

TRANSITIONING TO M-CMTS

Andrew Page, John T. Chapman, and Carmelo Iaria of Cisco Systems, Inc.
Glenn McGilvray of Scientific-Atlanta, Inc.

Abstract

The CableLabs® Modular Cable Modem Termination System (M-CMTS™) series of specifications introduces powerful new tools to enhance the architecture of the CMTS. Cable operators have several options for migrating from the current DOCSIS architecture to the M-CMTS architecture, and must evaluate the options relative to their individual success criteria to determine how best to migrate to M-CMTS.

To assist cable operators in formulating an M-CMTS migration strategy, this paper provides an overview of the options and highlights the advantages and disadvantages of each approach. The paper identifies and analyzes key criteria for selecting the M-CMTS Core and Edge QAM (EQAM) options and concludes with a recommended approach for migrating to an M-CMTS architecture.

1.0 INTRODUCTION

The Next Generation Network Architecture (NGNA) framework defined in 2004 by Comcast, Time Warner Cable and Cox Communications asserted that the DOCSIS® CMTS could and should be modularized. These cable operators had the foresight that next generation broadband competition was mounting in the form of xDSL, FTTH and other technologies. To remain competitive, they needed to radically increase the capacity of their DOCSIS networks, while reducing their costs and improving their operational efficiencies. They also needed to maximize the return on investments from their DOCSIS networks by

deploying a flexible infra-structure, capable of supporting DOCSIS 1.x, 2.0 and 3.0 features simultaneously, as well as transitioning all services to a converged, end-to-end IP network. The operators recognized that a modular CMTS architecture could fulfill these requirements and be the key ingredient of the next-generation cable network. In addition, with a modular CMTS architecture, best-of-breed solutions could be deployed, allowing operators to choose components from multiple vendors based on individual product strengths and the criteria deemed most important by each operator.

The M-CMTS series of specifications published in 2005 by Cable Television Laboratories (CableLabs) codified the CMTS architecture envisioned by the NGNA architects. The specifications clearly define the functional components of the architecture and the external interfaces supported by each component. As with most technical specifications, vendors were given the freedom to take different approaches to implementing the M-CMTS architecture. Cable operators must, therefore, evaluate each approach relative to their success criteria to determine which approach to follow. To assist cable operators in formulating an M-CMTS deployment strategy, this paper provides an overview of the available options and highlights the advantages and disadvantages of each option.

2.0 M-CMTS ARCHITECTURE

Prior to the release of the M-CMTS specifications, the DOCSIS CMTS was defined as a single chassis that integrates the DOCSIS MAC and PHY functions and all upper layer protocols. Herein, this type of

CMTS, is referred to as an Integrated CMTS (I-CMTS).

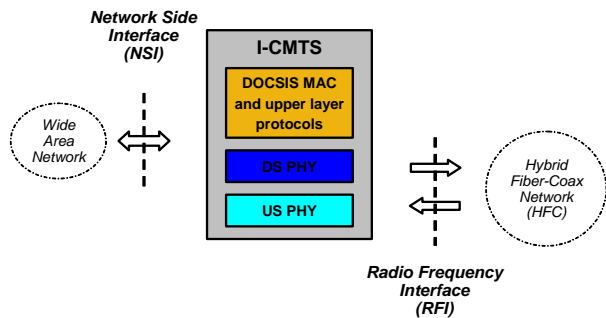


Figure 1: I-CMTS

According to Synergy Research Group, cable operators, worldwide, purchased more than \$2.0 billion of I-CMTS equipment from 2002 to 2005, including nearly 35,000 I-CMTS chassis. With such a large and valuable embedded base of I-CMTSs, cable operators must leverage these assets to their full extent as they transition to an M-CMTS architecture.

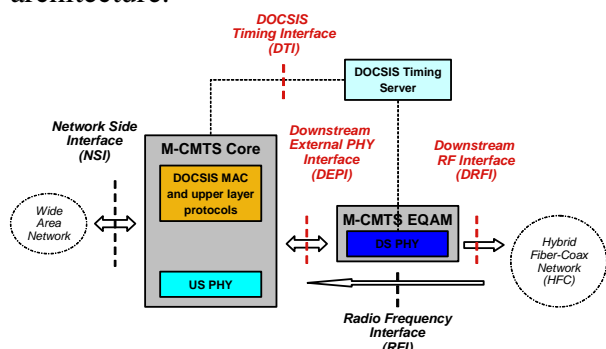


Figure 2: M-CMTS Phase 1

The M-CMTS architecture partitions the I-CMTS into a set of at least two physically separate components with standard interfaces between them. The M-CMTS Core was so named because it retains the core functionality of the I-CMTS, namely the DOCSIS MAC and upper-layer protocols (referred to collectively as Layer 2+). The M-CMTS Edge-QAM (EQAM) supports the downstream PHY function and the Downstream External PHY Interface (DEPI),

which is the interface and associated protocol by which the M-CMTS Core and EQAM communicate with each other. The M-CMTS EQAM must comply with the Downstream Radio Frequency Interface (DRFI) specification, which is an update to the original DOCSIS RFI specification and includes requirements for transmitting multiple QAM channels per physical RF port, a common practice with state-of-the-art EQAM technology. Both the M-CMTS Core and EQAM must comply with the DOCSIS Timing Interface (DTI) specification which defines the interface and associated protocol for synchronizing the M-CMTS components with a common timing reference.

