Remote PHY Going the Distance

A Technical Paper prepared for SCTE by

Marek Hajduczenia, PhD
Network Architect, Distinguished Engineer
Charter Communications
14810 Grasslands Dr, Englewood, CO 80112
marek.hajduczenia@charter.com

Glen Hardin
VP, Technical Fellow
Charter Communications
14810 Grasslands Dr, Englewood, CO 80112
Glen.Hardin@charter.com

John Chapman
CTO Broadband & Fellow
Cisco Systems
300 East Tasman Drive, San Jose, CA 95134
Glen.Hardin@charter.com

Table of Contents

List of Figures

1. Introduction

Why RPHY go the distance?

The answer to this question has a dramatic impact on the plant topology, network implementation, compute scaling and, most importantly, the operational performance of the overall system. One of the cost reduction opportunities is in the consolidation of HFC hub operations. HFC hubs can be consolidated, simplified, or even completely decommissioned by transitioning traditional integrated CMTS (iCMTS) platforms to virtualization solutions. This drives the opportunity for the complete hub-collapse and aggregating more customers into larger metro sites.

Charter has been evaluating different technologies with Remote PHY (also referred to as R-PHY, R PHY, or simply RPHY) being one of those. Charter also examined the distance question of RPHY design in detail from a theoretical and practical perspective. While there are general statements to that effect, the relationship between specific distances and performance has not been identified to date.

Spectrum Enterprise has been deploying vCMTS with subtended video cores and RPHY devices for years with many of its hospitality clients. The commercial application of RPHY has given Charter real-world data of actual deployments. This data is very relevant to the residential application for data services. Additionally, significant studies have been conducted using impairment generators in trial deployments, validating the impact of symmetric and asymmetric latency on the data throughput, video services, etc.

This paper explains the challenge at hand, identifying key technical risks and limitations, and looks at the lessons learned from the existing video centric RPHY deployments within the Charter network. Details on the internal studies with impairment generator are also presented, exploring the limits of RPHY architecture performance, providing a much more positive picture of the system-level performance.

2. RPHY Background

RPHY is the new and more future-proof HFC architecture. RPHY replaces the previous generation of analog fiber with digital IP based fiber and moves the function of generating RF spectrum to the fiber node.

2.1. A Brief History

The RPHY architecture was invented by John Chapman of Cisco in 2001 [1] [7] [8] [9] and was originally part of the Modular Headend Architecture (MHA). The main protocols required for the operation of RPHY include:

- **R-DEPI**, the remote downstream external PHY interface, a pseudowire that connects the CMTS Core to the RPD [2],
- **R-UEPI**, the remote upstream external PHY interface, a pseudowire that connects the RPD upstream to the CMTS Core [3],
- **GCP**, the generic control plane, that is the management protocol for the RPD [4], and
- **R-DTI**, the remote DOCSIS Timing Interface, that uses IEEE 1588 to manage the DOCSIS timestamp between the CMTS Core and the RPD [5].

DEPI, UEPI, and DTI were written in 2004. GCP was written in 2013 at a Pete's coffee shop. DEPI and DTI were standardized as part of MHAv1 at CableLabs in 2005 [10]. UEPI and GCP, along with a

"Remote" prefix, were standardized at CableLabs as part of CDOCSIS in 2014 [11] and MHAv2 in 2015 [12].

2.2. How RPHY Manages DOCSIS REQ-GNT Delays

All modern access systems, such as WiFi6, mobile 5G, PON, and DOCSIS, use a scheduled upstream. They all work in a similar manner. The DOCSIS system coupled with a RPHY system is shown in Figure 1.

The delays of this system were analyzed in extensive detail during the MHAv1 process. The delays, based on 2004 CMTS technology, are described in the MHAv1 DEPI specification, Appendix I, "DEPI and DOCSIS System Performance" [13] and in Appendix I of this white paper.

The first thing to be mentioned is that all the delays shown, apart from the three RPD queueing delays, are common to RPHY, RMACPHY, and I-CMTS systems. That is because they are DOCSIS delays. The REQ and MAP get special processing in the RPD that allows them to skip the data queue which provides very little delay. The RPD upstream data queuing is bringing a packet from a 100 Mbps DOCSIS interface to a 10 Gbps Ethernet interface, so it also has negligible delay.

Let's spend a moment looking at DOCSIS delays. Let's also assume that MAP messages are sent every 2 milliseconds on average. When a packet arrives at the CM, it is processed and then it has to wait up to a full MAP time for a REQ slot. When the REQ arrives at the CMTS, it is processed and put into a scheduler queue where it waits for up to a full MAP time for the scheduling process to initiate. A MAP is then built and sent to the CM.

The CM needs to receive the MAP in advance of when it needs to use it. The CMTS does this by creating MAP far enough in the future that the MAP will never arrive late. This is known as MAP advance. It includes engineering margin and can be as much as a MAP time. The MAP then has to play out and it can be up to a MAP time before the data packet is sent from the CM if it is at the end of the MAP. The CMTS then receives the data packet, processes it, and sends it out.

All in all, there are at least four MAP times chained together which creates a delay of 4 to 8 ms, in addition to PHY and queueing delays. Very fast CMTSs can get this time sequence down to 5 ms best case. On average, this delay is 12 ms, and if there is contention on the REQ channel, DOCSIS delays can quick add up to 50 ms. [15]

The additional delay from the RPD could be less than 0.25 ms, which is less than a 5% impact.

The summary of the delay analysis is that most of the request-grant delay in absorbed in the DOCSIS system itself. These measurements were shown in [15]. For systems less than the DOCSIS specification limit of 100 miles, the RPHY system adds less than 10% to the overall delay. In practice, since upstream delay is rounded off to MAP intervals, this additional delay may not even be seen in a short loop system.

The delay time in DOCSIS is much higher than RPHY, which makes RPHY nearly equivalent to DOCSIS.

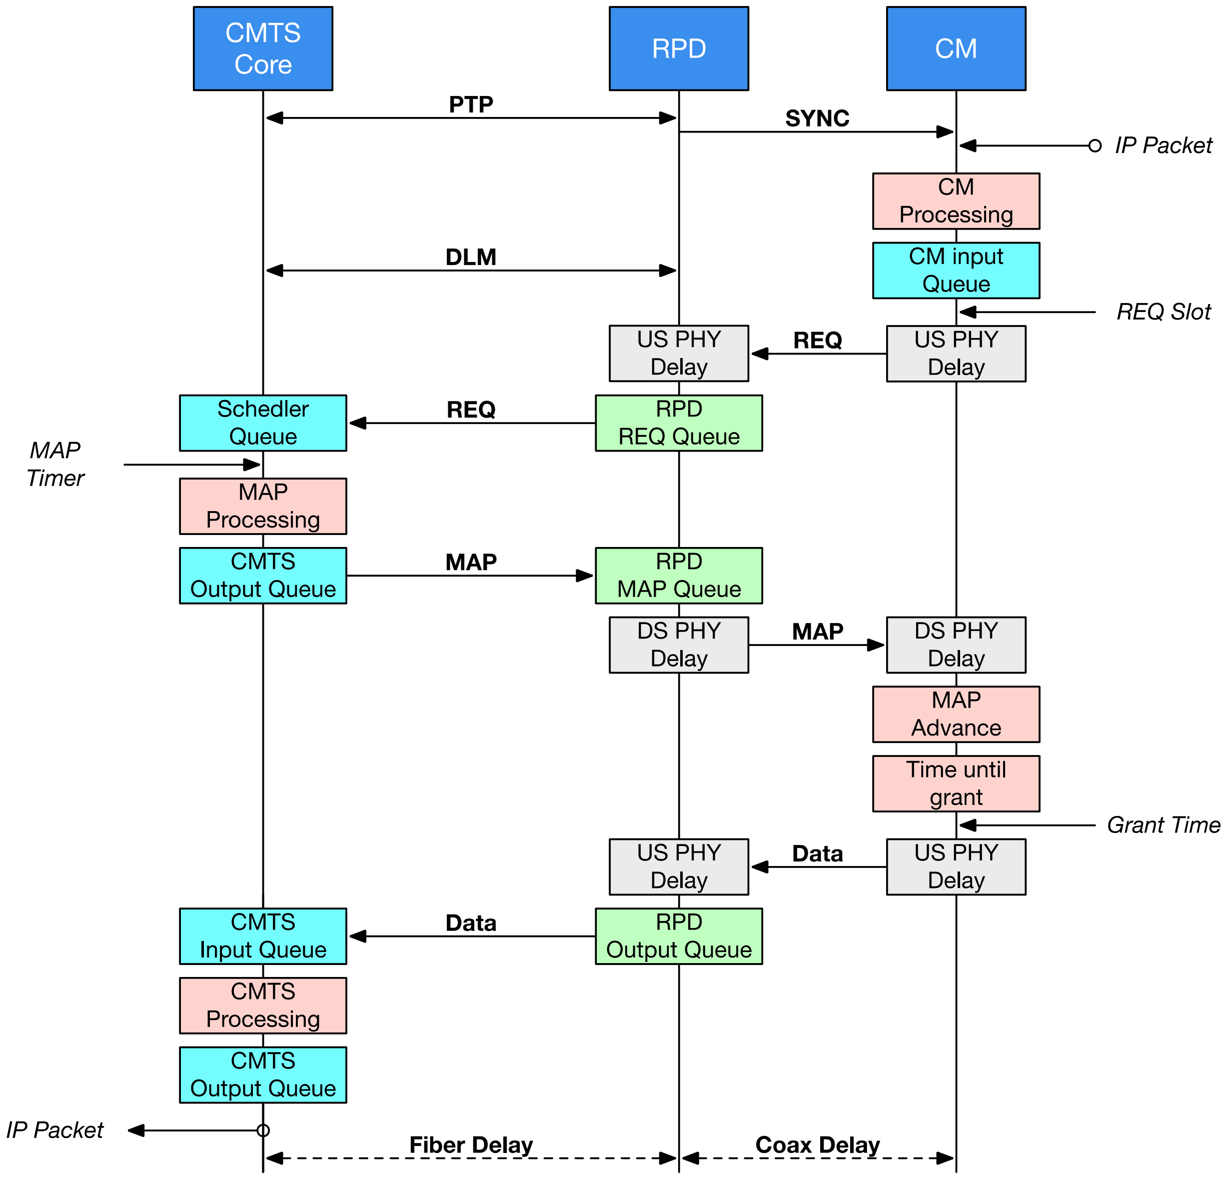

Figure 1 – RPHY and DOCSIS Delays

2.3. DOCSIS Latency Measurement (DLM)

The CMTS Core and the RPD are synchronized together to the reference PTP clock [16] [17]. This means that both the CMTS Core and the RPD are time-synchronized. The DEPI specification contains a special header operation called DLM that allows the timestamps to be recorded and compared at each point. The difference between the two is an accurate measurement of the delay of the fiber optic cable.

When the measured fiber delay is greater than the equivalent of 100 miles, the CMTS Core uses value from DLM to calculate a larger MAP advance time.

*It is DLM that allows the CMTS to make RPHY work with **almost any length of cable**.*

2.4. DOCSIS Predictive Scheduling (DPS)

DOCSIS predictive scheduling is a further enhancement that is not in the DOCSIS standard but is used by several CMTS vendors. DPS issues extra grants to active CMs, allowing these CMs to accomplish two objectives:

1. Send data packets from before receiving a schedule grant, or even before requesting upstream bandwidth.
2. Send request messages as a piggyback message and to not have to worry about contention in the system request slot. This will reduce both short-term latency and the long-tail latency.

