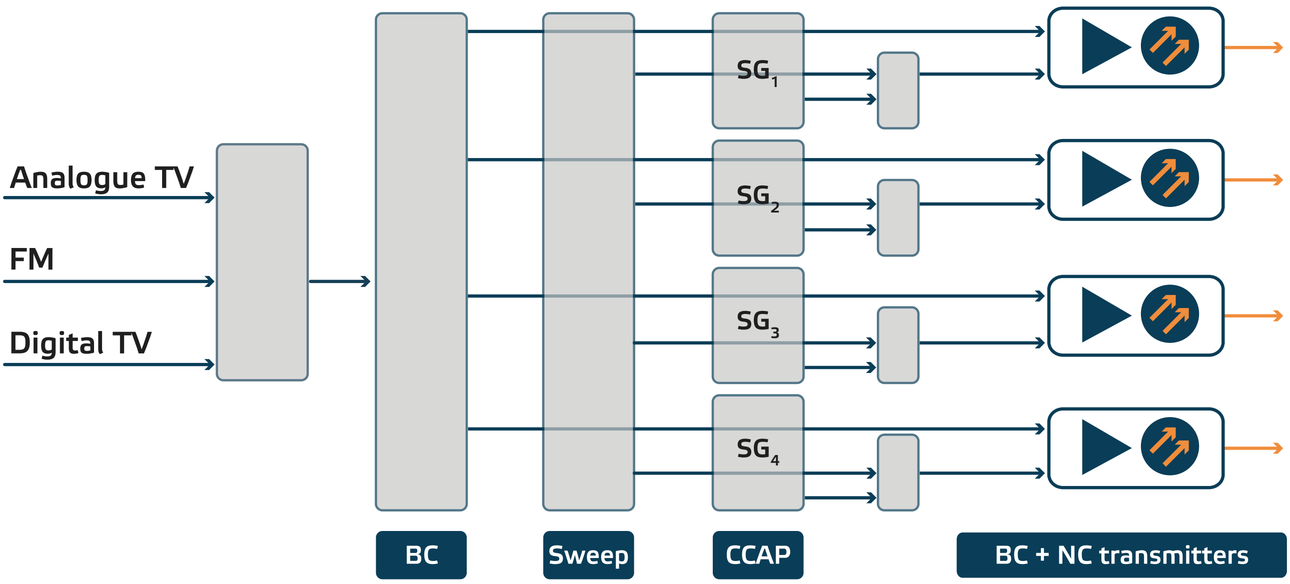

OPTIONS

Option 1:

The DOCSIS Core does everything

Option 1, illustrated in Figure 4, requires the DOCSIS Core to do everything that Video Headends used to do. In this article, the term 'Video Headends' means devices that manage, e.g. multiplexing, encryption and PSI/SI table insertion. For the sake of clarity, we assume that the DOCSIS Core is also the Principal Core, although Cablelabs's specifications do not require it [4]. A DOCSIS 3.1 capable Core most likely has 10 GbE ports that do not feed Remote PHY devices directly over point-to-point fibre links. First, one regional headend often feeds so many optical node areas that we need Ethernet aggregation switches between the DOCSIS Core and the Remote PHY to increase the number of Ethernet ports. Second, the 10 GbE capacity per RPD is most likely too much, at least at the beginning.

Could traditional Video Headend functions, such as local content reception, table insertion, multiplexing and conditional access for broadcast video, be performed by the DOCSIS Core? Based on market research, this seems to be a dream and not yet a reality. It would require massive integration exercises because of a plethora of different conditional access providers. Local video content reception at regional headends will complicate this as well because the local content is not often available over IP in a format that the DOCSIS Core would natively understand. Also, encrypting linear television services is a resource-hungry task and consumes processing capacity.

The approach shown in Figure 4 has several important consequences:

- Linear television burdens the DOCSIS Core, which now must manage L2TP encapsulation of all traffic.
- During software upgrades, the DOCSIS Core may interrupt linear television services.
- The cost of the DOCSIS Core may increase due to new pricing models, which are based on how much bandwidth subscribers consume [5].
- The DOCSIS Core should manage encryption, multiplexing and table insertion and it should be integrated with the existing legacy CA systems.

On the other hand, the approach offers a clear path to full IP delivery, which may be welcomed by many operators who see IP as a catalyst for access network convergence.

Figure 4: *The DOCSIS Core does everything.*

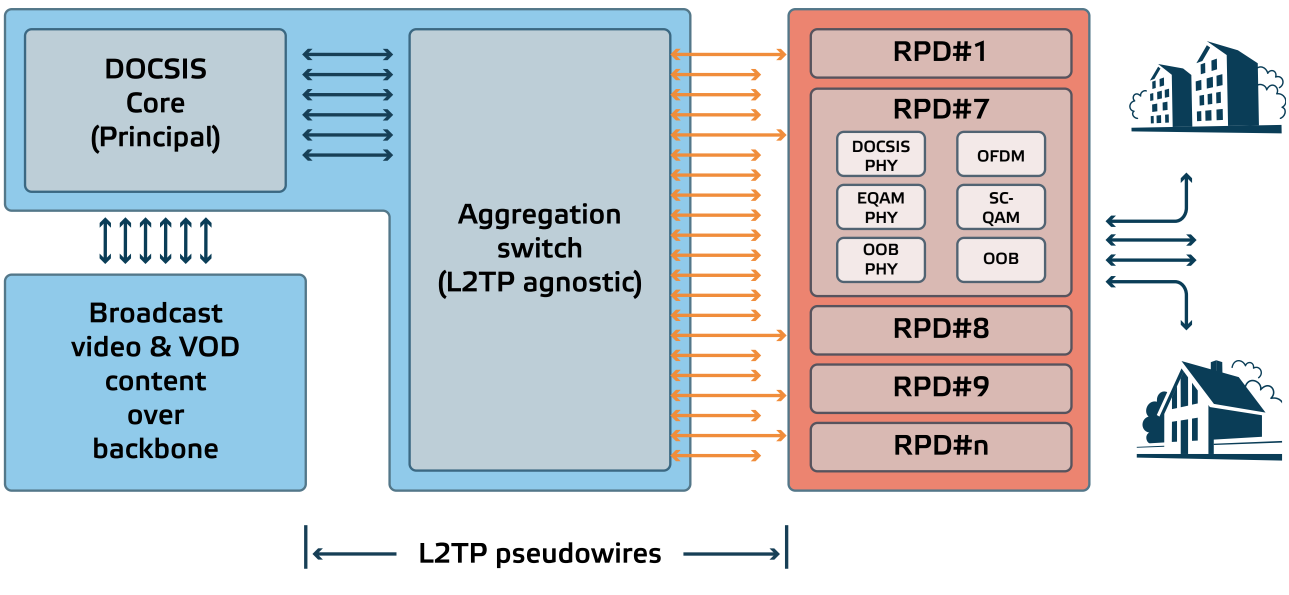

Option 2:

The DOCSIS Core handles L2TP encapsulation of video traffic

Option 2, illustrated in Figure 5, shows how the Video Headend sends all traffic through the DOCSIS Core. Unlike in option 1, the DOCSIS Core now just manages the L2TP encapsulation of linear television services, while traditional Video Headend functions are performed by a separate platform. Option 2 is more realistic than option 1 because the DOCSIS Core does not need to support features such as local content reception, multiplexing and conditional access.

