Data-Over-Cable Service Interface Specifications
Converged Cable Access Platform

Converged Cable Access Platform Architecture Technical Report

CM-TR-CCAP-V03-120511

ISSUED

Document Status Sheet

Document Control Number:	CM-TR-CCAP-V03-120511			
Document Title:	Converged Cable Access Platform Architecture Technical Report			
Revision History:	V01 - Released 12/22/10 as CM-TR-CMAP-V01-101222 V02 - Released 6/14/11 V03 - Released 5/11/12			
Date:	May 11, 2012			
Status:	~~Work in Progress~~	~~Draft~~	**Released**	~~Closed~~
Distribution Restrictions:	~~Author Only~~	~~CL/Member~~	~~CL/ Member/ Vendor~~	**Public**

Trademarks

Contents

Figures

Tables

This page intentionally left blank

1 SCOPE

1.1 Introduction and Purpose

This Architectural Overview Technical Report is intended to provide an introduction to the Converged Cable Access Platform (CCAP) architecture. This document describes the architecture and discusses the various CableLabs specifications that contain normative requirements pertaining to a CCAP. In addition, this document describes the architectural entities and interfaces that make up the implementation, as well as the protocols they support.

2 INFORMATIVE REFERENCES

This technical report uses the following informative references. References are either specific (identified by date of publication, edition number, version number, etc.) or non-specific. For a non-specific reference, the latest version applies.

[CCAP OSSI] Data-Over-Cable Service Interface Specifications CCAP Operations Support System Interface Specification, CM-SP-CCAP-OSSI, Cable Television Laboratories, Inc.

[DPoE ARCH] DOCSIS Provisioning of EPON Architecture, DPoE-SP-ARCHv1.0, Cable Television Laboratories, Inc.

[DPoE MEF] DOCSIS Provisioning of EPON MEF Specification, DPoE-SP-MEFv1.0, Cable Television Laboratories, Inc.

[DPoE MULPI] DOCSIS Provisioning of EPON MULPI Specification, DPoE-SP-MULPIv1.0, Cable Television Laboratories, Inc.

[DPoE OAM] DOCSIS Provisioning of EPON OAM Specification, DPoE-SP-OAMv1.0, Cable Television Laboratories, Inc.

[DPoE OSSI] DOCSIS Provisioning of EPON OSSI Specification, DPoE-SP-OSSIv1.0, Cable Television Laboratories, Inc.

[DPoE PHY] DOCSIS Provisioning of EPON Physical Layer Specification, DPoE-SP-PHYv1.0, Cable Television Laboratories, Inc.

[DPoE SEC] DOCSIS Provisioning of EPON Security Specification, DPoE-SP-SECv1.0, Cable Television Laboratories, Inc.

[DRFI] DOCSIS Downstream RF Interface Specification, CM-SP-DRFI, Cable Television Laboratories, Inc.

[DSG] DOCSIS Set-top Gateway (DSG) Interface Specification, CM-SP-DSG, Cable Television Laboratories, Inc.

[DTI] DOCSIS Timing Interface Specification, CM-SP-DTI, Cable Television Laboratories, Inc.

[EQAM VSI] Edge QAM Video Stream Interface Specification, CM-SP-EQAM-VSI, Cable Television Laboratories, Inc.

[ERMI] DOCSIS Edge Resource Manager Interface, CM-SP-ERMI, Cable Television Laboratories, Inc.

[IEEE 802.3ae] IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gbs Operation.

[IEEE 802.1AX] IEEE 802.1AX-2008 IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation

[IEEE 802.3ba] IEEE Standard for Information Technology-Specific Requirements - Part 3: 40Gb/s and 100Gb/s Ethernet, IEEE Std 802.3ba-2010.

[ITU-T G.8031/Y.1342] ITU-T Recommendations G.8031/Y.1342 SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS: Packet over Transport aspects – Ethernet over Transport aspects; SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS: Internet protocol aspects – Transport Ethernet linear protection switching, 2009-11-13

[L2VPN] Business Services over DOCSIS: Layer 2 Virtual Private Networks, CM-SP-L2VPN, Cable Television Laboratories, Inc.

[MHA TR] EQAM Architectural Overview Technical Report, CM-TR-MHA, Cable Television Laboratories, Inc.

[MULPI] MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0, Cable Television Laboratories, Inc.

[OSSI] Operations Support System Interface Specification, CM-SP-OSSIv3.0, Cable Television Laboratories, Inc.

[PHY] Data-Over-Cable Service Interface Specifications Physical Layer Specification, CM-SP-PHYv3.0, Cable Television Laboratories, Inc.

[RFC 2328] IETF RFC2328/STD0054 OSPF Version 2, April 1998.

[RFC 2453] IETF RFC2453/STD0056 RIP Version 2, November 1998.

[RFC 3209] IETF RFC3209, RSVP-TE: Extensions to RSVP for LSP Tunnels, December 2001.

[RFC 3376] IETF RFC3376, Internet Group Management Protocol, Version 3, October 2002.

[RFC 3810] IETF RFC3810, Multicast Listener Discovery Version 2 (MLDv2) for IPv6, June 2004.

[RFC 4023] IETF RFC4023, Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE), March 2005.

[RFC 4293] IETF RFC4293, Management Information Base for the Internet Protocol (IP), April 2006.

[RFC 4364] IETF RFC4364, BGP/MPLS IP Virtual Private Networks (VPNs), February 2006.

[RFC 4448] IETF RFC4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks, April 2006.

[RFC 4601] IETF RFC4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised), August 2006.

[RFC 4664] IETF RFC4664, Framework for Layer 2 Virtual Private Networks (L2VPNs), September 2006.

[RFC 4724] IETF RFC4724, Graceful Restart Mechanism for BGP, January 2007.

[RFC 4760] IETF RFC4760, Multiprotocol Extensions for BGP-4, January 2007.

[RFC 4761] IETF RFC4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling, January 2007.

[RFC 4874] IETF RFC4874, Exclude Routes - Extension to Resource Reservation Protocol-Traffic Engineering (RSVP-TE), April 2007.

[RFC 5036] IETF RFC5036, LDP Specification, October 2007.

[RFC 5120] IETF RFC5120, M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs), February 2008.

[RFC 5303] IETF RFC5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies, October 2008.

[RFC 5420] IETF RFC5420, Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol-Traffic Engineering (RSVP-TE), February 2009.

[RFC 5709] IETF RFC5709, OSPFv2 HMAC-SHA Cryptographic Authentication, October 2009.

[RFC 5711] IETF RFC5711, Node Behavior upon Originating and Receiving Resource Reservation Protocol (RSVP) Path Error Messages, January 2010.

[RFC 6020] IETF RFC6020, M. Bjorklund, Ed., YANG - A data modeling language for the Network Configuration Protocol (NETCONF), October 2010.

[RFC 6221] IETF RFC6221, D. Miles, Ed., Lightweight DHCPv6 Relay Agent, May 2011

[RFC 6241] IETF RFC6241, R. Enns, et al., Ed., NETCONF Configuration Protocol, June 2011.

[SEC] DOCSIS 3.0 Security Specification, CM-SP-SECv3.0, Cable Television Laboratories, Inc.

[TEI] TDM Emulation Interface Specification, CM-SP-TEI, Cable Television Laboratories, Inc.

2.1 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027;
  Phone +1-303-661-9100; Fax +1-303-661-9199; http://www.cablelabs.com
- The Institute of Electrical and Electronics Engineers, Inc., IEEE, Internet: standards.ieee.org
- Internet Engineering Task Force (IETF) Secretariat, 48377 Fremont Blvd., Suite 117, Fremont, California 94538, USA. Phone: +1-510-492-4080, Fax: +1-510-492-4001. http://www.ietf.org.

3 TERMS AND DEFINITIONS

This document uses the following terms:

Cable Modem Termination System	A headend component that provides the operator network side termination for the DOCSIS link. A CMTS communicates with a number of Cable Modems to provide data services.
Converged Cable Access Platform	A headend component that provides the functionality of a CMTS and an Edge QAM in a single architecture with greater QAM density and overall capacity.
Edge QAM	A head-end or hub device that receives packets of digital video or data from the operator network. It re-packetizes the video or data into an MPEG transport stream and digitally modulates the transport stream onto a downstream RF carrier using QAM.
Ethernet Passive Optical Network	A point-to-multipoint, fiber to the premises network architecture in which unpowered optical splitters are used to enable a single optical fiber to serve multiple premises.
Hybrid Fiber-Coax System	A broadband bidirectional shared-media transmission system using optical fiber trunks between the head-end and the fiber nodes, and coaxial cable distribution from the fiber nodes to the customer locations.
NETCONF	An IETF network management protocol that provides mechanisms to manipulate the configuration of a device. NETCONF executes YANG-based XML files containing configuration objects.
RF Combiner	Headend equipment that accepts multiple input signals and delivers a single output that is equal in phase and amplitude.
Service Group	A set of channels for a given service (e.g., Video On Demand, High-Speed Internet) delivered via a number of fiber nodes to corresponding subscribers of that service to a single subscriber device.
Switch Route Engine (SRE)	Management line card where NSIs are located and switching and routing take place.
YANG	A language used to model data for the NETCONF protocol. A YANG module defines a hierarchy of data which can be used for NETCONF-based operations, including configuration, state data, remote procedure calls (RPCs), and notifications.

