

Advantage Remote PHY

An Analysis of Remote PHY vs Analogue Deep Fiber Total Cost of Ownership

Introduction

Over the last 20 years, the cable access network's DOCSIS system and Hybrid Fiber Coaxial (HFC) plant have steadily evolved, supporting the access network's growth and viability as a leading Internet access technology, with more than 172 million broadband subscribers worldwide. The current iteration of the cable access network, the Converged Cable Access Platform (CCAP), is widely deployed and has begun to deliver the latest 3.1 version of DOCSIS. Each transformation of the DOCSIS system and HFC plant has delivered an increase in density and scale over the previous, enabling operators to cost-effectively address the exponential year-over-year bandwidth increases their customers demand.

Already on the horizon, the next major evolution in the cable access network is a shift to a Distributed Access Architecture (DAA). By moving access hardware from the headend to smaller hub sites or into the plant, DAA provides a number of benefits to cable operators, including reduced operational costs and bandwidth growth.

Although several options exist for implementing a DAA, this paper will examine the technological and business benefits of the standardized Remote PHY (RPHY) platform. We will compare an RPHY fiber deep deployment with an integrated CCAP and traditional HFC deployment to show how RPHY can be used to expand network capacity and reduce CapEx and OpEx.

Comparing RPHY fiber deep and integrated CCAP/HFC deployment

Why compare an RPHY fiber deep and an integrated CCAP/HFC deployment? Although fiber deep RPHY offers a number of advantages, it does represent a new deployment model that must be operationalized. Therefore, some operators might opt to continue splitting fiber nodes until they ultimately reach the same fiber depth as fiber deep RPHY. Theoretically, this approach would offer the same capacity expansion as fiber deep RPHY without the challenges of incorporating new technology and adapting to a new operational model.

This approach, which we will call analog fiber deep, is really just continuing down the current path of splitting fiber nodes into new segments and adding corresponding integrated CCAP capacity in the headends and hubs. Although not as flexible as RPHY fiber deep, analog fiber deep would enable a fiber deep or node plus zero architecture, but with analog modulated optics, that is, traditional HFC.

Background: RPHY deployment options

Before we begin our comparison of an RPHY fiber deep and an analog fiber deep deployment, let's walk through the options available in an RPHY DAA.

Separating the CCAP core and PHY functions, the RPHY architecture supports a number of deployment options. With considerations such as geographic network distribution, short-term and long-term goals, and anticipated subscriber and bandwidth growth rates, an operator might choose to deploy RPHY using:

- **An RPHY shelf:** A shelf can contain a few or many RPHY devices (RPDs). This deployment option allows a hub to contain only RPD shelves, while centralizing the CCAP core. Shelves can also be used as “port extenders” when paired locally with a CCAP core.
- **A business as usual fiber node location:** This deployment uses the existing or traditional fiber node location, typically passing several hundred homes, and incorporates a number of actives or amplifiers in line after the fiber node. An operator could deploy an RPD in an existing location or as part of a node segmentation effort.

- **A new fiber deep node location:** Because of its extreme segmentation, this deployment offers the most benefit from RPHY because each fiber deep location typically serves 50 to 70 homes, and there are no active components past the RPD/fiber node location. This deployment is also referred to as “node plus zero,” meaning zero active components after the fiber node.

A combination of these deployments is possible and even viable in many situations. For example, an operator could deploy both business as usual and fiber deep nodes when converting an existing node to RPHY. This type of deployment would enable the operator to roll out fiber deep only to the areas that need capacity relief. Alternatively, an operator could deploy a mixture of RPHY shelves and nodes in a given area. Whichever path is chosen, the flexibility of the RPHY deployment architecture enables operators to directly address capacity needs.

There are also several options for locating the core function of RPHY. Operators can distribute the CCAP core in hubs and headends, essentially using the same location where integrated CCAP is today, or they can choose to centralize the CCAP core in a city or even in a region.

Comparison criteria

To accurately compare OpEx and CapEx for RPHY fiber deep and analog fiber deep deployment, we will consider the costs associated with construction per cable mile, bandwidth per service group, segmentation, CCAP chassis density and scalability, equipment and HVAC powering estimates, headend footprints, and the RPHY CIN.