On the other hand, options 1 and 2 have common issues because, even now:

- Linear television burdens the DOCSIS Core, which now must manage the L2TP encapsulation of all traffic.
- During software upgrades, the DOCSIS Core may cause linear television service interruptions.
- The cost of the DOCSIS Core may increase due to new pricing models, which are based on how much bandwidth subscribers consume [5].

Figure 5: *The DOCSIS Core handles the L2TP encapsulation of video traffic.*

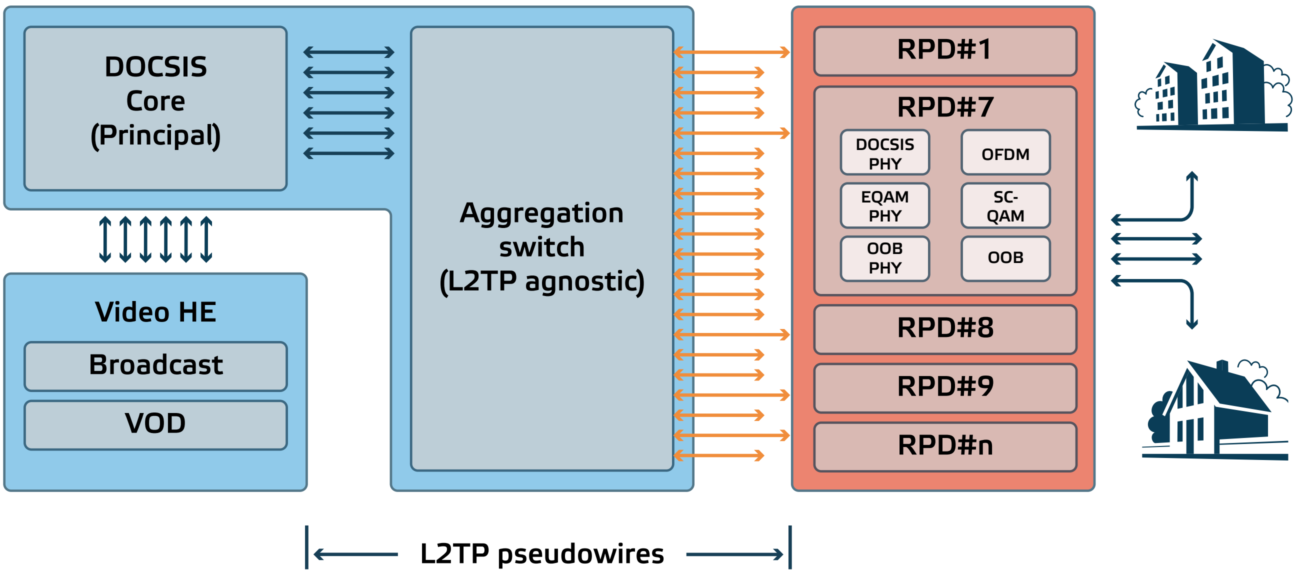

L2TP is comprised of two types of messages, control messages and data messages (sometimes referred to as "control packets" and "data packets", respectively). Control messages are used in the establishment, maintenance, and clearing of control connections and sessions. These messages utilize a reliable control channel within L2TP to guarantee delivery. Data messages are used to encapsulate the L2 traffic being carried over the L2TP session.

Unlike control messages, data messages are not retransmitted when packet loss occurs. The data message format for tunneling data packets may be utilized with or without the L2TP control channel, either via manual configuration or via other signaling methods to pre-configure or distribute L2TP session information.

[https://tools.ietf.org/ html/rfc3931#section-4.1.1]

Option 3:

Video content bypasses the DOCSIS Core

Option 3, illustrated in Figure 6, shows how services from separate Engines, in this case broadcast and VOD engines, are bypassing the DOCSIS Core. A L2TP agnostic aggregation switch delivers the L2TP-encapsulated video and data traffic to Remote PHY devices, which use OFDM and/or single carry QAM (SC-QAM) transmission for data traffic and SC-QAM for linear broadcast television. FM radio services can be converted to NDF channels, but it should be noted that the digitized FM band takes about 500 Mbit/s, and thus stresses the link between the FM to IP and IP to FM converters. Option 3 is less sensitive to changes in the DOCSIS Core, thus, e.g. software upgrades, do not cause linear television service interruptions. On the other hand, the system is a bit harder to manage since the video engines must communicate L2TP encapsulation instructions to the DOCSIS Core or else someone must manage the system manually. One of the benefits of option 3 is the DOCSIS Core bypass, the DOCSIS Core is not loaded with a linear television payload and the lower load thus improves cost efficiency. However, the video engines must be capable of managing a L2TP frame encapsulation, something that many Video Headend solutions are not used to doing. The L2TP-encapsulated frames are mandatory because the frame headers are used to address the selected SC-QAM channels in the Remote PHY; otherwise, the RPD will not know where the incoming video traffic should be routed.

