VIRTUAL EXPERIENCE
OCTOBER 11-14

2021 Fall Technical Forum
SCTE® • NCTA • CABLELABS®

Extended-CIN

A Remote Head-End Solution for Space Re-Allocation in CIN Deployment

A Technical Paper prepared for SCTE by

Deepa Phanish, Ph.D.
Technical Analyst
Cox Communications
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328
+1 404-664-8816
deepa.phanish@cox.com

Alan Skinner
Principal Engineer
Cox Communications
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328
+1 404-210-3122
Alan.Skinner@cox.com

John Huang, IP Engineer, Cox Communications

Igor Tavrovsky, Reliability Engineer, Cox Communications

Ernest Fabre, Design Engineer, Cox Communications

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Cox's Distributed Access Architecture (DAA) standard calls for deployment of a Converged Cable Access Platform (CCAP) chassis acting as Remote PHY core in every metro edge facility. To support the deployment of these chassis, we need substantial amounts of rack space, power, and HVAC at these critical facilities. In locations that have severe constraints we need an alternative solution to enable Remote PHY while avoiding highly expensive facility augmentations. In this paper, we explore our network design options to deploy a CCAP chassis non-locally in a host facility and to utilize our Converged Interconnect Network (CIN) to reduce the footprint at the remote edge facility. We discuss how to prioritize these solutions by their impact to service availability. The reliability analysis provides further insights into the design/decision thresholds for selecting a host site in a successful application of the remote head-end Extended (E)-CIN solution. We further discuss the implementation, limitations, and challenges of the E-CIN solution and assess impact to the business in terms of capacity planning and cost in comparison to the full-CIN solution. The outcome from this comprehensive analysis is very useful in deciding favorability of a given site as an E-CIN candidate.

1. Introduction

The access network has continued to evolve by leveraging new technologies – the recent one being Remote PHY– to meet the ever-increasing demands for bandwidth. Remote PHY is a distributed architecture that encompasses moving the physical (PHY) component from the traditional Cable Modem Termination System (CMTS) out to the edge, thereby extending Ethernet closer to the customer and providing the capability to support greater bandwidth. This architecture will ultimately enable cable operators to deliver Gigabit service tiers at a fraction of the cost of replacing the existing Hybrid Fiber Coax plant with fiber. In this context, CIN is a flexible, resilient, and extensible network that interconnects the CCAP core with Remote PHY Devices (RPDs). It is essentially the infrastructure that supports the distributed access and fiber-deep architectures of Remote PHY.

Implementation of the traditional CIN network requires deploying the CCAP core within the service provider's critical facilities. These CCAP chassis have significant space and power requirements that would need to be accommodated at those facilities. An example of the CCAP core is the Cisco cBR-8 shown in Figure 1. It has the below chassis specifications:

- Weight: 429 lb. (195 kg) maximum fully loaded
- Height: 13 RU (22.75 in / 57.78 cm)
- Width: 17.45 in (44.32 cm) with no rack mounts, 17.65 in (44.83 cm) with rack mounts
- Overall Depth: 28.075 in (71.3 cm)
- Operating temperature (nominal): 32 to 104ºF (0 to 40ºC) Sea Level

The Cisco cBR-8 router can be either mounted on the rack at the front or in the middle. Also, the router can be either mounted on a standard 19-inch wide four-post equipment rack unit or a two-post rack unit. It is also power intensive with the below requirements:

- Cisco cBR-8 Lifetime Facility Power Requirement: 9000 W
- Hardware Facility Power Requirement (D3.0): 7300 W
- Hardware Facility Power Requirement (D3.1): 7900 W
- Average fully loaded chassis between 4500 and 5200 W

Figure 1 – Cisco cBR-8 CCAP, 13RU Height

To deploy the CCAP chassis and digitalize a facility that is constrained on space, power, and HVAC, we would need cost intensive facility expansions/augments which can run into millions of dollars. In certain cases, an expansion may not even be a feasible option, requiring office move and re-design of the outside plant with new fiber builds. In this paper, we explore the concept of Extended (E)-CIN, which essentially means the digital CCAP is hosted in a separate facility than the site that houses the Remote PHY Aggregation switches (RPAs). In other words, E-CIN involves the de-coupling of R-PHY core resources (i.e., cBR-8 data & video cores) from the edge of the CIN. This can enable digitalization of facilities with physical space constraints and thereby recover rack space from decommissioning obsolete analog equipment, thus avoiding the regrettable spend on expansions. Although physical space constraint is the primary driver, E-CIN can also optimize network efficiency by consolidating host resources and serve as a conceptual proof towards the adaptation of future designs like "centralized CCAP" and/or "virtual CCAP".

In the rest of the paper, we refer to the facility where the cBR-8 is located as the "host" site, and the facility where the RPA resides as the "remote" site. Inherently, in most cases, E-CIN will increase the optical distance between RPDs and the CCAP core, which can add unique challenges, primarily regarding latency performance and network reliability. In section 2, we first consider all the possible network topologies and perform a reliability analysis to derive insights on the best choice of the host site for a given remote site. In section 3, we discuss the implementation of the designed solution. In section 4, we discuss the latency and throughput performance of E-CIN. In section 5, we discuss the financial impact to business for capacity planning and finally conclude in section 6.

2. Network Design

Important design questions on hand for E-CIN are:

1) What are the possible topological solutions for E-CIN?

2) How does one choose the most optimal solution and host site for a given remote site?

In this section, we first explore the topological possibilities by considering both L3 Internet Protocol (IP) and L1 optical Dense Wavelength Division Multiplexing (DWDM) layers. Secondly, we present an approach based on the analysis of network reliability to draw insights on the choice of host site.