CapEx estimates will be developed for both scenarios. CapEx will be estimated based on bandwidth and hardware scalability and includes the cable construction cost estimates. It should be noted that all cable construction and fiber node costs are incurred in the first evolution, and additional construction and node cost, are not incurred in the later cycles.

Background: sample hub

We will begin by defining some metrics for a typical, representative hub. The hub location was selected and analyzed, and the existing HFC plant evolved in steps to a “node 0” fully segmented environment. To provide an apples-to-apples comparison, we will assume that the CCAP core stays in the same headend or hub location as the integrated ‘CCAP.


¹ <https://marketintelligence.spglobal.com/blog/cable-s-broadband-boom-by-the-numbers>.

Our sample hub consists of ~38,400 homes passed with a mixture of commercial, residential, and Multidwelling Units (MDUs). The hub serves 70 fiber node locations of various sizes and varying densities per cable mile.

Based on the selection of sample nodes, the average node size and cable density across the hub are ~550 homes passed per fiber node and 147 homes passed per cable mile. (See Figure 1.)

Figure 1. About the sample hub

- 38,400 homes passed with a mixture of commercial (36%), residential (61%), and multi-dwelling units (3%) (mdu's).
- The sample hub serves 70 fiber node locations of various sizes and varying densities per cable mile.
- Average node size and cable density across the hub is 550 homes passed per fiber node and 147 homes passed per cable mile.
- Sample hub is comprised of 54% aerial and 46% underground miles.
- Fiber Node Density is High (63%), Medium (23%), and Low (14%).




Construction cost per cable mile

We used three representative nodes to perform a construction analysis. Costs were then extrapolated across the entire hub. Each of the three fiber node service areas uses short reach fiber extending from each existing fiber node location known as the parent node. Tap values, passive devices, and a small percentage of cable were replaced as required to achieve the node plus zero design.

Fiber was extended to the same average depth of 60 homes passed for both the analog fiber deep and RPHY deployments. Figure 2 depicts the average cost per cable mile and the total number of fiber deep node locations. Extrapolating the costs shown in the chart, we can calculate the total cost of cable construction as ~\$3.296M, or \$85.80 per home passed. Because the fiber depth and coaxial design is identical for the analog fiber deep and RPHY fiber deep deployments, the construction cost is simply added to the aggregate CapEx cost summaries for both approaches.

Figure 2. Fiber node location and construction costs



CapEx and OpEx summary for construction cost per cable mile

No differentiation between analog fiber deep and RPHY fiber deep deployments.


Bandwidth per service group

In the analysis, bandwidth is increased in steps. The initial bandwidth estimate used assumed a DOCSIS 3.0 environment without video convergence. The mid bandwidth estimate introduced a minimal amount of DOCSIS 3.1 OFDM loading. The long bandwidth estimate assumed the rebalancing of DOCSIS channel loading and the addition of converged MPEG broadcast and narrowcast video. The DOCSIS rebalancing factors in fewer DOCSIS 3.1 modems to begin with and an

increased number of DOCSIS 3.1 modems over time, justifying a conversion of some DOCSIS 3.0 downstream channels to DOCSIS 3.1.

Figure 3 depicts the channel loading in equivalent 6 MHz increments and the bandwidth per service group as a result of each step (initial, mid, and long). Note that at the long step, we have achieved approximately 1 GHz channel loading of 149 channels and ~5 Gbps of capacity per service group.

Figure 3. Channel loading and bandwidth scalability



CapEx and OpEx summary for bandwidth per service group

No differentiation between analog fiber deep and RPHY fiber deep deployments.


Segmentation

The segmentation analysis was initially developed using four fiber deep node locations per service group. Because the fiber depth and physical node size are the same for the analog fiber deep and RPHY fiber deep, we can assume that the service group size is initially four fiber nodes (60 homes passed), or 240 homes passed per service group. In the analog fiber deep deployment, each of the RF ports on the integrated CCAP chassis is split four ways, feeding each of the fiber deep nodes. In the RPHY fiber deep deployment, a virtual splitting and combining of four RPDs is used to achieve an equivalent service group size.

