New Converged Access Architectures for Cable Services

John Ulm, Gerry White
Motorola Mobility

Abstract

CMTS architectures have been evolving over the years, from the initial platforms with fixed upstream and downstream ratios to flexible integrated and modular CMTS. Today, we are seeing the emergence of a new generation of converged CMTS + UEQAM called CMAP. This paper reviews various CMTS/CMAP architectures including integrated and modular options. We then discuss migration strategies from today's silo infrastructure to a single converged platform. From here we investigate potential future evolutions as the industry prepares for all IP video delivery. We also look at the inclusion of application level services into these converged edge platforms to support new added value services.

INTRODUCTION

The authors have had the privilege of participating in the genesis of the cable broadband industry at LANcity and then author the first DOCSIS specifications. We have watched DOCSIS grow and change over the last 15 years. DOCSIS 1.0 enabled data services for cable operators to compete with DSL. DOCSIS 1.1 introduced QoS capabilities to enable voice services. DOCSIS 2.0 enhanced the upstream bandwidth while DOCSIS 3.0 enabled wider bonded pipes to withstand FTTP competition. Now, the next major impetus for DOCSIS evolution is IP video delivery.

IP video delivery over cable has been slowly gathering critical mass, but to date has always just been on the edge of taking off. Numerous architectures and other technical proposals such as M-CMTS and CMTS Bypass have been discussed over the years to enable IP delivery; but these have not received widespread adoption. With the proliferation of intelligent video capable devices such as gaming consoles, smart phones and tablets there is a renewed emphasis on IP video. When combined with the latest industry efforts around a converged next generation CMTS + EQAM platform known as CMAP; these may be the final ingredients needed for widespread IP video deployment.

Below we dive into the various CMTS & CMAP architectures. Then we look to how to enable existing equipment to be integrated with CMAP equipment to provide a smooth migration from early IP video deployments to IP video everywhere. Finally, we delve into the future evolution of CMAP and the integration of application level services into these converged edge platforms.

THE EVOLUTION OF DOCSIS ARCHITECTURES

In the early days of DOCSIS, CMTS systems tended to have fixed upstream to downstream ratios since the early generation silicon required a tight coupling between upstream and downstream. The price per CMTS channel then was measured in tens of thousands of dollars. These early CMTS systems enabled broadband data service and voice services to get off the ground and become well established. As internet traffic continued to skyrocket, people in the industry started to look at ways to significantly reduce the cost of DOCSIS downstream channels so that DOCSIS could expand to meet this growth. These efforts eventually led to the Modular CMTS (M-CMTS) and DOCSIS 3.0 specification work at CableLabs.

M-CMTS

While DOCSIS 3.0 emphasized larger data pipes via bonded channels, the M-CMTS effort was focused on breaking the logjam on cost per DOCSIS channel. The first critical concept it introduced was the ability to decouple the downstream from the upstream. This would allow operators to add downstream capacity independent from upstream. In earlier systems with upstreams and downstreams bundled together, you paid handsomely for the upstreams even if you only needed downstream capacity.

The second important aspect of M-CMTS was the desire to ride the EQAM cost curves for the downstream PHY (i.e. QAM + Upconverter technology). EQAM costs were on the order of one tenth the cost of a DOCSIS channel at that time. So the Downstream External PHY Interface (DEPI) was developed as part of the M-CMTS work. However, this required a change in existing EQAM. They needed to include a DOCSIS Timing Interface (DTI) to support critical DOCSIS timing. This created a new category of EQAM devices called Universal EQAM (UEQAM).

Modular vs. Integrated

A number of the CMTS vendors recognized that the M-CMTS benefits could also be implemented inside their single integrated CMTS (I-CMTS) system. It was possible to decouple the upstream and downstream inside existing CMTS and apply the underlying technology that was driving down EQAM costs to the CMTS downstream cards.

A pictorial view of I-CMTS and M-CMTS is shown in Figure 1 which illustrates some of their key differences. From a block diagram level, the main difference for M-CMTS is that the downstream PHY has been moved over to the UEQAM. From a cost perspective, M-CMTS was intended to create a healthy vendor ecosystem with lots of competition.

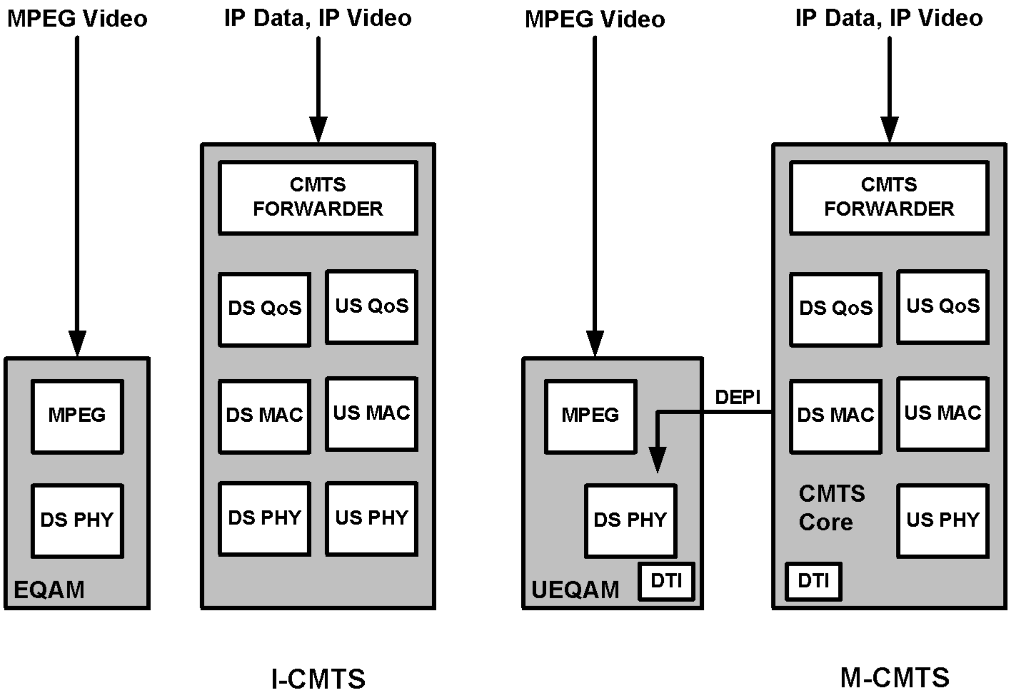

Figure 1 I-CMTS and M-CMTS Components

While multiple UEQAM products have come to market, this is balanced by the limited number of CMTS core vendors that chose to implement the DEPI functionality.

The M-CMTS approach has additional hidden costs. M-CMTS requires a DTI server that both the CMTS and UEQAM need to implement. There are also interconnect costs that may include an Ethernet switch. The costs are compounded as redundancy is factored into the system.

On the operational side, the I-CMTS has a clear advantage. It offers a single vendor system integration that leads to reduced operational costs. For M-CMTS, the operator needs to be the system integrator. The M-CMTS approach also took a longer time to get to market since it was a new multivendor interface. Early adopters went through numerous integration challenges.