2.1. Network Topology

The illustration in Figure 2 is representative of the CIN implementation within the Cox network. HUBC0x are the Metro hub routers controlling metro traffic, and DSRJ0x are the Distributed Service Routers connecting to the backbone network. The Remote PHY Aggregation switches (RPAs) aggregate the Remote PHY Devices (RPDs), and the Digital Physical Interface Card (DPIC) Aggregation switches

(DPAs) aggregate the connections from DPIC providing an ethernet network between the CCAP core and the RPDs. Note that the Host Site could also be the Regional Data Center (RDC).

For the problem on hand, we need to establish communication between the RPDs at the remote site and the CCAP within the host site. There are a couple options to consider:

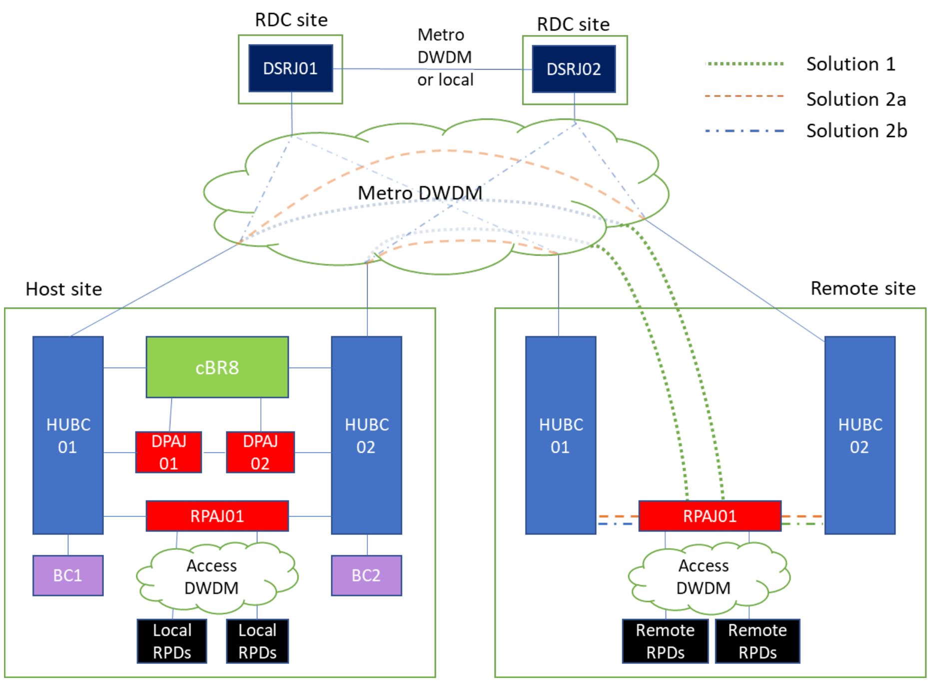

Figure 2 – Generalized Cox Metro Network Topology Showing Routing Options for E-CIN

Solution 1: One way to achieve this would be to leverage the DWDM metro optical network to transport traffic directly between the remote site RPA and the host site hub routers. In essence, the RPAs are re-homed manifesting as the expansion of the host site's access network.
Pros: This would be the least hop solution with low latency, and thus more reliable.
Cons: The solution is not scalable, as aggregation over the DWDM network is very expensive. Therefore, we do not discuss this solution any further in this paper.

Solution 2: The second possibility is to route the traffic via the hub routers. Traffic from the CCAP core hops via the hub routers in the host site before being transported over the metro optical DWDM network to the RPAs via the hub routers at the remote site.
Pros: Provides a scalable solution by aggregating traffic over the uplinks of the hub routers.
Cons: The solution has more hops, and consequently higher latency and lower reliability in comparison to the first solution.

- **Solution 2a**: The hub routers at the remote site can be directly connected to the hub routers at the host site over the metro DWDM network without passing traffic through the Distributed Service Routers (DSRs). This option creates a spur from the host site to the remote site, rather than utilizing the existing hub & spoke architecture. This is achieved by provisioning a primary wave and a secondary protecting wave between the hub routers at the two ends.
  Pros: Fewer hops in comparison to solution 2b.

Cons: Unless there is a pair of direct fiber between the two sites within the optical network topology, it results in a non-optimal aggregation over the DWDM network. Additionally, this alters the IP network topology from the standard (where uplinks connect back to the DSRs at RDC).

- **Solution 2b:** The hub routers at the remote site are not directly connected over to the hub routers at the host site, so the traffic instead hops via the DSRs over the metro DWDM network.
  Pros: Standardized metro topology in alignment with full CIN. Optimal aggregation of waves over the metro DWDM network.
  Cons: More hops and likely longer latency in comparison to solution 2a.

In summary, since scalability, DWDM aggregation and standardization are of top business priorities, our preferred solution is to reserve solution 2a for sites that have a direct fiber pair to another site. This implies we have a single node ring L1 optical network (L3 is secondary spur from primary set of hub routers), which would be the host in this case. For the more common multi-node ring topology of the L1 metro DWDM network (L3 topology is hub-and spoke from DSRs), solution 2b would be more preferential. Selection of host site in this case requires further analysis.

2.2. Reliability Analysis

Separating the CCAP core in one facility from the edge components in another facility increases the number of network elements, and specially fiber distance between the customer premise and the headend. This compels an evaluation of service availability to validate the Service Level Agreement (SLA) requirements, as well as draw insights into the impact of CCAP location on the network reliability.

2.2.1. Problem Definition

In a metro DWDM network, although mesh topologies exist, it is more common to have the ring topology to uplink the hub routers to the DSRs with a primary wave and a secondary protection wave. Considering the worst case - What is the impact on service availability as we vary the CCAP location?