Figure 6: *Video content bypasses the DOCSIS Core.*

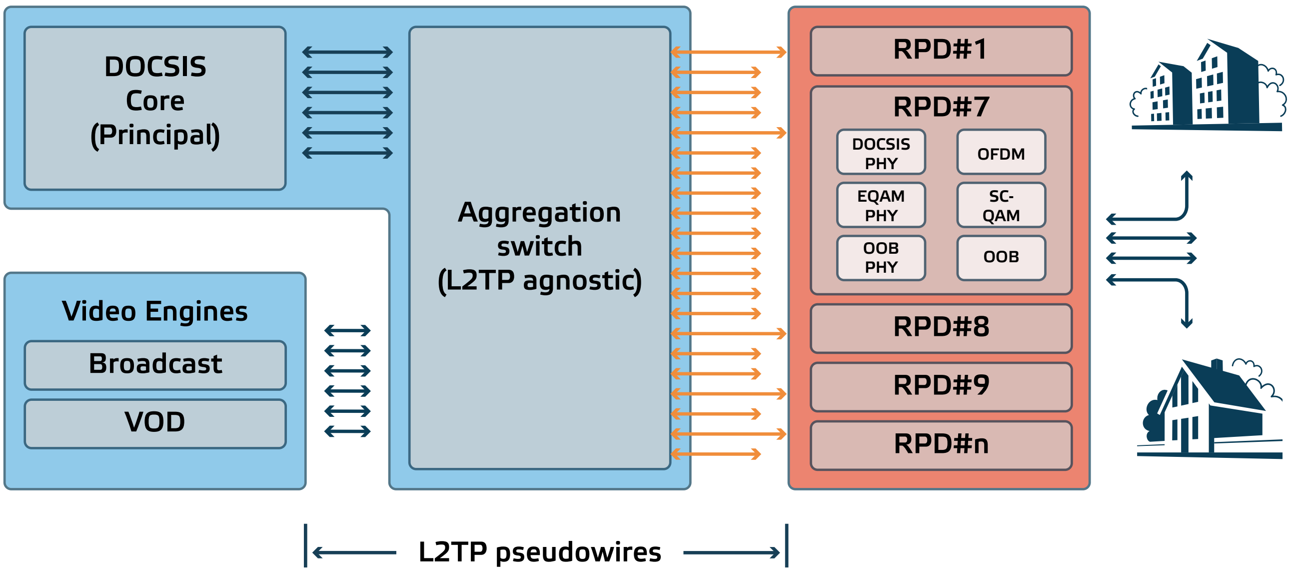

FM BAND: To support the 20.5 MHz FM band, 87.5 to 108.0, 25.6 MHz of spectrum is required, assuming a 1.25x oversampling for the guard band. This has been rounded down from 25.625 because the 25.6 MHz frequency is a convenient multiple of 5.12 MHz. The CCAP Core (DOCSIS Core in this white paper) sends 10-bit I/Q symbols packed into DEPI frames. The symbols are sent in I/Q pairs with the 10-bit I sample followed by the 10-bit Q sample. [CM-SP-R-OOB-I08-171220]

512 Mbit/s = 2* (25.6 Msymbols/s * 10 bit/symbol)

Option 4:

Video content bypasses the DOCSIS Core, Video Engines are converted to Video Auxiliary Cores

Option 4, shown in Figure 7, is similar to option 3. However, now the previously introduced Video Engines become Video Auxiliary Cores. These Aux Cores have data plane and control plane components, and the AUX Cores can manage a subset of RPD resources. While the option brings new features to the Video Engines, it also facilitates separate management of broadcast and data services. Therefore, traditional video services and data services (DOCSIS) can be managed individually by two teams. Otherwise, the pros and cons of option 4 are the same as those of option 3.

Figure 7: *Video Engines are converted to Video Auxiliary Cores.*

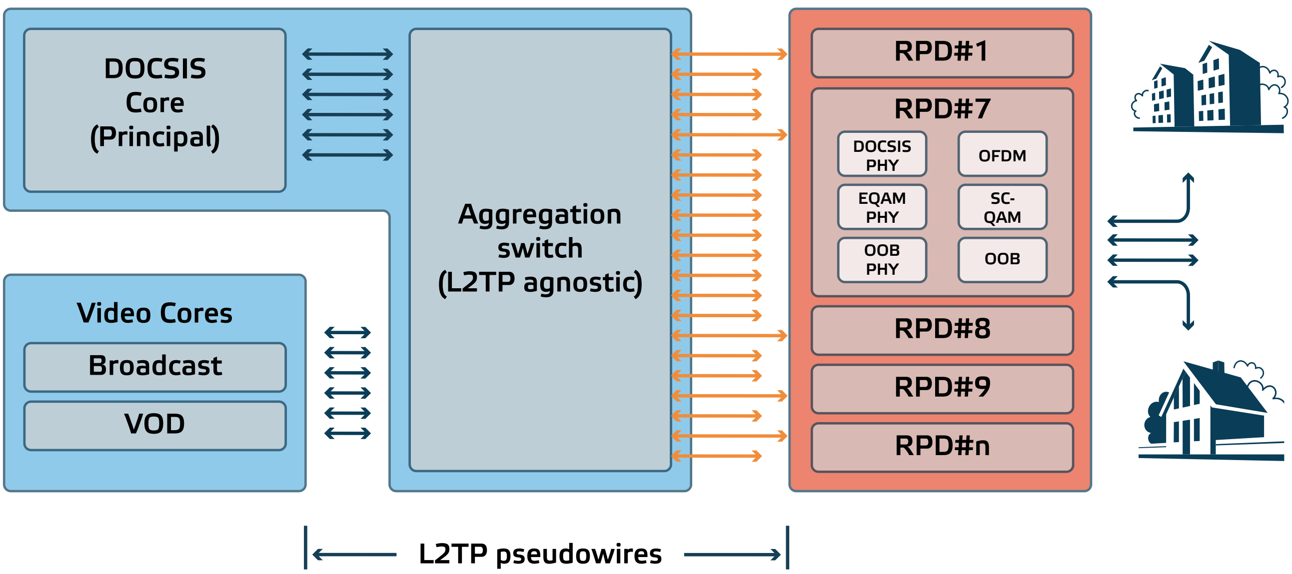

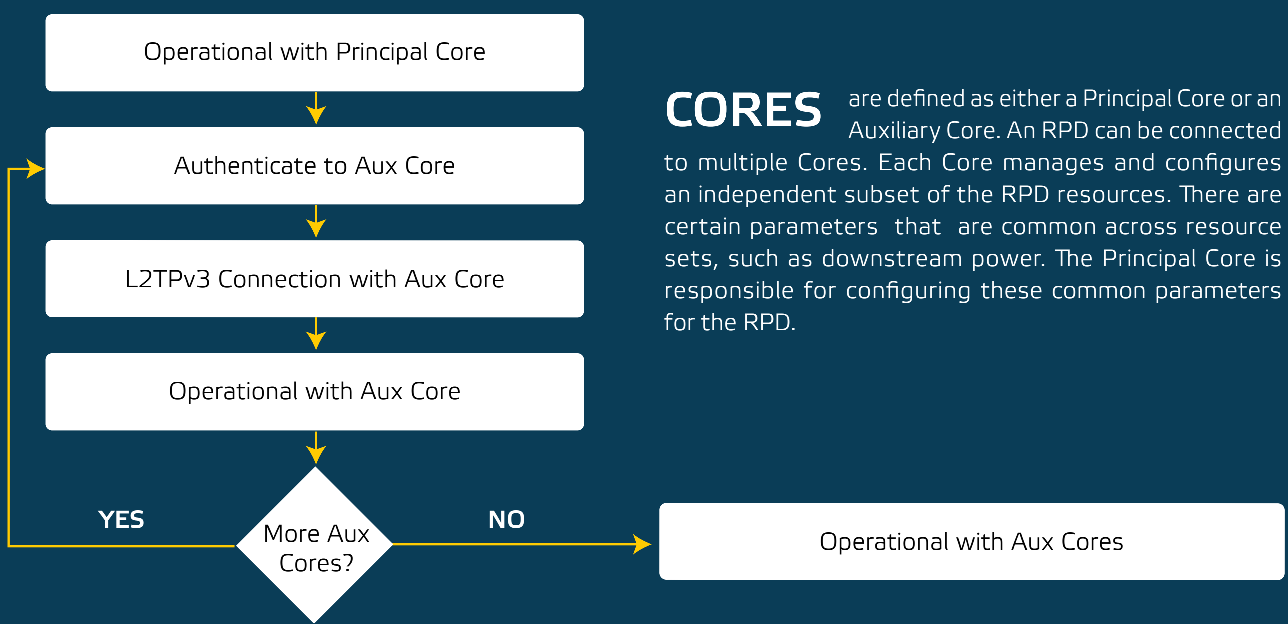

CORES are defined as either a Principal Core or an Auxiliary Core. An RPD can be connected to multiple Cores. Each Core manages and configures an independent subset of the RPD resources. There are certain parameters that are common across resource sets, such as downstream power. The Principal Core is responsible for configuring these common parameters for the RPD.