The upstream PHY function can either be supported by the M-CMTS Core or in a physically separate component. However, the current release of the M-CMTS specifications do not define a separate upstream PHY component or the protocol by which the M-CMTS Core and upstream PHY component would communicate with each other. The upstream PHY component and interface [referred to herein as upstream receiver and Upstream External PHY Interface (UEPI)] will be addressed in a future release of the M-CMTS specifications. Of course, vendors will have the freedom to implement the upstream receiver as a standalone device or integrate it with the M-CMTS EQAM.

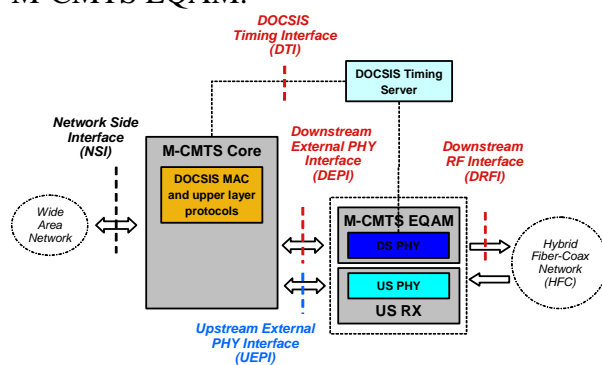


Figure 3: M-CMTS Phase 2

The fact that DEPI has been standardized, while the upstream receiver and UEPI are not yet standardized, implies a phasing of M-CMTS solutions. Such phasing should be considered when formulating an M-CMTS deployment strategy. For the purposes of this paper, we define two M-CMTS phases, as shown in Table 1.

	M-CMTS Core	M-CMTS EQAM	Upstream Receiver
Phase 1	Layer 2+ Upstream PHY	Downstream PHY	(none)
Phase 2	Layer 2+	Downstream PHY	Upstream PHY

Table 1: M-CMTS Phases

In Phase 1, the M-CMTS Core performs all functions of an I-CMTS, except for the downstream PHY function. In this phase, the upstream ports are integrated in the M-CMTS Core, while the downstream ports are supported in the EQAM. In Phase 2, the M-CMTS Core is purely a Layer 2+ component. An upstream receiver is introduced to support the upstream PHY function.

3.0 MIGRATING TO AN M-CMTS ARCHITECTURE

The M-CMTS architecture splits the I-CMTS into an M-CMTS Core, EQAM and (optionally) upstream receiver. There are two approaches to transition from an I-CMTS to an M-CMTS solution:

- The first approach (which we will refer to as Option A throughout the text) is to add the required interfaces to the I-CMTS so it can operate as an M-CMTS Core and utilize an external

EQAM and (optionally) upstream receiver.

- The second choice (which we will refer to as Option B through the text) is to add the required interfaces to the I-CMTS so it can operate as an M-CMTS EQAM and (optionally) upstream receiver, and install a new M-CMTS Core.

This section provides an overview of these options and discusses the key criteria for assessing the M-CMTS Core product choices (section 4.0) and EQAM product choices (section 5.0).

In addition to the M-CMTS Core and EQAM implementation options, cable operators may also have options for implementing DTI when building an M-CMTS solution. The DTI system is based on a client/server architecture. The DTI client must be implemented on both the M-CMTS Core and EQAM components, while the DTI server theoretically can be hosted on the M-CMTS Core, on the EQAM, or on a separate computing platform. However, in this paper, we assume a new DTI Server is required in all cases.

3.1 Transitioning the I-CMTS to an M-CMTS Core (Option A)

Since the M-CMTS Core retains most of the functionalities of an I-CMTS, the most logical approach to building an M-CMTS solution is to evolve the existing I-CMTS into an M-CMTS Core. With this option, the migration can be gradually accomplished in two phases: Utilize a low-cost, external M-CMTS EQAM to perform the downstream PHY function (in Phase 1) and then add an upstream receiver to extract the upstream PHY function (in Phase 2).

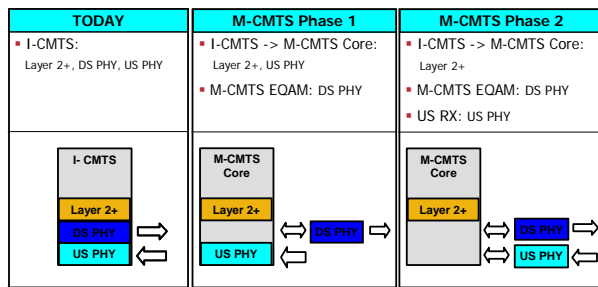


Figure 4: M-CMTS Migration Option A

Executing this option involves the following high-level tasks.

Repurposing the I-CMTS into an M-CMTS Core (Phase 1):

1. Upgrade the I-CMTS to M-CMTS Core
 - a) Add new line card(s) that support the downstream MAC function and Gigabit Ethernet (GE) ports to the existing I-CMTS.
 - b) Add/upgrade the I-CMTS hardware to comply with DTI.
 - c) Upgrade the I-CMTS software to support DEPI, DTI and other M-CMTS-related features.
2. Install a new M-CMTS EQAM(s) that is compliant with DEPI, DTI, and DRFI (or upgrade an existing video EQAM’s hardware and software to comply with the M-CMTS specifications).
3. Install a new DTI Server.

Repurposing the I-CMTS into an M-CMTS Core (Phase 2):

1. Enhance the M-CMTS Core to support external upstream PHY
 - a) Add new line card(s) that support the upstream MAC function and GE ports to the existing I-CMTS (if not

supported on the downstream MAC line card).

- b) Upgrade the I-CMTS software to support UEPI.

