Constructing a Convergence Lab

Lessons Learned From Building a Converged Network at CableLabs

A Technical Paper prepared for SCTE•ISBE by

Matthew Schmitt
Principal Architect
CableLabs
Pleasanton, CA
303-661-9100
m.schmitt@cablelabs.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

Cable operators are in the midst of deploying or preparing to deploy a variety of distributed access architecture (DAA) approaches. When this is done by aggregating multiple child nodes together onto a single point-to-point (P2P) coherent optics link, it has the effect of pushing very high capacity fiber-based Ethernet links deep into their networks. Beyond improving their DOCSIS®-based residential broadband services over coax, it also opens up numerous new business opportunities—such as mobile fronthaul and backhaul, business ethernet, remote passive optical networks (PON) and more—along with the ability to converge all of these different technologies onto a single network infrastructure.

CableLabs® has been working jointly with our members and technology vendors to develop the tools and technologies needed to make this vision a reality. However, while it's one thing to develop the technology, it's quite another to build devices based on it, and still another to integrate it together into a working whole. While CableLabs is not in a position to manufacture those devices, we are well positioned to bring them all together in a new network infrastructure convergence lab.

This convergence lab will allow us to demonstrate how these various technologies can be brought together to operate over a single physical network, showcasing the work that's already been done, and learning the implications of operating them together. It will also create a platform on which to test out new technologies going forward to see how they integrate together into a converged whole.

This paper goes through some of the background on DAA and P2P coherent optics technology, reviews the process we went through to build the lab, any lessons learned thus far, and the current status of what is an ongoing process of continuing to grow and develop this lab.

2. DAA and P2P Coherent Optics Background

2.1. Pre-DAA Hybrid Fiber-Coax (HFC) Networks

The typical access network architecture deployed by cable operators prior to the advent of DAA approaches employed a Hybrid Fiber-Coax (HFC) approach to reach their customers with video and data services. A simplified version of this architecture is shown in Figure 1 below.

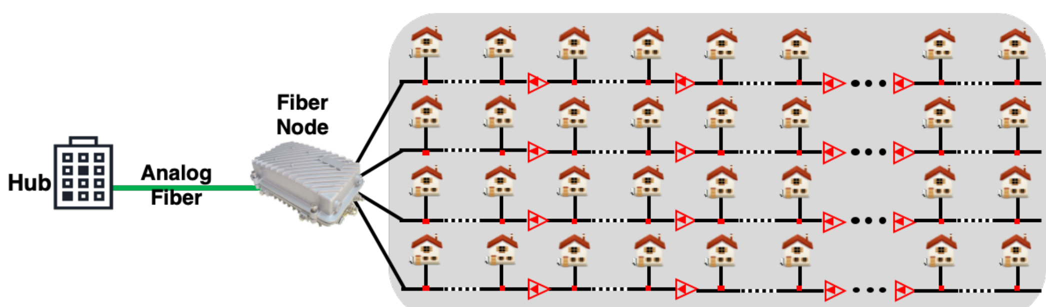

Figure 1 – Simplified Cable HFC Architecture

In this example, downstream radio frequency (RF) signals for data and video are generated in a Hub facility, converted from electrical to optical using an analog laser, transported over fiber to a Fiber Node at which point the signal is converted back to electrical, which is then transmitted over a coaxial cable network to end customers. In the upstream a similar set of operations occurs, originating at the customer premises and terminating at the Hub. The distance from the Hub to the Fiber Node is typically 20-80 km, although longer distances are possible; the distance from the Fiber Node to the end customer is generally just a few km.

When the capacity requirements of the group of customers sharing a portion of the network (a Service Group) exceeds the capacity of a Fiber Node, operators will typically split a Fiber Node into multiple Nodes. Each of those Fiber Nodes will require a separate link back to the Hub, but will often share a single fiber pair using Dense Wavelength Division Multiplexing (DWDM) to save costs. Each new Service Group also requires new equipment at the Hub site, which can create pressure on space, power, and cooling.

2.2. A Fundamental Shift in Cable Access Networks

In response to these pressures, as has been covered in previous papers and presentations, many cable operators have begun a fundamental shift in their networks by moving the RF generation previously performed in the Hub out into the network. This allows them to convert their existing fiber network from analog optics to digital optics and to utilize Ethernet networking across the fiber portion of their cable plant. The effect is to create a deep fiber Ethernet network.

This approach is what we refer to as a Distributed Access Architecture (DAA), because functions that were previously centralized are now distributed. A simplified version is shown in Figure 2 below. On the surface it looks very similar to Figure 1: homes are still connected to coax, which terminates at a device that is connected to fiber, which in turn is connected to the hub. However, under the surface things are dramatically different: the fiber is carrying Ethernet data over digital optics rather than modulated RF signals over analog optics, and the RF is generated at a Remote PHY Device (RPD) or Remote MAC/PHY Device (RMD) directly connected to the coax without the electrical-optical-electrical conversion that occurred previously.

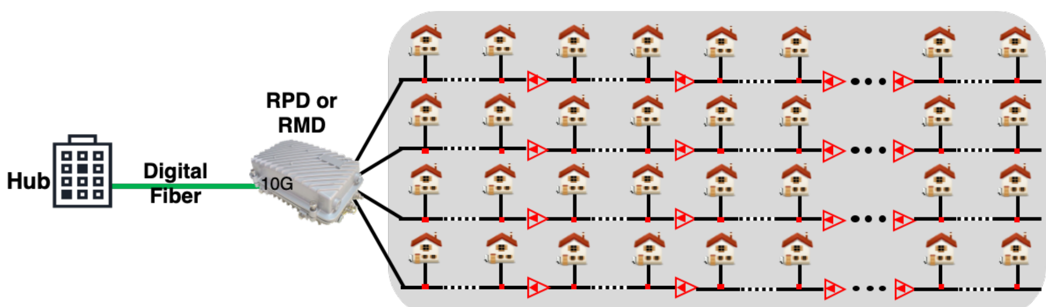

Figure 2 – Simplified Distributed Access Architecture

This change has a number of advantages:

- It reduces the equipment requirements for Hubs;
- It allows cable operators to utilize lower cost off-the-shelf digital optics for the fiber link instead of specialized analog optics; and
- By avoiding electrical-optical-electrical conversions, the signal to noise ratio (SNR) of the electrical signal on the coaxial cable is significantly improved.