Option 5:

Traditional broadcasting

Option 5 is a safe choice, but causes a misstep with respect to achieving full IP delivery. Unlike the other alternatives, option 5 will not help if regional headends are heavily occupied and running out of space. Nonetheless, the existing legacy systems can remain because they make option 5 the only alternative if broadcasting analogue television channels is a necessity. As Figure 8 illustrates, the Remote PHY will receive traditional broadcasting over a data fibre (RF overlay) or over a separate fibre. The Remote PHY housing includes a separate traditional fibre node, and the linear television transmission does not burden any data devices (DOCSIS Core, Aggregation switch, Remote PHY). Thus, software upgrades for the DOCSIS Core, switches or Remote PHY do not cause linear television interruptions. While the option [6] could be judged old-fashioned, it is a cost-effective and robust alternative. In theory, this alternative supports a return-path RF overlay as well, but the arrangement is at least exotic.

Figure 8: *Traditional broadcasting.*

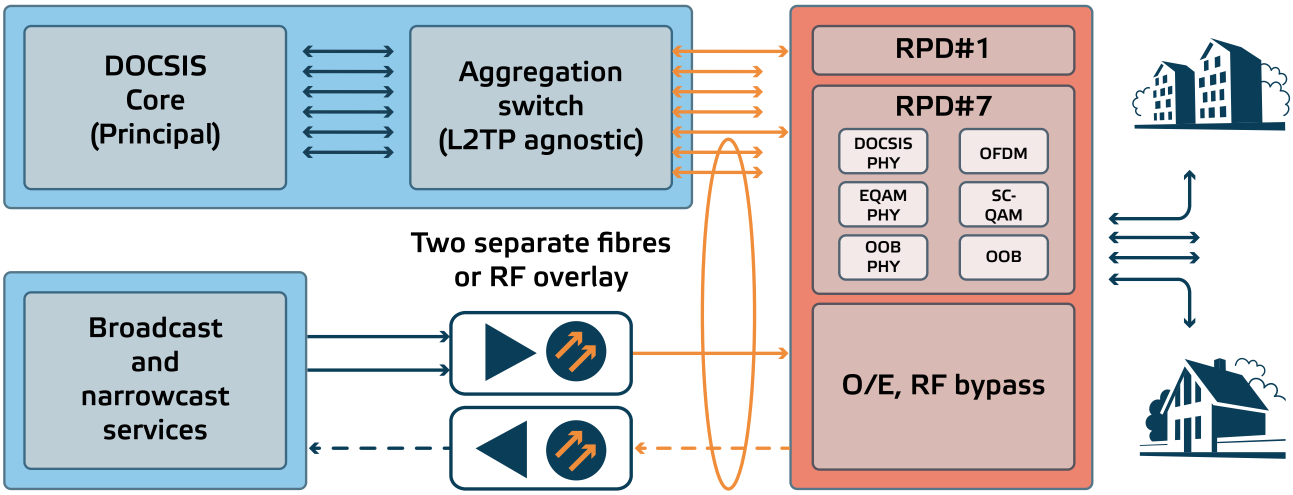

> **The only alternative if broadcasting analogue television channels is a necessity.**

Options compared

While Figure 9 compares the different options, namely their pros and cons, selecting the best alternative is impossible. However, we will make a more comprehensive and practical comparison in the 'Conclusions' section. Please note, the quantifiability of the factors 'cost effectiveness' and 'easiness of integration work' should be questioned. Operators have different starting points, which influence these factors. However, the author included them in the framework because the framework is meant to offer a template for operators to compare the pros and cons of different options from their own perspective. In the figures, 'cost effectiveness' refers to CAPEX, and, since real discounted product prices are not publicly available, an accurate cost comparison becomes impossible. The author would favour using TCO calculations instead of CAPEX, thus including OPEX is preferred when operators perform cost comparison calculations. Similarly, 'easiness of integration work' should be considered by knowing the starting point; thus, the comparison made in this article is, at most, only a coarse-grained one. The pros and cons are not of equal weight and the correct decision depends on many matters. It depends not only on the matters that we just discussed, but also on several other challenges, which we discuss next.