2. Install a new external M-CMTS upstream receiver or add upstream-only PHY line card(s) to the M-CMTS EQAM that supports DTI and UEPI.

The I-CMTS is the centerpiece of the DOCSIS network today. It possesses the intelligence for managing all DOCSIS features and services. The success of a cable operator’s DOCSIS service offerings hinge upon the performance and reliability of the I-CMTS. Evolving the I-CMTS into an M-CMTS Core retains the value of the I-CMTS and avoids the risks associated with introducing a new M-CMTS Core.

In Phase 1 of this approach, cable operators are able to use the existing upstream ports on the I-CMTS as part of the M-CMTS solution. In fact, there is no change to the upstream functionality of the I-CMTS in Phase 1. Cable operators can gradually transition to M-CMTS with lower risk of service interruption.

This M-CMTS migration approach is a direct application of a fundamental premise of the M-CMTS architecture: the cost of growing the downstream channel capacity of a DOCSIS network can be dramatically reduced by extracting the downstream PHY function from the I-CMTS and utilizing external EQAM devices that have been cost-optimized for price-sensitive video-on-demand (VoD) applications. Although the M-CMTS EQAM must support new features and functionality specific to the M-CMTS architecture, it is very similar to the video EQAM. Therefore, equivalent pricing is expected. The use of external EQAM devices are not only cost-effective, but also present

cable operators with an opportunity to deploy purpose-built products in their DOCSIS networks from the leaders in the EQAM market, and even use the same EQAM device to support both video and DOCSIS applications.

3.1 Transitioning the I-CMTS to an M-CMTS EQAM (Option B)

The second option for building an M-CMTS solution is to install a new M-CMTS Core chassis and evolve the existing I-CMTS into an M-CMTS EQAM and upstream receiver.

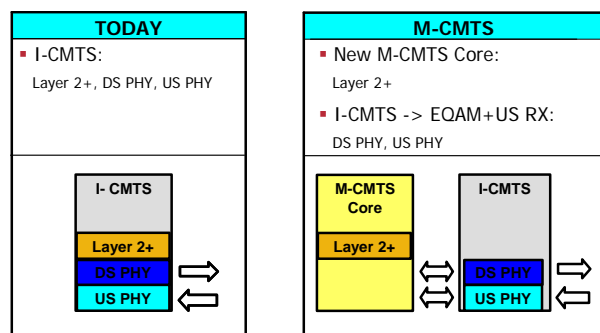


Figure 5: M-CMTS Migration Option B

Executing this option involves the following high-level tasks.

Repurposing the I-CMTS into an M-CMTS EQAM

1. Install a new M-CMTS Core chassis with DOCSIS Layer 2+, DEPI, UEPI and DTI support.
2. Repurpose the I-CMTS into an M-CMTS PHY Layer chassis
 - a) Add downstream-only PHY line card(s) with DRFI-compliant RF ports to the existing I-CMTS (or utilize the existing line cards in the I-CMTS to support the downstream PHY function)

- b) Add upstream-only PHY line card(s) to the I-CMTS (or utilize the existing line cards in the I-CMTS to support the upstream PHY function)
- c) Add/upgrade the I-CMTS hardware to comply with DTI
- d) Upgrade the I-CMTS software to support the DEPI, UEPI and DTI protocols and other M-CMTS-related features.

3. Install a new DTI Server.

Considering that the EQAM and upstream receiver are the low-value components of the M-CMTS architecture, and are the targets for cost optimization, repurposing the existing I-CMTS into an M-CMTS EQAM and upstream receiver devalues the cable operator's investment in the I-CMTS. Introducing a new M-CMTS Core forces the cable operator to re-purchase the high-value Layer 2+ function of the I-CMTS and could introduce risk if the new M-CMTS Core is unproven in a DOCSIS network.

To independently scale the downstream and upstream channel capacity in this approach requires new PHY line cards in the I-CMTS. Considering the advantages of a standalone EQAM product (as described in section 5.0), it is probably more prudent to invest in the standalone EQAMs than in additional PHY line cards for the I-CMTS.

4.0 M-CMTS CORE CHOICES

Making the right choice for the M-CMTS Core is a fundamental step to successfully migrate from today's I-CMTS to an M-CMTS architecture and meet the strategic imperatives that drive such a migration.

In review, the two options for implementing an M-CMTS Core are:

- Option A: Evolve an existing I-CMTS to become the M-CMTS Core
- Option B: Choose a completely new M-CMTS Core

This section evaluates the pros and cons of these two options. Underlying assumptions are: (1) The cable operator has successfully deployed High-Speed Data (HSD) and Voice over IP (VoIP) DOCSIS services, and (2) Throughout the transition to an M-CMTS architecture, the operator must sustain or grow these services and offer new services.

Each option will be tested against the following key evaluation criteria:

- Support for existing DOCSIS features and services.
- Scalability for existing DOCSIS services.
- Support for new DOCSIS services (next 3-5 years).
- Capital budget impact (next 5 years).
- Support existing DOCSIS features and services

The first objective in any network architecture migration is to implement the migration with minimal or no impact to existing services. For this reason, when migrating from an I-CMTS to an M-CMTS architecture, it is critical to ensure that the M-CMTS Core can support all DOCSIS features and services required by the cable operator.

Primarily, this means that all Layer 2 and Layer 3 features currently supported by the I-CMTS must be supported by the M-CMTS Core so that the same level of service, security and performance can be offered to subscribers. At the same time, and perhaps more critical to the success of the migration, it

also means that the cable operator must be able to deploy and manage the M-CMTS Core without significant changes to back-office systems and operating procedures.