That SNR improvement allows devices based on the DOCSIS 3.1 specifications to operate at higher modulation orders, which increases total network capacity.

2.3. The Need for Aggregation

Each RPD or RMD typically has a 10 gigabit-per-second (Gbps) port, which provides sufficient capacity for generating a full spectrum of digital video and DOCSIS 3.1 signals. Therefore, if Fiber Nodes were simply being replaced with RPDs or RMDs on a 1:1 basis, all that would be required to connect them to the Hub would be a single 10 Gbps optical link.

However, a key driver for the move to DAA approaches is to enable more Service Groups with fewer customers on each one. Figure 3 shows a simplified version of a scenario where the single service group from Figures 1 and 2 is split into multiple smaller service groups, each serviced by a separate RPD or RMD (each of which we refer to as a child node), and aggregated together at an aggregation node (AN) onto the existing fiber. The AN sits where the Fiber Node was previously.

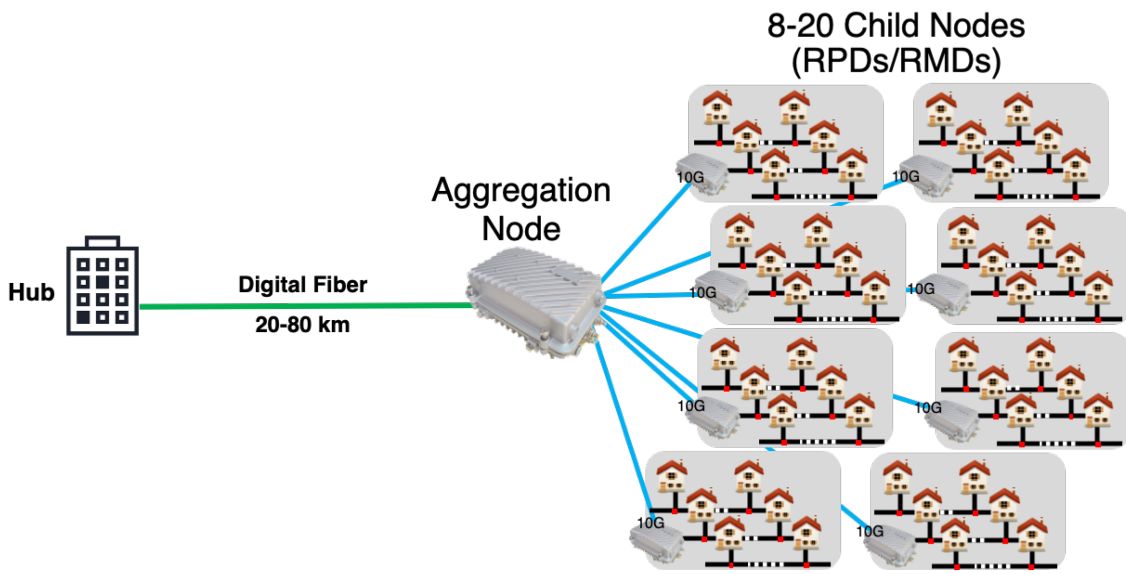

Figure 3 – DAA with Multiple Child Nodes

One straightforward means of building this approach would be to passively aggregate 10 Gbps optical links from each child node using DWDM, with the AN containing a passive mux/demux. This technology is readily available and well understood, and so represents a logical initial approach.

However, it has its limits, since at a typical spacing of 100 GHz for each signal cable operators will be limited to about 48 of these 10 Gbps links over a single fiber pair. And, the cost goes up linearly with every 10 Gbps link that is added.

2.4. The Case for Coherent

Architects at CableLabs postulated that utilizing a P2P coherent optics link operating at 100 Gbps or 200 Gbps per wavelength might prove advantageous over an approach using multiple 10 Gbps optical links with DWDM as described above. One end of the coherent optics link would be terminated at the Hub facility, and the other end at a device located in the AN called a Coherent Termination Device (CTD); the CTD would also terminate multiple 10 Gbps links from the Child Nodes, which could utilize low cost grey optics operating at 10 km or less. The aggregation could be performed in multiple different ways: at layer 1 with a muxponder; at layer 2 with a switch; or at layer 3 with a router.

Conceptually this looks exactly the same as what was shown in Figure 3 above, with the following exceptions: the AN has a CTD in it which forwards traffic at layer 1, 2, or 3 rather than containing a DWDM mux; the links connecting the CTD to the Child Nodes use low cost grey optics (instead of higher cost colored optics); and instead of carrying those 10 Gbps signals all the way back to the Hub (which requires longer reach, higher cost optics), they are aggregated onto one or more P2P coherent optics links for transport back to the Hub.

CableLabs staff conducted an analysis to compare the costs associated with each approach for supporting multiple child nodes in the same footprint as an existing fiber node. We found that the key determination for which approach would cost less depended on the number of child nodes. If the number of child nodes being aggregated together is relatively small, the DWDM approach was more cost effective, because the cost of the P2P coherent optics transceiver and the CTD is shared across only a few child nodes. However, as the number of child nodes increases that dynamic flips, because the cost of the DWDM approach goes up linearly with each child node due to the need to add a pair of higher-cost tunable 10 Gbps transceivers, whereas with the P2P coherent optics approach the cost for each additional child node is much smaller because only very low cost grey optics transceivers are added.

Further, while there are already cases where the P2P coherent optics approach was lower cost, we also saw potential opportunities to bring that cost down further, which would lower that crossover point.