4 ABBREVIATIONS AND ACRONYMS

This document uses the following abbreviations:

AES	Advanced Encryption Standard
ANCP	Access Node Control Protocol
APC	Angled Physical Contact
ARP	Address Resolution Protocol
ASM	Any-Source Multicast
AWGN	Additive White Gaussian Noise
BGP	Border Gateway Protocol
BSoD	Business Services over DOCSIS
CATV	Cable Television
CBR	Constant Bit Rate
CCAP	Converged Cable Access Platform
CLI	Command-Line Interface
CM	Cable Modem
CMTS	Cable Modem Termination System
CPD	Control Point Discovery
CPE	Customer Premises Equipment
CSA	Common Scrambling Algorithm
DBG	Downstream Bonding Group
DCS	Downstream Channel Set
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DLC	Downstream Line Card
DPIC	Downstream Physical Interface Card
DPoE	DOCSIS Provisioning of EPON (Ethernet Passive Optical Network)
DRFI	Downstream RF Interface
DS	Downstream
DSG	DOCSIS Set-top Gateway
DTI	DOCSIS Timing Interface
ECM	Encryption Control Message
ECMD	ECM Decoder
ECMG	ECM Generator
ECMP	Equal-Cost Multi-Path
EOAM	Ethernet Operations, Administration and Maintenance
EoC	Ethernet over Coax
EPON	Ethernet Passive Optical Network
EQAM	Edge QAM
ER	Edge Router
ERM	Edge Resource Manager

ERMI	Edge Resource Manager Interface
FFT	Fast Fourier Transform
Gbps	Gigabits per second
GigE	Gigabit Ethernet
GMQ	General Membership Query
GRE	Generic Routing Encapsulation
HSI	High-Speed Internet
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IPDR	IP Detail Record
IPTV	IP Television
IS-IS	Intermediate System To Intermediate System Protocol
L2	Layer 2
L2VPN	Layer 2 Virtual Private Network
L3	Layer 3
LACP	Link Aggregation Control Protocol
LAG	Link Aggregation Group
LAN	Local Area Network
LDP	Label Distribution Protocol
LDRA	Lightweight DHCP Relay Agent
LSP	Label-Switched Path
MAC	Media Access Control
M-CMTS	Modular CMTS
MCX	Micro Coaxial
MEF	Metro Ethernet Forum
MHz	Megahertz
MIB	Management Information Base
MLD	Multicast Listener Discovery
MPEG	Moving Picture Experts Group
MPLS	Multiprotocol Label Switching
MPTS	Multi-Program Transport Stream
MSO	Multi-System Operator
MTA	Multimedia Terminal Adapter
MVPN	Multicast Virtual Private Network
ND	Neighbor Discovery
NNI	Network to Network Interface
NSI	Network-Side Interface
OAM	Operation, Administration, and Management
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band

OSPF	Open Shortest Path First Protocol
OSS	Operations Support System
OSSI	Operations Support System Interface
OTT	Over-the-Top Voice over IP
OUI	Organization Unique Identifier
P2MP	Point-to-Multipoint Communication
PCMM	PacketCable Multimedia
PCR	Program Clock Reference
PE-CE	Provider-Edge – Customer-Edge
PEG	Public, Education, and Government channels
PHY	Physical Layer
PIC	Physical Interface Card
PIM-DM	Protocol Independent Multicast – Dense Mode
PIM-SM	Protocol Independent Multicast – Sparse Mode
PON	Passive Optical Network
PPM	Parts per Million
PSTN	Public Switched Telephone Network
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RA	Router Advertisement
RF	Radio Frequency
RFoG	Radio Frequency over Glass
RIP	Routing Information Protocol
RS	Router Solicitation
RSVP	Resource Reservation Protocol
RSVP-TE	RSVP – Traffic Engineering
SC	Subscriber Connector
SCTE	Society of Cable Telecommunications Engineers
SDV	Switched Digital Video
SIP	Session Initiation Protocol
SNMP	Simple Network Management Protocol
SPTS	Single Program Transport Stream
SRE	Switch Route Engine
SRM	Session Resource Manager
SSM	Source-Specific Multicast
STB	Set-Top Box
TEI	TDM Emulation Interface
TFTP	Trivial File Transfer Protocol
TLS	Transparent LAN Service
UCH	Universal Cable Holder

UML	Unified Modeling Language
VBR	Variable Bit Rate
VLAN	Virtual Local Area Network
VOD	Video on Demand
VoIP	Voice over IP
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VSI	Video Stream Interface
XML	Extensible Markup Language

5 CCAP ARCHITECTURE GOALS, BENEFITS AND OVERVIEW

5.1 Fundamental Goals of the CCAP

The Converged Cable Access Platform (CCAP) is intended to provide a new equipment architecture option for manufacturers to achieve the Edge QAM and CMTS densities that MSOs require in order to address the costs and environmental challenges resulting from the success of narrowcast services. The CCAP leverages existing technologies, including DOCSIS 3.0, Modular Headend Architecture, and current HFC architectures; and also can include newer ones, such as Ethernet optics and EPON (Ethernet Passive Optical Network).

The CCAP provides an alternative approach to the implementation of converged video and data services described in the Modular Headend Architecture (MHA) Technical Report (i.e., Modular CMTS with Universal Edge QAM). Similar to MHA, the CCAP provides sharing of QAM channels for different narrowcast services, but adds the capability of sharing broadcast QAM channels.

The key functional goals for CCAP include:

- Flexible use of QAM channels for the various services offered by MSOs, enabling modification to the number of QAMs using MPEG transport stream-based services (e.g., for VOD, SDV, etc.) versus DOCSIS-based services (e.g., HSI, voice, video over IP, etc.) over time though a single configuration point.
- Individually configurable assignment of QAM channels to various service groups, such that it would be possible to have HSI/voice service groups, VOD service groups, and/or SDV service groups overlap in different ways without requiring that these service groups be identical.
- Efficient implementation of separate sets of QAM channels for narrowcast and broadcast applications, such that QAM channels for narrowcast services can be individually implemented for each RF port, and QAM channels used for broadcast services can be shared among the RF ports in each downstream line card (DLC).
- Simplification of the RF combiner network by providing all QAM channels for all digital services from a single RF port, only leaving certain legacy functions for RF combining.
- Option to add content scrambling – both standardized and proprietary (e.g., PowerKEY™, DigiCipher®, etc.) – without requiring special-purpose hardware, such that a CCAP from any vendor can optionally implement the appropriate scrambling mechanisms without increasing the complexity of the platform.
- A transport-agnostic network architecture allowing implementation of EPON and other access network technologies natively within the CCAP. The CCAP will be expected to support additional access technologies and higher capacity uplink interfaces in the future with pluggable or otherwise replaceable components, allowing upgrade to a new access technology via installation or replacement of access modules.
- Modularization of the software environment, allowing upgrades to be applied to specific services without impacting other services. This partitioning also helps to ensure that software issues in the implementation of a given service do not necessarily impact other partitioned services.
- Significant operational improvements, including environmental efficiencies (e.g., reduced space, power consumption, and heat dissipation), implementation of functions such as upstream health monitoring, continuous wave carriers for plant amplifier biasing, and many other operational enhancements.

5.2 CCAP Benefits

5.2.1 Service Multiplexing Flexibilities

The CCAP provides efficient implementation of Edge QAM (EQAM) blocks by implementing separate sets of QAM channels for narrowcast and broadcast applications. QAM channels for narrowcast services are individually implemented for each RF port, while QAM channels used for broadcast services are shared among all the RF ports in each downstream line card (DLC). The number of narrowcast and broadcast QAMs supported on each RF port is flexible.

The CCAP provides the ability to map narrowcast and broadcast QAMs in different combinations to specific downstream ports via configuration. The same narrowcast video QAM can be mapped to multiple downstream ports, allowing for overlap of SDV and VOD service groups. This allows an operator to create service groups on a decoupled, service-by-service basis and effectively deal with service group inequalities.

The CCAP can configure any QAM in a given CCAP RF port for DOCSIS or Edge QAM applications. This allows all QAMs for a given service group to be generated from a single RF port. By allowing the configuration of any QAM on any RF port for broadcast, SDV, VOD, or DOCSIS, the CCAP provides the ability to transition to next generation video services (via DOCSIS or a new access technology) as necessary.

5.2.2 Bandwidth Capacity and Density Gains

The CCAP is designed to greatly increase the capacity of a single edge device, delivering all narrowcast and broadcast services via the downstream RF ports deployed (15-20 downstream ports on a small chassis; 40-60 on a larger chassis). The CCAP is expected to support multiple 10, 40, and/or 100 GigE interfaces with the ability to support a downstream capacity of over 150 Gbps. On a typical downstream line card in a large CCAP chassis, this traffic will be utilized by up to 12 downstream RF ports. QAM channels are flexible across service types. Each downstream port is capable of supporting up to 158 QAM channels. For example, a port could have 64 narrowcast QAMs and 96 broadcast QAMs. Each upstream port is capable of supporting a minimum of 4 DOCSIS RF upstream channels, with 6 channels being a recommended implementation. Support for higher data rates is possible with implementation of newer access technologies, such as EPON via DPoE. The ratio of downstream to upstream ports is variable, with an expected ratio of 2 upstream ports for every downstream port.

5.2.3 High Reliability and Redundancy Capabilities

Given the scope of each RF port providing all services for a given service group, it is important that the operation be highly reliable. Therefore, a CCAP is expected to support redundancy for critical components; however, even in a non-redundant configuration, the CCAP should provide sufficient up time. This, coupled with management of service group size, allows the size of failure groups to be reduced.

The CCAP is designed with a "wire once" approach: physical interface cards (PICs) implement the upstream and downstream physical interfaces, allowing replacement of line cards without impact to the cabling. N+1 redundancy allows line card replacement without impacting services for longer than the failover time and without the need to re-wire upstream and downstream connections. This reduces mean time to recovery for the CCAP.

The CCAP is designed such that software upgrades can be performed against a specific functional module, allowing an upgrade to a specific service that does not impact other services on the CCAP. Previous versions of the software images are available in local storage, allowing simplified regression to the last software image in the case of recovering from a failed software upgrade.

The high availability of the CCAP network-side interface (NSI) includes two technical components: network path redundancy and connections with multiple links. The CCAP incorporates multiple NSI ports, allowing path redundancy in connectivity to the MSO's regional network. While many redundant connectivity topologies are feasible, this document focuses on two examples. In the first example, the NSI ports of a CCAP are interconnected with multiple links to two Edge Routers (ERs), as shown on Figure 5–1. In the second example, the CCAP NSI ports are attached with multiple links to a single ER, as shown in Figure 5–2. Such a scenario may be possible if the ER platform provides robust redundancy features, resulting in sufficiently high availability.

In either case the redundant network paths are comprised of independent components, thus avoiding the possibility of a single point of failure.