DOCSIS IPTV Bypass Architectures

Competition between I-CMTS and M-CMTS helped drive down DOCSIS channel costs quickly. However, EQAM costs continued to drop at an equally quick pace. As the cost of DOCSIS channels headed towards $1,000, the EQAM channel costs drove towards $100. This led to the development of a DOCSIS IPTV Bypass Architecture. This was first published in 2007 [1]. Other variants of CMTS Bypass have appeared over the following years.

While CMTS Bypass approaches promised more cost effective delivery of IP Video, it was optimized for Constant Bit Rate (CBR) video using multicast delivery. It has several drawbacks when considered for more general purpose delivery. There are issues which have prevented wide scale deployment including unicast delivery, VBR, DOCSIS 3.0 bonding, mixing data & video in the same channel and privacy.

THE DAWN OF A NEW ERA

NGAA + CMAP Highlights

In 2009, Comcast expanded its Next Generation Access Architecture (NGAA) initiative to include other MSO’s and vendors in an effort to define the requirements for a converged CMTS, EQAM + PON product. This would be called the Converged Multiservice Access Platform (CMAP). Several papers on CMAP may be found in the references [2, 3, 4].

CMAP provides a highly integrated system with all services in a single box. It supports voice, data and IP video services along with legacy MPEG video, digital broadcast and commercial services. A major tenet of CMAP is that all narrowcast services for a given serving group are delivered from a single CMAP RF port, greatly collapsing the RF combining network. This leads to significant Head End space savings and extensive operational savings. Hence, the CMAP needs to function as both a CMTS and an EQAM. The CMAP based solution offers significant operational advantages. It provides a single management point for the entire system, including video, voice and data services.

This level of integration is now possible thanks to the on-going advances in silicon technology. CMAP targets 64 narrowcast channels per port with up to 96 digital broadcast channels shared across multiple ports. The goal is to have downstream blades which are based on EQAM technology. The same downstream technology that is driving EQAM costs lower will also drive CMAP downstream blade costs as well.

Given the large number of downstream channels in a CMAP device, another key CMAP attribute is its ability to reassign any channel to be either a legacy MPEG video or a DOCSIS data channel. For example, an operator could initially deploy CMAP with 80% of the downstream channels being legacy video; then over time adjust the mix until they reach 100% IP delivery. This will be one of the key CMAP features that enable the transition to IP Video services.

With the downstream PHY technology increasing density tenfold, the rest of the DOCSIS solution including CMTS forwarder and downstream QoS must scale as well. DOCSIS traffic in a typical CMTS often needs significant processing. This includes classification to map packets to a service flow followed by complex traffic scheduling. IP Video traffic has some key characteristics such as large packet sizes and fixed delivery intervals that will allow the CMAP to process these packets with lower overhead than the HSD traffic on today's CMTS. This has sometimes been referred to as a "DOCSIS Lite" approach. Effectively, the bypass concepts are now being implemented inside a single box rather than across several boxes. This along with new generations of network processor silicon means that the processing overhead of IP Video will be essentially the same as legacy MPEG video transport.

CMAP will enable a "standards" based solution. All IP traffic moves transparently over the DOCSIS infrastructure. No special control protocols are needed for IP Video transport in contrast to bypass options which require a proprietary protocol at the EQAM and Cable Modem Gateway. With a generic DOCSIS solution, you can simultaneously support multiple IPTV protocols over the same infrastructure. You also have flexibility to change over time as new protocols such as adaptive streaming take hold. A CMAP solution is ideally suited for generic unicast IP video delivery such as Over The Top (OTT) video or any other managed IP video traffic coming from the "cloud".

The CMAP solution also allows you to take advantage of the statistical gains from large DOCSIS bonding groups and VBR video delivery. CMAP can enable bonding groups of 16 or 32 channels that deliver up to 1Gbps downstream service rates. Because of the large bonding groups, the CMAP solution can mix high speed data with the IP video and provide full DOCSIS services over a single pipe. This provides for a very flexible network transport since it appears as a single IP pipe to the system. This is another feature that enables the IP Video migration.

Integrated and Modular CMAP Overview

The CMAP architecture supports both an integrated (I-CMAP) and a modular approach (M-CMAP). We will look closely at these two CMAP approaches and discuss their pros and cons in rolling out an IP Video delivery system. The major components of both are shown in Figure 2 below.

The I-CMAP components look almost identical to the I-CMTS shown earlier. The ability to process MPEG transport streams is added to the downstream MAC block. This is the additional capability required to implement EQAM functions. The other key addition is a common CMAP OSSI component that enables multiple devices to appear as a single configuration and management entity. While the I-CMAP functional component blocks are the same as I-CMTS, the I-CMAP capacity is tenfold larger than today's CMTS. Given the similarity of component integration to today's CMTS, it is expected that most existing CMTS vendors will opt for this approach, at least out of the gate.

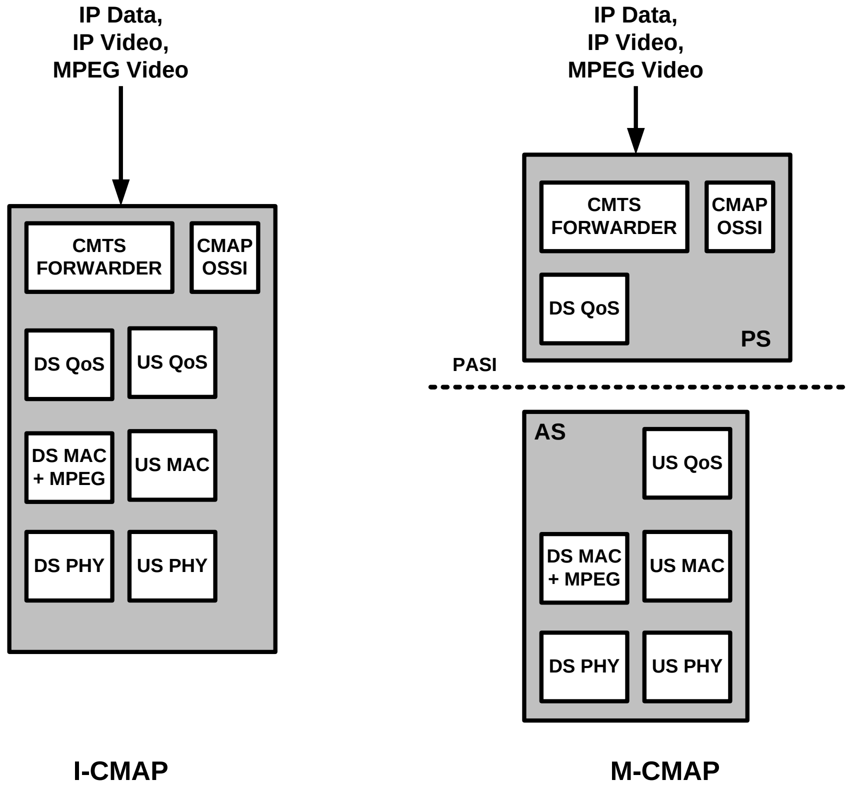

Figure 2 I-CMAP and M-CMAP Components

The M-CMAP approach partitions the CMAP functions into two pieces: a Packet Shelf (PS) and an Access Shelf (AS). A comparison of Figures 1 and 2 shows that this is a radically different partitioning than M-CMTS. The Packet-to-Access Shelf Interface (PASI) specification will become a new CableLab's standard to define the communication between the PS and AS.

