THE NEW BROADBAND EDGE:
SETTING THE STAGE FOR THE "*CONVERGED SERVICES ARCHITECTURE*"

John Treece
Juniper Networks, Inc.

Abstract

Since the introduction of the first cable modem termination systems for MCNS support in the mid to late '90s, the access network has continued to challenge Moore's Law: Rapidly advancing technologies, developed to meet the growing needs of cable operators worldwide. The industry now finds itself at the inception of a new and even more complex phase/era. It includes, arguably, the most dramatic and potentially rewarding transformation the cable access network has ever seen.

INTRODUCTION

The genesis of the modular CMTS (M-CMTS) was originally conceived during the cable industry's Next Generation Network Architecture (NGNA) effort. This new architecture works to de-couple the historically integrated platforms, most notably the PHY and MAC layers.

By de-coupling the distinct functions of these devices, suppliers can focus on core competencies, thus creating greater innovation and expedited roadmaps.

At the same time, cable operators benefit from this increased innovation, and gain a more cost effective and flexible approach to scaling the last mile infrastructure and introducing new services. As an example, a near-term benefit with the introduction of these specifications will allow operators to leverage the same QAM modulators used today for MPEG video transport, for data -- essentially creating a "universal edge device." Longer term benefits include implementing a platform that will allow MSO's to migrate to an all-IP environment with the ability to support quad-play services.

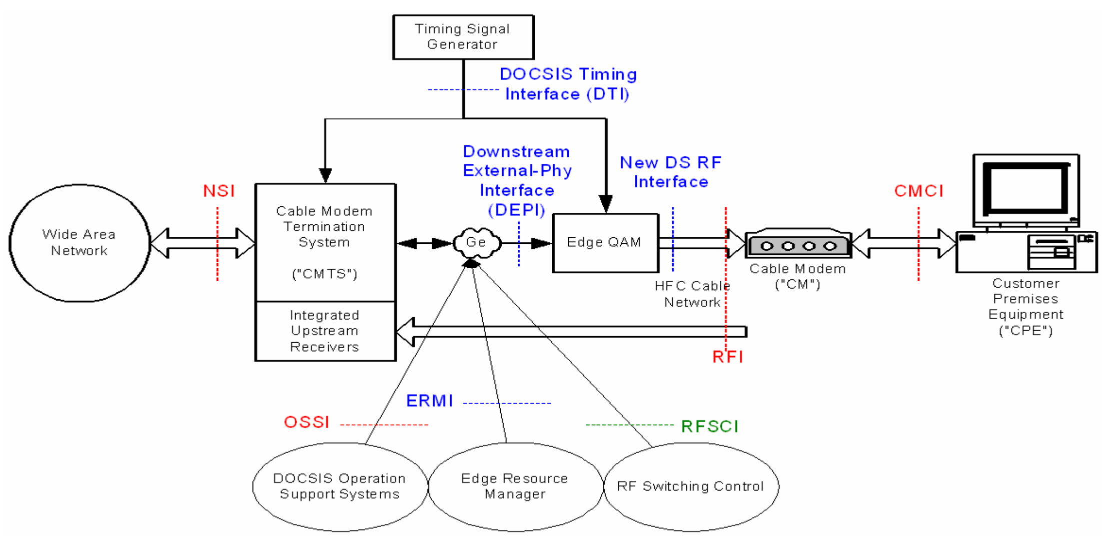

Figure 1

ARCHITECTURE

MSOs have invested upwards of $85 billion dollars over the last 30-40 years on the last mile infrastructure, and they are now starting to focus their capital, both intellectual and economic, on making better and more intelligent use of this existing infrastructure. One could argue that the value of the HFC networks is in its transparency. In other words, it's simply an analog signal pipe. It doesn't care about DOCSIS or QAM for that matter, and becomes a huge advantage for operators as you simply have to change the end points to run new and profitable applications. What this means however is, although we have this great and flexible HFC architecture, really the capabilities and the efficiency and the value of the HFC are dictated by the intelligence of the systems on either end.

This is evident by the emergence of many new initiatives such as, Digital Simulcast, Switched Digital Broadcast, the deployment of Next Generation On-Demand architectures, the move towards downloadable security, new ad insertion technologies, and finally new access network initiatives such as the modular CMTS.