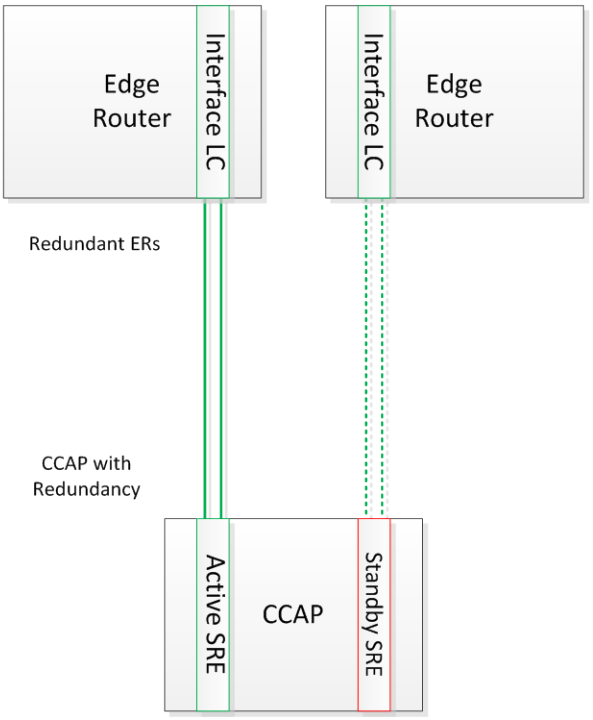

Figure 5–1 - CCAP and Two Redundant Ers [1]

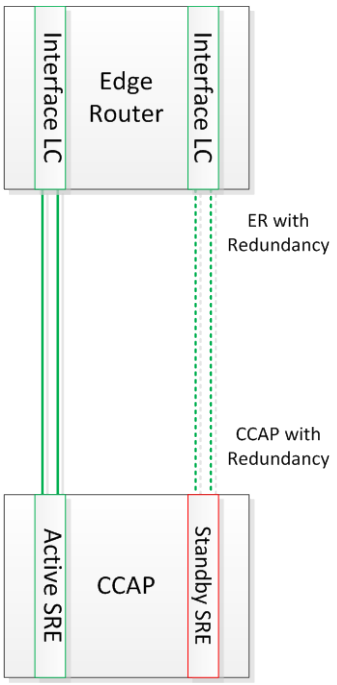

Figure 5–2 - CCAP and a Single Redundant ER [1]

[1] Both Figure 5–1 and Figure 5–2 depict connections between the CCAP and ERs terminating on the standby Switch Route Engine (SRE) line card. The diagrams are drawn under the assumption that the CCAP is capable of operating NSI ports even when the SRE is in the standby mode.

In both examples, the redundant connections between the CCAP and ER(s) may be active to allow the CCAP and the ER(s) to participate in neighbor protocol information exchange. In addition, the subscriber traffic may be distributed (load balanced) between the redundant paths, or the network elements may be configured to give a preference for one of the redundant connections. The dotted links shown in Figure 5–1 and Figure 5–2 are meant to represent connections that may carry subscriber traffic, or provide backup capacity, depending on the operator's preferences.

Each of the connections between the CCAP and the ER may comprise multiple links. Bundling of multiple links may be necessary to scale the aggregate connection bandwidth when the capacity of single link is insufficient. When deployed with excess capacity, multilink bundles can provide additional protection against failures of an individual link. Two standards-based bundling protocols may be utilized, Link Aggregation Control Protocol (LACP), defined in [IEEE 802.1AX], or Equal-Cost Multi-Path routing (ECMP).

Link aggregation, also defined in [IEEE 802.1AX], is used as a means to provide high availability. In a case that one (or more) of several links in the link aggregation group (LAG) fails due to equipment or connectivity failure, the remaining links provide connectivity at a lower throughput. Operators may use other mechanisms on top of LAG to provide redundancy for high availability, for example Ethernet path redundancy.

Another use of link aggregation on the CCAP NSI can be as a means to linearly increase NSI data throughput capacity. The capacity is increased by assigning more links of the same capacity to the same "aggregated" interface.

LACP is used to provide automatic LAG configuration. Links can be automatically assigned to a LAG if they are discovered to interface between the same two devices and have the same capacity. Nevertheless, static LAG configuration can be used without the use of LACP.

The LACP link failure detection mechanism will work even with a LAG interface belonging to a standby SRE. This ensures that the status of links is known even on a Standby LAG, so that failover can happen in a more predictable and reliable way.

Ethernet Operations, Administration, and Management (OAM), according to the ITU-T G.8031 over ITU-T Y.1731 standard ([ITU-T G.8031/Y.1342]), can be used for path failure detection and failover operation. While LACP detects failures on a link-by-link basis, it cannot propagate through intermediate switches. Ethernet OAM operates on a "bundled interface" level and can propagate through intermediate switches. Therefore, Ethernet OAM provides protection on a higher level: the path level. Interfaces between the ER and the CCAP in this scenario can be single links or LAGs and they can pass through separate intermediate switches.

5.2.4 Configuration and Management Simplifications

The CCAP will allow configuration of both CMTS and EQAM functions from the same configuration interface. CCAP configuration will move away from SNMP-based configuration and focus instead on the processing of XML configuration files that hold the configuration details for all services on the CCAP. The CCAP utilizes a common object model, which standardizes the configuration objects across vendor implementations, and therefore simplifies the management of device configuration. Local storage and versioning of configuration files aids rapid recovery of services when a primary component has failed and been replaced by an offline component. The CCAP is also expected to support traditional command-line interface (CLI) methods of configuration, as well as next-generation configuration protocols such as NETCONF [RFC 6241].

Management of CMTS and EQAM functions are also consolidated on the CCAP. The CCAP implements all MIBs required in [CCAP OSSI], which includes appropriate CableLabs, SCTE and IETF MIBs.

5.2.5 Rack-Space and Power Reduction

One of the key benefits of the CCAP is to achieve significant environmental efficiencies. To that end, Figure 5–3 and Figure 5–4 demonstrate an example of the space and power savings achieved by deployment of the CCAP in a typical system.

Figure 5–3 depicts a typical installation in a headend consisting of the various digital services, including broadcast, SDV, VOD, and HSI equipment, plus the corresponding combiner and lasers/receivers.

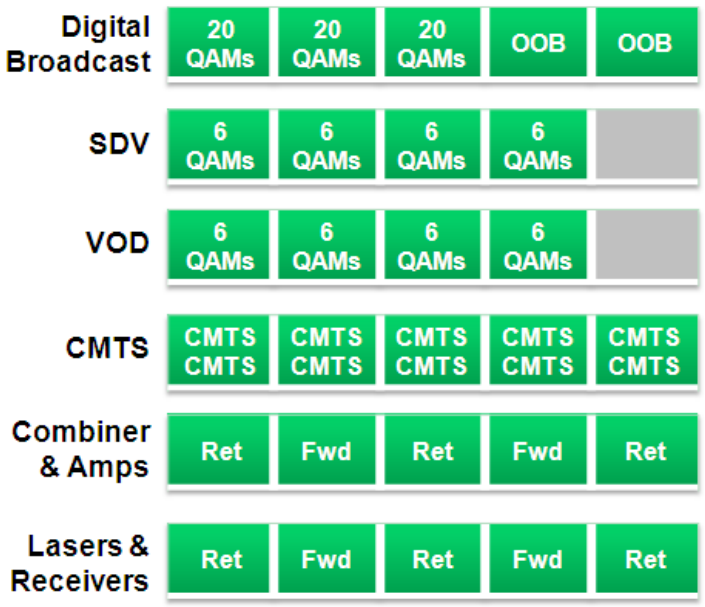

Figure 5–3 - Typical Headend Space Usage

The example shown in Figure 5–3 is intended to serve a typical population, combined in such a way as to result in 160 HSI service groups, and 120 VOD and matching SDV service groups.

Considering typical CMTS and Edge QAM equipment available today, this service group configuration would require about 10 CMTS chassis and about 4 racks for VOD and SDV, each containing 6 Edge QAM chassis configured for 64 QAM channels, each at a density of 4 QAM channels per RF port. The digital broadcast lineup is composed of 60 individual QAM channels, plus the corresponding out-of-band equipment.

Figure 5–4 depicts the analogous installation when considering the deployment of equivalent CCAP equipment in a medium-sized chassis.

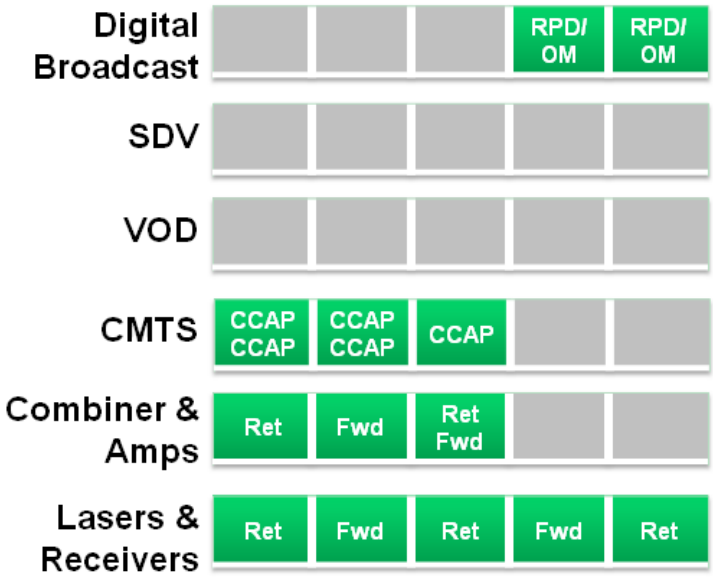

Figure 5–4 - CCAP Deployment Space Usage

Figure 5–4 shows the following:

- Given that the CCAP chassis would have twice the density of a typical CMTS, only half the number of CCAP chassis are required (compared to CMTS chassis), resulting in equivalent space savings.
- Additionally, the CCAP chassis, in its basic implementation, includes all the necessary QAM channels for supporting the VOD and SDV services. Therefore, no additional equipment is needed to support these functions, resulting in significant additional space savings.
- Given that the CCAP also supports sufficient broadcast QAM channels, the space previously allocated to the broadcast equipment is no longer needed, further contributing to space savings.
- Finally, it is estimated that half of the space allocated to the combiner network would be saved, resulting in even further space savings.

With all this taken into account, as much as half of the space previously required is needed for deploying the CCAP. Moreover, given that the CCAP can serve twice as many narrowcast QAM channels as the previous architecture could, the depicted CCAP scenario actually results in even greater space savings, providing twice as much capacity in half the space.

In addition, a cursory analysis of the difference in power consumption, assuming typical power draw for existing equipment and the expected power consumption for the CCAP, yields an estimated power savings of greater than 50%. And, this is taking into account the use of 32 QAM channels in the CCAP, or 2 times the capacity indicated in the original typical deployment.

With the decrease in equipment and power necessary to support this density of QAMs, the amount of heat generated by the equipment is also reduced, resulting in cooling savings in the headend.