At a high level, the PS contains all of the L3 and higher functions including classification, routing and traffic management for the downstream. The PS also includes the DOCSIS Downstream QoS and a large part of the DOCSIS control plane. This partitioning will lend itself to implementation by traditional router vendors, a number of whom participated in the CMAP specification development. One PS may control several separate AS.

The AS will primarily contain the PHY and lower level L2 MAC functions. The AS also contains the DOCSIS upstream QoS and the remainder of the DOCSIS control plane. Unlike M-CMTS, the AS will integrate both the upstream and downstream components. This eliminates the need for an external DTI. This split also has the benefit that the AS can be physically remote from the PS. The PS could be at a centralized Head End serving multiple AS, with some of the AS located in remote hubs. Many of today's EQAM vendors will look to AS products as they migrate to this new converged EQAM + CMTS world.

While the components in Figure 2 give a good representation of the data plane split in M-CMAP, it does not touch on the control plane. The DOCSIS 3.0 MULPI spec (MAC and Upper Layers) assumed an integrated approach like that used in I-CMAP. In

PASI, DOCSIS control and state machines have been split between the PS and AS. In some cases, they are shared between the two. A key tenet of M-CMAP is that the PS is the sole focus for the operator for configuration and management. The entire M-CMAP system appears as a single management entity with the PS coordinating with the multiple AS below it.

Modular vs. Integrated . . . Again

Once again, the cable industry is setting itself up to support two competing architectures: integrated vs. modular. Below is a discussion on several topics on how modular stacks up against an integrated approach for CMAP.

Expanded Vendor Ecosystem – One of the key rationale for a modular approach is expanding the vendor ecosystem. This is certainly true in the CMAP world as M-CMAP will enable both traditional router vendors and EQAM vendors to participate in this new converged CMTS + EQAM market. Having both architectures will lead to a healthy vendor ecosystem with the corresponding competition and innovation. However, an expanded vendor ecosystem by itself is not an indication of one approach being better than the other.

Best of Breed Components – In general, one of the advantages of a modular architecture is that you can mix and match to pick the "Best of Breed" for each individual component in the system. In M-CMAP, you split CMAP into a routing & packet processing engine and an Access shelf. The Best of Breed philosophy is only as good as the selections within each component. As we witnessed in M-CMTS, there was no choice selection when it came to the CMTS core. An unknown question for M-CMAP will be how many router vendors implement the Packet Shelf. If you do not get all of the major routing vendors, then you may not get your Best of Breed choice for routing.

In the CMTS world, some of the most complex components deal with the DOCSIS control plane and the upstream technologies. The existing major CMTS vendors have invested over a dozen years and hundreds of man-years of effort into DOCSIS interoperability testing and certification waves. If the current CMTS vendors opt for I-CMAP, then the Best of Breed DOCSIS and upstream implementations will not be available in an M-CMAP system.

M-CMAP does allow for other access technologies to be easily introduced. This may be well suited for introducing EPON technologies. In today's PON world, most OLT are L2 or simple L3 devices that would be well suited to the PS/AS split in M-CMAP. An EPON AS also does not need to worry about the complexities of splitting the DOCSIS control plane. This is a much simpler partition. Later down the road, the modular split may also enable other wireless technologies like WiFi hotspots or maybe 4G wireless technologies.

Rate of Innovation – Often a modular architecture enables vendors to focus on their area of expertise. This can offer the best likelihood for innovation provided the innovation rests completely within that modular component. The modular approach may actually inhibit innovation if the innovation must cross component boundaries such as PASI. The desire of an operator to maintain flexibility in selecting vendors often leads to the "least common denominator" of features being deployed. If a PS vendor and AS vendor combine to provide added value, then this is simply a two vendor integrated system rather than a truly modular one.

For I-CMAP, it is easier and faster for a single vendor to innovate in an integrated system without the need to coordinate changes with external partners and execute PASI interoperability tests. This may lead to better feature enhancements from integrated vendors. More innovation is possible when you own the entire solution.

Scaling from small to large hub sites – With an M-CMAP split, it's possible to make a simpler AS to fit into small hub sites. The PS based on traditional router platforms does not easily scale down to small hub sites. This leads to an M-CMAP scenario where a single PS supports multiple small hubs with remote AS. This is especially important as operators look to support small remote hub sites in a "lights out" manner. With the PS remotely located at a central Head End, these remote hubs can now be managed from a single location. This may turn out to be one of the most appealing attributes of a modular approach.

Outside of these small remote hubs, how does Modular stack up as we scale to medium and large scale systems? Conventional wisdom may say that the AS can achieve higher densities than I-CMAP because it has less functional components inside. I-CMAP must implement the CMTS Forwarder and downstream QoS functions above and beyond the AS. For medium size systems, it is not clear if this advantage will be enough to outweigh the cost/space/power of adding a PS into the system. So, the modular approach may be more suitable for very large sites where the overhead of the PS can be distributed across many AS.

However, there can be hidden costs with M-CMAP as we scale to very large sites. Every AS needs sufficient network resources to connect to the PS. These need to be redundant links as well. Some head end sites may need to connect multiple AS through an Ethernet switch network in order to conserve PS ports or to simplify head end wiring. Thus, a M-CMAP approach adds the cost of many additional 10G or even 100G network ports to the PS, AS and external switch that are above and beyond what is needed in the I-CMAP approach.

In reality, scaling is more dependent on individual designs than on integrated or modular split. Innovative I-CMAP designs can scale just as well or better than M-CMAP. With today's ASIC silicon technologies, an intelligent network processor can implement the CMTS Forwarder and downstream QoS functions in roughly the space that an AS requires for an internal switch to interconnect its PASI to its downstream components. We envision that I-CMAP densities will rival AS densities.

The overall system costs of an I-CMAP design may turn out to be noticeably lower as well. As mentioned above, the M-CMAP system will have added costs for the network interconnect and switches between the PS and AS. There is also the factor of the CMAP packet processing costs. Historically, the cost per port drops as functions are pushed to the edge of the network; as seen by comparing core router with edge router port costs. This begs the question as to whether it is better to push the CMAP packet processing to the edge in an I-CMAP or to implement it in a PS within a core access router? It remains to be seen which can achieve the most cost effective results.

Operational Complexities – One of the most important focuses of CMAP is the operational simplification. This is a major goal of CMAP. From a configuration and management perspective, the PS will make the system appear as a single CMAP entity.

Thus the M-CMAP can be managed and configured as simply as the I-CMAP system.

Beyond configuration and management, we need to consider other operational impacts. With I-CMAP, the operator has a single vendor support model. One vendor is completely responsible for the system and provides a single contact point for troubleshooting and diagnostic support. Overall, there are fewer components in the system which therefore has overall less complexity.