Using P2P coherent optics provides additional benefits as well: it can coexist with existing signals using DWDM, which in turn also opens the door to add additional P2P coherent optics links in the future. The net result is a 10-20x increase in available capacity for the fiber that cable operators have already deployed into the range of terabits per second, extending the lifetime of their existing network investment dramatically.

2.5. The Pathway to a Converged Network Infrastructure

In addition to the benefits identified above, moving to an architecture utilizing P2P coherent optics opens the door to new business opportunities by enabling a converged network infrastructure.

A DOCSIS network is, in effect, a very deep Ethernet network optimized for running over an HFC infrastructure. However, the only points at which you could access that network are at cable modems attached to the coax portion of the plant: accessing the fiber portion of the network cannot be done via DOCSIS, and instead was only possible by using separate fibers or separate wavelengths over the existing fiber.

However, let’s assume you’ve deployed a P2P coherent optics architecture using a network switch in the CTD. That CTD has become an Ethernet connection point that can tap directly into very high capacity fiber links. Any service that can operate with an Ethernet connection can connect at that point, including DOCSIS services, PON services, mobile xhaul services, business services, etc. All of them can be carried over the same Ethernet based fiber network, aggregated together via the CTD.

The result is a converged network infrastructure, as shown in Figure 4 below.

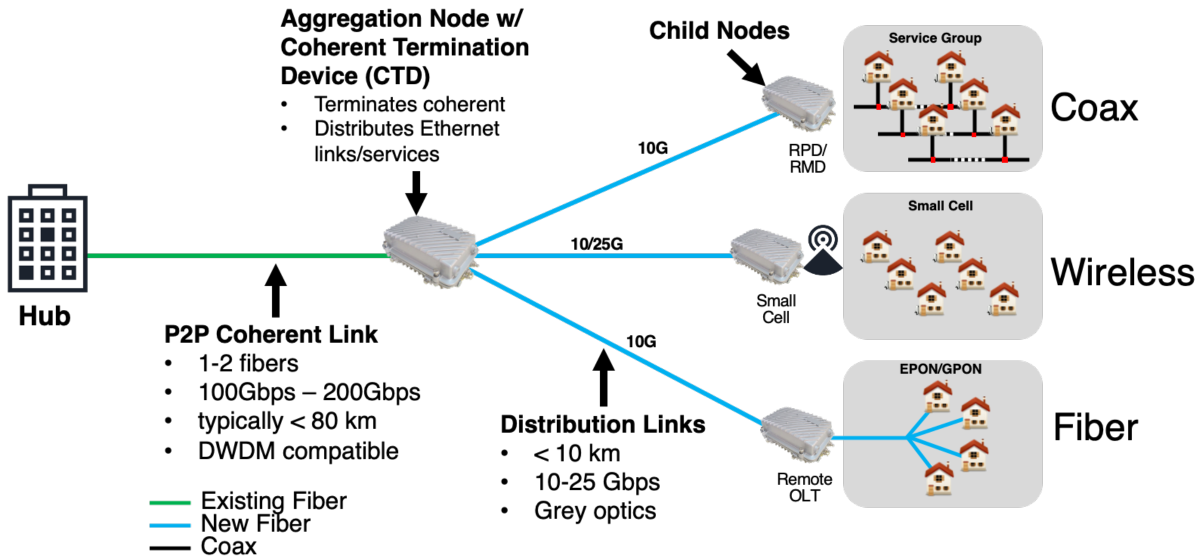

Figure 4 – Converged Network Infrastructure

With this type of architecture approach, the opportunities are limited only by our imaginations.

In order to help enable this future, CableLabs worked with our members and manufacturers to develop a series of specifications that define requirements for coherent optics transceivers operating at 100G and 200G. The intent was to provide manufacturers with guidance regarding the requirements for operating on a cable operator network, as well as to promote interoperability, which in turn helps promote scale and competition, thereby reducing cost. The specification for 100G per wavelength operation was initially released in 2018, and the specification for 200G per wavelength operation was initially released in 2019.

3. The Plan

3.1. Why Build a Convergence Lab

Another advantage of P2P coherent optics is that there is existing equipment that can be leveraged now, with more on the way. Interoperable 100G coherent optics transceivers that are compliant to the CableLabs specification requirements already exist, as do a variety of switches, routers, and muxponders that they can operate in, some of which are temperature hardened. Prototypes of interoperable 200G transceivers are expected this year as well.

However, we realized that a key piece – a fully weather hardened Aggregation Node with a CTD that could operate in it – was missing.

The plant architectures utilized by many cable operators around the world favor a clamshell type device that is fully weather hardened at the AN location. Unfortunately, that device doesn't yet exist and needs to be developed. In order for manufacturers to decide to invest the money to do that, they need to understand what requirements their customers have, and confirm that there's sufficient interest from those customers.

One way in which CableLabs is working to address these needs is by working with our members to define common requirements for a CTD.

Another way in which we realized we might be able to help was by demonstrating the capabilities of this new architecture, showing not only that it can support existing use cases exceptionally well, but even more importantly that it can support a broad range of other applications and services that could provide new revenue opportunities.

Thus, the idea for building a network infrastructure convergence lab was born. It would be a platform on which we could demonstrate how a variety of different services could coexist: DOCSIS broadband over coax; mobile fronthaul and backhaul; FTTH via PON; and business ethernet services. A showcase for products built to comply with a wide range of CableLabs specifications. And a test bed for validating how new products and technologies could integrate into that same infrastructure.

3.2. Finding Partners

In order to make that a reality, a key requirement would be obtaining the equipment—ideally from a variety of manufacturers—in order to make it work. So began a quest to find partners in this endeavor.