5.2.6 RF Combining Simplifications

Deployment of the CCAP simplifies the RF combiner network by providing all QAM channels for all digital services from a single RF port, only leaving certain legacy functions for RF combining. Rather than having to rewire the physical plant to make service group changes, the QAM content of a downstream RF port can be changed via the CCAP configuration interface. Downstream of the CCAP, legacy out-of-band, analog channels, and maintenance streams (balance, sweep) are the only things that need to be combined into the CCAP output.

5.2.7 IP Router Integration

A CCAP system will support IP routing and forwarding, including IP multicast proxies and forwarding, and IP service proxies from a CMTS, a DPoE System, and the hub router and switches that currently provide IP multicast functionality for video broadcast and SDV.

The CCAP operates as a single system for the purposes of IP forwarding; however, the QAM resources of the CCAP are also exposed to external VOD or SDV servers in order to allocate the QAM resource to MPTS video streams. As a result, the CCAP can operate as both a core when edge QAMs are controlled internally or, when implemented in a modular fashion, operate as an Edge QAM resource for the purposes of integration with VOD and SDV servers.

In the future MSOs will require MPEG-TS transport capability with a migration path to MPEG over IP (IP/UDP or IP/TCP). Such services will presumably be operated with IP over DOCSIS (and/or IP over EPON with DPoE). The CCAP will be required to support a wide range of IP applications during the expected lifetime of such a platform.

5.3 Supported Services in CCAP

This section discusses the various services supported by the CCAP on an HFC system. Some of these services can also be supported by a CCAP for an EPON deployment.

5.3.1 Video EQAM Services

Video services supported by the CCAP include digital video services that are supported today by existing broadcast and narrowcast Edge QAMs (EQAMs). These services include digital video delivered as: 1) broadcast digital video; 2) switched digital video; and 3) video on demand. The CCAP is not intended to support modulation of analog

video. Analog video signals are combined external to the CCAP with the output digital QAM signals of the CCAP prior to the input into the downstream fiber optic transmitter.

5.3.1.1 Broadcast Digital Video

Broadcast digital video services refer to the programming delivered in a channel lineup to subscribers in common, as opposed to just a particular subscriber or to a particular node. Sometimes, broadcast video services are referred to as linear broadcast services because of the time-linear nature of the broadcast, whereby the programming operates on a regular schedule that is not under the control of the viewers themselves (e.g., no native ability to pause, rewind, or fast-forward the program).

Broadcast digital video services are typically delivered to the hub site in "pre-packaged" multiplexes (a "broadcast lineup") that require minimal processing beyond local encryption, modulation, and upconversion.

The broadcast digital video services typically contain retransmission of over-the-air broadcast channels as well as programming supplied from various programmers delivered by satellite to a cable headend. Certain broadcast digital video services may also require local digital program insertion for advertisements targeted for particular ad zones, but the CCAP is not required to perform or support local ad insertion. Some broadcast lineups may require basic MPEG-2 "add/drop" type multiplexing to accommodate local public, education, and government (PEG) channels. For PEG channels, the CCAP is required to multiplex one or more SPTSs into an existing MPTS (Multi-Program Transport Stream) and perform MPEG table management (e.g., Program Association Table, Program Map Table).

Although digital music services delivered over a cable television (CATV) system are principally audio services, they are also included here as a broadcast digital video service, since the programming usually includes accompanying still pictures or video in the background along with descriptions of the musical track and artist.

5.3.1.2 Switched Digital Video

Switched digital video (SDV) is classified as a narrowcast video service, as opposed to a broadcast service. Although the SDV content is delivered in common to multiple subscribers, the target subscribers are only those of a particular SDV service group that corresponds to the SDV QAM channels delivered to one or more fiber nodes. SDV transmits channel programs to only those service groups where there is a subscriber viewing the channel. Bi-directional data transmission is a key element to the support of SDV since STBs that are SDV-capable will indicate that they are tuned to, or are tuning to, a particular channel.

With an SDV service, the system delivers to viewers programs that operate on a regular schedule and are not under the control of viewers. Thus, in that sense, SDV is a linear video service.

Broadcast video services and SDV services can coexist on the channel lineup transmitted on a given node. Programming considerations, analysis of viewership of particular programs, and available downstream frequency channels feed into decisions about which particular programs to include on downstream QAMs delivering SDV.

5.3.1.3 Video on Demand

A video on demand (VOD) service is defined as a video service delivered to a specific subscriber, generally in response to a real-time request for pre-packaged MPEG content. Thus, VOD is a narrowcast video service. VOD requires interaction between the subscriber STB and the program control system across a bi-directional CATV system. In addition, VOD typically means that the user has the ability to make a video selection and have control over its playback (e.g., pause, rewind, fast forward). In this sense, VOD is not a linear video service.

5.3.2 DOCSIS Services

Services enabled by the CCAP and other equipment supporting the interfaces specified in the various DOCSIS specifications include:

- High-speed Internet (HSI)
- PacketCable voice over IP (VoIP)
- Transparent LAN Service (TLS) over DOCSIS L2VPN
- Next generation video services delivered via DOCSIS

These services are discussed in the following sections.

5.3.2.1 High-Speed Internet

Interfaces for cable modems (CMs) and cable modem termination systems (CMTSs) have been defined in the DOCSIS specifications. Multiple generations of DOCSIS specifications exist in support of high-speed Internet services. The DOCSIS 1.0 specifications provide basic broadband Internet connectivity for one or more devices in the home. Among other things, they include the ability to rate-limit (cap) a particular customer's data rate to a cable operator selected value. The DOCSIS 1.1 specifications provide improved operational flexibility, security, and quality-of-service (QoS) features that support high-quality digital voice, interactive gaming, and commercial service level agreements (SLAs). The DOCSIS 2.0 specifications include increased upstream reliability and throughput for symmetric services. The DOCSIS 3.0 specifications provide a number of enhancements, most notably channel bonding, support for IPv6, and support for next generation video services. Channel bonding provides cable operators with a flexible way to significantly increase speeds to customers, with compliant devices supporting up to at least 160 Mbps in the downstream and at least 120 Mbps in the upstream.

The CCAP platform fully incorporates the functionality of DOCSIS 3.0 CMTSs – which includes backward compatibility for all previous DOCSIS generations – but with a much higher density.

5.3.2.2 PacketCable VoIP

The VoIP service discussed in this section refers to VoIP as provided by equipment supporting the interfaces specified in various PacketCable specifications. In this context, PacketCable VoIP provides a voice service that has voice quality, call features, and reliability that is expected from a primary line telephony service. Therefore, so-called over-the-top (OTT) VoIP services are not included in this category and can be considered as one of many applications that can be delivered under an HSI service. In addition to containing all the functionality of PacketCable 1.0, PacketCable 1.5 extends the PacketCable residential voice capability with capabilities such as fax and modem support, analog trunking for PBXs, and Session Initiation Protocol (SIP) for session management within and among PacketCable networks.

At a very high level, the PacketCable 1.5 architecture contains three networks: the DOCSIS HFC Access Network, the Managed IP Network, and the PSTN. The CMTS provides connectivity between the DOCSIS HFC Access Network and the Managed IP Network. Both the Signaling Gateway and the Media Gateway provide connectivity between the Managed IP Network and the PSTN. In the CCAP architecture, the CCAP performs the CMTS functions that are involved in support of PacketCable functionality today.

5.3.2.3 Transparent LAN Service

The Transparent LAN Service (TLS) discussed in this section refers to the functionality defined in the DOCSIS L2VPN specification. TLS is also sometimes referred to as a "Metro Ethernet" or "Carrier Ethernet" service, such as that specified by the Metro Ethernet Forum (MEF). TLS allows businesses to extend their Layer 2 Ethernet networks (LANs) across the connectivity cloud of a core network that could be by itself an Ethernet, IP, or MPLS (Multiprotocol Label Switching) network.

The DOCSIS L2VPN specification defines functionality required for the CM and CMTS to support TLS. It defines the CM interface providing TLS service on the subscriber side, and the CMTS interfaces on the core network connection side as well as the DOCSIS RF interface. The required functions reuse a lot of existing DOCSIS functionality and provide some necessary extensions. The DOCSIS L2VPN specification defines support for both point-to-point Ethernet connections (E-LINE type of service in MEF terminology) and multipoint-to-multipoint Ethernet connections (E-LAN type of service in MEF). It also provides support for non-multiplexed services (EPL – Ethernet Private Line) as well as services multiplexed on one CM port (EVPL – Ethernet Virtual Private Line). It has a robust support for QoS enforcement /guarantees and security/encryption for this type of business services.

In the CCAP architecture, the CCAP performs the CMTS functions as defined by the L2VPN specification. In particular, it performs necessary per-L2VPN flow packet encapsulation on the NSI, QoS enforcement, downstream encryption and other functions.

5.3.2.4 Next Generation Video Services

The CCAP is expected to support the managed delivery of IP video for next generation video services.

Managed IP video delivery has certain characteristics that can make the handling of that traffic by the CCAP less resource intensive. For example, the traffic flows have the following fundamental characteristics:

- **Large Packets:** IP video traffic consists predominantly of large packets on the order of 1300 to 1500 bytes. In order to sustain a given data rate, the CCAP needs to forward fewer packets per second than it otherwise would if the traffic flow consisted of small packets.
- **Long Sessions:** video sessions are generally long-lasting in comparison to internet sessions. Long duration sessions have the effect of minimizing session and flow setup and teardown overhead in the CCAP. Typical session duration is expected to be greater than 10 minutes.
- **Lower number of sessions:** video sessions are generally high bitrate sessions in comparison to internet sessions. High bitrate sessions require a lower number of sessions per QAM. The typical number of video sessions per QAM is expected to be lower than 40.

These characteristics of video traffic flows might allow a CCAP to be more cost-effectively sized with respect to the performance capabilities of a system than a system that is sized for full DOCSIS traffic of various sized packets on all data QAMs.

5.3.3 DOCSIS Provisioning of EPON

The CCAP is expected to support EPON interfaces for high speed Internet. To facilitate the interoperability of EPON with the CCAP, the CCAP will support the entire DOCSIS Provisioning of EPON (DPoE) suite of specifications. The DPoE system includes DOCSIS Provisioning functions and IP routing (forwarding and service) functions that are common with DOCSIS (RF) functions.