Let's take a deeper look into the traffic path over the metro core in case of solution 2b. The traffic path over an optical ring is shown in Figure 3 for both the steady state and fail-over cases. Considered here is a ring with 18 sites (A to R) spanning over a total mileage of ≈1300km with a single RDC location, akin to one of the largest metro rings on the Cox network. In steady state, the traffic from backbone starting at the DSR in the RDC traverses the shorter side of the ring to the CCAP location, again on the shortest path, returns to the RDC before finally reaching the customer site where RPDs are located. The flow is similar during failover, except that it traverses the longer sides of the ring. In the example shown in the figure, A is the RDC site, E is the site where CCAP head end is located, and C is the customer site of interest. It shows the primary path of traffic flow, and the secondary path when a failover occurs, say on the event of a fiber cut between sites B and C. For this exercise we vary the CCAP and the customer site over the ring and perform a reliability analysis of the metro network.

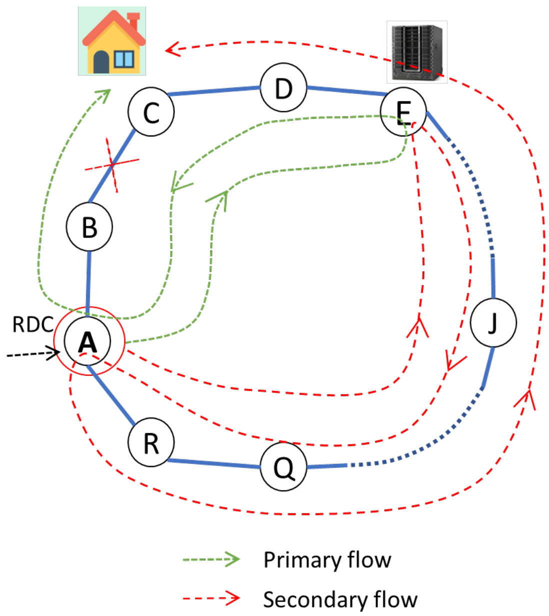

Fiber between locations	Fiber (kms)	Metro Fiber (kms)	Longhall Fiber (kms)
A -> Rs	30.7	24.56	6.14
R -> Q	65.5	13.1	52.4
Q -> P	111.7	11.2	100.5
P -> O	91	9.1	81.9
O -> N	86	12.9	73.1
N -> M	67.9	13.6	54.3
M -> L	55.5	5.6	49.9
L -> K	93.5	14.1	79.4
K -> J	87.9	17.6	70.3
J -> I	111.2	83.4	27.8
I -> H	70.7	28.3	42.4
H -> G	124.4	24.9	99.5
G -> F	58.3	43.8	14.5
F -> E	31.5	4.8	26.7
E -> D	99.8	15	84.8
D -> C	66.3	10	56.3
C -> B	63.8	12.8	51
B -> A	53.5	21.4	32.1
Total	**1,369.20**	**366.2**	**1,003**

Figure 3 – Protected Traffic Paths on a Metro Ring with the Fiber Distances

2.2.2. Modeling and Simulation

For reliability analysis we used the ReliaSoft BlockSim package allowing the creation of reliability block diagram and discrete repairable system event simulation. The fiber distances between sites for the case study is shown in Figure 3. Outside plant fiber mileage is distinguished between metropolitan and long-haul by their urban density for analysis since optical impairments (disruption of signal in the fiber optical link) largely depends on human activities, and they have different failure rates. The failure rates are modeled with lognormal distribution computing the means and variance for long haul and each metro individually. The customer service availability is then modeled with the primary and secondary path elements in a parallel system with all intermediate routes and optical infrastructure at the module level with built in redundancies.

Disregarding the "last mile" OSP infrastructure within the access network, we have simulated the service availability for customers at different sites on the ring as we vary the CCAP locations over the ring as well. For this comparative analysis, the Mean Time to Repair (MTTR) is assumed constant and equal to 4 hours. The results are shown in Figure 4. It can be observed that for a given CCAP location, network reliability is highest when the headend is local and decreases as the CCAP is moved farther away. Secondly, for all sites, the network reliability when the CCAP is located at the RDC site A is almost as high as when it is local. This is apparent from the figure as the line plot for "CCAP in A" is an envelope for the rest. It is due to the nature of traffic flow - locating the CCAP at the RDC does not add any additional fiber distance to the traffic path. Hence, the RDC would be the second choice, next only to local, from the perspective of network reliability. Further, if the RDC cannot host the equipment, then the facility closest to it needs to be considered.