Utilizing relationships we had developed in putting together the CableLabs P2P Coherent Optics specifications, as well as relationships that were developed as a part of building the P2P Coherent Optics interoperability demo shown at SCTE Expo 2019, and also relationships from the work done with CableLabs on Distributed Access Architectures (DAA), we reached out to a number of manufacturers with a fairly simple pitch to build the network shown in Figure 4 above in a lab at CableLabs. The benefit for manufacturers: feedback regarding how their devices operate in the network, visibility to our members when we report on the activity, and visibility to anyone that visits CableLabs and tours our lab spaces.

As a part of these discussions, we proposed to develop the lab in four phases.

The first phase would be focused on the main optical link from the "Hub" to the CTD, since all other applications and services would run over that infrastructure.

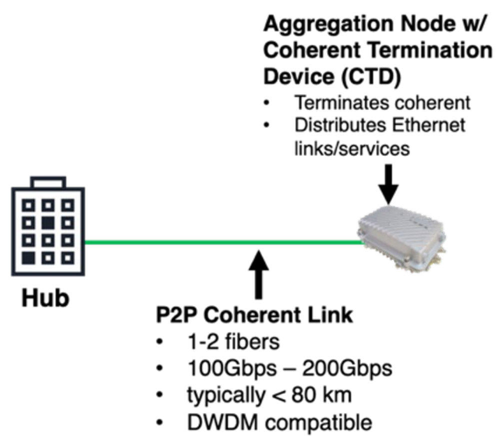

Figure 5 – Phase 1: P2P Coherent Optics

Ideally it would include equipment from multiple manufacturers so that we could demonstrate interoperability, and multiple transceivers operating simultaneously. That would demonstrate the ability of the solution to scale up as needed with multiple wavelengths operating at the same time. Playing the role of the CTD in the AN would be temperature-hardened equipment that would be suitable for deployment in a street cabinet (since clamshell devices were not available).

As a part of this phase we would also connect traffic generators directly to the CTD in the AN as well as in the Hub, which serves the dual purposes of validating the link and showing how a business ethernet service might work.

The second phase would be focused on building a Distributed Converged Cable Access Platform (CCAP) Architecture (DCA) network that would operate over the top of the P2P coherent optics link serving as the backbone of the architecture, since this would likely be the first, primary application for cable operators.

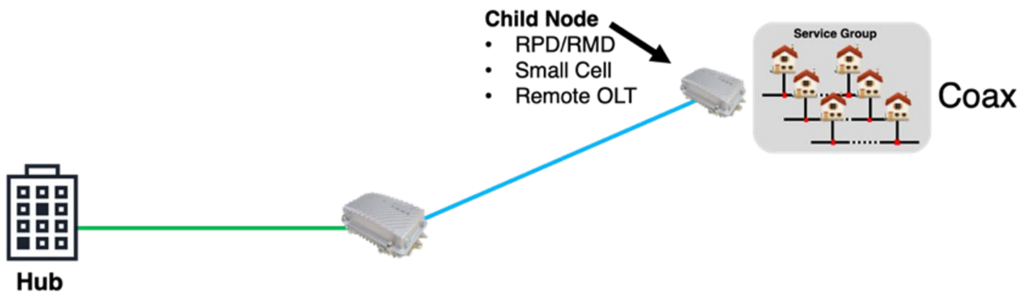

Figure 6 – Phase 2: Distributed CCAP Architecture

The intent was to start with the Remote PHY Devices (RPDs) that cable operators are starting to deploy, working with a CCAP Core in the Hub location.

The third phase would then look at layering in wireless solutions, which present an immense new business opportunity for cable operators, particularly with the rollout of 5G services that—in fronthaul applications in particular—have high capacity and low latency requirements, while also requiring higher densities than are typically the case today due to the use of higher frequency spectrum.

Figure 7 – Phase 3: Wireless

Within this, there would be opportunities to demonstrate the use of both direct fiber connections as well as DOCSIS links for backhaul and fronthaul transport of mobile traffic.

The fourth phase would incorporate a remote PON solution, with the remote OLT being either a standalone device connected via a short fiber run to the CTD (as shown below) or a device that connects directly to the CTD (such as an OLT in an SFP+ form factor).

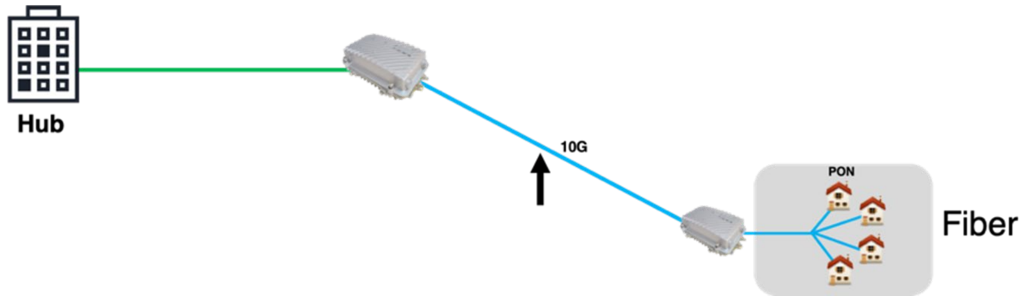

Figure 8 – Phase 4: PON

Note that any of the follow on phases after Phase 1 could be done in any particular order depending on equipment availability, and additional phases to add equipment to address other use cases would be possible.

Therefore, our primary focus was on finding partners to supply equipment to support Phase 1, with a lesser focus on Phase 2, leaving the remaining phases as future work.

3.3. The Construction Plan

Another question we addressed during these discussions was how to go about actually constructing the lab setup.

As mentioned, at SCTE Expo in 2019 the CableLabs booth featured a P2P coherent optics interoperability demo which included contributions from a number of different manufacturers. To prepare for that demo, we had held a sort of mini-interop, where each of the companies providing equipment sent an engineer out to CableLabs, and we spent several days piecing the demo together and ensuring we could make it work. We then took the demo apart, shipped it to New Orleans, and rebuilt the same setup we had used in the lab with substantial help from some of those same contributors.