The M-CMTS architecture is comprised of several different components as shown in Figure 1. They include the Timing server which supplies a synchronous DOCSIS clock and timestamp for each component; the Universal Edge QAM, capable of supporting both the DEPI and the DTI Timing interface; the MAC Core which contains all legacy CMTS downstream components, with the exception of the PHY; and the Edge Resource Manager which allocates bandwidth for both native MPEG video and DOCSIS which is primarily used for VoD today.

This new architecture will not only allow for reduced cost in hardware, but also reduced operations cost, because of associated reductions in complexity. An immediate savings in hardware can be seen by comparing today's model to a "pay as you go" model. When introduced, the current generation CMTS added improved density, high availability capabilities, as well as other feature upgrades from the first generation CMTS. These served the needs of the primary services at the time, essentially data services. Today, if MSOs need an additional downstream or upstream port for capacity relief, tied to it are either an additional physical downstream port or multiple upstream ports -- that operators may *not* need.

In addition, the move towards true convergence which will pave the way to "all IP" transport is necessitating the requirement to couple together robust Layer 3 functionality and traditional DOCSIS processing. A common delivery mechanism such as IP will allow for greater innovation and the evolution of new applications that work to support the "anything, anytime, anywhere" mentality that is rapidly emerging in today's consumer.

Figure 1 represents a logical overview of this architecture taken from Cable Labs who produced the initial four interface specifications. They are; Downstream external PHY (DEPI), DOCSIS Timing interface (DTI), M-CMTS Operations Support System Interface (MOSSI), Edge Resource Management Interface (ERMI), and the DOCSIS Radio Frequency Interface (DRFI).

An additional advantage to de-coupling the MAC and PHY portions of the architecture is the ability for the universal edge device to evolve without the burden of being physically linked to the MAC processing components allowing suppliers to innovate and create solutions that were once unfeasible. Because the PHY portion of the network is relatively simple, operators may now look at different physical media that may provide increased efficiency, such as the use of optics. For example, this would eliminate the need for optical-to-RF and RF-to-optical transitions respectively that take place in the remote hub sites today.

FEATURE SUPPORT

With required support for new features such as IPv6, Channel Bonding, new advanced security algorithms, and advanced Layer 3 support, these new platforms will require industry-leading expertise in both DOCSIS and IP. Whereas the features of the CMTS were primarily service-driven in the past, they will now be required to adapt to the advancements in DOCSIS, and ensure interoperability with the complex routing schemes that operators are putting in place. These new networks are being designed to efficiently transport video, data, and voice traffic locally, regionally, and nationally.

This means that next-generation access platforms will require seamless interoperation with these new robust metro networks to provide better scale and flexibility. They'll also need to offer multiple operational configurations to meet the needs of the MSO. And, they'll need to do all of this in a more cost efficient manner.

Of all the new features introduced in DOCSIS 3.0, IPv6 and Channel Bonding in particular have the potential to act as a catalyst for growth and future opportunity. All access network components that support voice, video, and data, will be required to support the legacy IPv4 address space, as well as IPv6. With the number of cable modem users in North America approaching 30 million and the rapid deployment of new services such as VoIP, operators are finding it more and more difficult to procure public address space for the growing number of IP devices. At the same time, the private IP space used for management is also becoming more difficult to manage. With some OSS systems evolving to a more centralized approach, the use of overlapping federated address space becomes less of an option. For these reasons, it is conceivable to think that we will see infrastructure support for IPv6 introduced for some operators in the 2006-2007 timeframe. This is another example of the critical requirement for new access network technologies to not only support historical DOCSIS processing, but to have the ability to support such new features as IPv6.

The need for greater downstream capacity has been analyzed intensely over the past few years as applications have emerged that require more and more bandwidth. This development, combined with increasing competitive pressures, has made Channel Bonding a priority, as evidenced by its inclusion in the now pending DOCSIS 3.0 specification.

Currently, a single 6MHz channel is used with the capability of producing up to 38Mbps (using 256QAM modulation) minus overhead. Operators today have service offers ranging from 1.5-15Mbps downstream. Channel Bonding has the ability to take that 38Mbps and increase it to speeds over 100Mbps with room to grow. Bonding also affords some additional technical advantages such as statistical multiplexing gains within a bonding group. However, data is not the only driver that can take advantage of this increased capacity to the home. As MSOs begin the deployment of IP capable set-tops boxes for video services, the new devices are being developed to support DOCSIS and will have the ability to tune to bonded streams. This becomes increasingly important as high definition content, which requires nearly triple the bandwidth of a standard definition stream, continues to grow as consumer demand also increases.