Another consideration in assessing the ability of an M-CMTS Core to support existing DOCSIS features and services is the stability of the design. Considering that the primary (Layer 2+) functions of the existing I-CMTS do not change in the M-CMTS architecture, operators can be confident that the I-CMTS will be able to perform equally well in supporting DOCSIS features and services when operating as an M-CMTS Core. If the new M-CMTS Core platform that would be deployed under Option B is unproven in DOCSIS networks, this option adds risk to the operator's ability to maintain existing service levels and operational effectiveness.

Based on the above considerations, choosing Option A (Phase 1) as the first step toward the migration to an M-CMTS architecture reduces the risk of impacting existing DOCSIS services and is a safer approach.

Scaling Existing DOCSIS Services

If seamless support of existing DOCSIS service is a key decision criteria for the M-CMTS Core, migrating from the existing I-CMTS to an M-CMTS architecture would not make much sense unless the M-CMTS architecture could meet the increasing scaling requirements posed by the projected growth of the HSD and VoIP services.

Obviously such growth will depend on a number of factors which are very specific to each region of the world; there is no single rule that can be applied everywhere. Generally speaking though, the trend for the broadband cable market is toward much higher download speeds. HSD tiers of 10 to 30 Mbps in the downstream and from 1 to 5

Mbps in the upstream are expected to be quite common in the next 12 to 24 months, with a number of geographies in the 50 to 100 Mbps range for the downstream. At the same time, penetration rates of HSD service are expected to reach the 50% to 60% range as more affordable broadband tiers are offered to entice dial-up users to upgrade to broadband. What the 'H' in HSD will stand for in just a year or two will be dramatically different than what we're experiencing today.

On the voice front, more and more cable operators are now successfully deploying VoIP services. Bundling VoIP with HSD and video services has proven to be a driving force behind the continued growth of the VoIP subscriber base, as well as a powerful tool for reducing subscriber churn.

The success of these DOCSIS services translates into three fundamental requirements relevant to scaling existing services that cable operators must consider when it comes to the M-CMTS Core decision:

- Enhanced Forwarding Capability
- Strong Control Plane Scalability
- Support for key features such as channel bonding and IPv6

Each of these requirements will now be examined in detail.

Enhanced Forwarding Capability

Enhanced Forwarding Capability means that the overall forwarding capacity needs to be much higher than today's I-CMTS. If an I-CMTS is today providing 1 Gbps downstream capacity to a given serving area, a good benchmark for an M-CMTS Core is the ability to offer at least 5 Gbps to the same serving area, with the ability to migrate to 10 Gbps. Although there will be deployments that require a significant increase on the upstream bandwidth, it is anticipated that, at

least in the short term, the HSD services will remain predominantly an asymmetrical service. This is especially true in light of the fact that the content being accessed by the residential broadband subscribers is becoming richer and richer and that music and video downloads are becoming a standard on any Internet portal.

Control Plane Scalability

While the evolution of the HSD service offerings will dictate the forwarding capacity requirement for an M-CMTS Core, growth in broadband penetration rates, and increased adoption of VoIP among broadband residential subscribers, will require the M-CMTS Core to have much greater control plane scalability than today's I-CMTS. Projections show that an M-CMTS Core must be able to support the control plane traffic for at least 50,000 to 60,000 DOCSIS CPE devices, and ideally up to 80,000 devices.

Support for Channel Bonding and IPv6 Features

Last, but not least, a number of upcoming features such as channel bonding and IPv6 are critical to scale existing DOCSIS services. The channel bonding functionality is fundamental to enable cable operators to offer HSD services with maximum download speeds above 30 Mbps—the practical limit of a traditional DOCSIS 1.x/2.0 cable modem using a single downstream channel. IPv6 support will also become critical due to IPv4 IP address exhaustion as the number of subscribers and their connected devices grows, as new services are offered, and as users demand seamless connectivity in their home network environments.

The ability of an M-CMTS Core product to scale for existing DOCSIS services depends more on the product implementation than whether it is evolved from an existing

I-CMTS or is an entirely new product. Cable operators should carefully consider both options, but be aware of the fact that some I-CMTS products in the market today will be able to scale to meet the requirements listed above, while others will not. For example, an I-CMTS with a passive backplane can grow with the service requirements without introducing fixed bottlenecks in the system. An I-CMTS with a passive backplane architecture can therefore scale without a forklift upgrade. Scalability is achieved by simply adding the newest-generation DOCSIS interfaces, backhaul interfaces and/or routing engine as needed.

Support New DOCSIS Services Planned in the Next 3-5 years

If the ability to scale deployment of existing services is an important driver for the migration towards an M-CMTS architecture, it is equally critical that the new architecture be able to support new types of services that the operator plans to deploy over the next 3-5 years. The range of future services currently being investigated by cable operators is quite wide and depends on the particular broadband competitive landscape in each particular region and the business focus of each operator. In this section, we'll focus on a subset of such services where the choice of an M-CMTS Core is particularly important.

One of the major threats to the cable operator's core business is the move of large Telcos into the residential entertainment video market with both on-demand and scheduled broadcast video content offered over their broadband infrastructure. While the IPTV technology is still at the initial deployment phase, few have any doubts that it will open significant business opportunities and change the way residential users experience video services. We believe it will be critical for the long term success of cable operators to make their migration to the M-CMTS architecture,

keeping open the possibility to deliver Video over DOCSIS (VDOC).