Services deployed over RF or EPON will vary over time. While it may be appropriate, for example, to support a specific service on one platform at one time, a change in bandwidth demands or other service delivery requirements could easily justify the move of a service from one platform to another. The CCAP needs to be flexible enough to support both DOCSIS (RF) and EPON-based services. The CCAP will not just support a fixed allocation of services to RF or EPON interfaces, but will instead provide a platform that could be configured as 100% DOCSIS (RF), 100% EPON, or any mix of access technologies, where that mix is only limited by the access technology installed, the platforms capability to support EPON slots, and the overall throughput capacity supported.

High-speed Internet, PacketCable, TLS, and the next generation video services described in Section 5.3.2 can all be delivered via EPON.

5.4 CCAP Architectures

Two reference architectures are provided in this section: one showing the digital video delivery infrastructure, and the other showing the high-speed Internet infrastructure. The CCAP was designed to fully support both types of services simultaneously. A discussion of using the CCAP in an M-CMTS architecture is also discussed.

5.4.1 CCAP MPEG Video Headend Reference Architecture

Cable headends acquire video from various sources to be provided to the subscriber via the access network. In Figure 5–5, the dotted lines represent the video data while the remainder of the diagram represents control elements or flows within the MPEG video system.

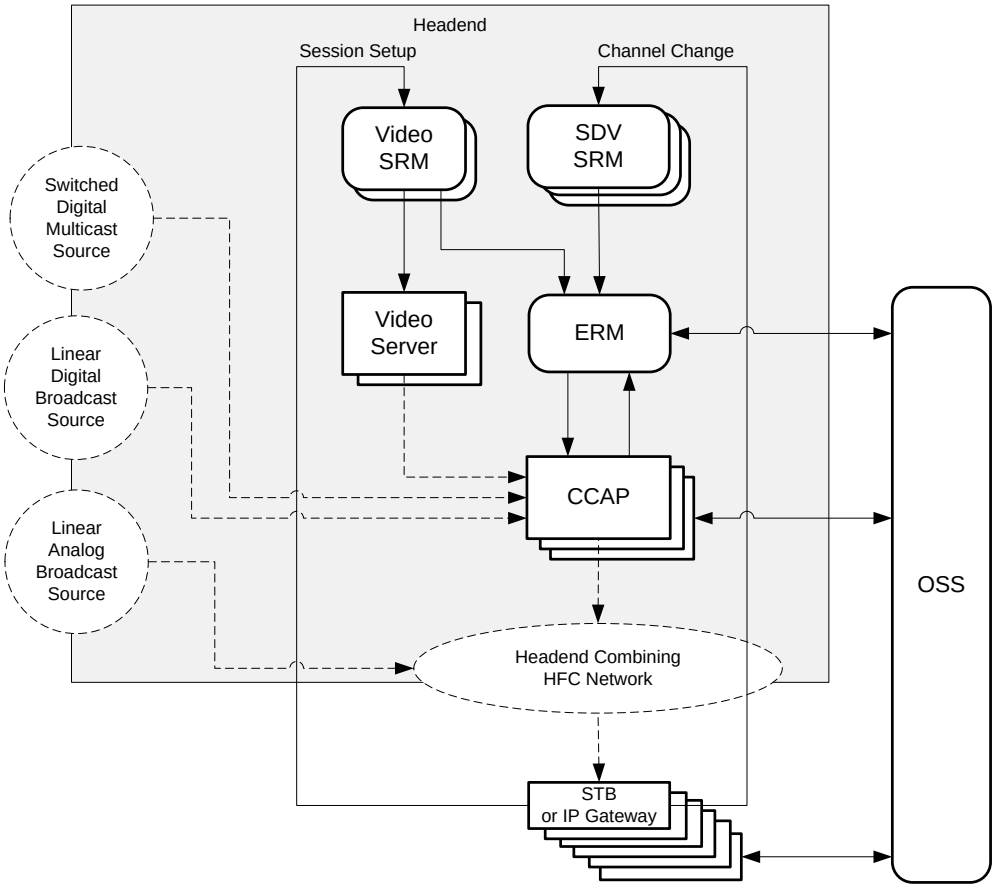

Figure 5–5 - CCAP Video Headend Reference Architecture

The CCAP has one or more ingress interfaces and multiple RF QAM outputs. The CCAP accepts input MPEG SPTSs or MPTSs transported via UDP/IP (multicast or unicast) over Ethernet, and multiplexes these input programs into an output MPTS that is then modulated and transmitted out one of the QAM RF outputs.

Digital video that is not broadcast continuously to service groups is controlled by the interaction of a service-specific client application on the STB, signaling to service-specific session resource managers (SRM) to request receipt of a video stream. When the STB client requests a stream, the SRM must acquire the necessary resources that allow the stream to be transported from source to destination. When the CCAP is first deployed, it is expected that VOD streams will be delivered as unicast, constant bit rate (CBR) SPTSs. The CCAP will be directed to route them to a particular QAM port and RF frequency, either via a UDP port mapping scheme, where the UDP port defines the RF port and frequency, or under the control of the Edge Resource Manager (ERM). In either case, the CCAP will be required to multiplex incoming SPTSs into MPTSs, and perform QAM modulation, up-conversion, and in some cases content scrambling. Variable bit rate (VBR) SPTSs may be supported in future releases.

Typical SDV architectures involve a centralized pre-processing of linear video streams into CBR SPTSs, a process known as "clamping". The SPTSs are then routed to a network encryption device that applies appropriate content security and multicasts the streams onto the network. An SDV SRM will have knowledge of the SPTS multicast groups that correspond to each service and will communicate this to the ERM. When a STB client sends a request to the network to view a particular service, the CCAP that is able to reach that STB client will be instructed by the ERM to join the multicast of that service and route it to a particular QAM. Therefore, the requirement on the CCAP is to respond to the ERM request and multiplex the encrypted SPTS into a transport stream with additional SPTSs and perform QAM modulation and upconversion.

Linear digital broadcast video is routed to the CCAP in "pre-packaged" multiplexes (a "broadcast lineup") that require minimal processing beyond content scrambling, modulation, and upconversion. Each downstream line card

on the CCAP will support at least a single broadcast lineup, sharing it among the service groups the line card serves. The CCAP can have multiple broadcast lineups per line card which are distributed across the service groups that the line card serves.

The ERM functional component is used to manage the use of transport bandwidth on the HFC network – the QAM channel and associated downstream frequency – out of the CCAP. The SRM uses the ERM to find a CCAP RF output having sufficient bandwidth and connectivity to the STB service group (serving area). In order to acquire the necessary RF/QAM bandwidth and CCAP resources needed to transport the stream to the service group, the SRM requests an ERM component function to allocate the bandwidth to the session manager. The ERM component provisions the CCAP to prepare it to receive the stream and direct it to the appropriate RF output using the allocated MPEG program number. The ERM can be either a controller interface outside of the CCAP or can be implemented internally to the CCAP.

The STB receives the QAM channel by tuning to the proper frequency, and can decode a single MPEG program from the MPTS. The STB is also responsible for providing the decoded A/V stream to the subscriber output device (i.e., monitor or TV) for presentation.

5.4.2 CCAP Data Reference Architecture

The CCAP performs all DOCSIS functions in the way that a traditional CMTS platform does. In addition, PON can be deployed on the CCAP to manage commercial HSI traffic. The following diagram, Figure 5–6, illustrates how HSI streams flow through the CCAP and the network to DOCSIS and PON devices and back.

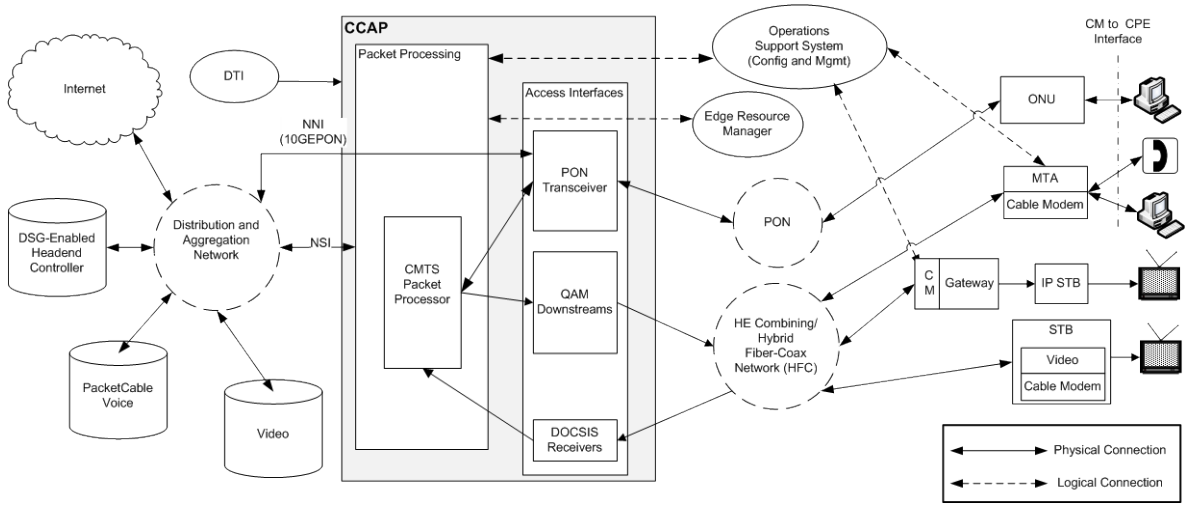

Figure 5–6 - CCAP Data Reference Architecture

The CCAP receives Internet content, video, PacketCable voice data, and [DSG] data through the NSI from a distribution and aggregation network, via one or more physical interfaces. Individual NNI (Network to Network Interface) may be provided for HSI business services using a 10G EPON interface.

The CCAP has many of the functions of a traditional provider edge (PE) router. Since the CCAP supports multiple protocols, it also cannot be purely IP. In particular, the growth of Ethernet services (provided directly to customers) and use of Ethernet for managing logical IP networks (IP-VPNs), whether in-house or for customers, necessitates the usage of pseudo-wire transports over MPLS.

The CCAP performs all of the MAC-layer functionality and all the initialization and operational DOCSIS-related processes. The MAC-layer functionality includes all signaling functions, downstream bandwidth scheduling, and DOCSIS framing. The CCAP creates the DOCSIS QAMs for the service groups the CCAP serves, and these QAMs are modulated and output as either downstream QAMs on the downstream line cards in the CCAP, or QPSK on the downstream PON interfaces.

The CCAP will deliver data and voice per the DOCSIS 3.0 and PacketCable 1.5 or PacketCable Multimedia (PCMM) specifications.