For this effort, we proposed to do essentially the same thing: have each of the participating companies send out an engineer to work together over the space of several days to install the equipment and get it all working. In some ways it would be even easier than what we did the previous year, because we wouldn't have to tear it down and ship it out to be rebuilt; we could build it right where it stay and be used. If any issues came up, we could deal with them on the spot.

It seemed a straightforward, easy approach for getting the lab up and running quickly.

3.4. The Response

After making our pitch, a number of manufacturers agreed to provide equipment either on loan, as a donation, or at a substantial discount. Specifically, the following list of companies (in alphabetical order) offered the following pieces of equipment:

- II-VI (formerly Finisar): a C Form-factor Pluggable 2 – Analog Coherent Optic (CFP2-ACO) transceiver module
- ADVA Optical: a network switch that supports coherent optics, targeted for the "Aggregation Node" location, along with coherent optics transceivers
- Ciena: two network switches that support coherent optics, one each for the "Hub" and "Aggregation Node" locations, along with coherent optics transceivers

- Edge-Core: a whitebox network switch that supports coherent optics, targeted for the "Hub" location, including support for CFP2-ACO modules
- EXFO: a pair of traffic generators capable at running multiple 100G streams of data simultaneously
- Lumentum: a CFP2-ACO coherent optics transceiver module
- Vecima: a Remote-PHY Device (RPD) and associated management station

With verbal agreements in place, we began working on all the necessary paperwork to allow these transactions to take place, making arrangements for equipment to be shipped out, etc.

All was going relatively smoothly, if perhaps a bit slower than we'd all have preferred: getting executive approvals, legal approvals, and documents signed always seems to take longer than you'd expect. Still, given that the actual construction of the lab should only take a matter of days with engineering support on site, some amount of delay wouldn't prevent us from completing the lab in a reasonable amount of time. Enough so that I felt comfortable submitting a request to SCTE to create this paper.

Then the Covid-19 pandemic hit with full force, blowing up many of our plans.

4. The New Reality

4.1. Restricted Access

Like many companies, as the rapid spread of the virus became clear, CableLabs instituted a work from home policy. Specifically, unless being in the building was absolutely essential to the core operation of the company, we were all to work from home. Further, travel between our California and Colorado offices was halted, and visitor access to the building was cut off.

This obviously presented several problems for the lab construction plan. First, I'm located in California, but the lab was to be built at our main office in Colorado; therefore, I would be unable to work on building the lab myself. Second—and most significantly—we would not be able to have any outside engineers visit the facility, meaning we would not be able to rely on their expertise to get their equipment operational. And finally, since this work didn't necessarily count as essential to the core operation of the company, even staff located in Colorado were unable to enter the building to work on the lab.

However, at that point in time, there was still hope that these restrictions would only be temporary: after a month or two of a tight lockdown things would get better and we'd be able to travel and access the labs again. Besides which, it was taking longer than expected to get the equipment agreements signed, equipment shipped, etc. So, we had time to spare and hoped that by waiting things out we would still be able to build the lab as planned (albeit perhaps a bit later than originally envisioned).

Therefore, we continued to work on arranging equipment shipments to our facility in Colorado so that it would be ready to go as soon as virus restrictions were lifted.

4.2. Equipment Arrivals

Thanks in no small part to the contributions and efforts of the various companies that had partnered with us on this effort, who had continued to work through the pandemic and their own restrictions, the needed equipment did indeed start to arrive. However, since there was no one there to receive, unpack, and work with the equipment it began to pile up, as shown in Figure 9 below.

Figure 9 – Packages piling up

In fact, as shown above, the boxes were starting to clog up one of our hallways. As soon as one of our local team members was able to re-enter the office, we found a more efficient stacking arrangement, as shown in Figure 10.

Figure 10 – Tower of boxes

So now we had much of the equipment we needed, but still no way to put things together until our level of access to the building changed.

4.3. A New Hope

While the situation with the pandemic did not improve nearly as much or as quickly as we might have hoped, it did improve enough that CableLabs was able to relax some of our building access restrictions. Visitors to Colorado were still out of the question, but local staff would be permitted limited access to the building for short periods of time when arranged in advance.

Among those that would be given some access to the building were the members of the CableLabs Optical Center for Excellence, the team that we have working on next generation optical technologies: because they can't bring their experiments home, they needed access to the optical lab in order to continue their work. If they could spare some time to lend a hand, given their expertise, they would be perfect for setting up the lab.

Fortunately for me, they are some of the most helpful, generous people you'd ever want to meet, and they agreed to help me out and work on setting up the equipment there on site without complaint. The caveat was that they also had to make progress on their own work, and so would have only a limited amount of time each week to help out. But I was grateful for any help I could get.

A wave of outreach to our partners who were supplying the equipment to request remote support to our on-site team followed; they were all supportive, and so a series of introductions were made to connect those that would be supporting the work remotely with the team on the ground.

Our prospects for building the lab were looking up.

4.4. Everything takes longer in a pandemic

The local team was able to access the lab starting in mid-July. They began by unpacking equipment from one of our partners and trying to get it setup and running with remote engineering support.

What I didn't yet recognize—but would soon be brought face-to-face with—was that everything takes longer in a pandemic. In particular, when you can only access the lab two days a week (in order to limit exposure), and are relying on remote support, problems that might've been resolved in hours or a day if people were on-site and dedicated to the task can take a week or more.

And it seems like the simplest problems are the ones that catch you off-guard and cause the most problems.

For example, we discovered when we unpacked a network switch from one of our partners that it required a 220V power supply, something we didn't expect—since all the equipment we'd worked with up to this point worked on 110V—but admittedly should have asked about in advance. After several calls and emails, we discovered that we did in fact have 220V power supplies in the building at CableLabs, and members of our Kyrio lab team were able to set one up for us.