With M-CMAP, the operator has a multi-vendor support model. The operator must become the system integrator. This is especially critical while the modular technology is still very immature. Troubleshooting problems across multiple vendors significantly increases complexity and is substantially more challenging. The M-CMTS DEPI spec is an order of magnitude simpler than the CMAP PASI spec and we can see how long it took to iron out all the DEPI interoperability issues. The multivendor integration will also be complicated in that one or more of the vendors participating may have never received DOCSIS CMTS certification before.

Another important operational impact is managing the availability of features and bug fixes and coordinating software updates. With a single I-CMAP vendor, the operator can be assured that a SW release has been tested as a complete system. With a multivendor M-CMAP system, the operator needs to coordinate the various SW roadmaps between the different vendors. If a bug requires a simultaneous patch on both the PS and the AS, the operator needs to ensure it works and coordinate the change. As the number of vendors increases and each vendor generates multiple software releases, the interoperability matrix the operator must manage to maintain a functioning system can become very large. M-CMAP becomes a more complex system to operate and manage over time.

Risk and Time to Market – A key item to consider when looking at I-CMAP and M-CMAP system is the risk and time to market impacts of each approach.

The M-CMAP approach needs to create a new Cablelabs spec called PASI. Following the spec creation, the testing procedures and certification infrastructure must be put into place at CableLabs. The entire ecosystem for M-CMAP needs to be created from scratch. Once the ecosystem is in place, the time to deliver new features will require additional time for PASI qualification before being ready for market.

The PASI spec also introduces new risks. It is partitioning DOCSIS functionality that has been operating in the field for a dozen years as a single entity. The downstream QoS and the DOCSIS control plane have become more complex in M-CMAP. On top of the split, new signaling is being introduced to try and manage a multivendor system as a single entity. The amount of industry effort and time to qualify PASI may match or exceed the multiyear efforts required for DOCSIS 1.0, 1.1 and 3.0 certifications and qualifications.

The one area where M-CMAP may reduce risk is in the area of new access technologies such as EPON. The PASI interface partition is well suited for the functions currently supported in EPON OLT.

The I-CMAP approach can leverage existing CMTS designs to get to market faster. I-CMAP devices will run through existing CMTS qualifications. You do not need to wait for the PASI ecosystem to get into place. Going forward, new features can then be added quicker with single vendor integration since you eliminate the added steps for multivendor interoperability and PASI qualifications.

The integrated approach will have fewer boxes leading to a simpler implementation. An example of this is the downstream QoS and DOCSIS control integrated in a single product. Thus, single Vendor integration can mitigate risk compared to complex multi-vendor interoperability around brand new standard.

IP Video Implications – CMAP will be a disruptive change in space, cost and power of a head end. It will enable the continued growth of legacy narrowcast video in the short term followed by the transformation from legacy MPEG to IP video. With its tenfold increase in QAM channels per port, CMAP will enable both the number of channels needed to roll out these services and the cost per channel to make this economically viable.

For the most part, I-CMAP and M-CMAP are identical in their ability to support IP video as described above. However, there is one key area that we need to investigate where they may differ, and that is in the area of multicast distribution. This topic was discussed in detail in reference [6].

From that paper we saw an example of an 80 service group system, For the I-CMAP system, ~2Gbps of multicast video needs to be delivered to each I-CMAP. For a M-CMAP system, the PS might perform the multicast replication. This would require ~50Gbps of bandwidth between PS and AS plus additional links for redundancy. If the AS performs the multicast replication, then QoS problems may be introduced if this traffic is mixed with other traffic that has already been groomed by the PS.

As operators look to make their transition plans to IP Video services, it will be critical that they understand the issues of unicast vs. multicast delivery and the resulting impact on M-CMAP vs. I-CMAP systems.

CMTS EVOLUTION CONTINUES

We do not envision that CMAP will be the end of the CMTS evolution. It will continue to morph and change to meet customer needs. The next section of the paper discusses some possible directions that the CMTS/CMAP evolution might take.

Mixed I-CMAP and M-CMAP Systems

While a lot of the previous discussions focused on comparing and contrasting modular with integrated CMAP, in actuality it makes sense to use the two together. In Figure 3 below we show an AS that is subtended to an I-CMAP device.

The I-CMAP device acts as the PS for the AS. Thus an operator could choose an I-CMAP initially for any of a variety of reasons like reduced risk, faster time to market, appropriate sizing for its head end. Then later when the operator needs to expand beyond the capacity of that I-CMAP, this can be done with the addition of an AS. For example, the I-CMAP may initially support 40 Service Groups. After splitting nodes, the operator could add an AS to the I-CMAP to support the additional 40 SG.

Instead of having two separate I-CMAP boxes to manage, this is configured and managed as a single box. This approach of an I-CMAP with an AS may also offer the best way for an operator to introduce new access technology such as EPON for business services by adding an EPON AS.