Figure 9: *Options, pros and cons.*

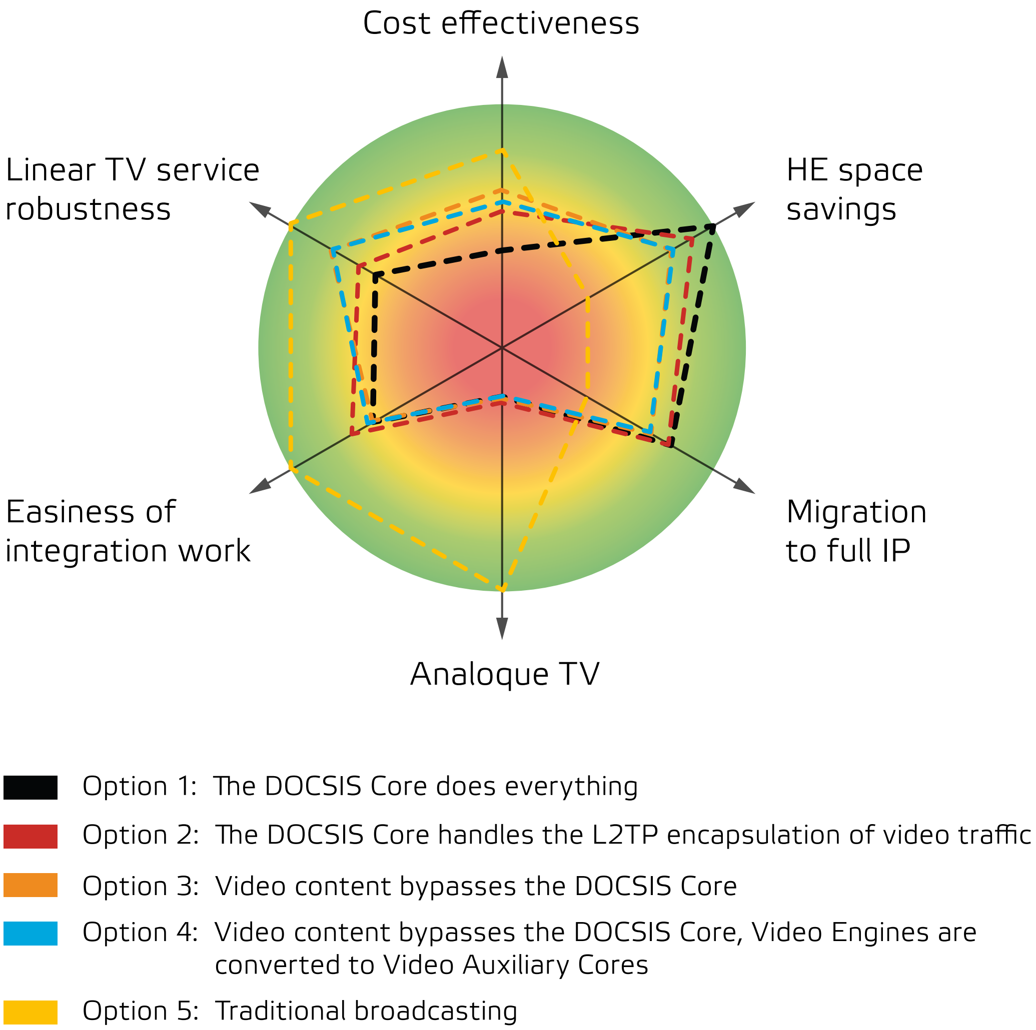

CHALLENGES

While deciding on the best option calls for strategical thinking, some practical and often operator-specific challenges require an operative mindset. One example of what operative people must consider is illustrated in Figure 11. Next, we discuss the challenges that the transition brings and their potential impact on the answer as to which options has the best fit in particular cases.

Challenge 1

Video delivery during the transition period

Renovating networks and rolling out distributed-access architectures takes time. In practice, the roll-out proceeds region by region, perhaps in the way shown in Figure 11. Rarely is the same street cabinet able to host an old fibre node parallel to a new Remote PHY node. Time is running out with respect to when the old fibre node can be switched to the Remote PHY node. When the incoming fibre is switched to the input port of the Remote PHY, the headend cabling must be switched as well, unless the operator has selected option 5. The transition period eases if the same Video Headend supports simultaneous video broadcasting and IP video delivery. Otherwise, consumers will experience long service breaks. Several decades ago, it was easier to switch off linear television during the night, but globalization has changed the landscape. Live content, such as sports originating from different time zones, might be broadcasted at night, making timing paramount. Although single service breaks may not be enough to create churners, these breaks result in dissatisfaction, pre-churn or even churn if customers are close to the tipping point [7].

Challenge 2

Local channels

Several content insertion points and local channels require attention. While some content can be processed and broadcasted locally, their management may cause challenges. In such cases, if the backbone between the regional headends and super headend has enough capacity, it might make sense to manage even local multiplexes centrally, as Figure 10 shows (the route B). This centralized approach is especially tempting if local channels are transcoded. But the transcoding of local channels is seemingly rare in real networks. To keep both avenues open (localized vs centralized), it will be important to have Video Headends that can manage multiplexing locally and also stream local channels over IP.