IMPLEMENTATION

Although MSO's are heavily engaged in the development during this innovative period, they are equally sensitive to protecting their investments and reducing stranded assets. Figure 2 shows a potential migration path that accomplishes the goals of introducing a flexible and scalable architecture, while re-purposing existing deployed assets.

The first diagram to the far left (Phase I) represents today's infrastructure with an integrated CMTS and separate QAMs for uni-cast video traffic. The second diagram (Phase II) shows the initial phase of implementation where the downstream PHY is separated from the legacy CMTS with the MAC core being introduced for Layer 3 and DOCSIS processing. In this phase, the downstream edge QAM device is now used for data, voice, and uni-cast video traffic. Since the upstream interface has yet to be defined, the upstream traffic is handled by the existing CMTS which is re-purposed to act as an upstream termination point.

During this phase, both upstream demodulators or burst receivers and MAC processing is handled within these devices. The final diagram (Phase III) represents a complete separation of MAC and PHY for both the downstream and the upstream.

Although formal work has yet to begin on the upstream interface, there still seems to be a division of thought relative to remote or centralized MAC layer processing on the upstream. Independent analysis continues to weigh the pros and cons of a MAC-less upstream PHY termination device.

Another interesting opportunity is the idea of consolidation of DOCSIS and IP routing functionality in the remote distribution hubs. Today, the CMTS is a consolidated MAC/PHY device that performs limited Layer 3 routing functions. These devices then route to a local aggregation router that usually sits within the same facility. However, the ability to combine robust Layer 3 routing capabilities with the DOCSIS processing can be an

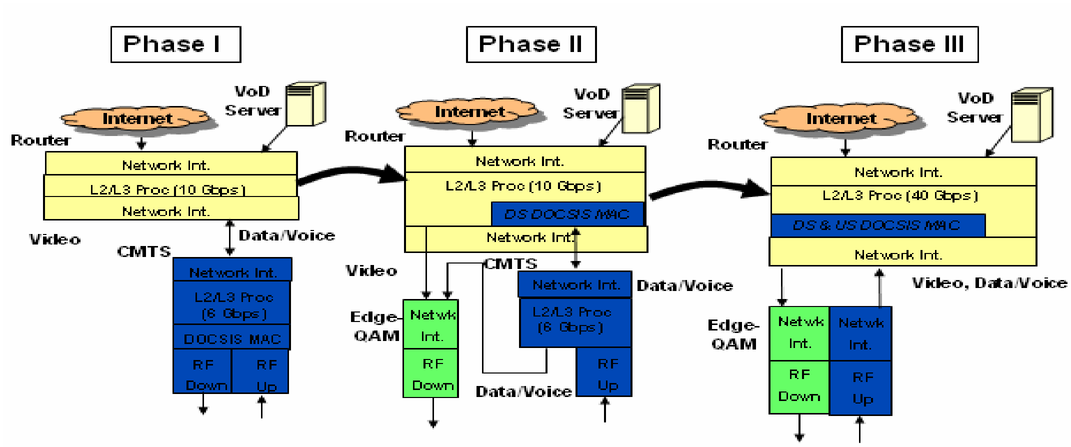

Figure 2

attractive alternative to today's architecture, particularly as operators move towards and all-IP environment. Because DOCSIS has reached a state of maturity within the industry, combined with the separation of the physical layer, it may now be possible to build a DOCSIS processing component or "blade" that lives within a traditional aggregation router. This option has the potential to reclaim valuable real estate within the headends, reduce power and environmental needs, and sets the stage for convergence with one device providing routing and processing for video, voice and data services.

This option provides a true solutions architecture with the ability to scale and evolve with the growing needs of today's consumer.

SUMMARY

Convergence has been a common topic of discussion throughout the last few years, however it seems now more then ever it is becoming a reality. Convergence has historically been categorized and discussed independently as; network, service, or in-home devices. However, it is now quite evident that convergence must and, in fact is, emerging within each of these domains, which is necessary to achieve the "anything, anywhere, anytime", content delivery objective. With this imminent evolution, it is imperative that the edge platforms provide the robust capabilities that will be needed to provide reliable performance for a plethora of services without being limited in their ability to grow with the network.