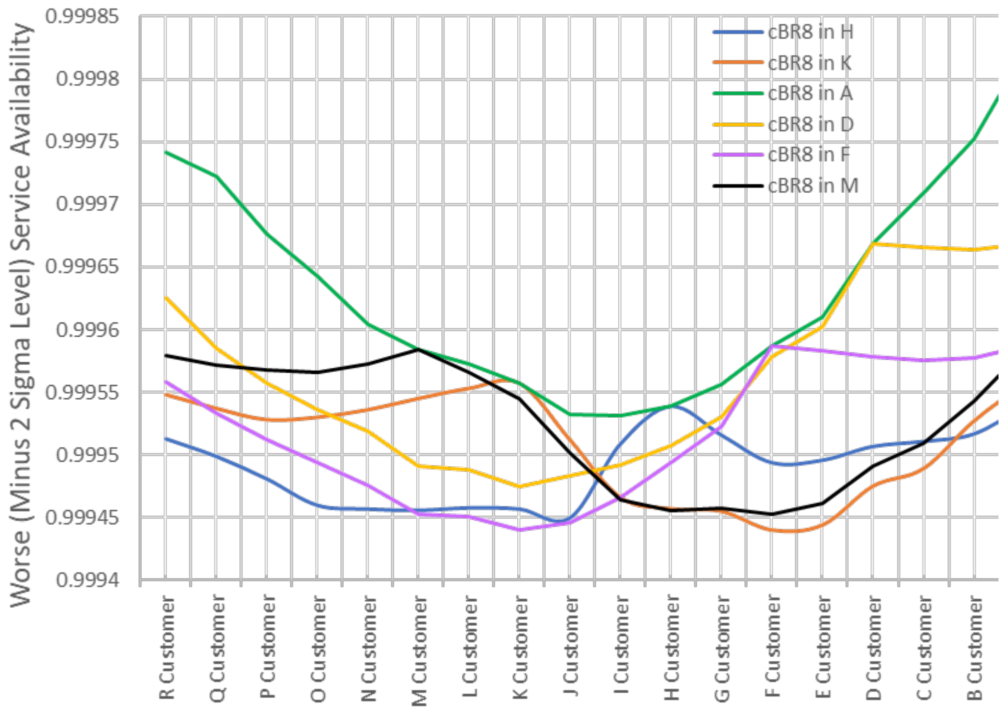

Figure 4 – Metro Core Service Availibity at Different Sites for Fixed CCAP Location

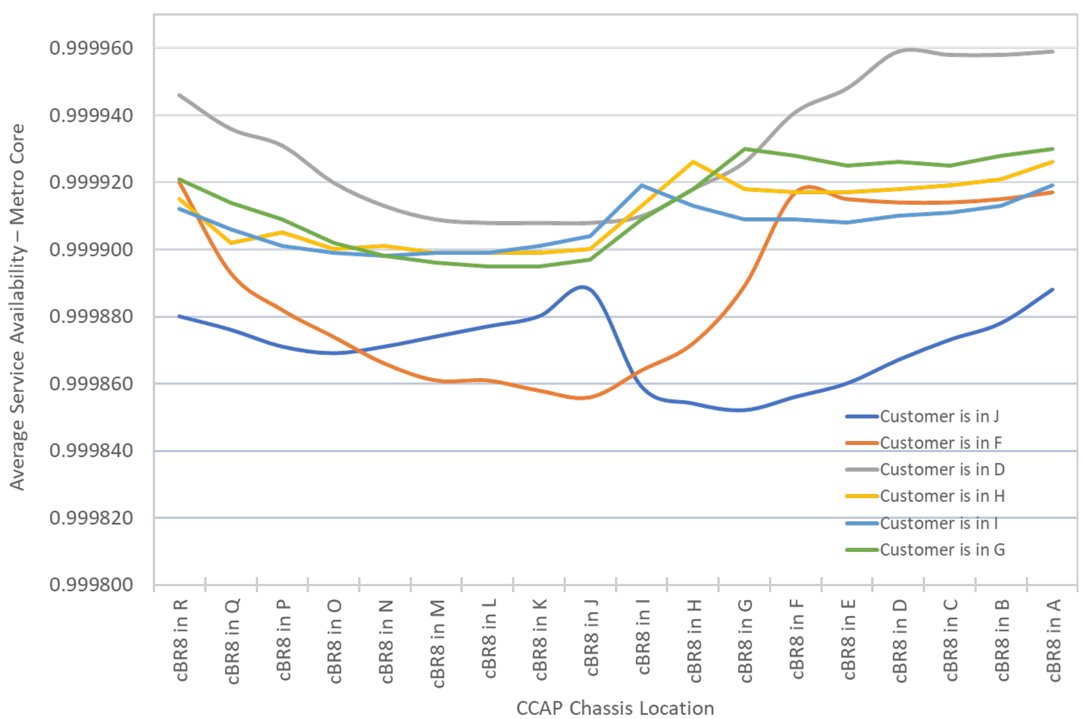

Figure 5 – Metro Core Service Availability at Specific Sites for Different CCAP Locations

In Figure 5, we can see that the inverse relation has a similar trend. It should be noted that the plots consider only the metro core network reliability and does not include the backbone and the Access OSP infrastructure. To qualify a host site, we need to verify that the absolute reliability for the entire network meets the SLA. Due to shape of the plot, it is possible for certain sites within the market to meet the requirements while certain others beyond a distance do not. In the above case study, we performed a detailed analysis including the "last-mile" R-PHY OSP infrastructure at Cox by considering the hardware, software, human factor, and commercial power outages. To be more accurate, we have modeled the MTTR with log-normal distributions with the log location $\mu = 3.3.4576$ and the log standard deviation $\sigma = 0.5287$. Similar distributions are used for RDC and hub site hardware with tailored parameters. This is a computationally very intensive model with over 1000 blocks simulating 5 years of operation and at least 5,000 iterations for acceptable convergence. For remote site F as example, the mean availability for CCAP location in RDC A resulted in 0.99964, and we can therefore qualify A as the host. We have observed a consistent drop in mean availability compared to Figure 5 by about 0.00028 for all host sites with the improved accuracy model.

2.3. Summary

In summary, directly linking the remote site RPA to the host hub routers (solution 1) is not scalable. To handle the RPD scale and to benefit from aggregation, the CIN traffic needs to hop via the remote site hub routers before passing through the DWDM network. Given the optical network topology, the remote site hub routers may have direct primary and secondary uplink to the hub routers at another site, if they have a pair of direct fibers between them. In this case, the second site would be the chosen host site. This is the "***Subtended hub-hosted***" solution. Generally, we have sites with optical nodes in the metro connected as a ring with one or two (split) RDC sites hosting the DSRs connecting over to backbone. In the standard metro IP topology with primary and secondary uplinks from the hub routers within the market to the corresponding DSRs, these uplinks are non-overlapping DWDM waves provisioned over the optical network on the primary and secondary fiber paths of the ring. Hence all traffic from a remote site hops over the RDCs before reaching the host site. Therefore, the RDC itself would be the primary choice for the host site in E-CIN. This not only maximizes reliability as shown in the above analysis, but also minimizes latency. This is referred to as the "***RDC hosted***" solution. In cases where the RDC facility does not meet the requirements for space, power, and HVAC, then the hub site closest to the RDC in fiber distance would be the next choice of preference for a host site. This is most generalized but least optimal, and it is called the "***Hub-hosted***" solution.