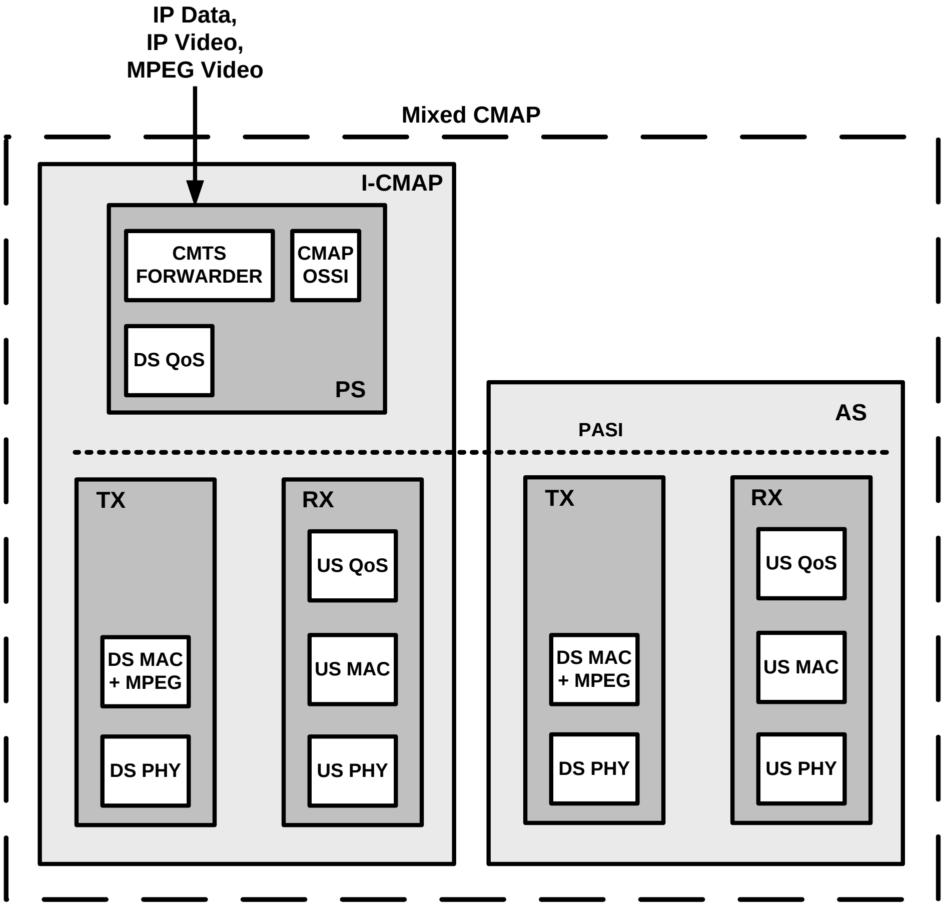

Figure 3 System with mixed I-CMAP and M-CMAP

Distributed CMTS Architecture

The CMTS architectures will continue to evolve as we witnessed when I-CMTS approaches evolved after the M-CMTS specs were written. The M-CMAP PASI specification was created to enable router vendors to become Packet Shelf providers and EQAM vendors to become Access Shelf providers. As discussed previously, this creates a brand new split inside the 750 page DOCSIS MAC and Upper Layer (MULPI) interface spec. This split was not based on any existing implementation and has associated risks with it. Existing I-CMTS vendors have developed systems over the years to independently scale the Transmit (TX) and Receive (RX) portions of CMTS. This work enables a different functional split than a pure M-CMAP System using PASI.

In figure 4 below, the AS components have been separated into separate TX and RX components. We shall call the interface between the PS and TX/RX components the *Distributed CMTS Interface*. The TX and RX components can be replicated as needed to increase downstream and upstream capacity while sharing a common PS.

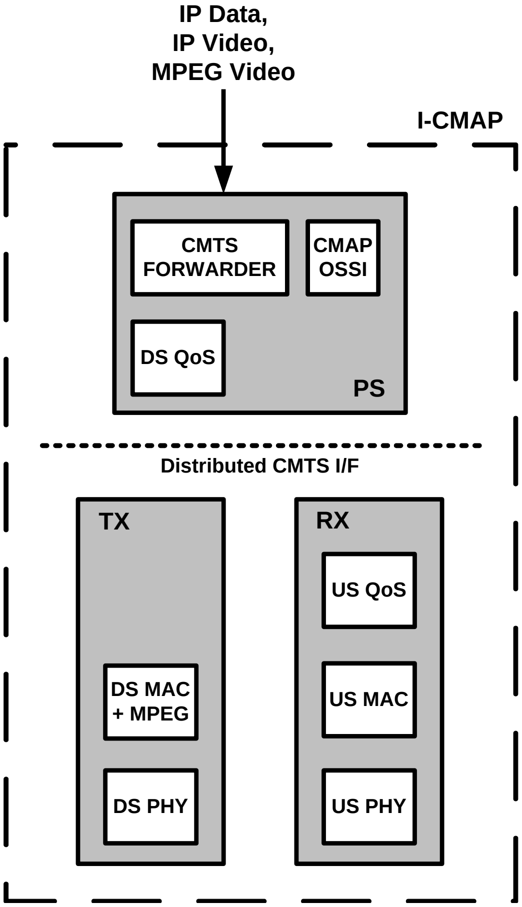

Figure 4 Distributed CMTS Architecture

The PS functions provide the common CMAP capabilities including packet processing, downstream QoS and OSSI. From the operator's perspective, the entire system still looks like a single I-CMAP box. But this alternative split in architectural components enables new possible configurations. The TX and RX components could be located in adjacent chassis. We shall refer to the chassis containing the PS component as the primary chassis while any chassis with just TX &/or RX components will be called a subtended chassis.

Distributed CMTS – Upstream Only

Previously, we discussed a mixed CMAP system where an AS was added to an I-CMAP. In the example stated, this allowed expansion from 40 to 80 Service Groups. Note that this requires a significant amount of interconnect between these components; perhaps as much as 100 Gbps. This could then be doubled if you added extra links for full redundancy.

With a Distributed CMTS system, you can attack the problem differently. Since the TX and RX components are separated, you could put all the RX components in a subtended chassis while the TX components stay closely coupled with the PS functions inside the I-CMAP. This is depicted in figure 5 below.

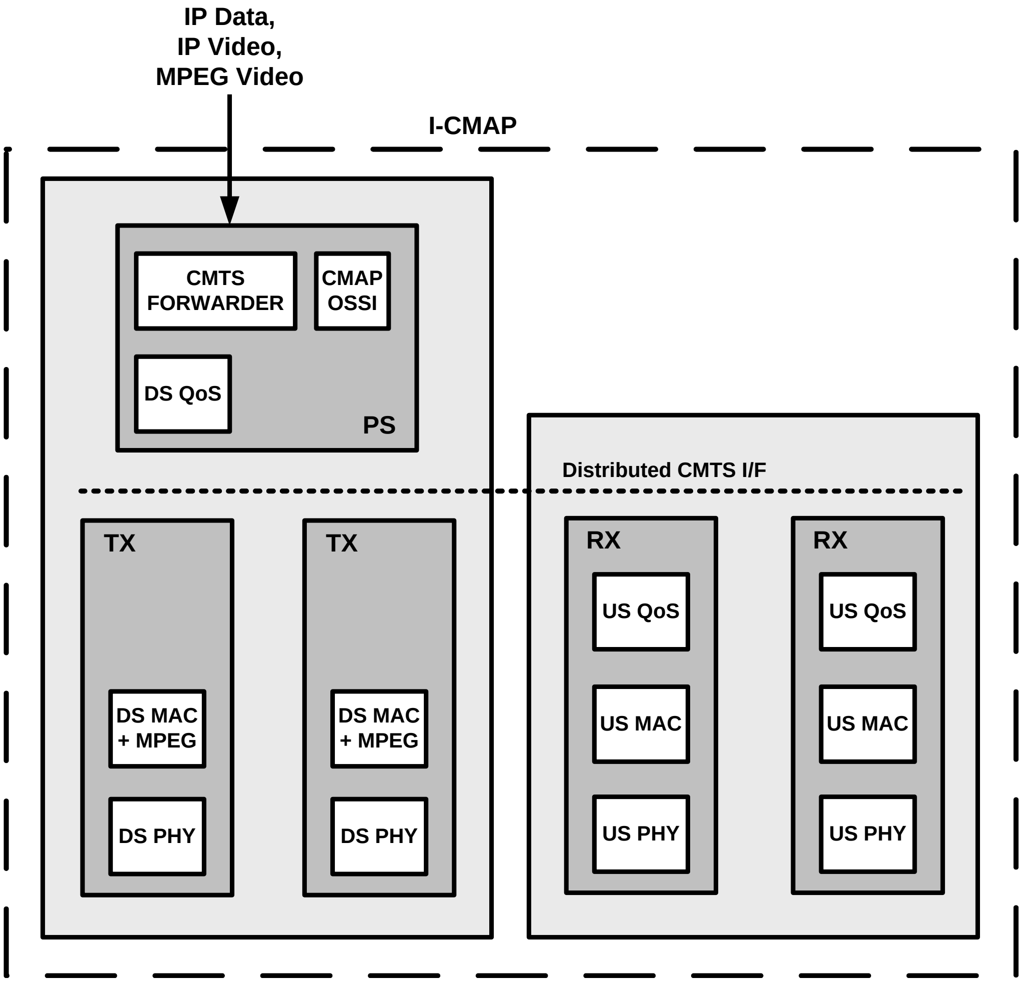

Figure 5 Distributed CMTS Architecture with Upstream Only

This partitioning of the architecture has a significant impact on the interconnection bandwidth required. In CMAP, the upstream bandwidth may be one tenth that required for the downstream. This means that the Distributed CMTS architecture in the above example would only need a single 10 Gbps link to support 80 upstream Service Groups. This compares to the 100Gbps plus redundancy that the M-CMAP approach requires as shown in the previous example.