Just one grant slot sent per MAP interval, or several intervals, to an active CM will low the upstream latency significantly.

DPS is a further performance enhancement that lowers the upstream latency to just a single delay of the cable length instead of tripling the cable length delay.

3. RPHY Service Architecture

RPHY is a type of distributed access architecture (DAA), whereby the physical DOCSIS layer is moved from the hub (typical location where CMTS is deployed) to the edge of access network, i.e., much closer to the end customers. In the RPHY architecture, the integrated CMTS (I-CMTS) is therefore divided into the physical layer, aggregation layer and routing/MAC layer, with each of these aforementioned layers mapped into new functional elements:

- the physical layer is implemented in the RPHY device (RPD),
- the aggregation layer is implemented using the IP/MPLS transport layer, and
- the routing/MAC layer is implemented using a virtual CMTS (vCMTS) software element running on a generic compute platform deeper in the carrier network.

This service architecture allows for the processing functions (schedulers, software management, bandwidth scheduling, etc.) to remain centralized, while the low-level DOCSIS physical layer processing is to be distributed in the most cost-effective manner, as shown in Figure 2. The use of the digital fiber-based transport IP/MPLS network between the centralized and distributed elements of this architecture substantially increases the reach of the DOCSIS vCMTS, eliminating the need for analog fiber deployment, analog fiber amplifiers, etc.

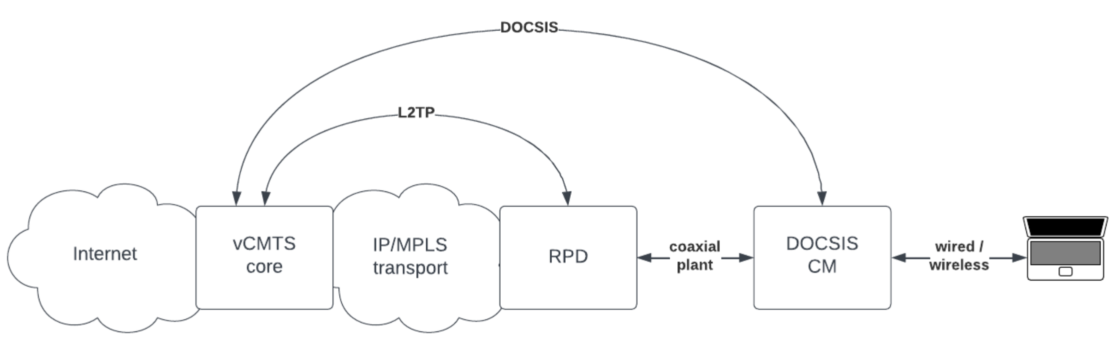

Figure 2: RPHY service architecture

The majority of low-level physical layer DOCSIS processing is now concentrated within the RPD at the network edge, which plays two important roles:

- In the downstream direction, the RPD converts IP-encapsulated data received from the vCMTS core into analog signaling and transmits the resulting RF towards individual end-user devices (DOCSIS CMs, video receivers, etc.)
- In the upstream direction, the RPD performs the reverse process, converting any analog signals received from individual end-user devices (primarily DOCSIS CM) into IP-encapsulated digital data that is then transmitted towards the vCMTS core.

The RPHY service architecture is highly flexible in which individual RPDs can be configured independently from other RPDs in the same or other serving groups. This facilitates delivering traffic to and aggregating traffic from customers sharing the same service lineup. Through an efficient use of unicast and multicast L2TP data tunnels between individual RPDs and a vCMTS core, specific content can be distributed to multiple locations at the same time, taking full advantage of underlying IP multicast capabilities of modern transport networks.

This architecture does, however, have a strong timing dependency, whereby the vCMTS and individual RPD must be tightly synchronized to permit centralized schedulers to properly stack transmissions from individual end-user devices attached to the RF portion of the network. At this time, given the existing level of technical maturity of this technology, the timing synchronization aspect remains the only technical challenge for any operator planning to deploy the RPHY service architecture. Careful design of the timing distribution is required, as discussed in more detail in the following sections.

3.1. Service Delivery Model

As indicated before, the service delivery model relies on a distributed CMTS platform where the service processing and routing functions are centralized within the vCMTS and all the DOCSIS-specific media conversion functions (PHY) are pushed to the edge of the network (RPD). All communication between individual RPDs and the vCMTS cores takes place on the IP/MPLS network core (when vCMTS cores are deployed in the regional/metro DCs) or the hub-level converged interconnect network (CIN, when vCMTS cores are deployed at the hub level).

In this case, the RPHY service architecture shares the packet-switched network infrastructure with other data services, thereby taking advantage of the economies of scale in carrier networks. The high-level service architecture is shown in Figure 3 where the vCMTS cores are located in a market-centralized facility (marked as data center (DC) vCMTS) or at individual hub sites (marked as hub vCMTS), and individual RPDs are deployed at one of the possible locations:

- Option 1 (hub location): At the hub, representing a direct 1:1 replacement of iCMTS platforms
- Option 2 (field location): In the field (using a strand mount or street cabinet format), replacing one of amplifier nodes and shortening the amplifier cascade
- Option 3 (on-prem location): At the customer premises, connecting to intra-building coaxial distribution network

Individual options outlined above present the evolution of residential DOCSIS deployment models (Option 1 and Option 2) as well as different application scenarios (typical residential deployments – see Option 1 and Option 2, and community deployments – see Option 3). Individual options are discussed in more detail in the following subsections.

The placement of the vCMTS cores is operator-specific and depends on several factors, including:

- Aggregation scale required to achieve proper economies of scale
- Availability of data center (DC)-class facilities in the network and their location
- IP/MPLS transport network design and its capacity
- Acceptable splash radius in case of any DC outages and general risk aversion levels

Note that the decision to use a centralized (vCMTS deployed in regional DC locations) versus distributed (vCMTS deployed in the hub) delivery model does not necessarily mean that only a single vCMTS deployment schema is present in the given network. For any operator, it is likely that both the distributed and centralized models will be used for different applications.

The existing residential high-density deployments translate more naturally into a distributed, CIN-based vCMTS deployment with a high number of RPDs replacing existing iCMTS DOCSIS ports at the given hub or replacing individual amplifier nodes (RPD Option 1 or Option 2).

It is also possible that both the centralized and distributed vCMTS architectures eventually converge to a certain degree, sharing the compute infrastructure but following their own service models to best fit the target customer population. Charter's enterprise services more naturally follow the centralized model, with a much lower number of strategically deployed RPDs (RPD Option 3) aggregated into a single vCMTS location per market.

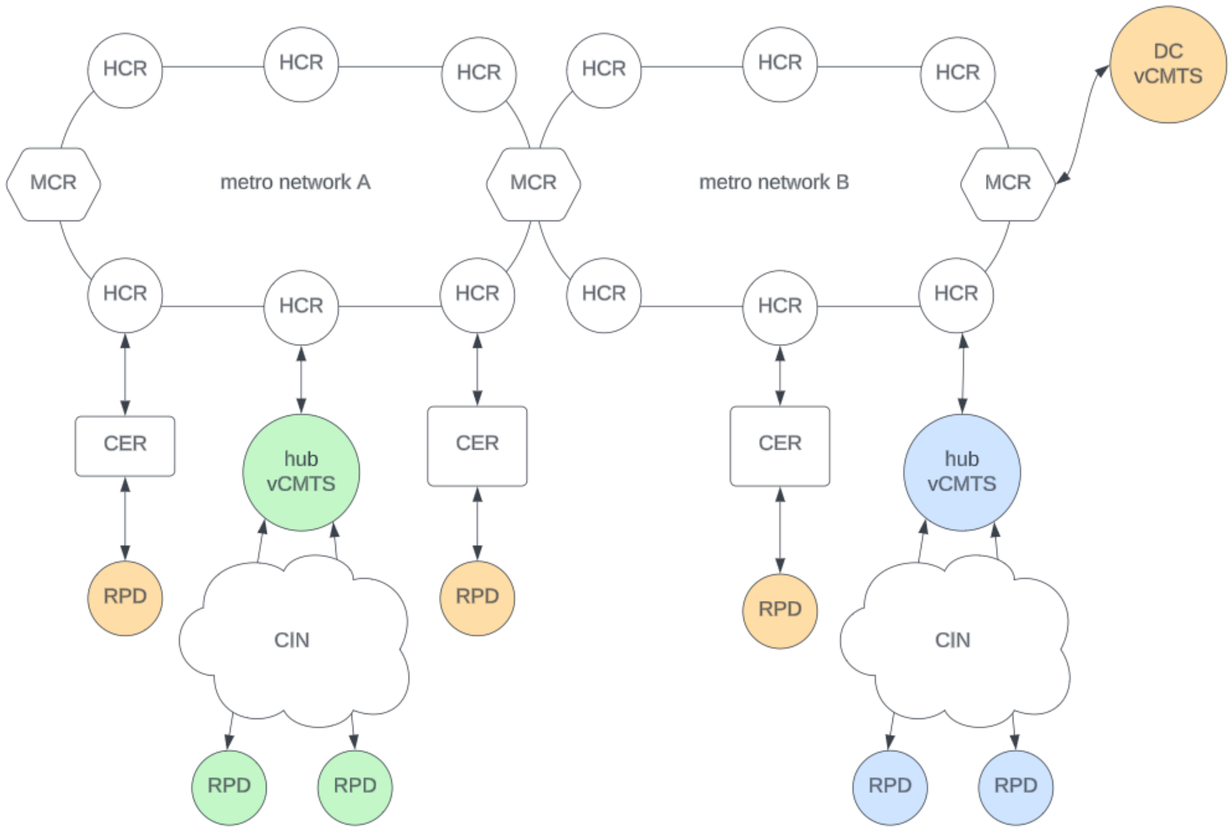

Figure 3: High-level service architecture used in the Charter network

3.1.1. RPD Placement: Option 1

This particular option represents a direct replacement for an iCMTS platform deployed today at the hub level, whereby the routing and MAC layer functions are extracted from the iCMTS chassis and re-deployed as the vCMTS either at the same hub or pushed deeper into the network, while individual RF iCMTS ports are replaced with RPDs connecting to field coaxial cabling, as shown in Figure 4.

While this option may seem a bit questionable in terms of operator advantages, it does provide a steppingstone towards the deployment of RPDs closer to the customer premises. This option presents a learning opportunity for field operations and network operation center (NOC) teams, extending the life span of the existing outdoor coaxial plant while transitioning towards a more scalable and distributed service architecture.