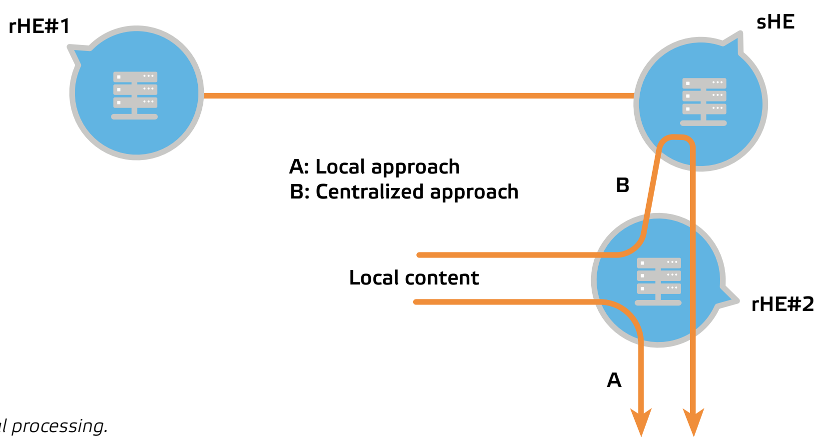

Figure 10: *Local vs. central processing.*

Figure 11: *Transition.*

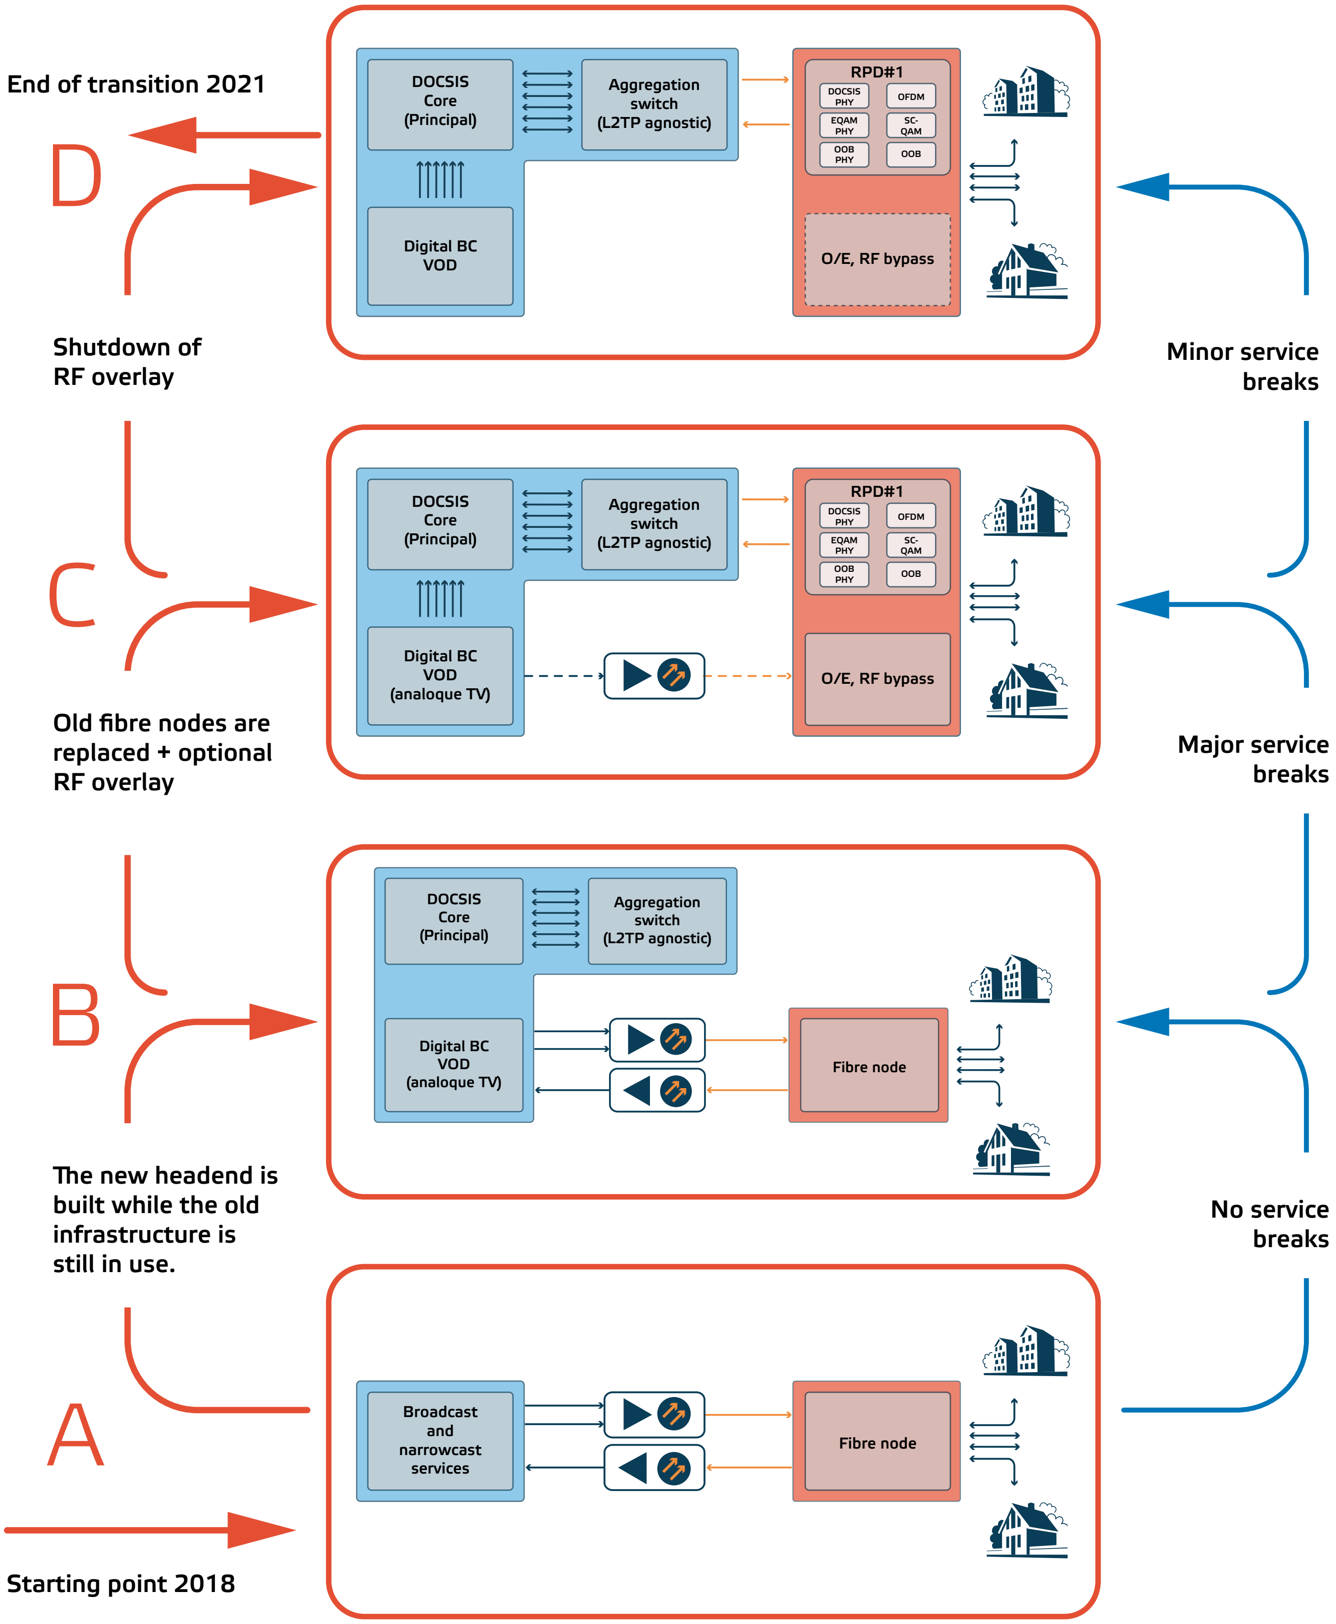

Challenge 3

Conditional Access Systems (CAS)

Before the introduction of the Simulcrypt framework, operators were locked in with particular CAS vendors. However, even nowadays the Simulcrypt framework does not standardize every detail and Video Headends must still be separately integrated with every conditional access system. Due to existing legacy systems, especially due to set-top boxes, many operators require that Video Headends work with several CA systems at the same time. Noting this becomes paramount if option 1 is of interest. While the requirement to support several CA systems is imperative for European Video Headend suppliers, it might be less obvious for many DOCSIS Core vendors.

Challenge 4

PSI/SI tables

A set-top box needs tables. Without tables, it does not know what to listen to, and without tables many elemental features, such as the EPG, are not available. But why do these tables constitute a challenge when operators implement distributed architectures? Because the NIT table tells the frequencies, constellations and symbol rates of different channels. But which device knows all that information? The DOCSIS Core can manage and thus know how different frequencies are used. However, it would be a bold assumption to think that the DOCSIS Core will create tables automatically and perhaps even dynamically. This raises several additional questions. Is the Principal Core able to inform other network elements about what kind of NIT tables should be created? Will the broadcast SC- QAM channel settings change, thus at times requiring NIT table modifications? Are NIT tables different across the network or can the same channel line-up cover the whole network?

Challenge 5

Multicasting

Two different multicasting alternatives may seem confusing:

A. Multicasting of L2TP-encapsulated video traffic between the regional headend and Remote PHY devices so that the RPD transmits the incoming multicast traffic over DVB-C QAM outputs. See Figure 12a.

B. Multicasting of L2TP encapsulated video traffic (OTT) between the regional headend and Remote PHY devices so that the RPD transmits incoming multicast traffic over DOCSIS QAM/OFDM outputs. See Figure 12b.

Alternative A:

The L2TP frames are easy to manage when the need is to build a point-to-point link between the DOCSIS Core and the RPD. However, when switches are placed between the Cores and the RPD, things may become complicated, as typically switches are L2TP agnostic. While the switches will see layers below the L2TP frames (Ethernet and IP), the RPD routes traffic based on the L2TP frames. So how will the L2TP-unaware switches know which RPD wants to receive particular multicast content? In practice, the Principal Core must ask the RPDs to send IGMP-join messages so that the switches learn which RPD wants to receive a particular multicast stream. In the above-mentioned scenario, the assumption is that all video and data traffic is in the same subnet. The switches should rather be L3 capable (i.e. simple routers) if video and data traffic are actually in different subnets. In this case, one DOCSIS Core port could be reserved for video services, while other ports are used for data traffic. In practice, this case requires that the RPDs receive content from several IP addresses.