A detailed description of possible VDOC architectures is outside the scope of this document, but it is relevant to highlight key aspects that impact the ability of an M-CMTS Core to support a VDOC service offering. In addition to the characteristics described earlier (such as Enhanced Forwarding Capacity, Control Plane Scalability and availability of features such as IPv6 and channel bonding), a number of fundamental requirements to enable VDOC services exist, including:

- Enhanced IP Multicast support (including Multicast QoS)
- Intelligent Multicast Routing to optimize video bandwidth allocation
- Ability to enable fast channel change
- Integration with Entitlement Servers to provide secure authentication
- Advanced Admission Control capabilities

When choosing an M-CMTS Core that can enable VDOC services, it is critical to carefully evaluate the product's DOCSIS, IP and Video features. If VDOC will expand the cable operator's service offering, it will be critical to continue to introduce innovation in the communications segment with services such as video telephony, video conferencing or integration with mobile services.

We will group these new services under the name of Next Generation Communications Services (NG-COMMs). Among the DOCSIS technologies that will enable NG-COMMs services, PacketCable Multimedia (PCMM) and PacketCable 2.0 are becoming increasingly important.

- Some of the functions being defined by PacketCable 2.0 include the capability for enhanced residential

VoIP services and video telephony, feature integration across service platforms, and mobility services.

- PCMM opens up the possibility for cable operators to launch very innovative services such as multi-party game playing and video-conferencing by extending the QoS mechanism beyond the VoIP application. Figure 6 presents the CableLabs PCMM architecture diagram.

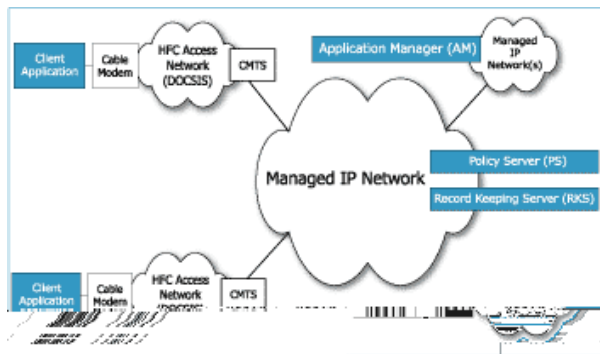


Figure 6: PacketCable Multimedia Diagram

A key aspect for the support of PCMM and PacketCable 2.0-based services is the ability of the M-CMTS Core to integrate well with the Application Manager, the Policy Server, and the Record Keeping Server. In the case where the M-CMTS Core is an evolution of an I-CMTS (Option A) that has already been through PCMM integration steps, this will translate into much greater service velocity than the case where such application servers need to be integrated with a new CMTS Core.

To summarize, some may argue that a newly developed M-CMTS Core platform (Option B) can support new services better than an I-CMTS since the I-CMTS was originally designed for today's services. However, before drawing conclusions, it is important to recognize the flexibility of an I-CMTS platform and its ability to evolve to meet future service requirements. In addition, the optimization of the underlying DOCSIS infrastructure and the integration of the

I-CMTS with advanced service architectures that has likely occurred during field deployment might prove a key asset to move faster towards the launch of such services.

Capital Budget Impact (Next 5 Years)

Perhaps the most pressing business driver for a cable operator to migrate to an M-CMTS architecture is the ability to add DOCSIS capacity and meet increasing bandwidth requirements at a fraction of today's cost. In this section, we consider how well the two migration options considered in this paper meet this.

It is important to realize that, with the separation of the downstream PHY into an external EQAM, the M-CMTS architecture has simultaneously achieved two key objectives. These objectives are:

- Separate the downstream PHY from the MAC functionality
- Separate the downstream PHY from the upstream PHY

The first objective is critical to leverage existing video EQAM technology for DOCSIS traffic, and therefore, reduce the cost of DOCSIS downstream ports. The second objective is as critical since it removes inefficiencies introduced by the fixed ratio between downstream and upstream ports in today's I-CMTS.

The impact on CapEx achieved through the first objective is easy to grasp by comparing the cost between a downstream port in a video EQAM and on a DOCSIS CMTS, as well as the expected decline.




Figure 7: CMTS DS Port vs. EQAM Port

To better understand the impact of the separation of downstream and upstream PHY functions, let's consider the following capacity sizing exercise: a cable operator wants to deploy 10 Mbps downstream and 2 Mbps upstream service in a network with 500,000 households passed. It projects to reach 10% take rate (50,000 subs). Assuming 200 kbps downstream and 40 kbps upstream are allocated to each subscriber, 10 Gbps in the downstream direction (50,000 * 200 kbps) and 2 Gbps in the upstream direction (50,000 * 40 kbps) will be required. This translates into 250 downstream and 200 upstream ports (assuming 40 Mbps per downstream and 10 Mbps per upstream port).

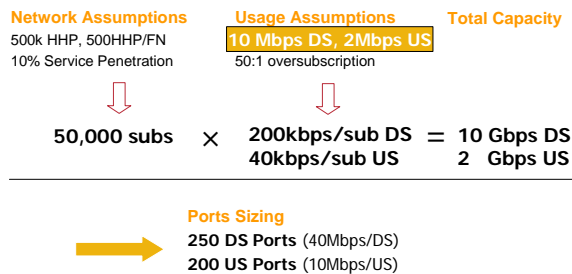


Figure 8: Example 1

Today downstream and upstream ports come packaged in one line card with 4 to 6 upstream ports per downstream. While the earlier traffic assumptions show the M-CMTS requirement to be approximately 1.25 downstream ports for every upstream port, this means that upstream capacity that is not required will be purchased, leaving the upstream ports under-utilized. For example, if the line cards come with 5 downstream and 20

upstream ports, 50 of such line cards will be required to meet the downstream capacity requirement, leaving the operator with 1000 upstream ports while the traffic requirement indicates that 200 are sufficient (20% upstream port utilization). If the line cards have 2 downstream and 12 upstream ports, this will mean that in order to purchase the 250 downstream ports required, 125 line card with 1200 upstream ports will be purchases (16.7% utilization). In reality, the HFC topology puts additional requirements beyond pure capacity in order to connect each fiber node to the CMTS. Although we've chosen to ignore this factor for simplicity, this doesn't impact the validity of the argument presented.