Another important benefit of this Distributed CMTS Architecture is the reduction in overhead costs for redundancy. The mixed I-CMAP + AS system requires TX and RX redundancy in both boxes. For a 14-slot I-CMAP product, it would implement 5+1 redundancy for its upstream and downstream cards for 20% overhead.

The Distributed CMTS Architecture example above only requires TX redundancy in the primary chassis, and only requires RX redundancy in the subtended chassis. Using the 14-slot example, the primary chassis could implement 11+1 TX redundancy while the subtended chassis implements 11+1 RX redundancy. This is less than 10% redundancy overhead. Thus the Distributed CMTS Architecture has cut the redundancy overhead by more than half.

MIGRATION STRATEGIES

CMAP systems promise a disruptive change in the space/cost/power of future Head Ends. But to get these benefits implies a fork lift upgrade of existing EQAM and CMTS equipment. Since the operators will be investing in EQAM and CMTS equipment up until the day CMAP systems are delivered, a critical question is how can operators leverage existing equipment and transition to a full CMAP system?

The CMAP spec team has partially addressed this by stating that initial CMAP systems should be capable of operating as a UEQAM downstream only product. With this reduced functionality, it is hoped that products come to market quicker to meet existing needs for legacy video expansion of VOD and SDV. Then later, this product can be upgraded to a fully compliant CMAP product. This migration story can apply to either an I-CMAP device or an AS.

I-CMAP + M-CMTS

The above strategy assumes that vendors provide a chassis that operates as an UEQAM today and is upgradable to CMAP in the future. This ignores a very large segment of dense UEQAM that will be deployed before the availability of these new chassis systems. Another strategy would be to use M-CMTS concepts combined with I-CMAP to leverage existing UEQAM. This is shown in the figure 6 below.

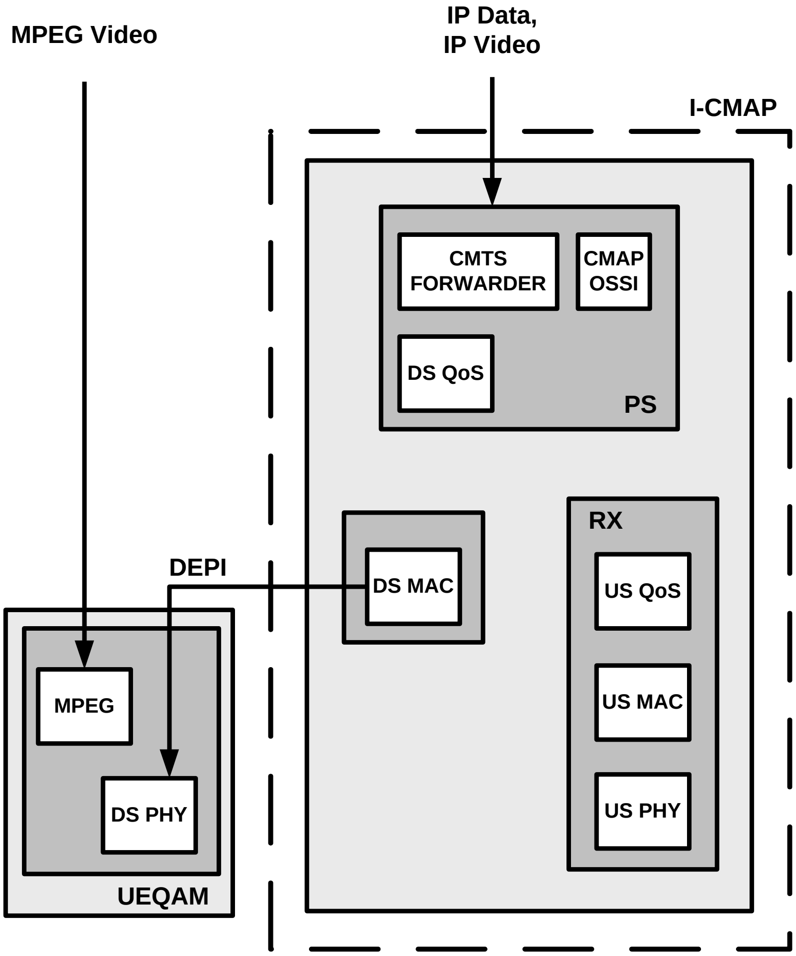

Figure 6 System with I-CMAP and M-CMTS

This M-CMTS + I-CMAP approach is not fully integrated from a configuration & management perspective. This is fine in the short term as legacy video is often managed as a separate silo from the data. But this could become a problem as operators start to converge into a single OSSI model.

In the above, note that the I-CMAP needs a downstream MAC component that then outputs over a DEPI interface to the UEQAM device. The other interesting aspect of the above split is that the legacy video (i.e. VOD and SDV) go directly to the UEQAM and bypass the PS. This seems a bit ironic since our early CMTS evolutions considered IP video bypass of the CMTS core. Now future evolutions could see legacy video bypass the CMAP PS.

Distributed CMTS Migration

A Distributed CMTS architecture can solve this problem in a different way. The UEQAM device may be enhanced to support the downstream MAC functions and become a TX component in the Distributed CMTS architecture. This is shown in figure 7.

By taking this approach, the enhanced UEQAM becomes a subtended chassis and the entire system is now managed as a single I-CMAP entity. This solves the long term migration problem to a single converged configuration and management system.

This approach is more cost effective by eliminating the need of the downstream MAC component from the primary chassis.