3. Implementation

In this section, we discuss the implementation of E-CIN and the necessary configurations specific to E-CIN deployment. Firstly, although general CIN routing guidelines apply, IP addressing needs to be tailored appropriately for successful operation. Secondly, supporting video service becomes more complex due to multiple channel lineups at the host core and requires additional configurations on the CCAP.

3.1. Networking

The host site will have the CCAP, Digital Physical Interface Card (DPIC) Aggregation switches (DPAs), and the Boundary Clocks (BCs). The remote site will have the "edge" CIN equipment: one or more RPAs, Remote PHY access DWDM equipment, and switches for local management/telemetry. However, remote sites will not have local boundary clocks for Precision Time Protocol (PTP); the boundary clocks at the host site serve the purpose instead.

In terms of routing, general CIN guidelines apply with appropriate updates to the IP space addressing and route advertisements to ensure reachability between the RPA, CCAP core and boundary clocks. Route advertisements specific to the Cox network is shown in Figure 6, and the same routing policy works for all solutions of E-CIN. In a typical Remote PHY network, a facility has two PTP boundary clocks and each is connected to a different router which in turn resides either on the "Hub1" or "Hub2" side of the network. BC1 is preferred by all RPDs on the CCAP, because the standard Remote DOCSIS Timing Interface (R-DTI) profile points to it. BC2 only comes into play if BC1 is unreachable. If the primary path between host and remote sites utilizes the Hub Router 1 path, then this standard CCAP configuration is used. If the optimal path between host site and remote sites uses the Hub Router 2 path, however, then BC2 is preferable because it is closer to the RPDs. In this case, a second R-DTI profile is configured on the CCAP and the RPDs in the remote site must use that profile instead of profile 1.

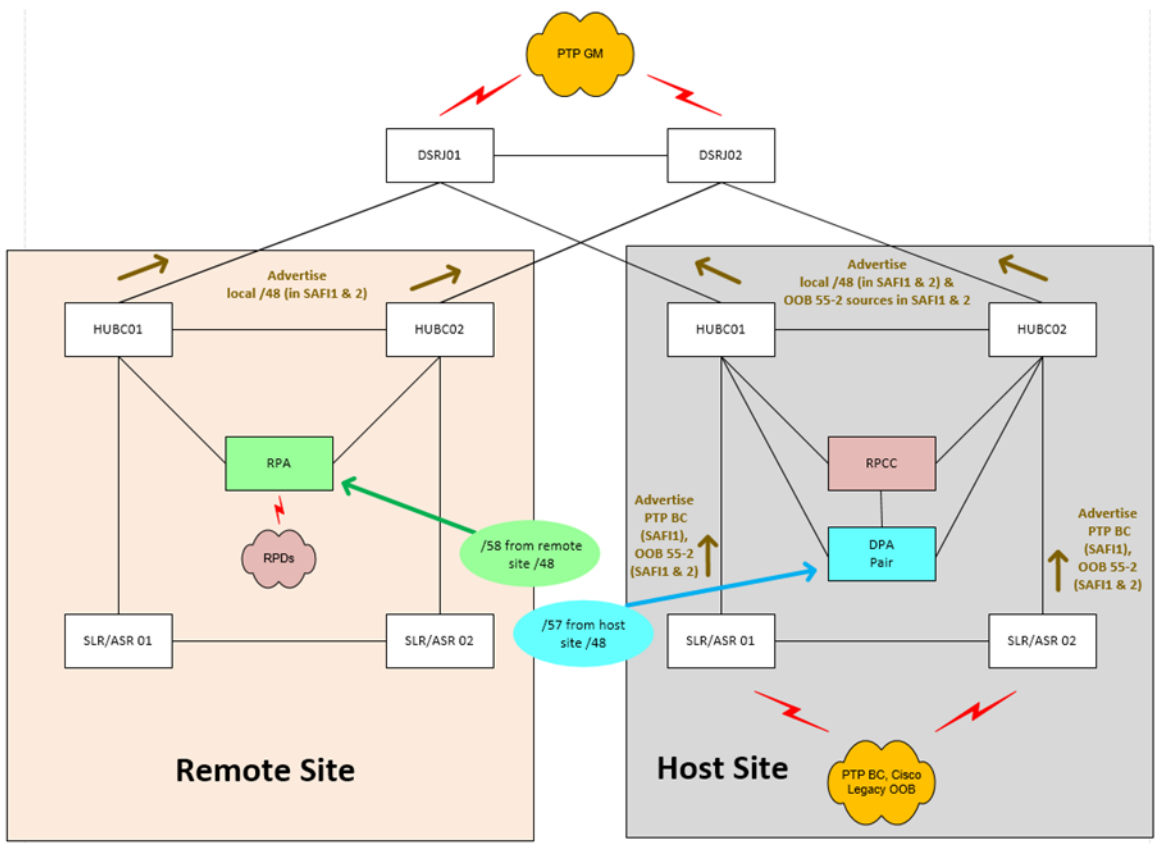

Figure 6 - Route Advertisements in the CIN Network

3.2. Video Support

One of the limitations to E-CIN is the increased operational complexity for supporting video services. In the case where the remote site has the same channel lineup, ad zones, DOCSIS Set-top Gateway (DSG) tunnels, legacy Out-Of-Band (OOB), and Public, Educational, and Government (PEG) channels as the host site, there may be no additional configurations required to support Quadrature Amplitude Modulation (QAM)-based video services including broadcast video and narrowcast video (Switched Digital Video (SDV) and Video on Demand (VOD)). However, if any of the above differ between the host and remote site, then care must be taken to configure the CCAP accordingly.