By breaking the fixed ratio between upstream and downstream ports, an M-CMTS architecture allows upstream and downstream capacity to be purchased independently so that CapEx can be optimized. In the earlier example, only ten 5x20 line cards are required to meet the upstream bandwidth requirements providing 50 downstream ports. The additional 200 downstream ports can be purchased independently by purchasing a MAC engine for the M-CMTS Core and a number of EQAMs for the downstream PHY interfaces.



Today's I-CMTS	M-CMTS Core + Edge QAM	
5x20 LC	5x20 LC +	
	24DS M-CMTS Edge QAM	
50 Cards	10 Cards	9 MAC Blades + 9 Edge QAM
250 DS	50 DS	200 DS
1000 US (20% utilization)	500 US (100% utilization)	-
7 Chassis	2 Chassis	

Figure 9: Example 2

As the example shows, significant savings can be achieved due to the separation of upstream and downstream ports in the

M-CMTS architecture. There is a key aspect of the above analysis that is important to stress. Due to the asymmetry of today's high-speed Internet service offering when compared with the available capacity on a DOCSIS line card, cable operators most likely have excess upstream capacity they've already purchased (upstream port under-utilization described above) for an I-CMTS deployment. With regard to this aspect, migrating an existing I-CMTS into a M-CMTS Core (Option A) clearly shows strong advantages from a CapEx standpoint. No additional upstream ports are required in addition to what is already available until those ports are fully utilized and the system is balanced with the addition of the required downstream capacity. For this reason, we expect Option A to provide the minimum impact to CapEx since it fully leverages the assets already purchased.

M-CMTS Core Conclusions

This section analyzes the key selection criteria a cable operator should consider when making a decision on the M-CMTS Core during a transition to an M-CMTS architecture. Table 2 summarizes important points.

Although a newly developed M-CMTS Core might look attractive when it comes to support for future DOCSIS services, this potential advantage must be balanced against the key downsides of the approach:

- Risk of impacting existing services if M-CMTS Core is unproven in DOCSIS networks
- Risk of delaying migration to M-CMTS if back-end integration effort is extensive
- Larger upfront CapEx negatively impacting cash flow

Choosing to migrate an existing I-CMTS into an M-CMTS Core might be the safest approach. The key is to choose an I-CMTS that can scale to support existing and new DOCSIS services, while keeping the changes to turn it into an M-CMTS Core at a minimum.

Criteria	Critical factors
Support for existing DOCSIS features and services	<ul style="list-style-type: none"> • Support all Layer 2 and Layer 3 features required for today's HSD and VoIP services • Integrate seamlessly in today's operations structure • Have the design stability of a field-proven implementation
Scalability for existing DOCSIS services	<ul style="list-style-type: none"> • Provide 5-10 Gbps forwarding capability • Support 50k-80k DOCSIS CPEs • Support Channel Bonding and IPv6
Support for new DOCSIS features and services (next 3-5 years)	<ul style="list-style-type: none"> • VDOC • NG Communications such as video telephony or video conferencing
Capital budget impact (next 5 years)	<ul style="list-style-type: none"> • Fully leverage existing capital expenditure in I-CMTS deployment

Table 2: Key Evaluation Criteria for M-CMTS Core

5.0 M-CMTS EQAM CHOICES

The EQAM is a key component of the M-CMTS architecture and is an important

factor in the success of an M-CMTS deployment strategy. This section contains an analysis of the key evaluation criteria to be considered in choosing which EQAM option best meets the operator's objectives for deploying an M-CMTS solution.

As explained earlier, cable operators have two fundamental options for deploying an EQAM as part of an M-CMTS solution. These two options are:

- Option A: Install a new standalone EQAM
- Option B: Repurpose the I-CMTS into an EQAM

The key criteria for assessing these EQAM options are:

- Cost optimization
- Implementation risk
- Vendor and product choice
- Density
- Scalability
- Flexibility
- Cost optimization

Cost is a primary consideration in assessing the options. The EQAM product must be designed for cost optimization—not just at initial deployment, but throughout the product life cycle. Product upgrades such as adding QAM channel capacity, supporting new features, implementing redundancy, and replacing components must be cost-effective.

EQAM costs are directly related to the power and signal quality of the outputs. Compared to the power and signal quality of today's video EQAMs, the DRFI specification raises the bar significantly. It will be a challenge for vendors to meet the DRFI requirements at the same price per QAM as video EQAMs. However, the leaders in the video EQAM market will be able to leverage a high-volume cost structure to optimize their

M-CMTS EQAM costs, and thus, will foster a competitive pricing environment.

By standardizing the interfaces and functionality of the EQAM, the M-CMTS specifications drive a common, well-defined feature set, which enables additional cost optimization in EQAM products. With a standalone EQAM, features like redundancy, scalability, density, manageability and reliability can all be targeted for the EQAM application.