The ERM functions in a similar way to the MPEG video session setup: to acquire the necessary RF/QAM bandwidth and the CCAP resources to transport the stream to the service group.

The outputs of the CCAP are combined with legacy OOB data (and possibly legacy analog video) in the headend combining network and are distributed through the HFC network or PON network, where they are received by the following devices:

- MTA/CM: Provides PacketCable voice services and DOCSIS HSI for telephony and personal computing applications.
- ONU: Terminates the PON traffic for business services applications.
- Gateway: Receives video via MPEG Transport and/or IP protocol for distribution within the home.
- Set-top Boxes: native QAM video and command and control data (either through DOCSIS or legacy out-of-band).

Upstream DOCSIS HSI and PacketCable voice traffic travel through the HFC or EPON and are received by either QAM upstream receivers or PON transceivers.

A DTI Server may be used to provide a common clock reference to synchronize to other TDM clock domains.

5.4.3 Modular Headend Architecture Functionality

To assist with the transition from traditional EQAM and CMTS devices to the converged implementation of the CCAP, a CCAP can also function as a Universal EQAM. This allows the CCAP to conform to the architecture defined in [MHA TR], allowing existing CMTS infrastructure to be leveraged while transitioning in CCAP devices. This approach will be essential during the transition from current CMTS/EQAM deployments to the deployment of the CCAP.

It is anticipated that devices implemented this way would be capable of operating in either Universal EQAM or CCAP mode and changing mode could require a software change.

5.4.4 Forwarding Mode Options

Non-Routing (NR) Mode is being described to explicitly state that the CCAP Reference Architecture is supportive of devices that operate in a routing mode, a non-routing mode, or both. In the case of a CCAP device operating in a non-routing mode, all the routing functions and routing protocols are performed in the ER, north of the CCAP device.

The motivating factors for NR Mode include:

1. Simplifying the development path for CCAP suppliers
2. Leveraging edge router routing capabilities, which provide an extended set of routing features compared to CCAP devices

The CCAP Reference Architecture diagrams shown in Figure 5–5 and Figure 5–6 are identical for both routing and non-routing modes. All physical and logical elements shown in the diagrams, as well as all the stated services (e.g., HSD, voice, VOD, SDV, etc.) are fully supported in both modes.

The high-level behavior change that defines the difference between routing and non-routing mode is the type of forwarding being performed in the CCAP device. Furthermore, all CCAP devices, regardless of forwarding mode, support DOCSIS and PacketCable specifications. The behavior description of non-routing mode is described in detail in section 7.8.

The forwarding mode is indicated in the IP-MIB ([RFC 4293]) by the ipForwarding and ipv6IpForwarding objects of ipv4GeneralGroup and ipv6GeneralGroup2 respectively, defined in [OSSI]. Both the ipForwarding and ipv6IpForwarding objects will report notForwarding(2) when the CCAP is operating in non-routing mode, and forwarding(1) when it is operating in routing mode. Furthermore, the operation mode of a CCAP supporting both routing and non-routing modes may be configured by these objects if the CCAP supports SNMP write operations.

6 SUMMARY OF DOCSIS SPECIFICATIONS AND APPLICABILITY

6.1 DOCSIS 3.0 Specifications

The following sections describe which DOCSIS 3.0 specifications are required for the CCAP, and the extent to which they apply.

6.1.1 MAC and Upper Layer Protocols Interface ([MULPI]) Specification v3.0

The [MULPI] specification defines the MAC layer protocols of DOCSIS 3.0 as well as the requirements for upper layer protocols (e.g., IP, DHCP, etc.). The CCAP is required to meet all CMTS requirements specified therein.

6.1.2 Physical Layer ([PHY]) Specification v3.0

The [PHY] specification defines the upstream physical layer requirements for hybrid fiber-coax systems that the CCAP must support. The CCAP is required to meet all of the requirements specified therein. In addition, the CCAP is also designed to be compatible with the European market. For a European CCAP, adherence to Annex B of [PHY] is required.

6.1.3 DOCSIS Security ([SEC]) Specification v3.0

The [SEC] specification defines security services for DOCSIS communications, providing the operator with the ability to secure the provisioning process of cable modems (CM) and protect cable modem users by encrypting traffic flows between the CM and the cable modem termination system; in this case, the CCAP. The CCAP is required to meet all CMTS requirements specified therein.

6.2 CCAP Specifications

6.2.1 CCAP Operations Support System Interface ([CCAP OSSI]) Specification

The [CCAP OSSI] specification defines new configuration interfaces based on a standardized, converged object model, supporting both EQAM and CMTS functions. This specification introduces and standardizes YANG-based configuration and NETCONF as a protocol to support the configuration and management of the CCAP. This specification augments the configuration and management requirements specified in [OSSI].

6.2.1.1 SNMP Requirements and Reporting Requirements

The SNMP requirements of the CCAP are based upon the requirements specified in [OSSI], but the CCAP does not implement all SNMP requirements specified therein. The CCAP is required to support SNMP v1 and v2, as well as at least 10 SNMP community strings with controlled access via access lists. The CCAP primarily diverges from the [OSSI] specification in its configuration methods. The CCAP is not required to implement SNMP as a configuration protocol; instead, the CCAP is configured through the processing of XML-based configuration files, the structure of which is defined in YANG instance modules.

The CCAP must support all standard event reporting mechanisms defined in section 8 of the [OSSI] specification. The CCAP is required to meet all IP Detail Record (IPDR) requirements specified in the [OSSI] specification. The CCAP is also required to support all IPDR service definitions defined in the [OSSI] specification.

6.2.1.2 CCAP Object Model

The [CCAP OSSI] specification implements object models for configuration, fault management, and performance management. These object models build upon models established in [OSSI].

The specification supports the use of YANG, allowing access to these object models via NETCONF [RFC 6020].

The CCAP implements a configuration object model based on existing DOCSIS 3.0 and DPoE configuration objects, extending them as needed to meet the advanced features of this platform.

All read-only MIB objects in the [OSSI] specification are implemented for fault and performance monitoring, but the read-write and read-create MIB objects are not mandatory for the CCAP. The non-mandatory status of SNMP for configuration allows these write/create MIB objects to be excluded.

6.3 Modular Headend Architecture Specifications

6.3.1 Edge Resource Management Interface ([ERMI]) Specification

The [ERMI] specification defines interfaces that are used by EQAMs, ERMs, and M-CMTS Cores. While the CCAP does not require that all of [ERMI] be supported, because the CCAP will interface with ERMs to dynamically control video and possibly DOCSIS QAMs, the following interfaces specified in [ERMI] are required:

- Registration interface to the ERM (ERMI-1)
- Control interface to the ERM (ERMI-2)

In addition to these interfaces, the CCAP is also required to implement switched digital video (SDV) as defined in the [ERMI] specification.

6.3.2 DOCSIS Timing Interface ([DTI]) Specification

The [DTI] specification defines the timing interfaces required for the DOCSIS M-CMTS architecture. While the CCAP must support stratum 3 clock accuracies, it is not required to implement the timing interface as specified in the [DTI] specification, although [DTI] is an acceptable implementation. If a DOCSIS timing interface is not implemented, then the external timing interface is required to support a method to lock itself to the internal DOCSIS clock to ensure traceability of the clock to the TDM hierarchy.

6.3.3 Video Stream Interface ([EQAM VSI]) Specification

The [EQAM VSI] specification defines the data plane requirements for receiving, processing, and transmitting MPEG transport streams in EQAMs. The CCAP implements all requirements in the [EQAM VSI] specification, with the following exceptions:

- Section 9 – Encryption and Encryption Interface: The CCAP may be implemented with its own content protection mechanisms, described in further detail in Section 7.2, Optional Content Protection of this report. In this case Section 9 is not required. The optional CCAP Scrambler is required to support payload input and payload output, as defined in [EQAM VSI].
- Section 12 – Input and Output monitoring: The CCAP implements robust MPEG transport stream monitoring.

6.4 Downstream RF Interface ([DRFI]) Specification

The [DRFI] specification defines the downstream radio frequency interface for EQAMs and CMTS; as such, the requirements specified in [DRFI] are required for the CCAP. While all [DRFI] specification requirements must be met by the CCAP, the CCAP does diverge from the [DRFI] specification in the following areas:

- Frequency accuracy: The CCAP requires a frequency accuracy of equal to or better than 5 ppm, 10 year aging over time and temperature.
- Port-to-port isolation: The CCAP requires a minimum port-to-port isolation of ≥70 dB from 50 MHz to 550MHz and ≥65 dB from 550 MHz to 1002 MHz.
- Frequency shift: The CCAP requires:
  - Carriers must not sweep across bands.
  - Ports must be muted when changing frequency.
  - Port output cannot be restored until the RF output is at the correct frequency and is stable.
- Output level adjustment: The [DRFI] specification contains general language regarding how the system should behave while changing output levels, which is required for the CCAP.

6.5 DOCSIS Set-Top Gateway ([DSG]) Specification

The [DSG] specification defines an interface and associated protocol that introduces additional requirements on a DOCSIS CMTS and DOCSIS CMs to support the configuration and transport of a class of service known as "Out-Of-Band (OOB) messaging" between a set-top controller (or application servers) and the customer premises equipment (CPE). The CCAP is required to meet all CMTS requirements in the [DSG] specification.

6.6 Business Services over DOCSIS Specifications

6.6.1 Layer 2 VPN ([L2VPN]) Specification

The [L2VPN] specification describes requirements for both CMTSs and CMs in order to implement a DOCSIS Layer-2 Virtual Private Network. The L2VPN feature allows cable operators to offer a Layer 2 Transparent LAN Service (TLS) to commercial enterprises. The CCAP is required to meet all CMTS L2VPN requirements and implement all relevant read-only L2VPN MIB objects. The CCAP is not required to implement L2VPN read-write and read-create MIB objects.

6.6.2 TDM Emulation Interface ([TEI]) Specification

The [TEI] specification defines a method for cable operators to deliver T1, E1 and NxDS0 emulation services that meet or exceed the quality requirement of applications that use such services. Implementation of TDM emulation is preferred for the CCAP, but no [TEI] specification requirements are mandatory.

6.7 DOCSIS Provisioning of EPON Specifications

A CCAP supporting EPON applications will be expected to implement the entire suite of DPoE specifications defined in the following sections.

6.7.1 DOCSIS Provisioning of EPON Architecture ([DPoE ARCH])Specification

The [DPoE ARCH] specification describes the architecture required for DPoE Networks.