Consider a Synamedia PowerKey market as an example. Legacy OOB is a Narrowband Digital Forward (NDF) signal delivered via multicast from a Kronback NDX source to a set of destination RPDs. The legacy OOB is assigned to a Digital Hub on the Explorer Controller. It is imperative that both the legacy OOB NDF signal and the DSG multicast flows to the CCAP are assigned to the same Digital Hub on the controller. Failure to do so will render Basic DSG Tuning Adapters inoperable. A CCAP hosting an E-CIN will have to be evaluated to determine whether the remote site is serviced by a different Digital Hub. If so, the CCAP will require a matching NDF pseudo wire for each set of DSG multicast flows.

For any multicast video transport streams sourced at the extended facility, those services will need to be routed back to the host CCAP so they can be ingested and mapped to an output like any other video source. Additionally, care must be taken when supporting multiple Ad Zones in a single CCAP. Multiple Ad Zones means that not all narrowcast video service groups are equal. This calls for additional coordination and operational processes. Before a narrowcast service group is associated with its first RPD, it must be first determined to which Ad Zone the RPD belongs, and then the correct zone should be associated with the service group. Further, any other RPD that gets associated with that service group must also belong to that same Ad Zone.

Extended CIN is limited to environments contained within a single market. There are several reasons for this, but the primary factor is video support. Because all video in a R-PHY network is ingested by the CCAP and then re-generated toward the RPDs, it is not practical (or even possible, in many cases) to support channel lineups and video encryption from multiple markets on the same CCAP. Within a market, when the channel lineup in an extended site differs from the lineup in the host site, the CCAP must be configured to support both, which can add additional complexity. Sometimes, it may not be advantageous, or even possible, to group some sites together. In such a case, multiple CCAP chassis would be needed at the host site to support E-CIN.

Here are some of Cox's best practices on CCAP configuration. Some of these could be relaxed or eliminated if a standalone video core were used in place of a converged data/video CCAP:

- No more than 6 full broadcast service groups per CCAP.
- No more than 12 total broadcast service groups per CCAP including PEG service groups.
- Only 1 Conditional Access System per CCAP. Sites tied to different set-top box controllers cannot be supported on a single CCAP.
- Except for very specific circumstances, only 1 main SDV lineup should be supported on a CCAP. Within the CCAP video config, its linecards are pointed to a SDV session server. The session server is associated with a main lineup. While it is technically possible to point different linecards at different SDV session servers, this should only happen for very small sites whose total capacity needs can be satisfied with 1- or 2-linecards on a single CCAP.

When a host facility is planned to support one or more extended facilities, it is advisable to consider the following options:

1. Support for all extended sites on each CCAP in the host facility. This means that all CCAPs are essentially created equal, and RPDs from any extended site can be moved to any host CCAP. This option is the most flexible but comes at a cost of complexity on the CCAP, and potentially reduced Data Over Cable Service Interface Specification (DOCSIS) Service Group capacity.
2. Segregation of CCAPs by serving footprint. This option would involve setting aside one or more CCAPs as being "E-CIN" hosts and mapping the extended site(s) to certain CCAPs. Doing this would reduce the amount of waste involved in the universal configuration, but at the expense of having to keep track of RPD mappings.

3. Cox is also considering a third option to consolidate video onto a standalone dedicated video core. In this scenario, only the DSG tunnel configurations on the Data CCAPs would be needed to host multiple remote facilities. The CCAPs could therefore be fully utilized for DOCSIS.

4. Performance – Latency, Throughput, and Distances

There are many factors that contribute to a customer's perceived latency. Only the contribution of the access network (CCAP to cable modem) is considered here. Figure 7 illustrates the DOCSIS request/grant cycle and the impact of E-CIN on the same. It is critical to have symmetric latencies from core to RPD and vice versa for both PTP and DOCSIS traffic for Remote PHY operation. It can be observed that the optical distance between the RPD and the CCAP core is the primary factor influencing this latency performance.

The latency between core and RPD is higher in an E-CIN environment primarily because of the additional fiber distances from the metro DWDM network. In a lightly loaded system, where congestion is not a factor, customer ping times to the CCAP are generally 5-6 ms in the best case. Those times start to go up almost linearly as you add distance between the core and the RPD as shown in Figure 8. Typically, a non-Extended hub site can serve RPDs as far as 100km away with only minimal impact to ping times. With Extended CIN, that distance from hub to RPD must be added to the distance from host facility to remote site, which could be as high as 1200km. Because upstream scheduling is done in the core and not in the node, a customer's best-case ping time to the CCAP increases linearly by 4X the latency added by Extended CIN because of the DOCSIS request/grant cycle. For example, if 400km (2ms) is added to the distance between the host facility and the remote facility, the customer will see an increase of 8ms in their ping times (best case).