The bandwidth per home passed was increased at each step by the addition of DOCSIS channels and video convergence, as explained in the bandwidth per service group section. After the maximum of 149 channels is achieved, the bandwidth per home is increased by reducing the number of fiber deep nodes per service group. In the analog fiber deep deployment, this is achieved by decreasing the split ratio on the integrated CCAP port from four nodes to two, and finally a single fiber deep node per service group (split 1 and split 2). In the RPHY fiber deep deployment, an equivalent service group size is achieved by incrementally deaggregating the switch port at the same rate.

Figure 4 depicts the service group segmentation of the five steps used to developed the CapEx and OpEx models within this study (initial, mid, long, split 1 and split 2).

Figure 4. Service group scalability and service group migration




CCAP chassis density and scalability

The CCAP chassis used for the analog fiber deep and RPHY fiber deep deployments have different service group densities, and therefore scalability across the same baseline hub is not the same. This reflects the flexibility in scaling that RPHY offers versus integrated CCAP, because a CCAP core can scale to maximize platform bandwidth capability. The CCAP core chassis assumes 112 service groups per chassis for the entire evolution of the five steps outlined in the segmentation section. In turn, the integrated CCAP (analog fiber deep)

chassis assumes an increase in port density in the final two steps from 56 ports to 84 ports per chassis. The 56-port chassis assumes the current density of the integrated CCAP chassis, and the increase to 84 ports per chassis is assumed in the final two steps by using next-generation integrated line-card hardware.

Based on the analysis defined in the segmentation section, Figure 5 depicts the chassis density and the number of CCAP chassis required for both deployments across all five migration steps.

Figure 5. CCAP chassis density and scalability



CapEx and OpEx summary for CCAP chassis density and scalability

Advantages of RPHY fiber deep deployment: The CCAP core count for RPHY fiber deep increases incrementally compared to the integrated CCAP chassis. These incremental changes are mainly caused by the CCAP core’s scalability with RPHY. It can support a higher number of service groups, resulting in fewer chassis deployed as compared to the analog fiber deep deployment. Fewer chassis mean less CapEx, less physical footprint, and lower power and cooling requirements, all resulting in a lower TCO.

Equipment and HVAC powering estimates

In the analog fiber deep deployment, a “pod” architecture was developed to simplify the operational transition from traditional HFC to analog fiber deep. The analog fiber deep pod is currently being used for real-world deployments.


Each self-contained pod is prewired and constructed with all of the forward and return optics, active RF

combining, and optical multiplexing to scale across one CCAP chassis. As noted in the CCAP chassis density and scalability section, each CCAP chassis initially serves 56 service groups composed of 224 fiber deep nodes (four fiber deep nodes per service group). Each pod consists of five 44 RU headend racks and assumes the use of Prisma II optics and upstream Enhanced Digital Return (EDR). Note: The equipment type is only specified to determine CapEx and OpEx (powering costs) for the analog fiber deep deployment.

A similar pod concept was also developed for the RPHY fiber deep deployment. Each RPHY pod uses a single CCAP chassis with integrated optics, switching, and optical multiplexing. Because the RPHY chassis has a greater service group density, and all forward and return optics are integrated into the chassis, the RPHY pod has lower power consumption than the analog pod.

Figure 6 depicts the aggregate power consumption of both pods, including HVAC estimates for each deployment. The OpEx estimate assumes \$.12/kWhr, or ~\$1,051 per kW/yr, for power costs.

Figure 6. Headend power consumption and annual powering costs




CapEx and OpEx summary for equipment and HVAC powering estimates

Advantages of RPHY fiber deep deployment: With an annual powering cost of \$104,700, the consumption rate of the analog fiber deep deployment is expensive. With the RPHY fiber deep deployment, the savings add up to a staggering \$64,000 a year, assuming no change in power costs as the years progress. This factor alone is one of the main reasons service providers are preparing to transition to RPHY fiber deep architectures.