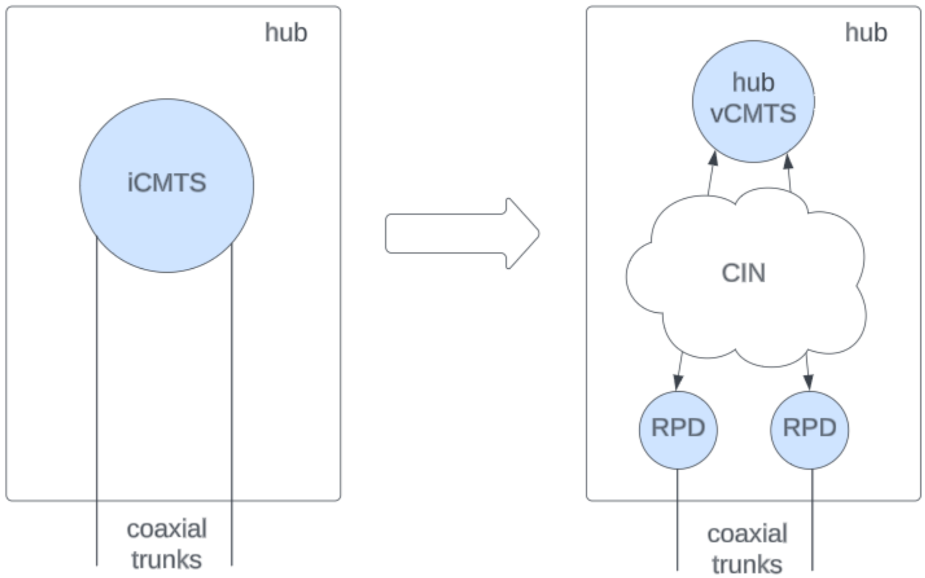

Figure 4: Option 1: Migration from iCMTS to vCMTS with hub-based RPDs

In the hub-based option, RPDs require only short distance digital optical connectivity to the hub network (CIN), resulting in none of the optical reach challenges characteristic of Option 2, as discussed in the next section. Also reducing the dependency on IP/MPLS network for connectivity.

3.1.2. RPD Placement: Option 2

Once the RPHY service architecture elements have been fully integrated into a regular network operational model, individual RPDs can be pushed further out into the field. The selection of which of the RF amplifiers in the node+N architecture gets replaced depends heavily on operator goals. A node+X (where X < N, i.e., the amplifier cascade becomes shorter) architecture eliminates some of the challenges associated with maintaining a large amplifier cascade (noise amplification and accumulation, pulse shaping, powering, management and everyday operations), but faces all the limitations of existing coaxial plant builds. The advantages and disadvantages of node+X architectures have been discussed in detail in literature already. Figure 5 shows an example of a migration from iCMTS with node+3 coaxial architecture into vCMTS with node RPDs and node+1 coaxial distribution network, providing all the advantages of shortened amplifier cascade and the use of digital feeding fibers.

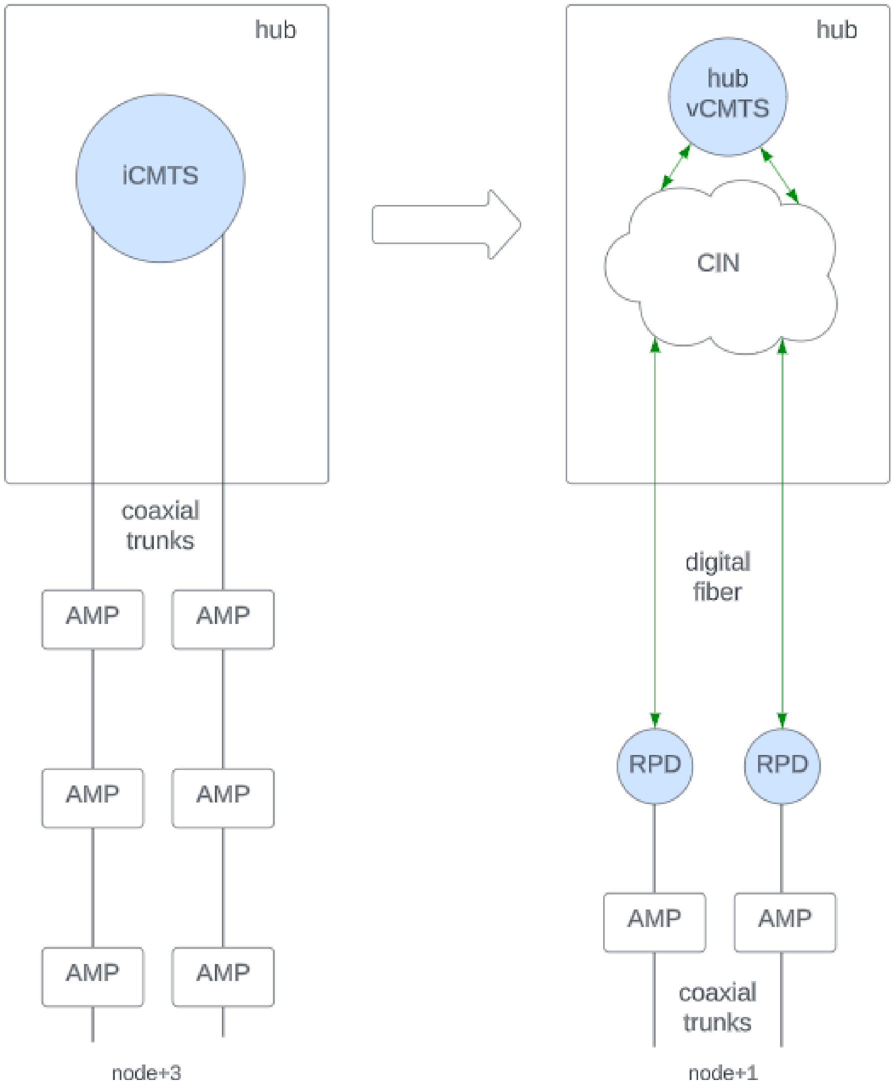

Figure 5: Option 2: Migration from iCMTS to vCMTS with node-based RPDs

It is important to note, though, that irrespective of the placement of the RPD (i.e., how many amplifiers remain in the coaxial distribution network), the RPD must be fed with a digital fiber using either a more common single mode fiber (SMF) pair with respective digital optics or a less common single SMF with the so-called bidirectional (commonly referred to as BiDi) optics. Where analog fiber is available, it can be easily converted into digital fiber by replacing active electronics at the ends of the fiber pair(s).

A single fiber path may be further reused for many RPDs by using standard DWDM optics, putting up to 20 bidirectional wavelength pairs on the same fiber path and using the methodology commonly used today in the enterprise-grade access networks. Depending on operator preferences, fixed DWDM or various types of tunable DWDM optics may be used.

Additionally, the bandwidth requirements per RPD also play an important role in the RPD deployment decisions. As DOCSIS 4 supports 10 Gbps data rates, the digital RPD backend must also support compatible data rates. With future planned DOCSIS 4.0 deployments targeting 10/1 Gbps data rates, a single RF port RPD must support at least 10G Ethernet. RPDs supporting two or more RF ports must obviously support higher aggregate data rates on their digital backend.

Intensity modulated optics are distance-limited, with typical 10G Ethernet SMF optics operating up to 80 km and with specially designed long-reach variant achieving 120 km. Along with the increase in the data rate, the viable distance decreases though, with 25G Ethernet optics typically limited to 40 km, 50G Ethernet optics limited to 20 km and 100G Ethernet optics rates at 10 km. This problem can be solved through the use of coherent optics, supporting the reach of hundreds of kilometers, though at the cost of increased optical module size and power consumption.

Considering the cost-sensitive aspect of the RPHY service architecture and the evolutionary character of the RPHY deployment, it is expected that 10G Ethernet backhaul is the most viable and cost-effective option for RPDs today. The use of longer reach and higher power 120 km optics is possible, though that might present additional thermal challenges, especially in the case of RPDs deployed in street cabinets.

3.1.3. RPD Placement: Option 3

Option 3 represents an RPD deployment model specifically tailored to community services, whereby an RPD is deployed at the customer premises (a basement attachment point, telecommunication cabinet, etc.) and feeds in the intra-building coaxial network. A large number of student dorms, hotels, multi-tenant buildings, etc., are still wired with coaxial cabling, representing the lowest cost access medium to individual residential units, eliminating the need for costly and time-consuming indoor rewiring, as shown in Figure 6.

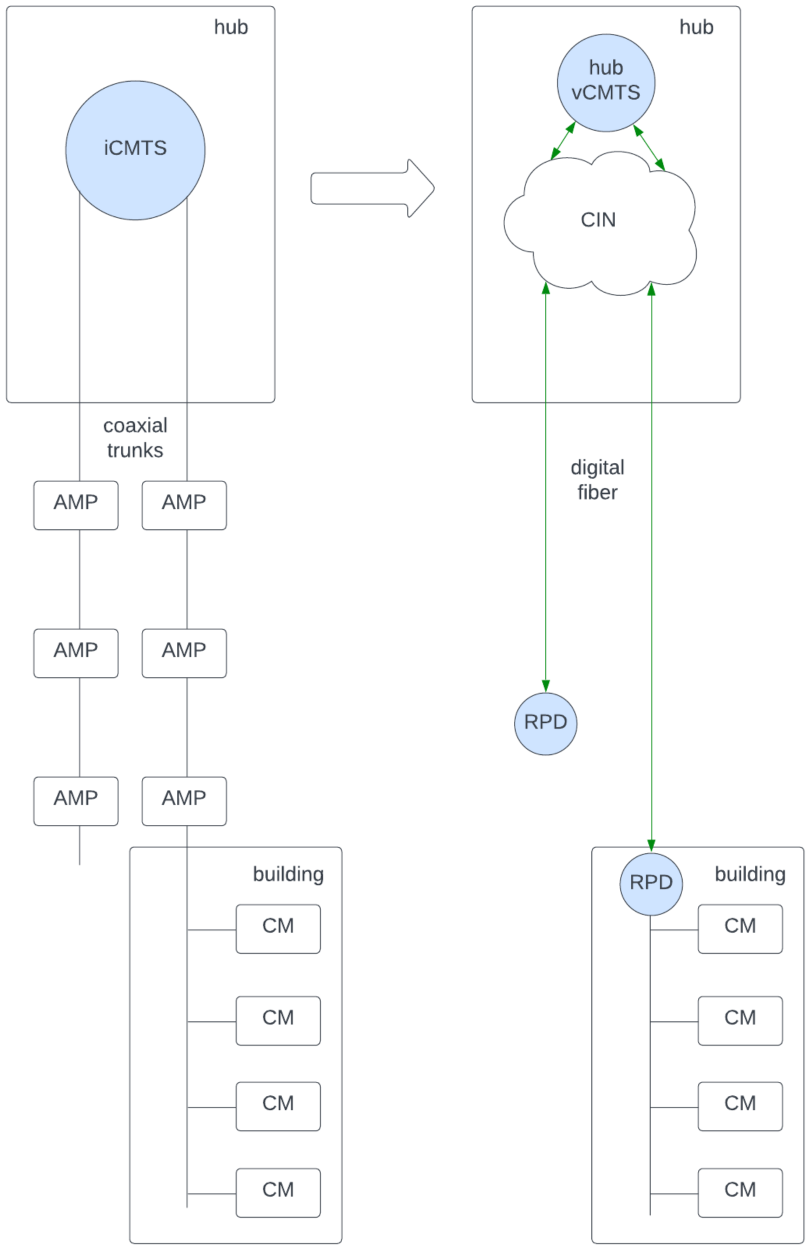

Figure 6: Option 3: In-building RPD-fed coaxial distribution plant

The challenges with the optical transport to individual RPDs as outlined in the description of Option 2 are still applicable in this case but may be circumvented to a certain degree through the deployment of a network interface device (NID), such as an aggregation router, feeding the RPD, as shown in Figure 7. In this case, the NID may implement more advanced coherent optics for extended reach and/or capacity in the digital section of the fiber (marked green in Figure 7), allowing for the use of lower cost gray optics in the RPD (marked red in Figure 7) and limiting the distance between the NID and the RPD to typical in-building scope of a few hundred feet at most. In the majority of deployments, the NID and RPD are physically collocated, allowing for the use of much cheaper and reliable short reach multi-mode optics.