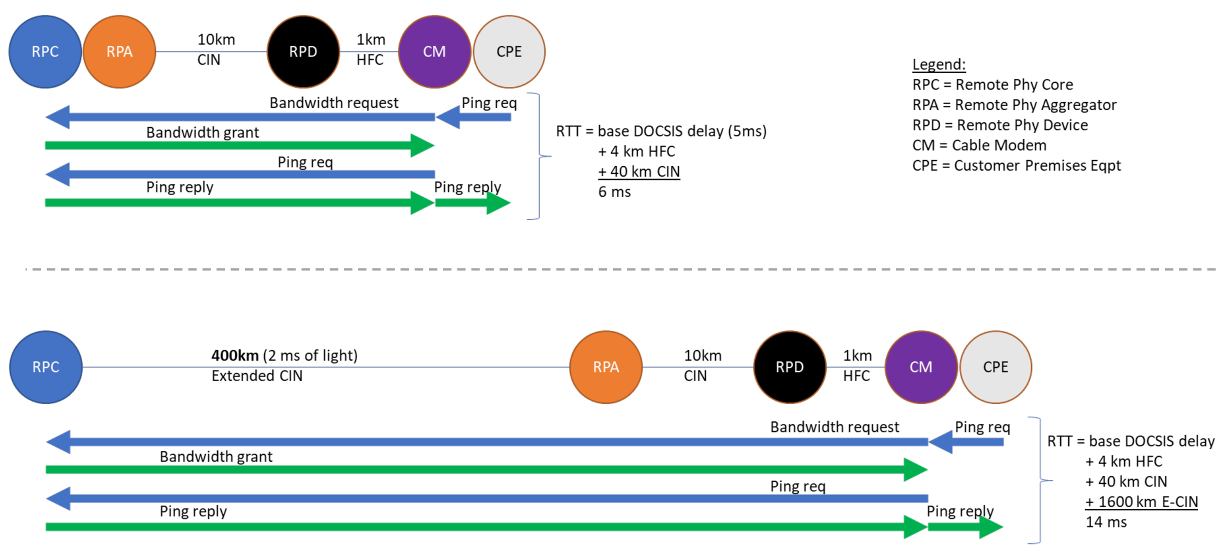

Figure 7 – DOCSIS Ping Time from CPE to CCAP in Regular vs Extended CIN

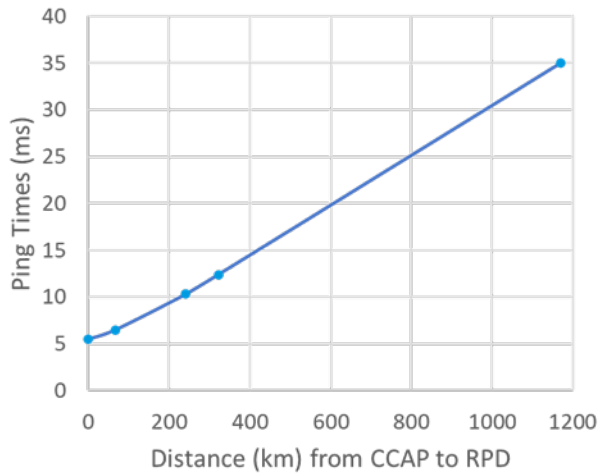

Figure 8 – Latency with Distance

Since DOCSIS is a timing-dependent technology, it is essential to maintain optimization, both in terms of preferring the shortest path (in steady state) and ensuring symmetrical (forward & return) traffic flow. The degradation in latency performance can be mitigated by ensuring that all Remote PHY traffic takes the optimal path from core to RPD, and from RPD back to core, without being load-balanced across multiple paths. This may require change in the Internet Gateway Protocol (IGP) metrics of the metro IP layer to default to the shortest path. Only during a failover should the suboptimal path come into play.

Field testing has shown that customers should be able to achieve full downstream and upstream throughput, even on the Gigabit tier, at distances up to 320km. Field experimentation is currently in progress to determine the exact distance limitation to preserve full Gigabit download. In the very worst-case scenario, which occurs during a path failover, at 1200km, downstream throughput was inconsistent and did not reach the usual gigabit-class downstream. The modems stayed online and continued providing service in a degraded state. However, use of secondary path needs to be minimized for achieving better throughput as well, in addition to latency. Limited field testing has shown that during a failover from short to long path, or from long to short, RPDs and cable modems generally stay online. Nevertheless, large scale testing would be needed to fully validate whether customers could see an interruption in service while their modems re-register.

5. Impact to Business

5.1. Capacity Planning

Depending on the E-CIN topology, the design may involve some “double-back” traffic across certain segments/links of the metro network. Therefore, bandwidth requirements need to be adjusted carefully. Let’s consider the generic hub hosted design for discussions in this section.

Traffic flow between Remote PHY core and the Extended CIN edge over metro core is unique in that it is deterministic and not load balanced. All the traffic – multicast, forward unicast, return unicast, and PTP – must always take the same path, for every RPD. This ensures traffic routes through the shortest path and there is symmetrical forward and return path latencies, which are necessary to meet the timing requirements for DOCSIS. Figure 9 shows how this can be achieved by designing the IGP metrics appropriately. The red path shows load balanced traffic, whereas the green path shows the L2TPv3 tunnel with a preference on HUB1 side of the network. The optically shorter path, either via HUB1 or HUB2, needs to be engineered as the preferred default and the other as the failover.