6.7.2 DOCSIS Provisioning of EPON MEF ([DPoE MEF]) Specification

The [DPoE MEF] specification describes the provisioning and operations required to support Metro Ethernet Forum (MEF) Ethernet Services in DPoE Networks, which use EPON as defined in 802.3ah and 802.3av.

6.7.3 DOCSIS Provisioning of EPON MAC and Upper Layer Protocols ([DPoE MULPI] Specification

The [DPoE MULPI] specification defines the MAC and upper layer protocols for DPoE Networks. The MAC in DPoE Networks is EPON.

6.7.4 DOCSIS Provisioning of EPON Operations Administration and Maintenance ([DPoE OAM]) Specification

The [DPoE OAM] specification defines the interface used for conveying management information between a DPoE System and DPoE ONU.

6.7.5 DOCSIS Provisioning of EPON Operations Support System Interface ([DPoE OSSI]) Specification

The [DPoE OSSI] specification identifies requirements for the adaptation or additions to DOCSIS specifications that are required to support DPoE Networks related to the Operations Support System functional area.

6.7.6 DOCSIS Provisioning of EPON Physical Layer ([DPoE PHY]) Specification

The [DPoE PHY] specification identifies requirements for the EPON PHY for the adaptation or additions to DOCSIS specifications that are required to support DOCSIS Provisioning of EPON.

6.7.7 DOCSIS Provisioning of EPON Security ([DPoE SEC]) Specification

The [DPoE SEC] specification identifies recommendations for the adaptation or additions to DOCSIS specifications that are required to support DOCSIS Provisioning of EPON (DPoE).

6.8 Summary of DOCSIS Specification Applicability

The following table summarizes the level of adherence the CCAP must have to the DOCSIS specification.

Table 6–1 - DOCSIS Specification Adherence

	Specification											
Device	**MULPI**	**PHY**	**SEC**	**CCAP OSSI**	**ERMI**	**DTI**	**VSI**	**DRFI**	**DSG**	**L2VPN**	**TEI**	**DPoE**
CCAP	M	M	M	M	P	O	P	M	M	M	O	M*
ERM	NA	NA	NA	NA	M	NA	NA	NA	NA	NA	NA	NA

* Required if EPON is implemented on the device

M = Mandatory
P = Partially Required
O = Optional
NA = Not Applicable

7 CCAP FEATURES AND CAPABILITIES

7.1 Service Multiplexing Capabilities

7.1.1 CCAP Service Groups

For the purposes of the CCAP, a service group is defined as a set of channels in a given service delivered via some number of fiber nodes to the corresponding subscribers of that service, provided by one or more CCAP ports. One of the concepts tied in with service groups is the reachability of certain signals from the CCAP to fiber nodes to the subscribers on those nodes, and likewise, the reachability of return path signals from fiber nodes to the CCAP. Thus, a service group is defined as a set of channels to or from a set of subscribers.

A downstream service group has traditionally been defined as a set of downstream channels carrying a particular service that reaches a specific set of optical nodes. The sizes of the service group for each service on the network, in terms of fiber nodes or actual subscribers, may be independent of each other. For the discussion below, it is assumed that the output of one port on the downstream line card (DLC) is associated to one fiber node. In addition, the number of channels provided for a particular service on each node differs for various service types available to subscribers.

The notion of a service group is important to understand some of the key functions of the CCAP. The following sections provide an overview of the different types of service groups and how they relate to each other.

7.1.1.1 Broadcast Service Groups

A broadcast service group consists of linear digital broadcast video channels corresponding to an advertising zone or a channel lineup with local or regional PEG channels. A broadcast service group typically spans more than one port on a DLC, possibly spans all ports on a DLC, and may even cross multiple DLCs.

7.1.1.2 Switched Digital Video Narrowcast Service Groups

An SDV service group consists of a number of downstream QAM carriers that are configured for video services for switched digital video applications. The SDV service group is a configured element of the CCAP and may span more than one RF port on a DLC.

7.1.1.3 VOD Narrowcast Service Groups

A VOD service group consists of a number of downstream QAM carriers that are configured for video services for video on demand. A VOD service group may be configured to be smaller than a DOCSIS service group, due to frequency re-use configurations and the number of QAMs needed for a specific serving area. The VOD service group may span more than one RF port on a DLC.

7.1.1.4 DOCSIS Narrowcast Service Groups

A DOCSIS downstream service group today is typically configured to be comprised of one to four optical nodes, depending on node size, service penetration, and data traffic load. There is some expectation (not necessarily a rule) that a DOCSIS service group would correspond to one port on a DLC.

7.2 Optional Content Protection

The CCAP provides an option to accept incoming transport streams that have had network encryption applied to keep them protected as they traverse the network of the cable operator, remove that network encryption, and then, based on the encryption mode specified for the content, apply the appropriate conditional access (CA) encryption for downstream transmission. The CCAP Decryptor and CCAP Scrambler are the two core optional functions of the CCAP, daisy-chained as shown in Figure 7–1, to dynamically process a large portion of the input payload, according to the requirements of each content protection session.

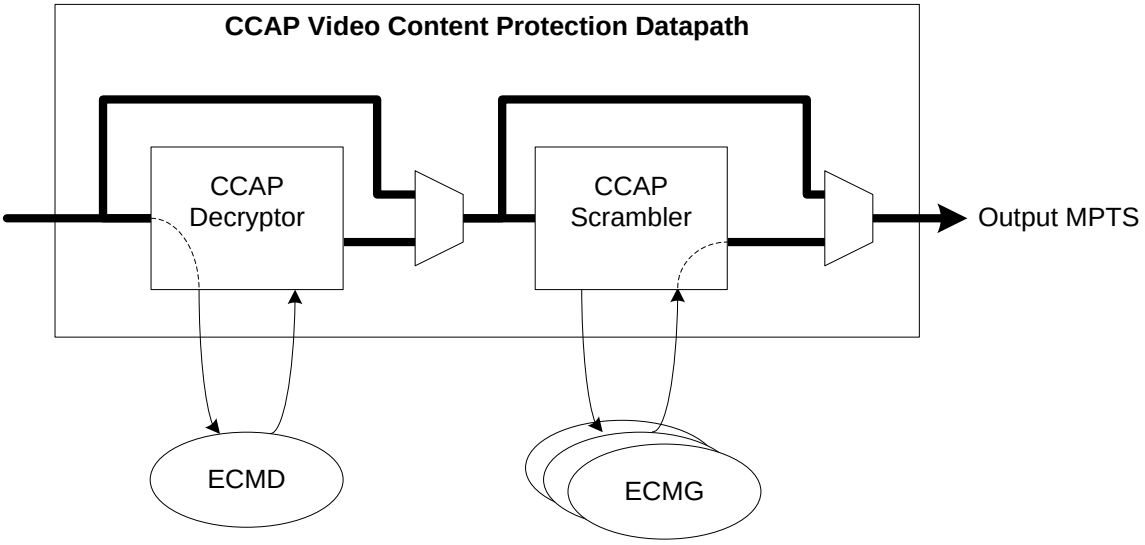

Figure 7–1 - Optional CCAP Video Content Protection Overview

The CCAP content protection data path is designed to support all of the decryption and content protection needs for the following existing and future video cable services:

- Video on demand
- Switched digital video
- Linear Digital Broadcast

The CCAP Decryptor and CCAP Scrambler functionally, when selected, reside on the downstream line card.

7.2.1 Network Decryption

The optional CCAP Decryptor component works in association with one external ECM decoder (ECMD) to remove the encryption layer used to secure the content distribution within the operator inner network. The control words and copy control information are retrieved from the incoming ECM to allow decrypting the associated payload, if required. The CCAP Decryptor supports 128-bit AES decryption.

7.2.2 Access Encryption

The CCAP Scrambler is a simple scrambling engine; the conditional access intelligence resides in the ECM Generator (ECMG). The CCAP Scrambler is under the control of one or more external ECMGs and applies the encryption layer required to secure the content distribution to the subscriber CPE devices, based on the ECMG-provided access criteria. In Simulcrypt operation, the Scrambler provides the same control word and access criteria to all ECMGs, which in return generate their matching ECMs. In non-Simulcrypt operation, the Scrambler receives both control words and ECMs from the selected ECMG, based on the provided access criteria.

The CCAP Scrambler supports both DigiCipher and PowerKEY conditional access systems, with support for the following encryption algorithms:

- DES (Data Encryption Standard)
- CSA (Common Scrambling Algorithm)
- 128-bit AES

7.3 QAM Replication

In order to simplify integration of the CCAP into existing systems, the CCAP is expected to implement a QAM replication feature. The purpose of this feature is to allow an operator to create logical service groups on a decoupled, service-by-service basis. This will provide the ability to replicate narrowcast video (SDV and VOD) QAMs across multiple ports on a given line card (not necessarily across all of the line cards in the chassis).

To effectively deal with service group inequalities, the CCAP will share a given set of SDV or VOD QAMs to other ports on the line card to form service groups with unique sets of HSI QAMs.

This is illustrated in the Figure 7–2; note that each DLC port has a unique HSI group of QAMs (not depicted).