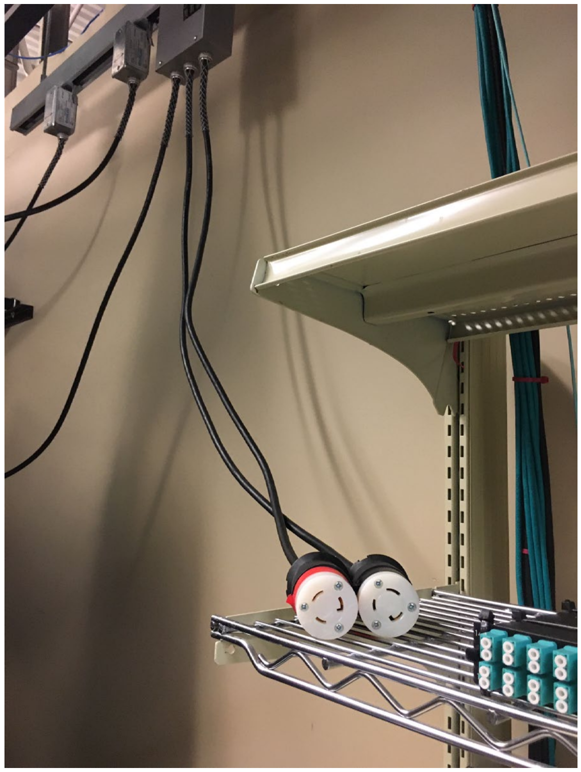

Figure 11 – 220V power

We initially thought we had overcome that issue with relative ease and minimal impact, but quickly realized that we weren't going to get off that easy: the power supply we had at CableLabs used an L6-20 receptacle (rated for up to 20 amps), whereas the plug for the box we were trying to install and setup used an L6-30 plug (rated for up to 30 amps).

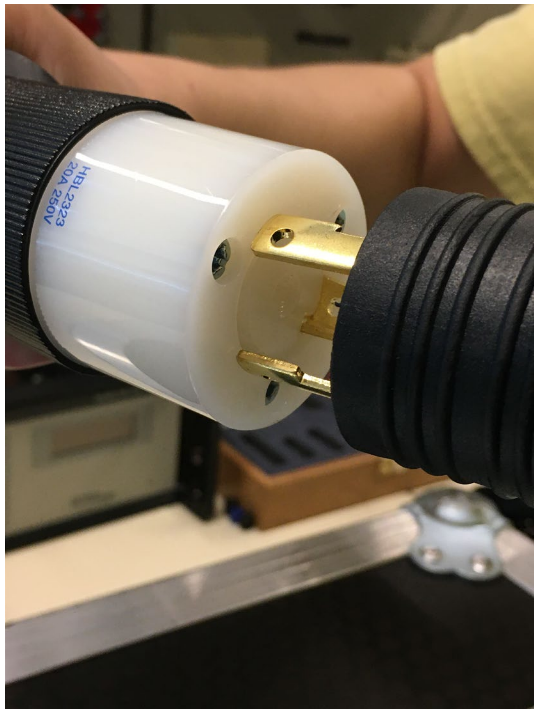

Figure 12 – The plug doesn't fit

We inquired with our partner regarding an alternate power cord, but were told that what we had was the only one they provided. We asked if the box would ever draw more than 20 amps, and were told that it probably wouldn't, although they didn't want to commit to that at that time. As a result, no one was comfortable simply using an adapter or replacing a plug. Instead, we started to look into purchasing a new 30 amp power dongle for our rail system, which we found would take at least 6 weeks.

This left us at a bit of a stand-still on this particular device until we eventually made two critical discoveries: while the plug on one end of the power cord that came with this network switch was specifically a 30 amp plug, the connector on the other end of the power cord was in fact only rated for 20 amps; and buried deep in the technical specs of the network switch we were using, we discovered that the switch itself had a max draw of just 9.3 amps. With that information in hand we became comfortable with the idea of using an adapter to connect the power cord to our 20 amp power supply, and quickly ordered one. And while that took just a few days to arrive, it also had to wait for the next window when someone would be in the lab, further delaying the day by which we could power up that device.

As I said, under these unusual circumstances, the simplest things become multi-week delays.

The power plug issue was just one example; each device setup ran into their own simple yet time consuming issues.

For one of them, the issue was licensing. For this particular manufacturer, they assumed that devices would always be connected to a live internet/network connection. Therefore, they had been designed to pull their license from over the internet, without which it wasn't operational.

However, in our case, we were setting these up in a lab without an external network connection; and while that might eventually happen, it certainly was not the case for our initial setup. There was an option to create a download for us to use, although we'd have to do it through their customer portal in which we weren't set up. It didn't get setup correctly the first time around, given it was a non-standard way of doing things. In the end we got it all working, but only after multiple emails, phone calls, and aborted attempts.

For another network switch we were attempting to setup, the issue wasn't power or licensing, it was transceiver compatibility.

This partner had provided us with a network switch under the assumption—which we shared—that one of the coherent optics modules we already had on hand would work. And initially it did, as we were able to power up the switch and have it recognize the transceiver module. All seemed fine. However, we then discovered that we couldn't set the frequency on which we wanted it to operate. It would only operate on a single fixed frequency, even though we knew the module was capable of tuning to different frequencies. Eventually the problem was traced to a compatibility issue: the switch would only support controlling the coherent transceiver module if it were also from the same manufacturer.

Our partner who had supplied the switch graciously offered to provide us with a pair of loaner transceivers, and moved to make that happen as quickly as possible. However, quickly in this case means getting a loaner agreement drawn up, having it reviewed by one legal department, having those changes reviewed by another, and finally getting it signed and processed. And then of course, waiting for shipping.

All of our issues were ones that could be overcome given time, but all of which also took longer simply because of the unique situation we found ourselves in with the pandemic.

5. Today and Tomorrow

5.1. Current Status

As of the writing of this text, here's where we stand today.