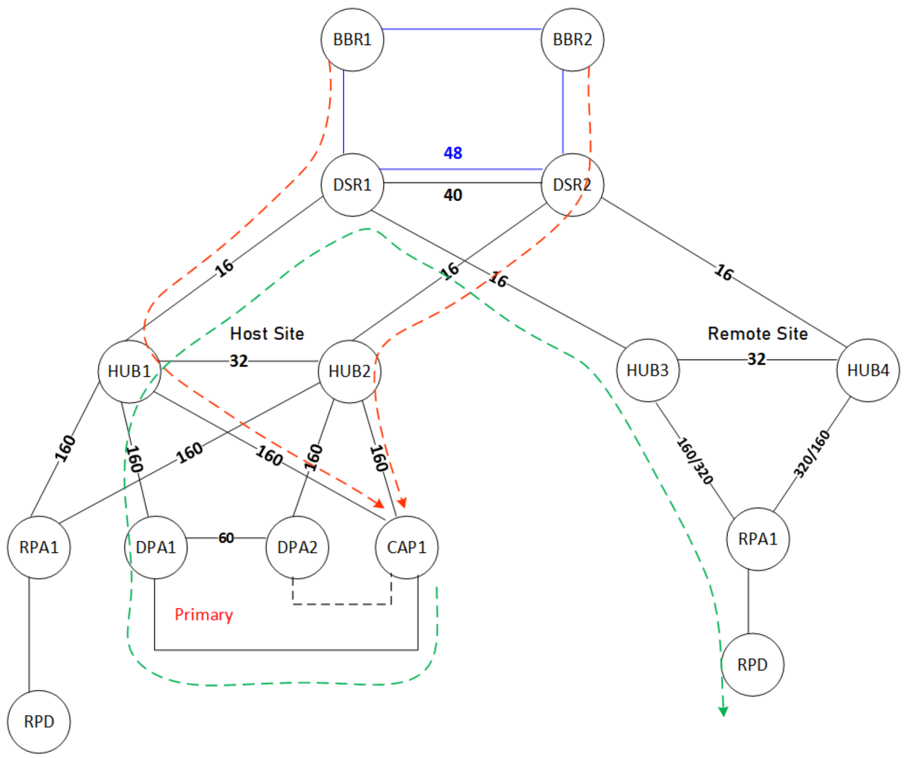

Figure 9 – Traffic Engineering with ISIS Metrics

In planning capacity for the metro network, the CIN component of the traffic needs to be monitored separately. For the metro-core links, we have the below capacity calculations (All traffic rates would be 95th percentiles in Gbps):

Host hub router to DSR uplink capacity = $\lceil max(Td_H + TCIN_R/2\, , M * (Tu_H/2 + TCIN_R/2))\rceil$

Remote hub router cross bar capacity = $\lceil Td_H/2 + M * TCIN_R \rceil$

Remote hub router to DSR uplink capacity = $\lceil max(Td_R\, , \frac{M}{2} * (Td_R + TCIN_R))\rceil$

Remote hub router cross bar capacity = $\lceil Td_R + TCIN_R \rceil$

Where Td_H is the total downstream traffic to the hub routers at the host site, Tu_H is the total upstream traffic from the hub routers at the host site, Td_R is the total downstream traffic to the hub routers at the remote site, $TCIN_R$ is the CIN component of the total downstream traffic to the hub routers at the remote site, and M is the bandwidth margin on the router uplinks. It is a good practice to set M=1.5, which means that the fill rate is below 66.66% in steady state operation. This margin is required to ensure healthy operation, since tunneled traffic - unlike load-balanced - would otherwise operate at 100% rather than 50% in steady sate unable to absorb temporary spikes.

Under special circumstances such as split-RDCs where DSRs in different locations, a single fiber cut could force traffic over the DSR crossbar, in which case, it needs to be augmented as well.

5.2. Cost Benefit

From a cost perspective, E-CIN has a wide range of possibilities. On one end, in cases where we could save expenses on facility expansion it could benefit the company, whereas, on the other end, where traffic rates are significantly large, E-CIN could be more expensive than regular CIN. Cost benefit is generally at best where facility augments can be avoided. Secondly, if resource sharing is maximized, remote site can share the CCAP chassis, its processor, and spare line cards; the boundary clocks; and the DPAs at the host. This could result in additional savings. However, depending on the video configuration, it may not be possible to fully leverage resource sharing, especially in cases where dedicated video core may be necessary.

E-CIN also comes with an additional expense – it needs more bandwidth capacity and hence growth expenditure on the metro core as discussed in the previous section. Bandwidth augments are not only necessary to address the double backing of CIN traffic, but also the tunneling effects viz., no load balancing, and operational margin. The main component here is the transport network cost, which depends on the traffic rates, optical network platform, transponder wave density, and its topology. Unless host and remote site traffic rates are very low, augments may be necessary on the uplinks of either one or both the hubs. We may need to typically add 2-4 wave augments on either end. Therefore, the financial impact for each remote and host site pair candidate needs to be assessed on a case-by-case basis to precisely estimate and compare costs.

6. Conclusion

Extended-CIN provides a novel technology to geographically decouple the Remote PHY core and the edge on already existing CIN infrastructure. It is useful reducing the footprint on a facility by consolidating core resources and therefore avoiding expensive augments. However, it can add unique challenges, primarily regarding network reliability, latency performance and operational complexity for video support. These risks should be minimized by optimally choosing the design and host sites with additional case specific analysis. It would be preferable to reserve E-CIN for cases where optical distances between the remote site's edge and the host site's core are relatively low and the cost benefit is high. As part of the future work, we will be evaluating the significant and applicability of E-CIN as newer technologies such as Remote MAC-PHY and virtual CCAP arise.

Acknowledgements

The authors would like to thank Deependra Malla, Jason Cole, Bill Wall, and James Stockdill at Cox Communications for useful discussions.

Abbreviations

BC	boundary clock
CCAP	converged cable access platform
CIN	converged interconnect network
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
DPA	dpic aggregation switch
DPIC	digital physical interface card
DWDM	dense wavelength division multiplexing

DSG	docsis set-top gateway
DSR	distributed service router
E-CIN	extended-converged interconnect network
IGP	internet gateway protocol
IP	internet protocol
MTTR	mean time to repair
NDF	narrowband digital forward
OOB	out-of-band
OSP	Outside plant
PEG	Public, educational, and government
PHY	physical
PTP	precision time protocol
QAM	quadrature amplitude modulation
RDC	regional data center
R-DTI	remote docsis timing interface
RPA	remote phy aggregation switch
RPD	remote phy device
SDV	switched digital video
SLA	service level agreement
VOD	Video on demand

Bibliography & References

Cable Television Laboratories, Inc., 2015. *Distributed CCAP Architectures Overview Technical Report*. Cable Television Laboratories, Inc

Chapman, J.T., 2013. *DOCSIS Remote Phy*. White Paper, Society of Cable Telecommunications Engineers.

Malla, D., 2021. *Modernizing Cox Communication's Access and Aggregation Network for PHY Deployment.* Society of Cable Telecommunications Engineers.

Datasheet, cBR-8 Converged Broadband Router, Cisco