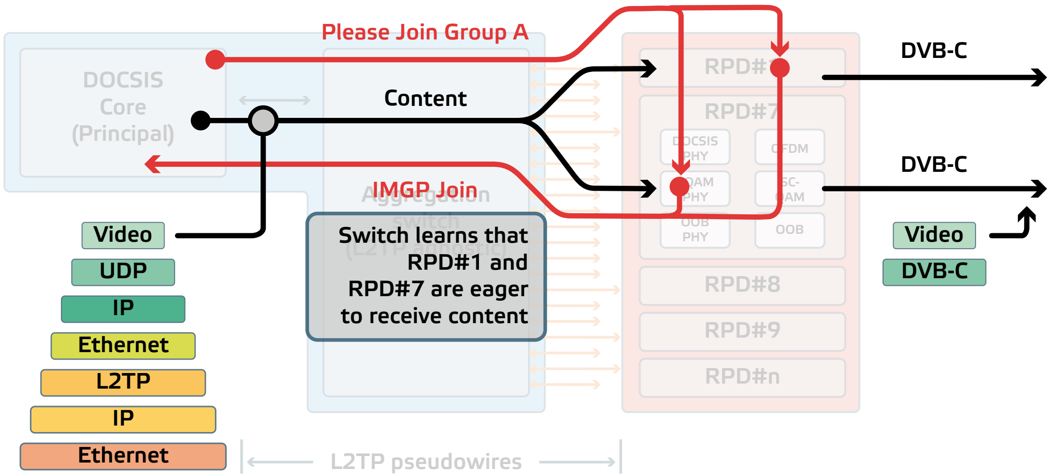

Figure 12a: *Multicasting.*

Alternative B:

At the time of writing this article, the author is not aware of any DOCSIS Core that could transmit the same DOCSIS IP multicast traffic (e.g. OTT video) across different downstream service groups as that of a single stream. In practice, the DOCSIS Core actually copies the traffic and sends it as many times as the stream is needed in different downstream service groups. This copying stresses the link capacity between the DOCSIS Core and aggregation switches. Also, there is a tendency towards smaller service groups, making the copying a more serious issue. In theory, this issue could be solved by a workaround configuration, where all Remote PHY devices would receive the same L2TP-encapsulated stream. However, the Remote PHY, according to the standard [8], does not process multicast join messages (IGMP join) sent by the customer premise equipment. So, the RPDs cannot filter traffic based on this information either. Even if such filtering would be possible, the arrangement would cause challenges in managing available QAM/OFDM capacity after the RPDs, as capacity should be managed by the MAC layer. Therefore, although the above-mentioned workaround configuration would save network capacity between the DOCSIS Core and aggregation switches, it would stress the network after the Remote PHY devices because all incoming traffic would be forwarded via the RPD. In practice, at least at the moment, it is simply better to accept that IP multicast is actually not a real multicast-exceeding downstream service group; rather, it works inside single downstream service groups.

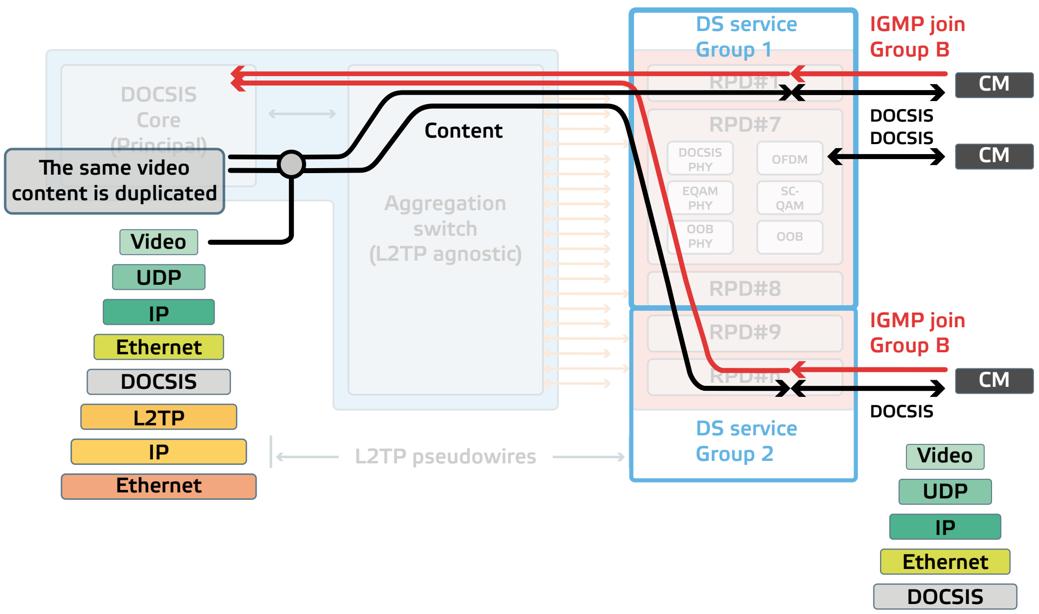

Figure 12b: *Multicasting.*

CONCLUSIONS

The shift to distributed-access architectures will not be painless. This paper hopefully covered many topics worth noting when operators think about how to master video transmission over next-generation networks. Figures 13 a–e recap options and challenges, and they will serve as guidance when operators think about which option is the best way forward. As previously mentioned, please note that the measurability of the factors 'cost effectiveness' and 'easiness of integration work' should be questioned.

Option 1: The DOCSIS Core does everything

Video delivery during the transition period	Linear TV and DOCSIS system changes must be managed at the same time.
Local channels	Does the DOCSIS core support local reception?
Conditional Access Systems (CAS)	Is the DOCSIS core integrated with CAS?
PSI/SI tables	Does the DOCSIS core generate PSI/SI tables?
Multicasting A	DVB-C multicasting requires special attention.
Multicasting B	IP Multicasting across downstream service groups is virtually impossible.

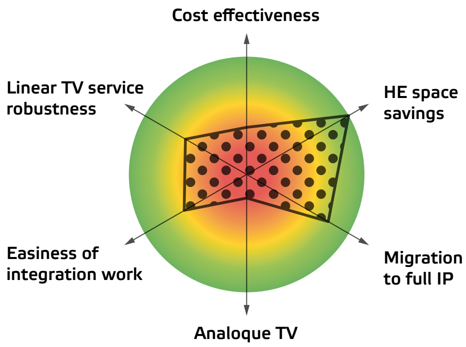

Figure 13 a. *Challenges and options: Option 1.*

Option 2: The DOCSIS Core handles the L2TP encapsulation of video traffic

Video delivery during the transition period	Linear TV and DOCSIS system changes must be managed at the same time.
Local channels	A separate Video HE is used for local reception.
Conditional Access Systems (CAS)	No changes required
PSI/SI tables	Video HE/engine/core can generate tables but which system manages dynamic changes? Is the NIT table the same across the whole network?
Multicasting A	DVB-C multicasting requires special attention.
Multicasting B	IP Multicasting across downstream service groups is virtually impossible.