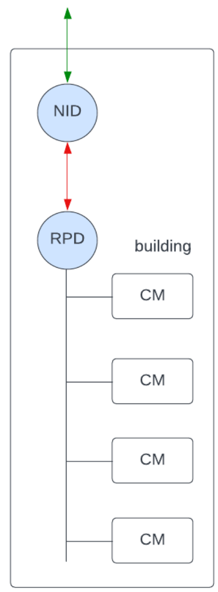

Figure 7: In-building RPD deployed with NID

3.1.4. vCMTS Placement: Centralized

In the centralized vCMTS deployment model, individual vCMTS cores are deployed in the regional or national data centers, as shown in Figure 8. In such a scenario, individual RPDs are deployed following one of the three aforementioned options, and the resulting L2TP sessions are then transported over the IP/MPLS transport network, covering metro, regional and national network segments, as applicable.

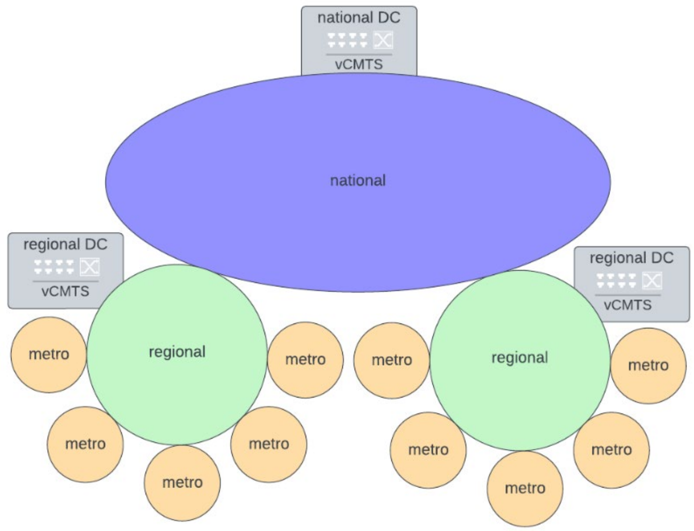

Figure 8: Regional and national DC locations for centralized vCMTS placement

The centralized vCMTS deployment model is best suited for low and medium density deployments of the RHY architecture, minimizing the vCMTS core cost and providing a deployment seed for this particular technology. Down the road, when the service density increases and a much larger number of vCMTS cores are required to serve the customer population, individual customers can be moved seamlessly to one of the distributed vCMTS units located closer to the network edge.

One of the drawbacks of this deployment model is the need to carry all the L2TP sessions from individual RPDs back to the centralized vCMTS cores, and thus aggregating a large volume of traffic at DC locations. Along with the increase in the number of supported RPDs, the aggregate volume of traffic ingressing and egressing the DC may become excessive, requiring transition to the distributed vCMTS deployment model covered in the following section.

Note that this deployment model does not affect the traffic flow in the network with any significance. Most networks are built around the hub-and-spoke architecture, carrying all Internet-bound traffic towards POP as effectively and quickly as possible. With the centralized vCMTS code deployment, such Internet-bound traffic is carried in L2TP sessions towards the target vCMTS core within the DC, where it is then decapsulated, and egresses the DC towards the typically co-located POP.

3.1.5. vCMTS Placement: Distributed

The distributed vCMTS placement model caters specifically to high-density customer populations, where a large number of vCMTS cores deployed at the hub limit the logical distance individual L2TP sessions have to travel over the IP/MPLS transport network to just the local CIN. L2TP is a non-zero overhead transport protocol and a very high volume of customer traffic transported over L2TP does have an impact on the network capacity. In the upstream direction where individual customer datagrams are typically small, the L2TP encapsulation results in a higher overhead when compared to the downstream direction.

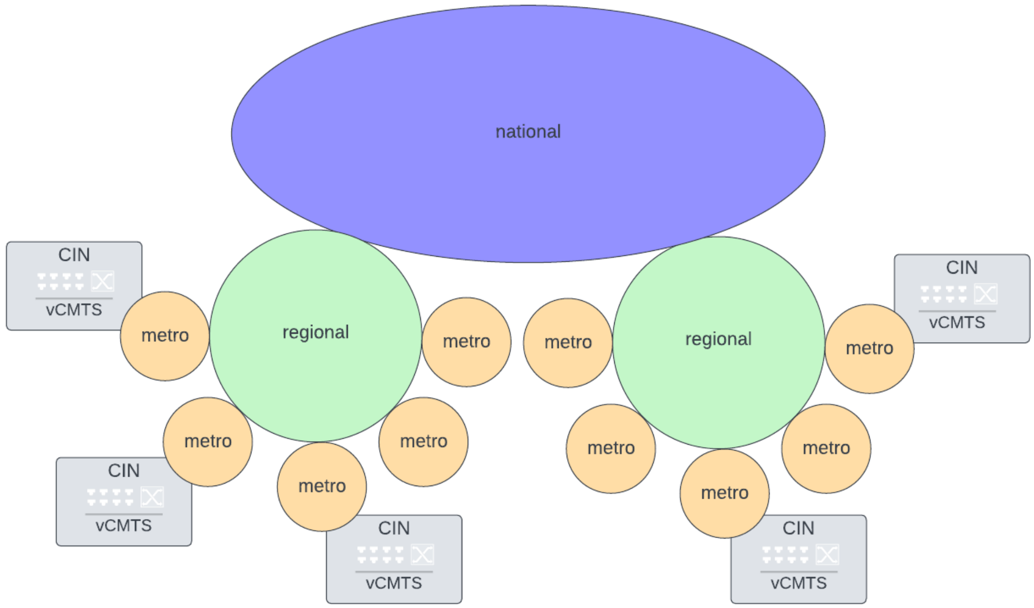

Figure 9: Distributed vCMTS placement

3.2. PTP Architecture

The Precision Time Protocol (PTP), standardized in IEEE Std 1588, defines an interoperable protocol for the distribution of phase, frequency and/or time of day information across various types of communication networks, including a range of specific network profiles. At Charter, PTP is distributed using the unicast IP encapsulation option (compliant with ITU-T G.8275.2 profile), traversing the IP/MPLS core network between the GPS-fed PTP Grandmaster clocks (GMC) attached to the metro core routers (MCRs), vCMTS cores collocated within the DC or local hubs (depending on the market and the delivered services), and RPDs deployed using one of the three options discussed before.

PTP traffic further traverses hub core routers (HCRs) and boundary clocks (BC) deployed where individual PTP domains need to be segmented to minimize the impact of packet delay variation (PDV) on the resulting PTP service quality. The high-level view of the aforementioned PTP architecture is shown in Figure 10.

Each metro network is equipped with two redundant GMC units, providing protected PTP sources for all PTP clients deployed within the given metro network. The PTP architecture comprises two layers, core PTP network (shown in red in Figure 10) and distribution PTP network (shown in green in Figure 10), providing a very high number of PTP clients without taxing GMC units with direct client connections.

Through the proper IP prefix allocation and reachability limitation via access lists and/or prefix visibility scoping, only PTP BC units are allowed to connect to GMC units within the given metro network. Individual PTP clients are permitted to connect to PTP BC units only, protecting GMC units from supporting thousands of PTP clients in large-scale deployments.

Additionally, PTP traffic is prohibited from traversing backbone networks (interconnecting individual metro networks, not shown in Figure 10), effectively constraining PTP traffic exclusively to the given metro. This self-imposed limitation increases slightly the number of GMC units that need to be deployed but constrains PDV to much smaller values typical to metro networks, resulting in improved PTP performance. Moreover, backbone routers are typically more challenging to optimize for PTP transfer, even in the presence of properly designed CoS/QoS policies.

Furthermore, to minimize PDV for PTP traffic and optimize its performance across all metro networks, updates to network QoS/CoS policies are needed, prioritizing PTP traffic over any customer traffic by configuring individual routing platforms to put PTP traffic into real-time queues.

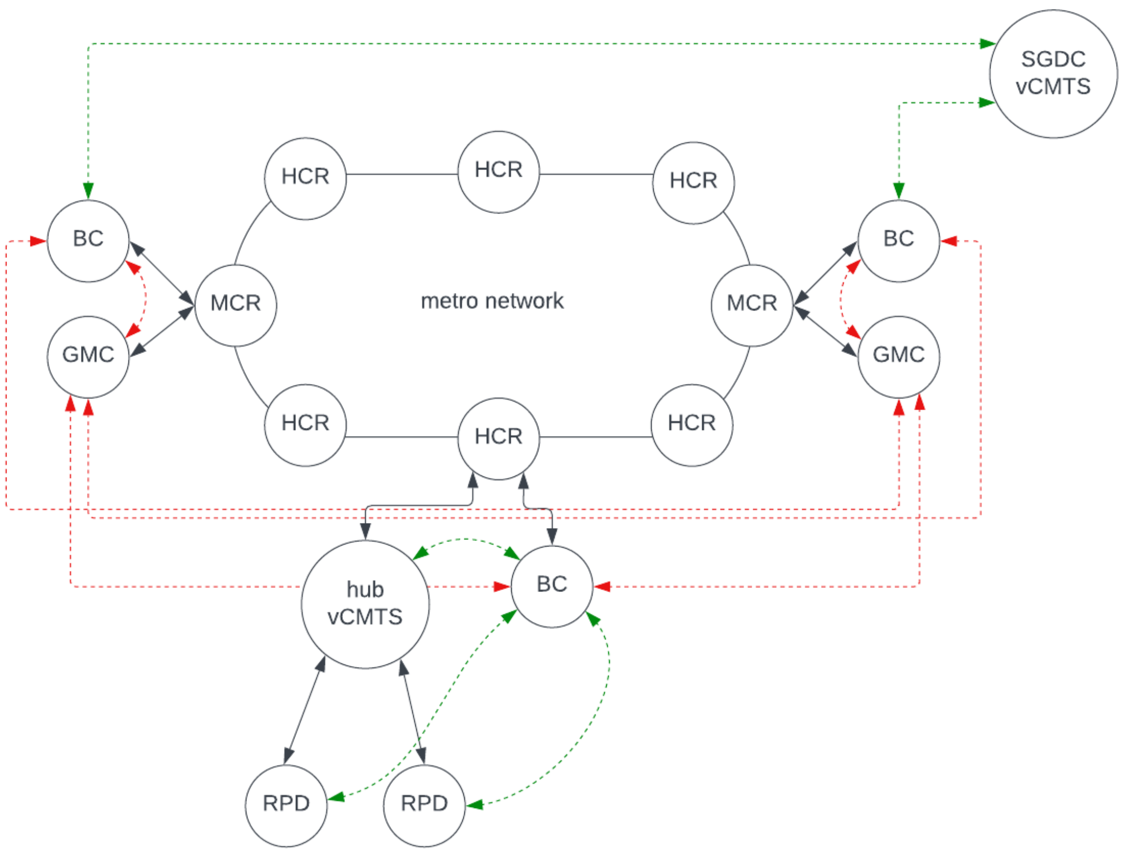

Figure 10: PTP distribution architecture

PTP traffic is critical for the proper operation of data (digital) services delivered using the RPHY architecture, where individual RPDs must be tightly synchronized (in terms of time of day as well as frequency, with phase synchronization not required) with the vCMTS core they are connected to. The synchronization is required, since the DOCSIS scheduler is hosted within the vCMTS cores deployed within DC locations, while the RPDs are deployed outside of DC locations, with the distance between vCMTS cores and individual RPDs ranging from a few miles to a few hundreds of miles in practical deployments. The said distance is primarily constrained by PDV that PTP traffic is subject to, as well as the network design and segmentation.