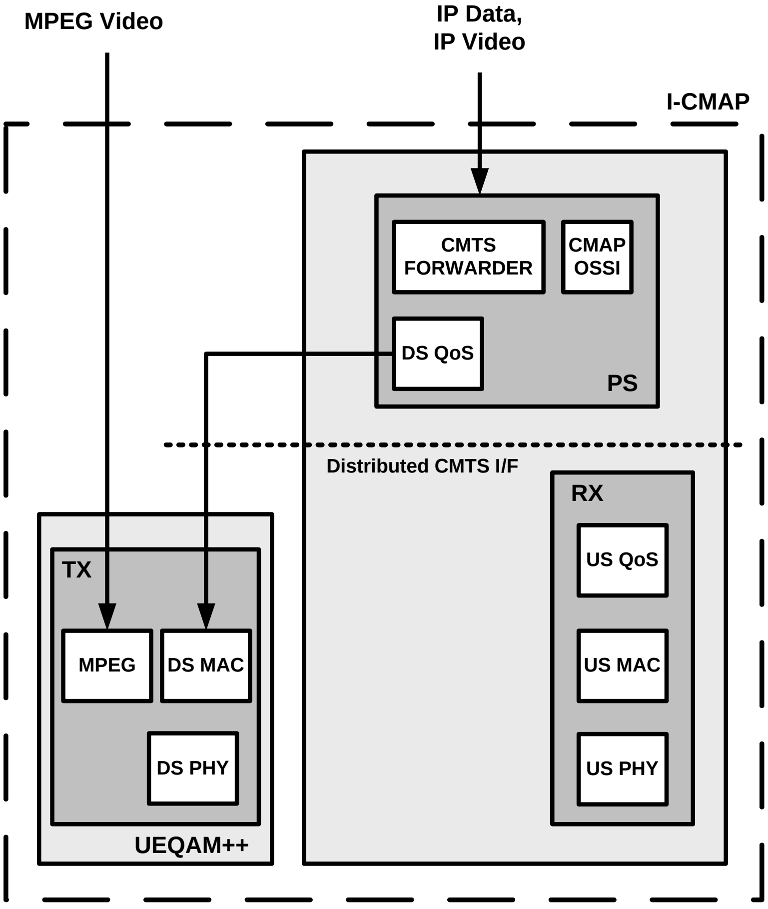

Figure 7 Distributed CMTS Architecture with Upstream Only

APPLICATION LEVEL SERVICES AND VIRTUAL ENVIRONMENTS

To date, CMTS systems have focused on providing data transport and offering L2 and L3 capabilities to service providers. This has enabled traditional data and voice services. As the CMTS evolution continues, we will start to see more capabilities added to these converged edge platforms. They will extend beyond data transport and begin to reach up to the application layer.

Offering basic data services to residential customers allows service providers limited opportunity for differentiation. In order to better compete, the provider should be able to offer additional services layered on top of the data transport. Examples of these include subscriber based services (rather than CM or STB based), business services, application acceleration and caching.

Application level services are typically delivered in a general purpose server environment rather than in an embedded platform. There are however situations when these services must be implemented in the embedded platform, such as when they need access to the data stream or need to have a packet by packet impact on how data should be forwarded.

Trying to determine ahead of time which application level features to embed in a system is to some extent an exercise in frustration. By the time the system goes into operation the applications have changed. In order to resolve this “catch 22” like situation the embedded platform needs to offer a development environment similar to that found on servers, so that functions can be developed quickly and added to the platform post deployment. This has not been practical in the past. Infrastructure systems must provide very high levels of availability and to preserve this have a restricted environment in which any hardware or software changes are the preserve of the system vendor and must be extensively tested prior to deployment.

Recent technology developments offer a potential resolution to this conundrum. The development of multi-core CPUs, hypervisors and virtual machine environments enable infrastructure equipment to offer an application programming environment which appears similar to a general purpose server but which can be isolated from the rest of the system. This can provide a sandbox in which application code can run but with the protection that any failures are limited and do not propagate through the rest of the system. Standard interfaces can similarly be used to enable the addition of hardware processing with filtering engines in the mainstream data forwarding path identifying and diverting packets for specialized processing.

While this appears to offer a utopian future in which equipment vendors, service providers and third party application developers can all add applications and even hardware to in service platforms there are of course several caveats. The applications are in the embedded platform because they need to influence the platform behavior to a greater or lesser extent. The equipment vendor must provide interfaces to enable this but must do it in a fashion so that they can be confident that operation will not be compromised for the rest of the platform.

This introduces the requirement for isolation between application level services and on any impact they have to the data and control planes of the system. The platform must evolve into a virtual environment where multiple service contexts (software and hardware) run in parallel with no knowledge of other contexts and more importantly with no access to them or impact on them. These service contexts must be capable of extending beyond the application sandbox into the rest of the system so that in effect a set of virtual systems is supported. References 7 & 8 describe academic and commercial work along these lines.

The potential value of an application service delivered in this manner can be illustrated by the following example. A cable MSO may offer WiFi hot spot service using DOCSIS cable modems (CM) for backhaul. The hot spots may be open to subscribers from two different wireless service providers (SP1 and SP2) who have negotiated different SLA's with the operator. Subscribers from SP1 may simply receive best effort service with a maximum rate limit while subscribers for SP2 receive an enhanced service set and higher rate limits.

To differentiate between the types of subscriber the system cannot use the CM address (as for conventional DOCSIS services) as it is a shared resource. Thus the subscriber must be identified with a wireless service provider and placed in the appropriate context. In the SP2 context the subscriber may then have access to services such as higher data rates, enhanced QoS, video caching or even local video processing hardware. This identification and service logic may be provided by SP2 and loaded into the converged edge platform.

BUSINESS SERVICES

Business services are a good example of the need to provide a set of virtual systems. Each business customer should have the impression that their service runs in its own virtual context and that this context is not vulnerable to disturbance from other business or residential contexts which share the common resources. For business service delivery the context is not simply a virtual application space but must include the control and forwarding planes. This can be delivered as VLANS, MPLS or L3 VPNs or as a full function virtual router depending on needs.

Next generation converged edge platforms may be deployed primarily as a residential service platform but must be capable of providing a range of business services as well. These services are typically delivered in compliance with well defined service level agreements (SLAs) and place additional requirements on the base platform. An SLA may define bounds on features such as downtime, packet loss, latency and jitter. The platform must not only deliver business services in compliance with these SLA's but must be able to monitor service delivery to confirm compliance. Thus to provide SLAs to each customer and continue to offer residential services the ability to provide virtual contexts and isolation between them is essential.

SERVICE CONTEXT ISOLATION

In order to maintain commercial SLA's the CMAP platform must isolate high value business traffic from the unpredictable loads created by best effort residential services. Ideally this isolation will be provided "end to end" and should be supported for both the data and control planes of the platform.

Data Plane

In the data plane, isolation is typically implemented with three components as the packet moves through the system. Packet classification on reception, internal buffering and queuing while it is stored and scheduling for transmission as it leaves the system.

Classification

Incoming packets must be classified to determine the service context in which they will be processed and the forwarding treatment they will receive in this context. The classification system must be flexible as core and regional networks use multiple options to segregate business traffic. These can vary from basic packet priority, through layer 2 virtual LANS to sophisticated MPLS and layer 3 virtual private networks (VPNs). The classification system needs to recognize each of these mechanisms in order to identify packets for priority treatment, mark them as such and hand them off to the queuing system.

Queuing & Buffering

The queuing system receives packets from the input classifiers and places them into buffers while they wait for transmission. The system is located at a point in the network where there is a very significant discontinuity in interface speeds. Fiber based interfaces to the regional networks operate at data rates which are three orders of magnitude greater than the HFC interfaces. Thus buffer space in the system must be large enough to accommodate packet bursts from the regional networks which can overwhelm HFC capacity. At busy times the buffers may be filled with a backlog of business service and best effort traffic. In order to maintain the SLA when buffers overflow low value best effort packets should be discarded rather than high value business service packets. Overflow in one service context must be contained within that context and not impact other contexts operating in parallel.