Implementation Risk

This section highlights some of the potential risks of the EQAM options that might impact the decision on which M-CMTS migration strategy to choose. For example, an I-CMTS runs on a large, complex code base that supports features which are superfluous to an EQAM. Repurposing an I-CMTS into an EQAM is therefore a more difficult task, and thus riskier than implementing an application-specific standalone EQAM. Considering that an existing I-CMTS that is repurposed to perform as an M-CMTS EQAM is also likely to perform as an M-CMTS upstream receiver, the product development effort will be more complex (and thus riskier) than developing a standalone EQAM.

Implementing both the EQAM and upstream receiver functions on the I-CMTS also adds risk to achieving interoperability with the new M-CMTS Core. This risk is even greater if a pre-standard version of UEPI is implemented by the M-CMTS Core and EQAM.

Vendor and Product Choice

The M-CMTS specification opens the door to the DOCSIS downstream channel market for EQAM vendors and significantly increases the total addressable market for their EQAM products (possibly by as much as twofold or

higher). For this reason, it is expected that several EQAM vendors will take this opportunity to introduce standalone M-CMTS EQAM products.

Cable operators who choose Option A (standalone EQAM) will have the freedom to choose EQAM products from multiple vendors, including the EQAM market leaders which to date have not been able to offer products for the DOCSIS network. The freedom to choose from multiple vendors ultimately fosters competition and aligns with the M-CMTS objective of using cost-effective EQAMs.

Enabling multiple vendors to offer M-CMTS EQAM products also results in greater product choice for cable operators. A standalone EQAM can be implemented with either a stackable or chassis form factor, giving operators the freedom to choose the best form factor for their particular environment. A stackable EQAM is likely more suitable in small hub sites where the growth can be accommodated with a few EQAMs, whereas a chassis-based EQAM is likely more suitable in large headend and hub sites where high growth is anticipated. The form factor of the existing I-CMTS may not be suitable where the operator intends to deploy the M-CMTS EQAMs. Of course, operator preference dictates which form factor is most suitable. Having EQAM products available in both form factors is necessary to fulfill market requirements worldwide.

Density

The dwindling available rack space in a typical cable headend or hub site and the ongoing migration from analog to digital services is escalating the priority of EQAM density (in terms of QAM channels per rack unit) as a key evaluation criterion for EQAM products. Recalling that a primary objective of the M-CMTS architecture is to enable

significant growth in DOCSIS downstream channel capacity, it is imperative that M-CMTS EQAM products provide the density required to support the projected growth within the cable operator's limited rack space.

As an example of the potential growth of DOCSIS downstream channel capacity, consider the case of a cable operator with a DOCSIS 3.0 network who wants to offer a HSD service with 100 Mbps maximum download speed. Such a service would have to be delivered on a four-channel bonding group (using 64 QAM Annex B modulation). This essentially matches the number of QAM channels typically installed for VoD service today. As new high-bandwidth services such as VDOC are offered, the number of DOCSIS QAM channels required can quickly surpass the number of narrowcast video QAM channels installed.

The earlier example illustrates the point that M-CMTS EQAMs must at least match the density of video EQAMs. Video EQAMs typically support up to 24 QAM channels per rack unit today, and are expected to at least double in density within the next 2-3 years. Ideally, M-CMTS EQAMs will achieve much greater density to allow for the dramatic growth in downstream channel capacity envisioned by the M-CMTS and DOCSIS 3.0 initiatives.

Standalone M-CMTS EQAM products will likely be based on or modeled after video EQAMs, and thus can be expected to offer similar QAM channel density. Stackable EQAMs are likely to achieve greater QAM channel density than chassis-based EQAMs because chassis-based products typically incorporate high availability features that are not normally found in products with a stackable form factor.

When repurposing the I-CMTS into an M-CMTS EQAM, the number of QAM channels supported by the downstream PHY card is the critical factor in determining the overall QAM channel density of the I-CMTS chassis. The overall QAM density of the I-CMTS must be comparable to the density of standalone EQAM products. Otherwise, the cable operator will essentially be wasting rack space by using the I-CMTS as an M-CMTS EQAM. The overall QAM density of the I-CMTS is particularly critical if the new M-CMTS Core will be co-located with the I-CMTS, and thus require additional rack space.

Scalability

There are numerous factors to consider in assessing an M-CMTS EQAM's scalability, including:

- Output power level and signal quality requirements (dictated by DRFI)
- Number of RF ports (may be limited by connector type)
- Number of QAM channels per RF port (block upconversion ratio)
- Number of line card slots (in modular designs)
- Backplane capacity per line card slot
- Total backplane bandwidth
- Power supply capacity
- Air flow/cooling capacity
- Number and speed of wide area network (WAN) interface ports
- Processor performance (for data and control plane traffic)
- System memory

It is beyond the scope of this paper to address each of the scalability factors listed above. However, some general observations about the scalability of the different types of M-CMTS EQAM products being considered are provided.

Due to their small form factor, stackable EQAMs are not scalable. Today's stackable video EQAMs are typically able to support up to 24 QAM channels, with the expectation of doubling in scale in the next 2-3 years. When additional QAM channels are required, additional stackable units are deployed in the network.

Scalability is important for chassis-based EQAM products, since these products occupy more rack space and have longer life spans than stackable products. Since the useful life of a chassis-based product is typically at least 5-10 years, the product must be designed not only for today's services and technology, but also in anticipation of tomorrow's services and technology. For example, a chassis-based EQAM should be able to provide the total system power required to support full loading with current-generation QAM modules, as well as full loading with next-generation QAM modules which will support twice the number of QAM channels. All of the scalability factors listed above should be considered in selecting a chassis-based EQAM.