PTP is not required for video services, given that typical RPDs are capable of operating in the so-called asynchronous mode, where there is only downstream traffic transmitted towards RPDs, requiring no synchronization between the RPD and the vCMTS core.

3.3. RPD Spectrum Plans

Each RPD can be configured independently with a specific RF spectrum plan, including any combination of digital (DOCSIS) and analog (video QAM) services. To keep things simple, a finite number of RPD RF spectrum plans are typically defined.

For Spectrum Enterprise applications, there are two primary RF spectrum plans in use: the low-split (LS) spectrum plan with Fiber Connect Plus (FC+) analog video services (see Figure 11) and the high-split (HS) spectrum plan with digital services only (no analog video, see Figure 12).

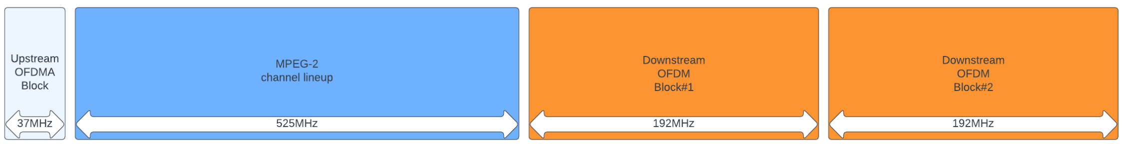

Figure 11: Low-split RF plan with FC+ analog video services

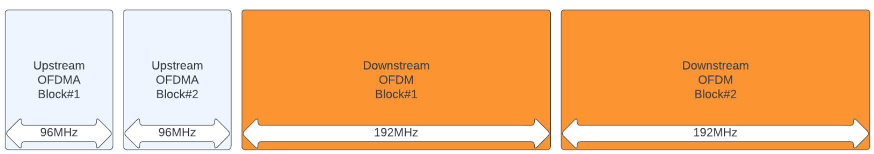

Figure 12: High-split RF plan with no analog video services

In comparison of a low split spectrum plan which delivers a complete analog video channel lineup (about 525 MHz worth) with MPEG2-encoded content, providing sufficient RF capacity for two downstream 192 MHz-wide OFDM blocks and a single reduced size (so-called *skinny*) upstream 37 MHz-wide OFDMA block, with no DOCSIS primary QAM channel. There is also no ATDMA in the upstream since all connected DOCSIS CMs are D3.1 compliant and there is no requirement for any backward compatibility at all. Apart from providing analog video services, this spectrum plan can theoretically provide around 2.5 Gbps downstream and around 180 Mbps upstream data rates, assuming 4k QAM constellations under sufficient modulation error ratio (MER).

The HS spectrum plan eliminates the analog video channel lineup, making much lower-frequency RF spectrum available for the upstream OFMDA blocks. Effectively, this spectrum plan provides sufficient RF capacity for two downstream 192 MHz-wide OFDM blocks and two full size upstream 96 MHz-wide OFDMA blocks. Therefore, this full digital spectrum plan provides around 2.5 Gbps downstream and around 1.2 Gbps upstream data rates, assuming 4k QAM constellations under sufficient MER (41 dB or more).

During testing, with properly engineered intra-building coaxial plant, much higher MER was achievable, reaching the average of 49 dB, as shown in Figure 13. Given the relatively well-controlled character of the intra-building coaxial plant, as well as much less challenging environmental conditions, maintaining the minimum MER required for 4K QAM constellation is not expected to be very complicated.

```
MER measurement for CM
         RxMER - avg(log(S/N)): 49.12 dB
          RxMER - log(avg(S/N)): 49.21 dB
--------|----------|------------------------------------------------------------
  RxMER |    Max   |         RxMER Histogram (Subcarriers)
    dB  |  Bitload |
--------|----------|--------10--------20--------30--------40--------50--------60-
  14 db |          |
  15 db |   16-QAM |
  16 db |          |
  17 db |          |
  18 db |          |
  19 db |          |
  20 db |          |
  21 db |          |
  22 db |          |
  23 db |          |
  24 db |          |
  25 db |          |
  26 db |          |
  27 db |  256-QAM |
  28 db |          |
  29 db |          |
  30 db |  512-QAM |
  31 db |          |
  32 db |          |
  33 db |          |
  34 db | 1024-QAM |
  35 db |          |
  36 db |          |
  37 db | 2048-QAM |
  38 db |          |
  39 db |          |
  40 db |          |
  41 db | 4096-QAM |
  42 db |          |
  43 db |          |
  44 db |          |
  45 db |          |
  46 db |          |
  47 db |          |
  48 db |          |*******
  49 db |          |****************
  50 db |          |**********
  51 db |          |*
  52 db |          |

* represents 100 subcarriers
```

Figure 13: MER for test location in production network

The application of the specific RF spectrum plan in the given customer location (site) depends on the requirement to support the analog video service, whereby customers using analog TV receivers are provided with the LS spectrum plan and the other customers are defaulted to the HS spectrum plan. The main difference between these two spectrum plans lies, obviously, in the lack of support for analog video channel lineup in the HS spectrum plan as well as higher upstream capacity for digital services. The downstream capacity for both spectrum plans was designed to be exactly the same, utilizing the full capability of existing off-the-shelf DOCSIS 3.1 cable modems.

4. RPHY Performance Measurement Scenarios

All of the results presented in this paper were collected in a live production network where a test customer site with two RPDs was deployed with distinct RF plans: RPD-A with the low-split (LS) spectrum plan with FC+ analog video services, and RPD-B with the high-split (HS) spectrum plan with digital services only (no analog video). A vCMTS core used for these tests was deployed in the regional DC (centralized deployment) model and was shared by a number of FC+ video-only RPDs.

In order to eliminate any testing variability associated with the use of public infrastructure (public speed test servers, for example) and the intermediate public Internet, all testing was conducted to an iPerf 3 server connected to the very same routing platform the vCMTS core is connected to. A 10 Gbps+ capable server, connected using four 10GE interfaces, was used to eliminate any potential bandwidth bottlenecks and provide an optimum path from the multi-chassis routing platform stack.

At the test customer premise, a 10 Gbps capable NID terminated the DWDM optical link and then fed the RPD via the inline impairment generator. The inline impairment generator is capable of adding IP-level latency (see Figure 14) and/or PDV to all or selected packets. For example, it is possible to add PDV to PTP traffic while not affecting other traffic (see Figure 15). This particular device is used in the test bed to emulate different options for network distance between the RPD and the vCMTS and observe the impact of the said network distance on the service performance (data throughput as measured using the iPerf 3 server).

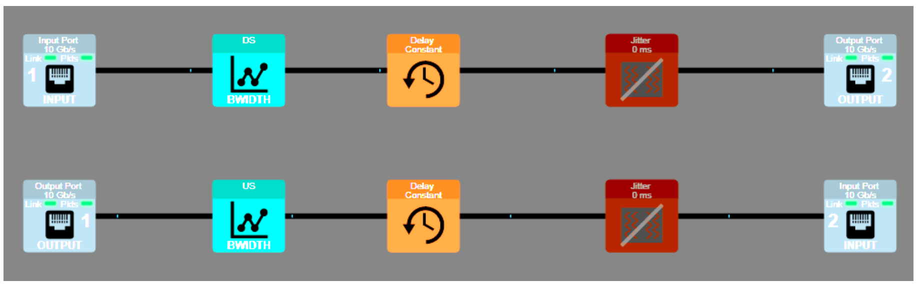

Figure 14: Impairment generator, latency / PDV added to data path to all IP traffic

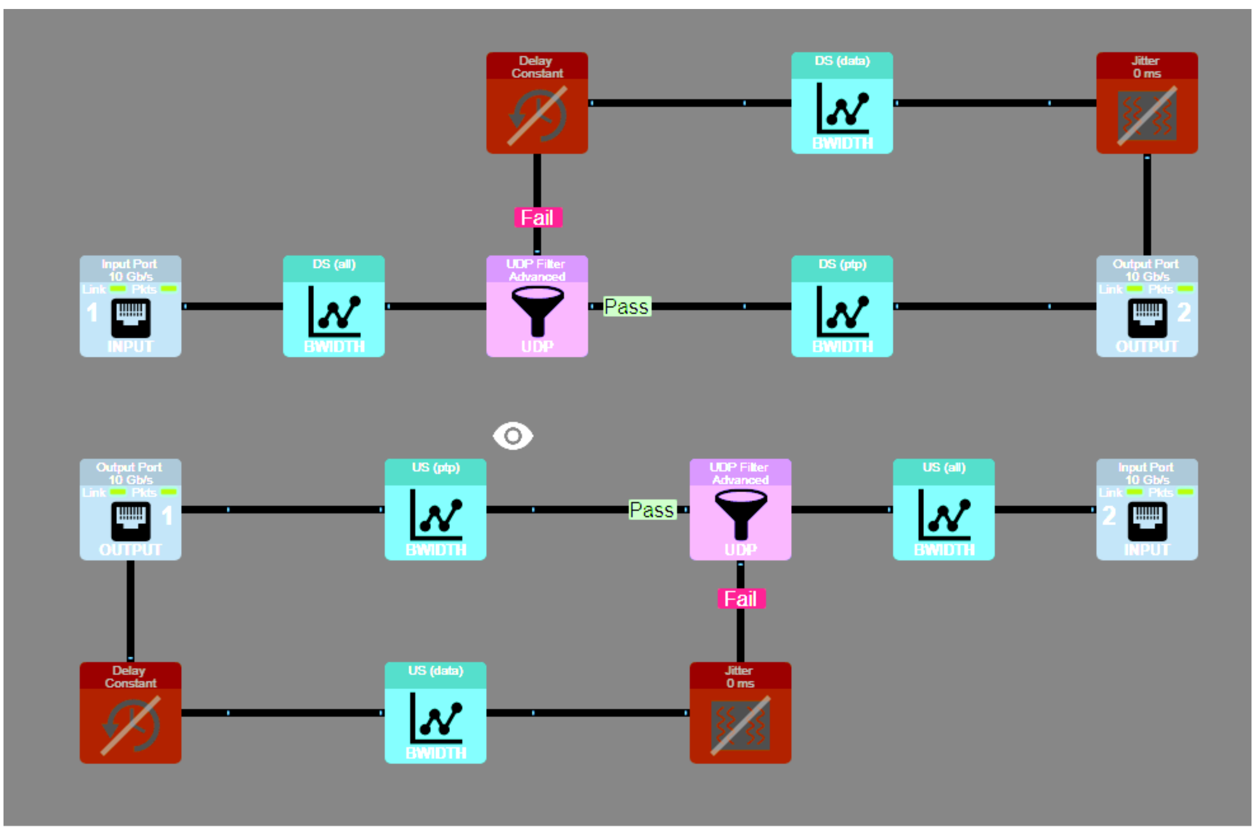

Figure 15: Impairment generator, latency / PDV added to data path to PTP traffic only

RPD A/B then fed individual coaxial distribution plants with connected DOCSIS 3.1 CMs. One of the DOCSIS CMs then fed a test laptop via 2.5GE connection to maximize the available throughput.