We now have both of the two Ciena switches up and running in our lab, with a transceiver in each unit talking to a transceiver in the other, as shown in Figure 13 below.

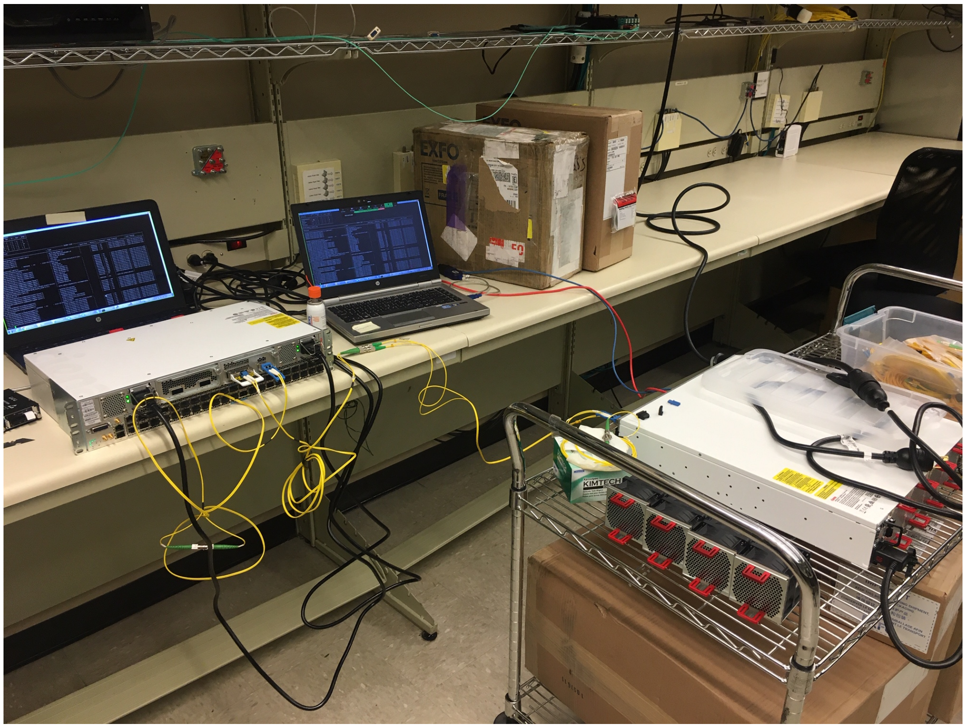

Figure 13 – Ciena to Ciena communication

The EXFO FTB-4 Pro traffic generators were recently unpacked and are now up and running; we're currently working on getting them setup to pass traffic across each of our switches.

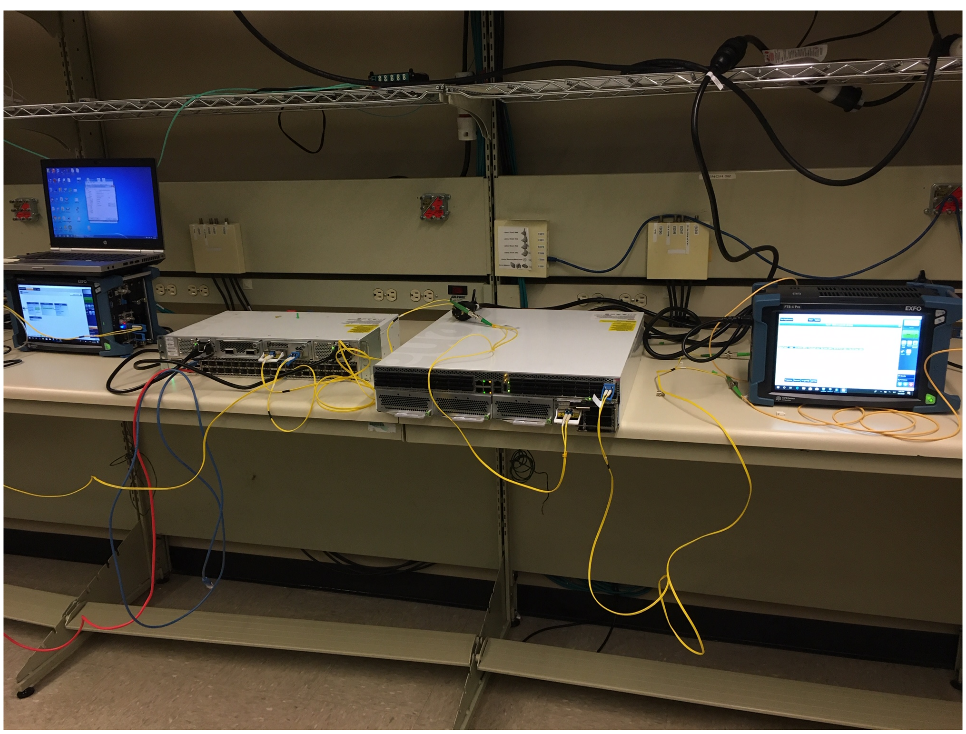

Figure 14 – EXFO traffic generators with Ciena switches

We've also gotten the ADVA switch up and running with in a loopback configuation.