Headend footprint

Based on the pod estimates explained in the equipment and HVAC powering estimates section, a headend footprint required for each deployment was developed and is summarized in Figure 7. Note that the footprint includes all of the hardware associated with each pod. In the RPHY fiber deep deployment, the Converged Interconnect Network (CIN) uses 100G and 10G switches for the “spine and leaf,” respectively. The CIN is described in the next section.

Figure 7. Analog fiber deep and RPHY fiber deep headend footprints



CapEx and OpEx summary for headend footprints

Advantages of RPHY fiber deep deployment: The headend space determines the number of equipment racks needed to support the deployment. The initial space required for analog fiber deep compared to RPHY fiber deep is 17 racks instead of 4 racks, a significant difference. As we grow the network and bandwidth, the number of racks for the analog fiber deep deployment scales to 45, while the RPHY fiber deep deployment only requires a total of 6 racks. Again, RPHY fiber deep provides operators with significant savings because headend space can be translated into a cost per square foot. Headend space is a large part of the deployment TCO, and often MSOs simply do not have the space available. RPHY fiber deep enables operators to constrain headend space and dramatically reduce cost per square foot.

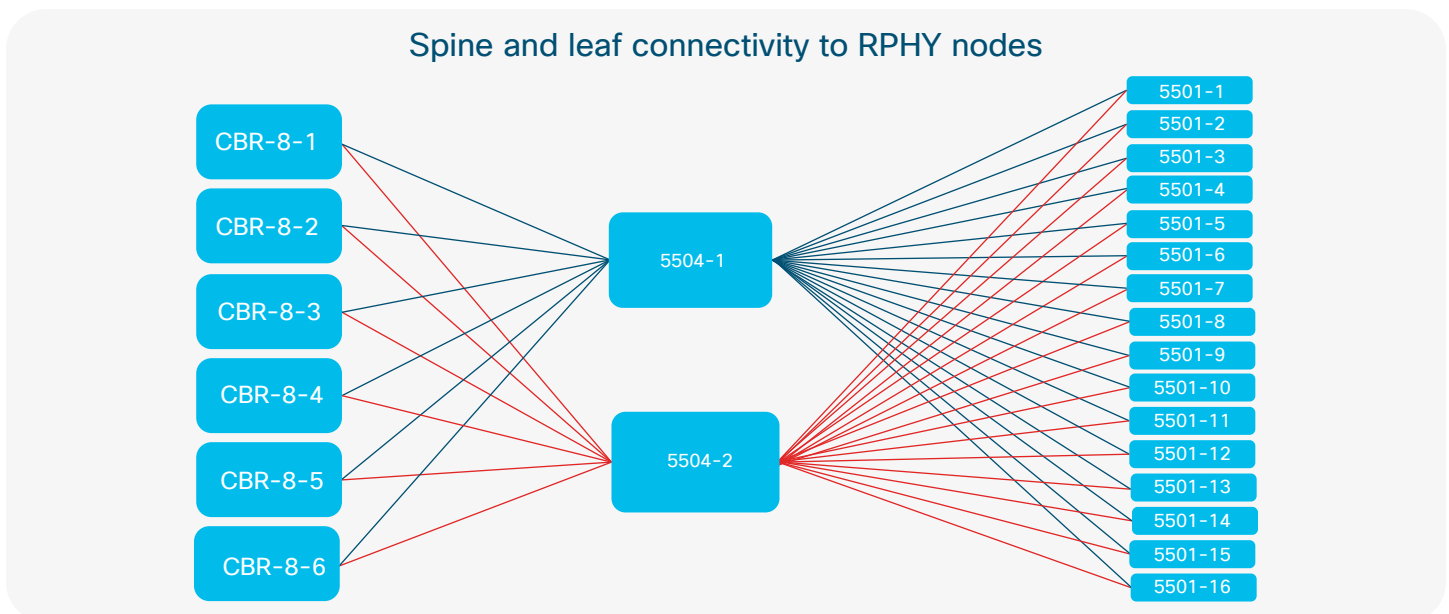
RPHY CIN

As noted in the headend footprint section, the RPHY fiber deep deployment includes the spine and leaf connectivity to compose the CIN. Each CCAP chassis is connected to the spine using 4x10G or 100G optics (depending on the evolution step). Each port on the spine is then correspondingly connected to the 40/100G