The resulting test topology is shown in Figure 16.

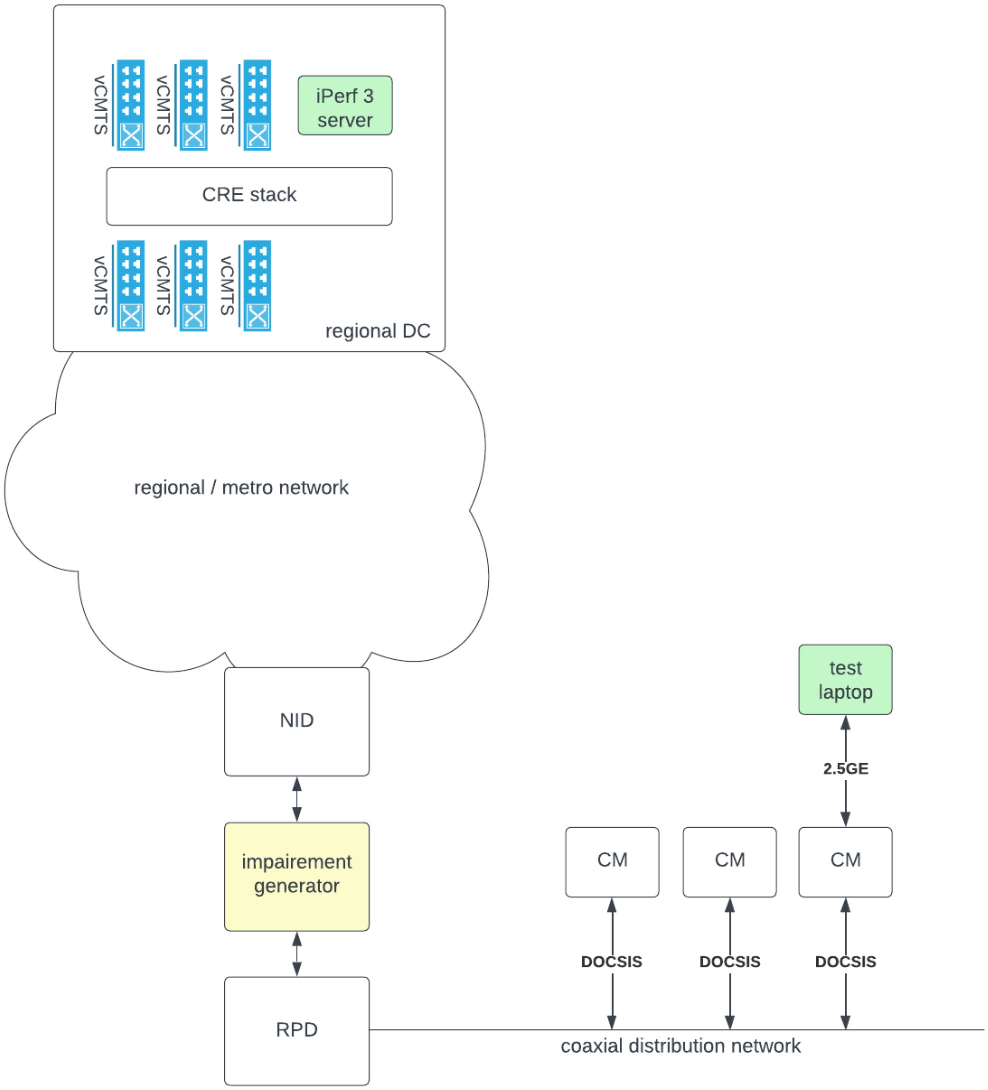

Figure 16: Test topology with in-line impairment generator

In each scenario described below, the service performance (throughput) was measured for a number of different extra latency points, corresponding to the distance between the vCMTS and the RPD equal to: 0, 2.5, 126, 158, 190, 221, 253, 285, 316, 348, 379, 506, 632 and 1264 miles. These individual distance points were derived from added one-way latency assuming the fiber propagation delay of 1.5ns/foot.. Measurements past the 1,300 mile mark were limited by the ability of the vCMTS to maintain communication with the RPD. Each measurement was taken at least five times and averaged to eliminate the potential impact of any single measurement outlier. Measurement at zero mile-added latency corresponds to inherent network latency (regional and metro) associated with the test data circuit.

Additionally, the impact of network latency asymmetry was examined. When enabled, the operation of ECMP may result in the downstream and upstream data streams taking different paths across the provider network, resulting in different latency. This problem is shown in Figure 17 , where the downstream direction (green path) is different than the upstream direction (red path), typically resulting in a different network latency. Asymmetric network latency does have an impact on the PTP, and specifically, on the synchronization precision since PTP is not able to compensate for the network latency asymmetry.

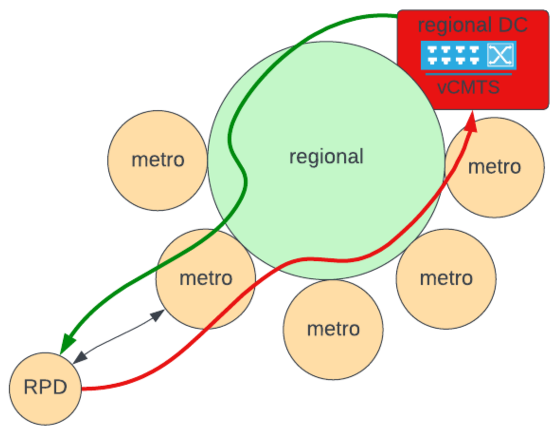

Figure 17: Asymmetric routing problem

Three scenarios are examined in each case: a fully symmetric (ideal) network, a network with 100 miles of asymmetry (around 0.8ms, a value that can be easily accumulated in the production network) and the worst-case scenario of 700 miles asymmetry (around 5.6ms) observed to date in the production network.

All the presented data rate figures represent L3 throughput and are inherently lower than the maximum L2 Ethernet data rate for the link between the test laptop and the DOCSIS CM (2.5 Gbps).

4.1. Impact of Network Latency on HS Services

The observed downstream and upstream service performance in the HS service scenario is shown in Figure 18 for scenarios with symmetric and asymmetric (100 miles and 700 miles worth of latency asymmetry) conditions.

The downstream direction performance is mostly independent from the network latency symmetry, showing the highest performance for the symmetric latency, and decreasing slightly along with the increase in the latency asymmetry. This is an expected behavior since the downstream direction does not require any scheduling and the observable throughput is only limited by the performance of the TCP itself. There is also very limited dependence on the network latency value (network diameter size), likely due to the TCP windowing as well as the use of parallel test streams filling the test circuit very efficiently.

The upstream direction exhibits much stronger correlation with the increase in the network latency, where with sub 400 miles of latency, the throughput decreases from ~1250 to ~1050 Mbps and then falls off much quicker, reaching around 550 Mbps for 1264 miles marker. The dependency on the network latency asymmetry is less expressed than in the downstream direction, likely due to the extra buffering available in the upstream direction along the data path, and especially in the DOCSIS CM.

3.2 Impact of Network Latency on LS Services

The observed downstream and upstream service performance in the LS service scenario is shown in Figure 19 for scenarios with symmetric and asymmetric (100 miles and 700 miles worth of latency asymmetry) conditions.

The observations for the HS covered in the previous section equally apply to the LS scenario, with the only major observable difference in the upstream data rates supported by the system. The decrease in the upstream throughput comes from the use of very limited upstream RF spectrum (single skinny OFMDA block) when compared to the HS scenario with two full OFDMA blocks.

3.3 Impact of PTP Path Pinning

In this scenario, the PTP traffic was isolated from the data traffic and was subject to symmetric network latency, while the data traffic was subject to the previously examined three network latency scenarios: symmetric latency, 100 miles of network latency asymmetry and 700 miles of network latency asymmetry. This scenario corresponds to the PTP traffic engineering design, in which PTP traffic is transported between RPD and the source clock in a bidirectional data tunnel, eliminating the side effect of ECMP operation in the network. While this approach does not eliminate the impact of the network latency itself, it helps manage the impact of network latency asymmetry on PTP traffic, eliminating any potential synchronization inaccuracies between the RPD and the vCMTS core.

Since the downstream transmission direction is not impacted by the RPD/vCMTS synchronization accuracy, only the upstream direction examined for the HS scenario is shown in Figure 20.

Comparing the upstream performance for the HS scenario without PTP path pinning (see Figure 18), LS scenario without PTP path pinning (see Figure 19) and both scenarios with PTP path pinning (see Figure 20), it is visible that the elimination of the network latency asymmetry for PTP traffic improves the throughput stability in the upstream direction. This makes, in turn, the upstream throughput less sensitive to the network latency on the data plane.

PTP path pinning does require additional network configuration, where pre-computed and symmetric data paths are built for PTP traffic between the PTP source (BC) and PTP client (RPD, vCMTS) over the IP/MPLS transport network. Various approaches may be taken as far as the implementation aspect is concerned, including selective MPLS, RSVP-TE, etc., and each method has its own challenges and scaling requirements.

In the simplest implementation, each RPD would have its own PTP path built across the transport network. While it is a simple approach, it is not scalable, since it would require thousands of PTP paths to be configured and maintained, presenting an additional operational challenge.

A more scalable approach builds a PTP path between the given hub and the clock source, aggregating all PTP traffic from all PTP clients served from the given hub. In this approach, a single PTP path can serve hundreds to thousands of PTP clients, depending on the scale of the given RPHY deployment.