The best way to meet the scalability requirements of an M-CMTS solution in a chassis-based EQAM is to deploy a new purpose-built chassis that addresses the scalability factors for current and future DOCSIS services over the duration of the product life cycle. A chassis originally designed as an I-CMTS will likely have significant shortcomings in several of the scalability factors listed above.

Flexibility

The M-CMTS architects took great care in defining the M-CMTS EQAM so that it could be seamlessly integrated with a traditional video EQAM, and thus, enable both DOCSIS and digital video applications to share a common EQAM device. As a result, a new

product concept emerged from the M-CMTS architecture, the Universal EQAM. A Universal EQAM is able to support traditional digital video services, as well as DOCSIS services, and enables cable operators to make more efficient use of network resources.

Both of the M-CMTS EQAM options being considered in this paper can be implemented as a Universal EQAM by integrating the video EQAM functionality with the M-CMTS EQAM functionality. However, the choice of M-CMTS EQAM options does have implications on the flexibility achieved with a Universal EQAM.

There are three critical factors that should be considered in assessing the overall flexibility of a Universal EQAM:

- Range of Services Supported
- Resource Allocation Model
- Granularity of Resource Sharing

Range of Services Supported

Since the interfaces and functionality of the M-CMTS EQAM is standardized, all Universal EQAM products that are compliant with M-CMTS specifications will be able to support the full range of standard DOCSIS services. However, since the video EQAM is not standardized, each video service requires unique features on the EQAM. Therefore, the range of video services a Universal EQAM can support will depend on the product's feature set. For maximum flexibility, a Universal EQAM must support video on demand (VoD), digital simulcast, and switched digital video (SDV), as these services are critical to the cable operator's success.

Resource Allocation Model

The ability to dynamically allocate shared EQAM resources to services as needed

greatly enhances the flexibility of a Universal EQAM. Dynamic resource allocation enables the most efficient use of EQAM resources, especially when the EQAM resources are shared by a wide range of services. To support dynamic resource allocation, a Universal EQAM must interoperate with a resource management system. Without dynamic resource allocation, EQAM resources must be statically allocated to each service, potentially stranding resources and affecting service availability.

Granularity of Resource Sharing

There are multiple levels of granularity for sharing Universal EQAM resources among services. Services may be able to share a Universal EQAM per chassis, per RF port, per QAM channel, or even share the same QAM channel. The flexibility of a Universal EQAM is maximized when each QAM channel can be shared by multiple services.

Considering the above factors for maximizing the flexibility of a Universal EQAM, and the fact that existing video EQAM products can not be directly integrated with an I-CMTS EQAM, the best approach for achieving flexibility is to deploy a standalone Universal EQAM which integrates the M-CMTS EQAM functions into a video EQAM product capable of supporting VoD, digital simulcast and SDV.

M-CMTS EQAM Conclusions

This section analyzes the key evaluation criteria to be considered when assessing the options for deploying an EQAM as part of an M-CMTS solution. Table 3 summarizes the key criteria and the critical factors discussed in the analysis.

In summary, for each of the key criteria identified, deploying a standalone M-CMTS EQAM has significant benefits compared to

repurposing the I-CMTS into an EQAM. The cost and scalability advantages of a standalone EQAM is of critical importance, considering that these are primary objectives of the M-CMTS architecture. In addition, offering greater choice and flexibility with lower risk are important operational benefits of the standalone EQAM option.

CONCLUSION

The Modular CMTS (M-CMTS) series of specifications published by CableLabs introduces powerful new tools to enhance the architecture of the CMTS, the fundamental building block of all IP-based services over the cable network. The high-level objectives of the M-CMTS architecture are to improve the flexibility of the current DOCSIS network, enable independent scaling of downstream and upstream channel capacity, and improve the overall efficiency of the cable network with best-of-breed solutions. The solutions allow cable operators to prepare their DOCSIS infrastructure to enter the new broadband era.

This paper presents various approaches cable operators can follow to migrate from the current DOCSIS architecture to an M-CMTS architecture. The high-level tasks for executing the migration are described for each approach, followed by a discussion of the options for implementing an M-CMTS Core and EQAM as part of an M-CMTS solution.

A thoughtful analysis of the key criteria for assessing each M-CMTS migration strategy leads to the recommendation that cable operators should follow a two-phased approach that builds upon the current CMTS installed base to implement the M-CMTS Core function, and utilizes new standalone Edge QAMs in the first phase and external upstream receivers in the second phase. This strategy meets the objectives of the M-CMTS architecture, while providing cable operators greater choice and flexibility with lower risk.

Criteria	Critical factors
Cost optimization	<ul style="list-style-type: none"> Designed to optimize cost High-volume cost structure Standard EQAM feature set
Implementation risk	<ul style="list-style-type: none"> I-CMTS code base Design complexity UEPI standardization Interoperability
Vendor and product choice)	<ul style="list-style-type: none"> Vendor choice with standalone EQAM Standalone EQAM can be stackable or chassis-based
Density	<ul style="list-style-type: none"> High priority feature Match video EQAM density
Scalability	<ul style="list-style-type: none"> Critical for chassis-based EQAMs Must scale with services and technology Requires purpose-built EQAM
Flexibility	<ul style="list-style-type: none"> Universal EQAM All video and DOCSIS services Dynamic resource allocation Granularity of resource sharing

Table 3: Key Evaluation Criteria for M-CMTS EQAM