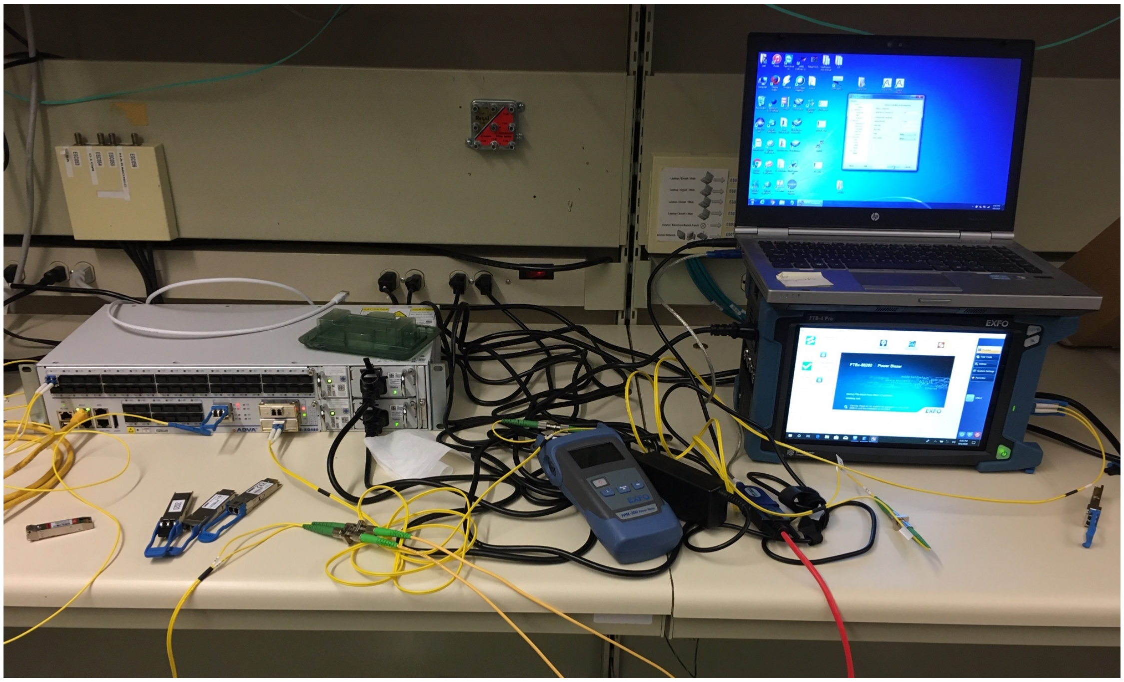

Figure 15 – ADVA switch with EXFO traffic generator

Other equipment is yet to be unboxed but will be soon.

5.2. Next Steps

In the short term, our plan is to get traffic running across the Ciena and ADVA switches, unpack and get the Edge-Core Cassini running with the CFP2-ACO modules that were provided to us by II-VI and Lumentum, and connect them all together using DWDM.

That will form the backbone of our lab to demonstrate network infrastructure convergence.

From there we'll be working on getting a Remote PHY implementation running over top of our optical network using the equipment provided by Vecima—in combination with a CCAP-Core already on-site—to demonstrate that service. That'll be followed by other phases as we're able, establishing that you can run multiple services over the same access network infrastructure simultaneously.

Beyond that, we plan to use this lab as both a showcase (it'll eventually be moved to a more visible location) and a working lab that we can use to test and demonstrate new components of a converged network. For example, we'll be interfacing it with equipment from other teams within CableLabs, such as those working on network virtualization, to create a truly converged network of the future. It's going to take a while to get there—much longer than I ever would have anticipated—but many of the building blocks are already in place, and more will be soon.

5.3. A note of thanks

While I may be the one writing this paper, given the restrictions on travel and building/lab access, I haven't been the one doing the work on the ground as I had hoped and intended. Instead, that's been done by a team of folks from multiple companies to whom I am greatly indebted:

- CableLabs
  - Haipeng Zhang
  - Jon Schnoor
  - Junwen Zhang
  - Mu Xu
  - Scott Bybee
  - Steve Jia
- Kyrio
  - Efren Torres
  - Mark Ambrozic
- ADVA
  - Clark Scott
  - Jack Yocum
  - Morgan Shaner
- Ciena
  - Craig McCoy
  - Darren McKinney
  - Fernando Villarruel
- Edge-Core
  - Jeff Catlin
- II-VI
  - John DeAndrea
  - Shan Esser
- IP Infusion
  - Srikanth Krishnamohan
- Lumentum
  - Ed Kirchoff
  - Ernest Muhigana
  - Jan Willem Poelman
  - Scott Swail
- Vecima
  - Colin Howlett
  - James Trueman

My sincere thanks to all of you, without whom none of what we have done (and will be doing!) would have been possible. And to anyone whom I may have inadvertently left off of this list, my apologies and deepest thanks.

This has been very much an industry wide effort, and it is my heartfelt hope that it will continue to be one going forward.

6. Conclusion

When I originally proposed this paper, the expectation was that I would be writing about what we had learned by successfully building a converged network infrastructure: that I would be describing why we had set out to build it, what had happened when we did, and hopefully about how we had demonstrated that these various services could operate effectively over a common access network infrastructure, along with perhaps a few issues the MSOs should watch out for in their own networks.

And while it did accomplish the first of those objectives, from there it instead evolved into a story regarding the very unusual and unique circumstances in which we all find ourselves, and the challenges that imposed in this particular situation. And a story about how we can push through those challenges when we all work together to overcome them.

Which to me is one of the key lessons to be learned here: that while everything takes longer in a pandemic, when we work as a team, we can push through those challenges and accomplish a lot.

This network infrastructure convergence lab is a work in progress, and in fact I hope and expect it will always remain one, with new equipment and technologies being incorporated for years to come. In fact, additional contributions are more than welcome, as we'd love to enhance what we can do in the lab and bring in as many different pieces of equipment for as many different services from as many different manufacturers as possible.

The lab may not be as far along as we'd like, but thanks to the contributions of multiple individuals we will get there , which I hope to be able to report on in the future.

Abbreviations

AN	Aggregation node
CFP2-ACO	C Form-factor pluggable 2 – analog coherent optic
CCAP	Converged cable access platform
CTD	Coherent termination device
DAA	Distributed access architecture
DWDM	Dense wave division multiplexing
G / Gbps	Gigabits per second
HFC	hybrid fiber-coax
P2P	Point-to-point
PON	Passive optical network
RF	Radio frequency
RPD	Remote PHY device
RMD	Remote MAC/PHY device