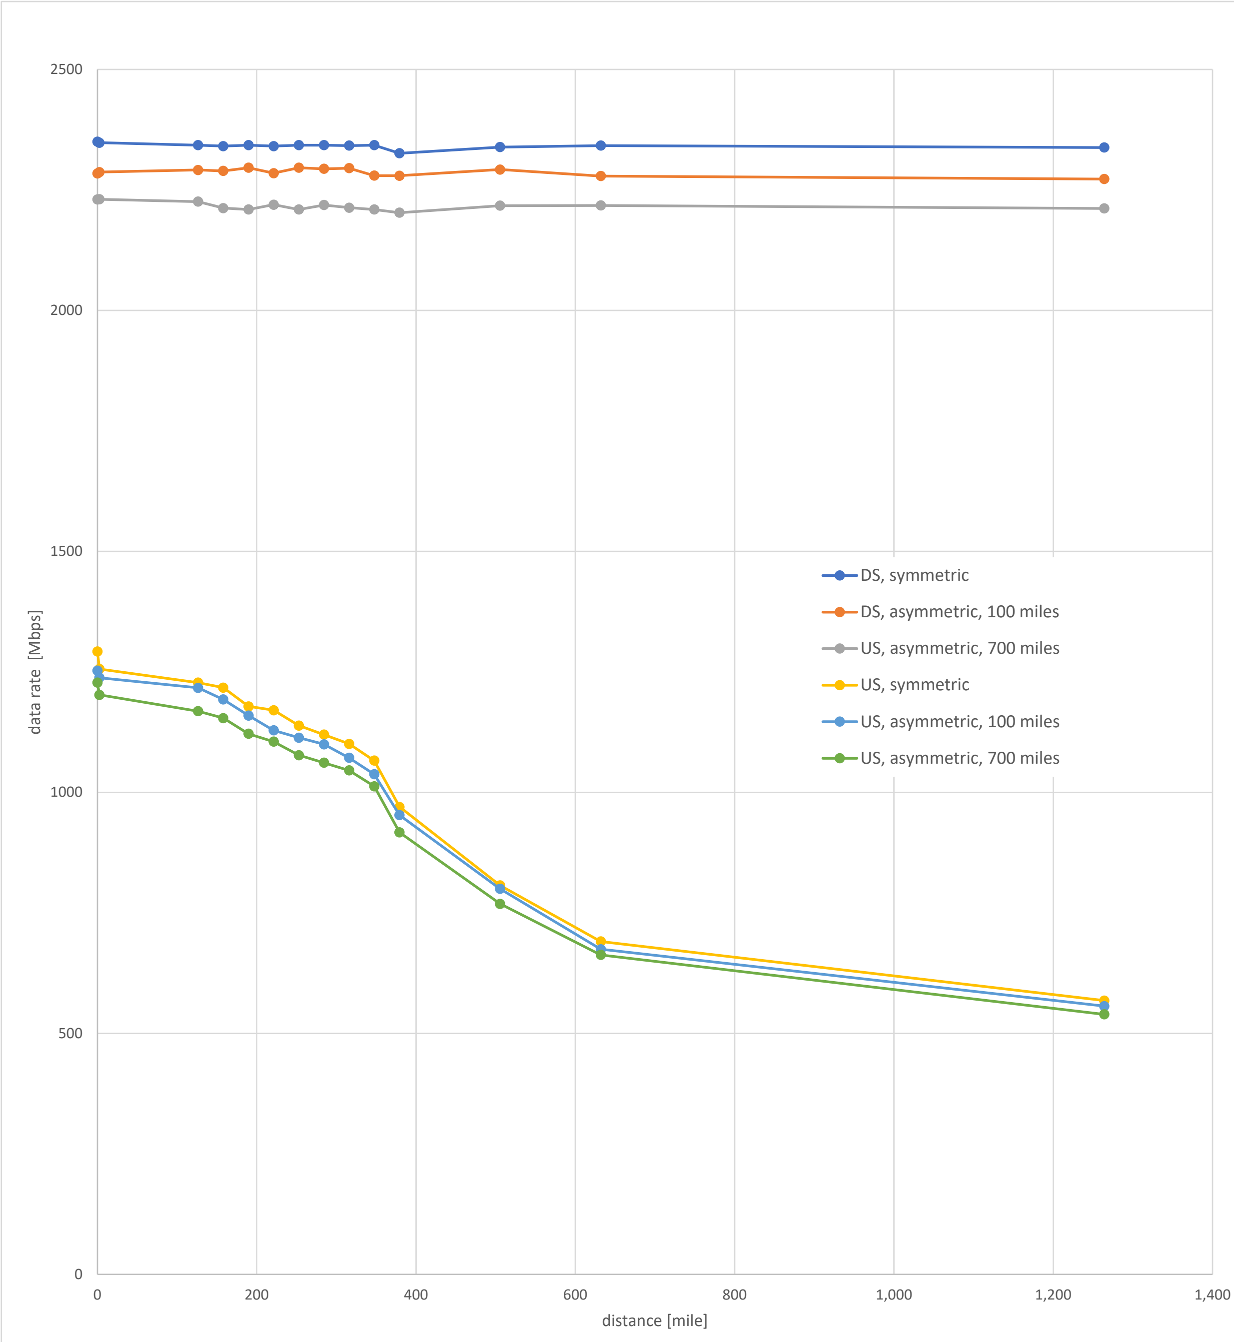

Figure 18: Upstream and downstream HS throughput with 0, 100, and 700 miles network asymmetry

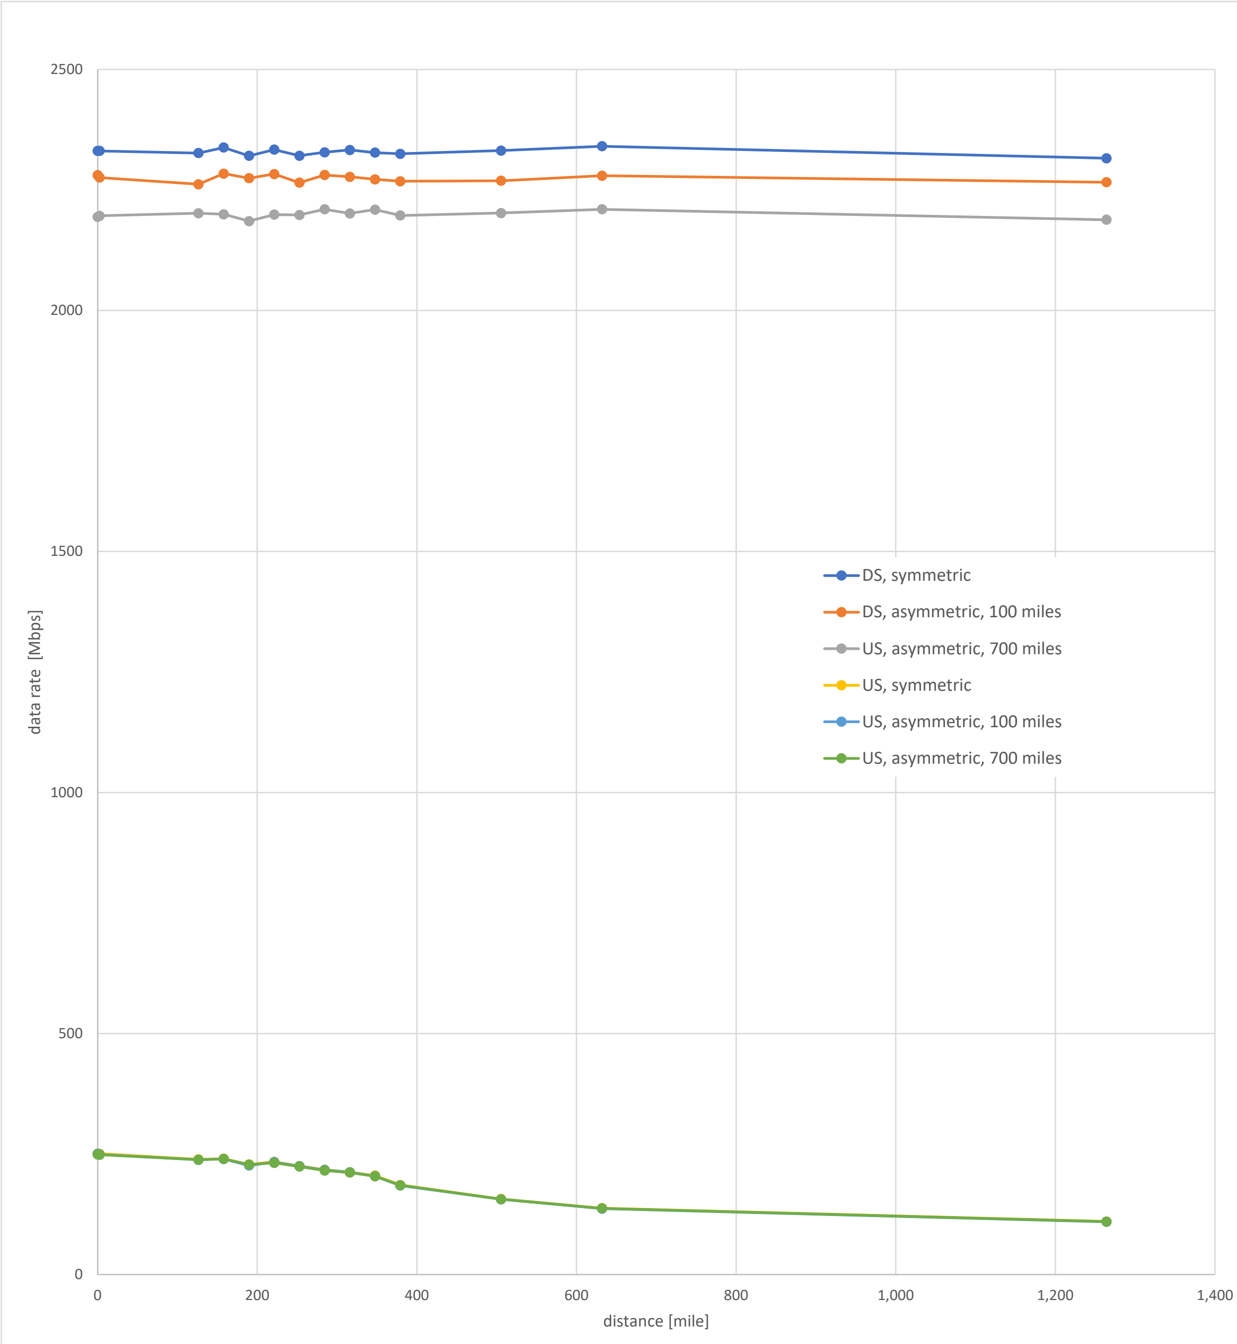

Figure 19: Upstream and downstream LS throughput with 0, 100, and 700 miles network asymmetry