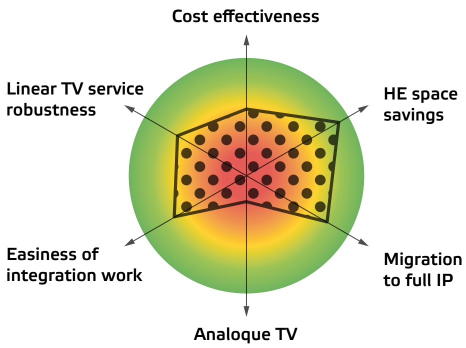

Figure 13 b. *Challenges and options: Option 2.*

Option 3: Video content bypasses the DOCSIS Core

Video delivery during the transition period	Linear TV and DOCSIS system changes can be managed separately.
Local channels	A separate Video HE is used for local reception.
Conditional Access Systems (CAS)	Is the video core/engine integrated with CAS?
PSI/SI tables	Video HE/engine/core can generate tables but which system manages dynamic changes? Is the NIT table the same across the whole network?
Multicasting A	DVB-C multicasting requires special attention.
Multicasting B	IP Multicasting across downstream service groups is virtually impossible.

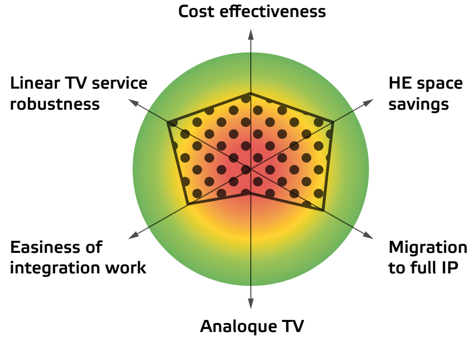

Figure 13 c. *Challenges and options: Option 3.*

Option 4: Video content bypasses the DOCSIS Core, Video Engines are converted to Video Auxiliary Cores

Video delivery during the transition period	Linear TV and DOCSIS system changes can be managed separately.
Local channels	A separate Video HE is used for local reception.
Conditional Access Systems (CAS)	Is the video core/engine integrated with CAS?
PSI/SI tables	Video HE/engine/core can generate tables but which system manages dynamic changes? Is the NIT table the same across the whole network?
Multicasting A	DVB-C multicasting requires special attention.
Multicasting B	IP Multicasting across downstream service groups is virtually impossible.

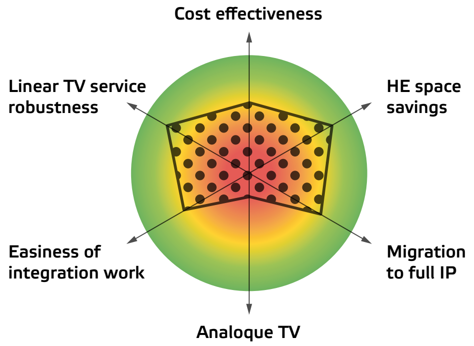

Figure 13 d. *Challenges and options: Option 4.*

Option 5: Traditional broadcasting

🟢	Video delivery during the transition period	No challenges
🟢	Local channels	No changes required.
🟢	Conditional Access Systems (CAS)	No changes required.
🟢	PSI/SI tables	No changes required.
🟢	Multicasting A	No changes required
🔴	Multicasting B	IP Multicasting across downstream service groups is virtually impossible.

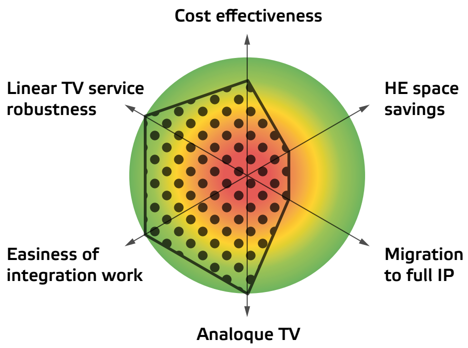

Figure 13 e. *Challenges and options: Option 5.*

LIMITATIONS

Our article in part ignores the option of separating the Principal Core and the DOCSIS Core, although it is certainly an attractive alternative. We ignored the option because we do not see it as impacting the decision regarding how video delivery should be arranged, although the option may increase the amount of integration work when a separate Principal Core, the DOCSIS Core and different Engines are integrated. In turn, the monetary benefits of making DOCSIS Cores into commodities and having a separate Principal Core may justify this 'divide and rule' approach.

We also ignored the option of including some, but not all, Video Headend functions within the DOCSIS Core. For instance, multiplexing could be arranged in the DOCSIS Core, while conditional access could be managed by other elements. These mixed approaches can be attractive in some cases, but they may complicate the management of linear television services.

The white paper focuses on European challenges, but most of them are valid in North America as well. However, we do not discuss Switched Digital Video here, as its importance in Europe is negligible. Also, the MACPHY alternative was ignored. Although deploying MACPHY-based distributed access introduces many challenges similar to those discussed in this article, certain details, for example IP multicasting, do differ.

Several options, such as option 1, might offer a smooth path to network virtualization. We ask industry practitioners to publish studies regarding the impact of virtualization on the attractiveness of different network migration options, as the topic was ignored by the author.

Video transmission is not the only challenge introduced by the shift to distributed access. We ignored important questions, such as which system orchestrates and manages the network, Video Headends, Engines and Cores when distributed access gets a foothold. What kind of powering requirements do new power-hungry remote PHY devices set for the network? How do you monitor HFC network performance when the traditional sweep is gone? Please, as we mentioned earlier, let us know what areas part 2 should cover. We need to pave the way in many interesting areas and your feedback will influence the order of the articles. We are eager to hear your feedback and receive questions through email. The article, despite its limitations, hopefully gave new tools and a framework for managing the Shift.

ABOUT THE AUTHOR

Arttu Purmonen, Vice President, Teleste Corporation

Before devoting his efforts to technical marketing, Arttu served as business director, product manager and engineer at Teleste Corporation. Arttu is an active content producer, and his interest lies in the latest network and video processing technologies. Understanding the customer perspective has always motivated Arttu, and his latest research work focuses on consumer retention rates, especially on the question of how operators can use intelligent network products to prevent pre-churn and churn. Arttu joined Teleste in 1997 and holds an MSc in technology; he is currently performing post-graduate marketing studies at Turku School of Economics.

Contact:
arttu.purmonen@teleste.com

The author thanks Teleste Corporation and the research team for their support.

The research team:
Jouni Karlsson (illustrations)
Päivi Vuori (editor)
Katja Kettunen (layout)

REFERENCES

[1]. https://www.teleste.com/video-headend/products/teleste-luminato (accessed April 2018)

[2]. https://www.teleste.com/node/21488 (accessed April 2018)

[3]. https://www.teleste.com/video-headend/products/luminato/luminato-application-module (accessed April 2018)

[4] https://apps.cablelabs.com/specification/CM-SP-R-PHY (accessed April 2018)

[5] https://www.acgcc.com/ciscos-docsis-new-licensing-model-realigns-economic-incentives/ (accessed April 2018)

[6] https://www.teleste.com/broadband-network/products/docsis-data-over-coax/ac9100-neo-rpd-node-remote-phy-module (accessed April 2018)

[7] https://www.teleste.com/node/21870 (accessed April 2018)

[8] https://apps.cablelabs.com/specification/modular-headend-architecture-v2-technical-report (accessed April 2018)

TELESTE CORPORATION
P.O.Box 323
FI-20101 Turku, Finland
www.teleste.com

P3I_Changes of video delivery in distributed architecture 05 2018