ports on the leaf switch. The RPHY fiber deep nodes are connected directly to the 10G leaf ports using point-to-point 10G DWDM transceivers. In Figure 8, you can see that the redundant CCAP LCs, spine hardware, and optical route diversity are designed into the network up to and including the 40/100G leaf ports. The primary and redundant optical routes and connectivity are also shown. Although it would be possible to include optical route diversity from the leaf network to the RPD, it was not included in this cost model. Even so, the level of hardware redundancy and route diversity provided on the RPHY fiber deep network far exceeds that of the analog fiber deep network.

Figure 8 shows the final network evolution (split 2). Note that the number of spine (2) and leaf switches (16) remains constant through each evolution, indicating that additional bandwidth is added by deaggregating the leaf ports, thus decreasing the number of homes connected to each port. Also, the use of a scalable, modular 100G switch means that capacity can be added to the spine network without any additional chassis hardware. Based on constructing a scalable CIN in the first evolution, the only additional hardware required is the requisite number of CCAP chassis. Figure 8 also shows the number of CCAP chassis scaling from 2 to 6 across the five evolution steps.

Figure 8. Spine and leaf connectivity to RPHY nodes



CapEx estimates

CapEx estimates were developed for both deployments. CapEx was estimated based on the bandwidth and hardware scalability described in previous sections and is inclusive of the cable construction cost estimates

provided in the construction cost per cable mile section. It should be noted that all cable construction and fiber node cost is incurred in the first evolution, and additional construction and node cost is therefore not incurred in the later cycles.

Figure 9. Total network CapEx and total cost per HHP



Conclusion

While the comparison in this paper between RPHY fiber deep and analog fiber deep is based on a representative but hypothetical hub, the OpEx and CapEx data clearly shows the cost effectiveness and advantages of an RPHY fiber deep deployment. Substantial OpEx and CapEx gains are made based on all criteria beyond those that simply define the logistics of the environment, including CCAP chassis density and scalability, equipment and HVAC powering estimates, headend footprints, and the RPHY CIN. Particular attention should be paid to the RPHY fiber deep deployment’s reduction in headend footprint requirements, because this can often negate the need for facilities expansion.

RPHY fiber deep: A pathway to savings

While this study compares the approach of an analog fiber deep versus an RPHY fiber deep deployment, RPHY technology can be used in other capacities to gain additional savings. For example, a Business As Usual (BAU) RPHY 2x2 node configuration could be placed at the existing node location, thus greatly reducing the CapEx required for fiber construction. This configuration provides the hub facilities with the same reduction in headend footprint and OpEx efficiency as the RPHY fiber deep network because of the elimination of analog optics and all associated active and passive RF combining. RPHY shelves could also be placed at hub locations, potentially reducing the space and power requirements even further, while providing a pathway to virtualization and eliminating plant upgrades altogether. Finally, there are some variations to the CIN that can provide incremental port and DWDM efficiency based on the placement and use of aggregation switching.

Cisco and RPHY

In May 2017, Cisco announced the availability of its Infinite Broadband solution, which uses RPHY technology to overcome the limitations of analog fiber and break through the HFC bottleneck. In addition to supporting bandwidth growth and reducing operational costs, the Cisco® RPHY solution enables cable operators to transform their infrastructures to simpler all-digital networks; decouple applications from the infrastructure; simplify management tasks; and create a platform for service delivery that is inherently flexible, scalable, and capable of supporting cable operator revenue growth and increased profitability.

An industry leader in cable access technology, Cisco is committed to developing standards-based solutions that allow cable operators to reimagine their networks and operating models. In fact, John Chapman, Cisco Fellow and Cable CTO, played a central role in the creation and ongoing development of the RPHY DOCSIS standard and invented the primary technologies that were the foundation for DOCSIS 3.0 and DOCSIS 3.1. The RPHY architecture and open-source remote PHY software program, now standardized by CableLabs, were his creations.