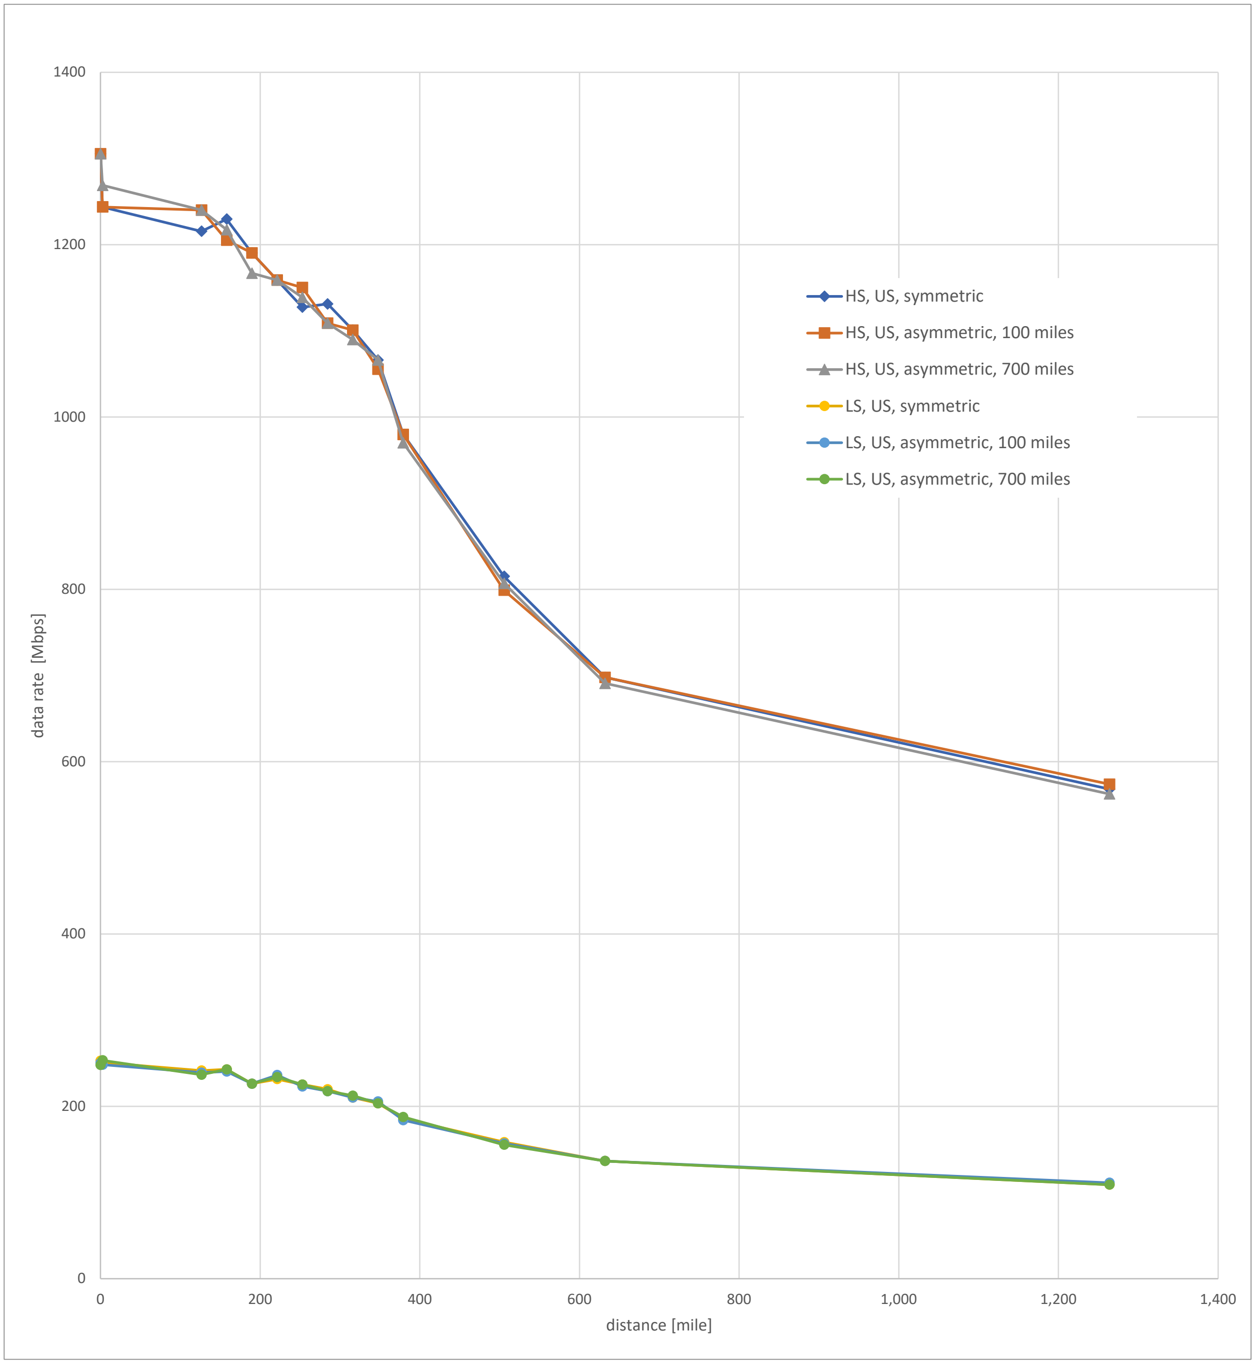

Figure 20: Upstream data throughput with PTP path pinning, HS service plan

5. Conclusions

The RPHY service architecture can operate at distances between the vCMTS and RPD, far exceeding typical DOCSIS distances supported by the latest generation of iCMTS. The downstream direction is not adversely affected by the increased distance between the vCMTS and RPD, while the upstream direction exhibits a drop in performance past ~400 miles of distance, likely related to the ability of the vCMTS to schedule the upstream direction effectively.

An effective service area (distance between the vCMTS and the RPD) of around 200 miles provides an opportunity to retire a portion of existing iCMTS units and collapse these DOCSIS plants into a much smaller number of vCMTS nodes. The exact savings will be largely operator-dependent and vary from market to market, driven primarily by the number of iCMTS ports being removed, customer density, availability of IP/MPLS fiber transport between individual hubs and the vCMTS location, as well as any additional work done on the outside coaxial distribution network.

When replacing existing iCMTS ports with RPD ports at the hub (Option 1), the perceived gains from the RPHY architecture are somewhat limited, representing only a disaggregation of the iCMTS platform itself. Along with the move of individual RPDs into the field and decreasing the size of the amplifier cascade, or the deployment of RPDs at the customer premises (Option 3), the aggregate data rates within the architecture will increase and the number of RPDs will grow as well, especially as individual serving groups are further split to increase their capacity.

There is also an observable improvement in the upstream direction performance when the PTP path pinning is used to eliminate the impact of network latency asymmetry on the vCMTS upstream scheduler.

6. Appendix I – Historical Calculations on REQ-GNT Delays

For reference, here are some delay calculations for RPHY done in 2004. These numbers are out of date and for lower speed interfaces. However, this is included to capture historical documents and to show how the inner structure of the CMTS works. These numbers can be combined with Appendix I of [13].

Remote PHY Delay Budget	EDCS-376514
Author Name, e-mail	John T. Chapman, jchapman@cisco.com
Date	July 20, 2004 (original: May 4, 2004)
Filename	remote-phy-delay-040720a.xls

Inputs:		
DS bit rate	40	Mbps
DS MPEGoIP packet size	1366	bytes
DS MAP Min Length	2000	usec
DS MAP Size	200	bytes
DS Early MAP Release	yes	yes/no
US bit rate	30	Mbps
US REQ packet Size	250	bytes
US Concatenation Size	8000	bytes
US Early REQ Release	yes	yes/no
Ethernet Switch Intrinsic Delay	25	usec
Ethernet Switch Hops	1	0 to n

Outputs:	**Today**	**Remote PHY**	**Delta**	**units**
REQ-GNT Delay Sub-total	**5107**	**5931**	824	usec
% increase of delay budget			16	%
Total With No Concatenation				
MAP Interval Time	2000	2000	0	usec
Rounded up to MAP intervals	**6000**	**6000**	0	usec
% increase of delay budget			0	%
Performance				
net PPS rate	167	167	0	PPS
US bit rate with 1500 byte packets	**2.0**	**2.0**	0	Mbps
Total With Concatenation				
MAP Interval Time	2133	2133	0	usec
Rounded up to MAP intervals	**6133**	**6133**	0	usec
% increase of delay budget			0	%
Performance				
net PPS rate	163	163	0	PPS
US bit rate at max concatenation	**10.4**	**10.4**	0	Mbps

REQ-GNT Delay Element (usec)	Today		Remote PHY	
	Sub-total	Total	Sub-total	Total
HFC Plant Round Trip Delay	1600	1600	1600	1600
CM Delay		267		267
CM intrinsic delay	200		200	
CM packet tx time (onto QAM)	67		67	
CMTS Delay		3240		
CMTS US PHY delay	200			
CMTS US REQ Queuing & Processing	1000			
CMTS MAP Advance Margin	1000			
CMTS Downstream PHY delay	1000			
CMTS MAP tx time (onto QAM)	40			
MAC Card Delay				2036
CMTS US REQ Queuing & Processing			1000	
CMTS MAP Advance Margin			1000	
CMTS Packet Encapsulation Delay			0	
CMTS switching intrinsic delay			25	
CMTS MAP tx time (onto GE)			11	
DS Remote PHY Delay				1738
Edge QAM intrinsic delay			400	
Edge QAM Queuing Delay			263	
Edge QAM Downstream PHY delay			1000	
Edge QAM MPEG-TS packet tx time			75	
US Remote PHY Delay				227
Remote US PHYdelay			200	
Remote US intrinsic delay			25	
Remote US packet tx time			2	
GE Switch Delay				63
GE Switch US instrinsic delay			25	
GE Switch US packet tx time			2	
GE Switch DS instrinsic delay			25	
GE Switch DS packet tx time			11	
REQ-GNT Delay Sub-total		**5107**		**5931**

Notes:

DS FEC delay of ~1 msec is included in the MAP advance time

MAP Size is not exactly 2ms. Instead, it is equal in length to an integral number of packets and IE slots.

Abbreviations

BC	boundary clock
BiDi	bidirectional
CAPEX	capital investment
CER	commercial edge router
CIN	content integrated network
CM	cable modem
CMTS	cable modem terminal system
CoS	class of service
DC	data center
DOCSIS	data over cable service interface specifications
DWDM	dense wavelength division multiplexing
FC+	fiber connect plus
GMC	grandmaster clock
HCR	hub core router
HD	high definition
HFC	hybrid fiber-coax
HS	high split
L2TP	layer 2 transport protocol
LS	low split
MCR	metro core router
MER	modulation error ratio
MMF	multi mode fiber
MPLS	multi protocol label switching
NID	network interface device
NOC	network operation center
OPEX	operational expense
OPEX	operational expense
PDV	packet delay variation
PHY	physical layer
PTP	precision time protocol
QoS	quality of service
RF	radio frequency
RPD	remote PHY device
RPHY	remote PHY
SMF	single mode fiber
vCMTS	virtual CMTS
iCMTS	integrated CMTS

Definitions

Downstream	Information flowing from the hub to the user
Upstream	Information flowing from the user to the hub

Bibliography and References

[1] John T Chapman, "Packet Fiber Node", US Patent 7,209,442, filed Jun 27, 2001, issued April 24, 2007 [link]

[2] John T Chapman, "Downstream Remote Physical Interface for Modular CMTS (DEPI)", US Patent 7,688,828, filed May 17, 2005, issued March 30, 2010 [link]

[3] John T Chapman, "Upstream physical interface for modular cable modem termination system (UEPI), US Patent 7,639,617, filed May 17, 2005, issued Dec 29, 2009 [link]

[4] John T. Chapman, "Generic Control Protocol (GCP)", US Patent 9246700, filed Sept 12, 2023, issued Jan 26, 2016 [link]

[5] John T. Chapman et. al, "Managing time offset and frequency drift in asynchronous docsis remote phy network environments (DTI)", US Patent 9,722,739, filed April 15, 2015, issued Aug 1, 2017 [link]

[6] John T Chapman, Alon Berstein, et. al., "Virtual Cable Modem Termination System", US Patent 9,369,339, filed Sept 30, 2013, issued Jun 14, 2016 [link]

[7] John T Chapman, "Modular CMTS Architecture", NCTA Spring Technical Forum, 2005 [link]

[8] John T. Chapman, "DOCSIS Remote PHY", *SCTE Cable-Tec Expo Fall Technical Forum*, Atlanta, GA, Oct 2013.

[9] John T. Chapman, "Remote PHY for Converged DOCSIS, Video and OOB", *NCTA/SCTE Technical Forum*, Los Angeles, Jun, 2014. [link]

[10] CableLabs, Modular Headend Architecture Specifications, 2005 [link]

[11] CableLabs, CDOCSIS Specifications [link]

[12] CableLabs, Remote PHY Specifications (MHAv2) [link]

[13] CableLabs, "Downstream External PHY Interface Specification", CM-SP-DEPI-I08-100611 [link]

[14] CableLabs, "DOCSIS 3.1 Physical Layer Specification", Section 1.2.1 Broadband Access Network, CM-SP-PHYv3.1-I20-230419 [link]

[15] John T. Chapman, Gerry White, Hang Jin, "Impact of CCAP to CM Distance in a Remote PHY Architecture", *SCTE Cable-Tec Expo Fall Technical Forum*, New Orleans, Oct, 2015. [link]

[16] IEEE Std 1588: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems; Institute of Electrical and Electronics Engineers, published 2019

[17] ITU-T G.8275.2: Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network, International Telecommunication Union, Telecommunication Standardization Sector, published 2022