Scheduling

When there is backlog in the system packets must be scheduled for transmission based on their "value" but while continuing to provide a reasonable level of residential service. Business traffic should have priority but not cause starvation for residential services. In order to achieve this, a relatively sophisticated scheduler is required for both upstream and downstream traffic, which should be capable of arbitration between services contexts and service differentiation within a context.

Control Plane

The previous section describes the operations required in the data plane to support business services. The control plane must provide the mechanisms to configure and manage these operations. This requires both relatively static provisioning mechanisms such as configuration files and command line interfaces and more dynamic mechanisms such as routing and session control protocols to establish and maintain virtual private networks. Just as the data plane is required to isolate business service traffic the control plane should provide similar isolation. Thus a system may provide multiple control plane contexts. Each context has its own static and dynamic control plane mechanisms so that the control plane for each major business service may be self contained. This avoids problems such as configuration errors from propagating between independent business service offerings.

SLA VALIDATION

To successfully provide business services an operator must not only be able to deliver them according to the SLA, but must also be able to measure this delivery and be able to provide metrics showing how the service is operating. These metrics are needed both for internal use by operations staff and also potentially for delivery to the customer in order to confirm service delivery. Thus the converged edge platform must measure, store and export a number of variables on a per customer basis. These will typically include average and burst data rates, service outages and service impacting errors such as packet loss. They will need to be measured in granular time intervals and exported periodically using a mechanism such as IPDR for off line compilation and analysis. As with data and control plane operations the monitoring and reporting must be cognizant of the service contexts.

UPSTREAM BANDWIDTH

Traffic for residential services is dominated by multimedia downloads and as a result is very asymmetric. Business services are typically much more symmetrical and require higher upstream data rates. Thus service providers need to offer higher upstream peak rates to business customers but more significantly must be prepared to support higher average upstream data rates.

All versions of DOCSIS since 1.1 provide the capability to isolate residential and business services (by defining service flows with specific QoS which can be scheduled independently). This ability to schedule the upstream traffic independently is only useful if sufficient bandwidth is available for both residential and business customers.

Fortunately current DOCSIS 3.0 systems have the potential to operate up to 85MHz. Using the full 5-85MHz spectrum along with the DOCSIS S-CDMA capabilities can enable up to 400 Mbps of upstream bandwidth once the operator needs it. Existing residential services can stay below 42MHz while multiple new 100 Mbps services can be offered in the 40 to 85MHz band. CMAP will enable future downstream rates approaching 1Gbps. These downstream rates will require increased upstream burst rates on the order of 200 to 300Mbps. Next generation CMAP devices must be capable of extracting all possible upstream bandwidth to meet these demands.

MULTIPLE ACCESS NETWORKS

Not all business services will be delivered over DOCSIS. Ethernet and EPON based access networks will be used to offer data rates beyond those practical in the DOCSIS network so that a range of business services can be offered using multiple technologies.

EPON is a point to multipoint technology providing similar services and operational issues to DOCSIS. In order to simplify the operation and provisioning of EPON services for cable operators, the DOCSIS Provisioning of EPON (DPoE) standard [reference 9] has been developed. With this in place integration of EPON into a CMAP like platform is a practical proposition.

While existing EPON deployments are typically based on a 1Gbps version, newer 10Gbps PON technology is quickly becoming feasible and cost effective. For many service providers, it will make sense to go directly to 10G PON services and obtaining a significant competitive advantage. Integrating 10G PON capability into a CMAP system will become a key differentiator in the near future.

Point to point Ethernet will continue to be used for the highest business customer tier. Next generation converged edge platforms can easily support Ethernet interfaces and can offer the same set of enhanced application level services over Ethernet as well as DOCSIS and EPON networks.

SUMMARY

Over the years, we have witnessed the evolution of CMTS architectures as services were added and bandwidth needs increased. From this, we saw the introduction of M-CMTS and competing I-CMTS systems. The move to IP video over cable will be the next major driver of CMTS evolution. It will drive systems to increase downstream port density and reduce costs per downstream channel.

The CMAP architecture is the next major step in this CMTS evolution and has widespread support from multiple operators and vendors. It is currently being standardized through CableLabs. CMAP supports both integrated and modular solutions. Either mode is capable of supporting IP video. Operators will have plenty of choice to select vendors in either camp.

However there are significant tradeoffs between the modular and integrated solutions which operators will need to consider carefully. The areas that need careful consideration include:

- Expanded Vendor Ecosystem
- Best of Breed components
- Operational complexities
- Scaling from small remote hub sites to large Head Ends
- Risk and Time to Market
- IP Video implications
- Migration Strategies
- Rate of Innovation

CMAP systems promise a disruptive change in the space/cost/power of future Head Ends. But to get these benefits implies a fork lift upgrade of existing EQAM and CMTS equipment. Since the operators will be investing in EQAM and CMTS equipment right up until the day CMAP systems are delivered, a critical question is how can the operator leverage existing equipment and transition to a full CMAP system? This will include solutions such as Access Shelves and I-CMAP that can initially operate as UEQAM and integration of M-CMTS concepts with integrated CMAP.

Will I-CMAP and M-CMAP be the end of the DOCSIS CMTS Architecture evolution? Definitely not. We envision that there will be a blending of the two systems where I-CMAP devices act as the PS for other AS. This could give the operators the best of both worlds.

Beyond that, innovation will keep moving forward and both I-CMAP and M-CMAP will morph to adapt. We gave one potential example with a Distributed CMTS architecture that could allow the downstream TX component to be separated and scaled independent from the upstream RX component.

Past this, these new converged edge platforms will need to start integrating application level services to empower the service providers and keep them competitive. This will include expanded business services and expansion into new access technologies, including wireless.

In any respect, this is an exciting time for the cable broadband industry as we begin a new seismic shift.

REFERENCES

[1] M. Patrick, J. Joyce, *"DIBA - DOCSIS IPTV Bypass Architecture"*, SCTE Conference on Emerging Technology, 2007.

[2] J. Salinger, *"Proposed Next Generation Cable Access Network Architecture"*, SCTE Conference on Emerging Technology, 2009.

[3] J. Salinger, *"Understanding and Planning CMAP Network Design and Operations"*, SCTE Cable-Tec Expo, 2010.

[4] J. Finkelstein, J. Salinger, *"IP Video Delivery using Converged Multi-Service Access Platform (CMAP)"*, SCTE Canadian Summit, 2011..

[5] J. Ulm, P. Maurer, *"IP Video Guide – Avoiding Pot Holes on the Cable IPTV Highway"*, SCTE Cable-Tec Expo, 2009.

[6] J. Ulm, G. White, *"Evolving Architectures for Cable IP Video Delivery"*, SCTE Canadian Summit, 2011.

[7] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner.OpenFlow: enabling innovation in campus networks, ACM SIGCOMM Computer Communication Review, 38(2):69–74, April 2008.

[8] OpenFlow switch Specification, www.Openflow.org

[9] DPoE™ Architecture Specificatio, www.cablelabs.com DPoE-SP-ARCHv1.0-I01-110225