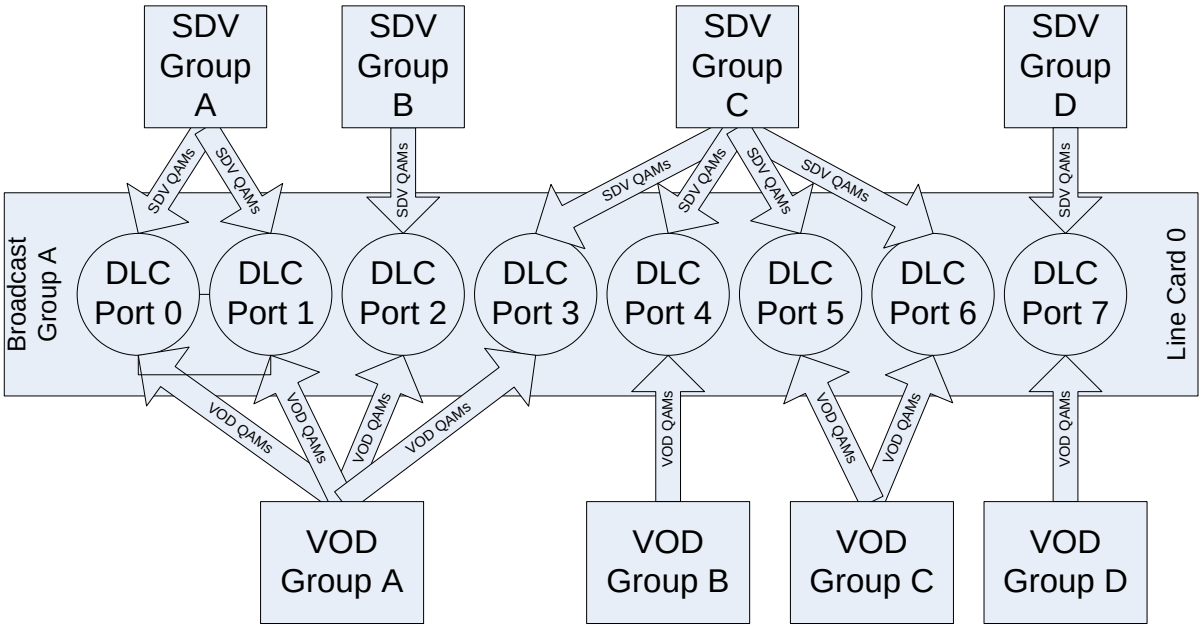

Figure 7–2 - QAM Replication

Note that the SDV and VOD QAMs are not necessarily coupled, but can be combined in various SDV/VOD pairs. In this illustration, each service group has a unique set of HSI QAMs.

The CCAP will be capable of replicating the contents of all narrowcast QAMs configured for native MPEG transport stream video services (e.g., VOD or SDV), and those configured for DOCSIS services, to a minimum of 3 other QAMs operating at the same frequency on other ports on the same DLC. The replication can be done to any ports on the same DLC and to any ports on different DLCs.

7.4 Spectrum Surveillance

Due to the nature of the CCAP system and its targeted services, the CCAP is placed at a critical location within the cable operator's network. As the primary bridge between the back-office network and the HFC plant, the CCAP is responsible for transmitting and receiving all of the signals on the HFC plant for MPEG-TS SDV, MPEG-TS VOD,

MPEG-TS Broadcast, DOCSIS High-Speed Internet, DOCSIS VoIP, and next generation video services. Many of these services require a fully-operational DOCSIS connection in the HFC return path for both data transport and for signaling/messaging transport.

To ensure the quality of the upstream, the CCAP is expected to be capable of monitoring the upstream return path to help identify:

- any return path RF issues that might negatively impact the performance of DOCSIS upstream channels, or
- proprietary signaling/messaging signals propagating in the upstream direction.

This upstream surveillance monitoring will provide the cable operator with up-to-date information on the RF quality of all of the upstream channels being used, and the entire upstream spectrum that is currently unused (in case the cable operator decides to utilize that spectrum for new upstream channels in the future). This upstream surveillance data will provide spectral information in the form of Fast Fourier Transform (FFT) outputs that could be used to quantify the magnitude of the Additive White Gaussian Noise (AWGN) on the upstream HFC plant, as well as the magnitude of ingress noise and impulse (burst) noise that might be present on the upstream HFC plant.

The CCAP will permit the cable operator to monitor the upstream spectrum without interrupting the transmission of DOCSIS data that is being simultaneously transmitted in the upstream direction. In addition, the CCAP will permit the cable operator to periodically schedule "quiet times" in the DOCSIS upstream channels to permit the RF noise to be successfully measured within the spectrum of each DOCSIS upstream channel. To ensure ease of use, the cable operator will be allowed to enable or disable this CCAP Spectrum Surveillance feature using simple configuration commands.

7.5 CCAP Configuration Management

The CCAP combines the functionality of an Edge QAM with a CMTS into a single platform designed to reduce operational costs and provide network flexibility. In order to provide operators the simplest path to deployment of CCAP with existing OSS systems, the goal for configuration and management of the CCAP is to treat the configuration of these very distinct platforms in a consolidated fashion.

This section will provide an overview of the various aspects of configuring a CCAP chassis including background on the choice of object modeling language chosen and a discussion of the various object models and their constructs.

7.5.1 YANG Data Modeling Language and XML Background

The configuration of a CCAP chassis can be accomplished using a variety of methods such as via a command-line interface, through file-based processing, or methods such as NETCONF or Web services. Underlying these configuration methods is a common configuration object model which defines the parameters that are to be configured. For the purposes of the CCAP, the configuration object modeling language is UML and the configuration data modeling language used is YANG.

YANG is a data modeling language for the NETCONF network configuration protocol that has been developed within the IETF to allow for modeling of configuration data, network element state and network events (see [RFC 6020]. Only the configuration data modeling portion of YANG is utilized. The option of using NETCONF to configure the CCAP is a new feature of the CCAP.

A CCAP configuration XML schema - derived from the CCAP configuration YANG data model is the basis for the creation of the XML file which will be used to configure the CCAP. The CCAP will parse the entire XML configuration file and process the configuration objects that are present in the file. Internal processing of configuration objects specified in the CCAP configuration file is vendor-specific.

7.5.2 Configuration Object Model

A CCAP Configuration and Management Object Model has been developed in UML to define the elements (objects) and their parameters (attributes) that will need to be represented in the YANG configuration data model and eventually, the configuration XML schema. The object model also defines the associations between objects.

One example of a CCAP object is the Downstream RF Port. The Downstream RF Port has its own attributes such as number (1…N where N is the number of ports on a line card), Administrative State (the status of the port), RF Mute

and Channel Power. The Downstream RF Port is also defined to be associated with a particular Downstream Line Card object and a set of Downstream Channel objects; these associations define the location of the Downstream RF Port in the chassis and the particulars of the QAM channels that will be carried on the Downstream RF Port.

7.5.3 Configuration Data Model

The CCAP YANG configuration data model is created from a direct translation of the CCAP configuration UML object model into a set of YANG modules. The YANG data model is constructed in a tree format using modules and sub-modules. The CCAP YANG modules can be used by a YANG translation tool to generate the XML schema used for the XML configuration file. Alternatively, the CCAP can validate the XML configuration file directly against the set of YANG modules via NETCONF or other methods.

7.5.4 CCAP Configuration File Processing

The CCAP is configured via the execution of an XML configuration file that is transferred to the file system on the CCAP. The CCAP parses the entire XML configuration file and processes the configuration objects represented in the file as a sequence of individual element operations. Individual element operations can succeed or fail; the CCAP will log unsuccessful operations.

Before a configuration file is applied to the CCAP, the CCAP performs several checks against the file, such as verifying that the configuration file is well-formed XML and that it validates against its schema. If the configuration file does not pass these checks, the CCAP will reject the file. The CCAP can also reject individual objects within the configuration file. In all rejection cases, the CCAP will log the rejection as an error.

The CCAP supports partial configuration, allowing the CCAP to not process objects that it either does not understand or that have invalid parameters, while continuing to process the objects that have no issues.

7.5.5 CCAP NETCONF-Based Configuration

The CCAP may also be configured via NETCONF as specified in [RFC 6241]. The CCAP uses standard NETCONF edit-config commands to execute XML-based configuration parameters. The XML configuration data can contain explicit "merge", "replace", or "delete" operation values at various nodes within the configuration tree provided. In this manner, a NETCONF command may contain merge operations for one branch of the tree, replace operations for another portion, and delete operations for yet another branch. The CCAP supports partial configuration, allowing the CCAP to not process objects that it either does not understand or that have invalid parameters, while continuing to process the objects that have no issues.

7.6 PON Configuration: DOCSIS Provisioning of EPON

Cable operators have recognized the value of including alternative access technologies into their network topology. This has included exploration and deployment of various wireless and passive optical solutions. The common characteristic for all of these alternative access technologies is that they have uncommon OSS models. Specifically, they do not look nor feel like a DOCSIS-oriented system to the operator that is responsible for deploying and managing services provided by this new technology. As such, there is value in mediating the interaction between these new access technologies and existing back-office tools, processes, and operator expectations.

The DOCSIS Provisioning of EPON (DPoE) specifications provide a service overlay of the DOCSIS and CMTS management framework on an IEEE 802.3ah/av EPON network. DPoE systems rely on the EPON MAC and PHY and the upper-layer DOCSIS protocols defined by CableLabs. The DOCSIS MAC and PHY do not apply.

7.6.1 The DOCSIS and DPoE Networks

The following diagram summarizes the primary systems and elements involved in a typical HFC-based DOCSIS network. For brevity, only a sample of the back-office systems used to provision, manage, authorize, and control the network are included.

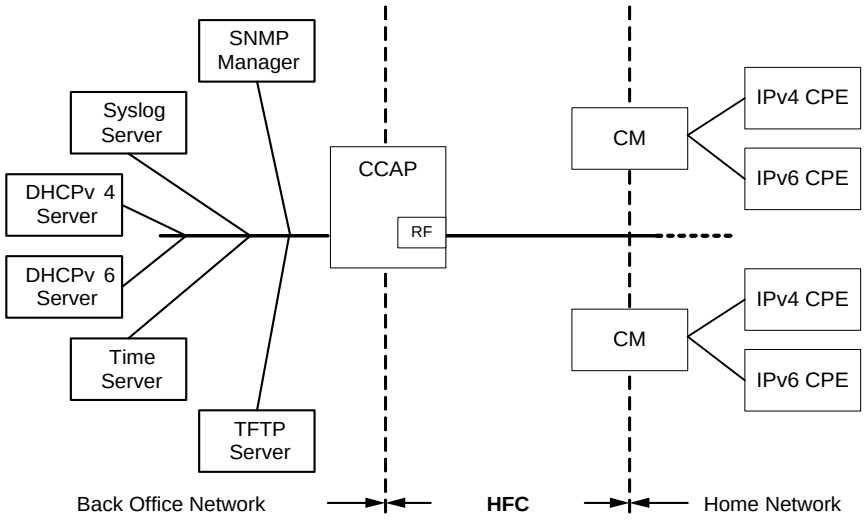

Figure 7–3 - DOCSIS 3.0 HFC Network Using CCAP

The DPoE specifications define a system that is analogous to a CMTS. The DPoE System need not be a single device, but instead could be a collection of devices that includes an Optical Line Termination (OLT), router, and DOCSIS emulation system. Collectively, these separate devices would be referred to as the DPoE System. The DPoE System provides the logical interfaces and protocol translations necessary to integrate EPON devices into the DOCSIS OSSI framework. This permits operators to take advantage of standard EPON functionality while retaining their investment in back office operations and systems, leaving CPE unchanged.

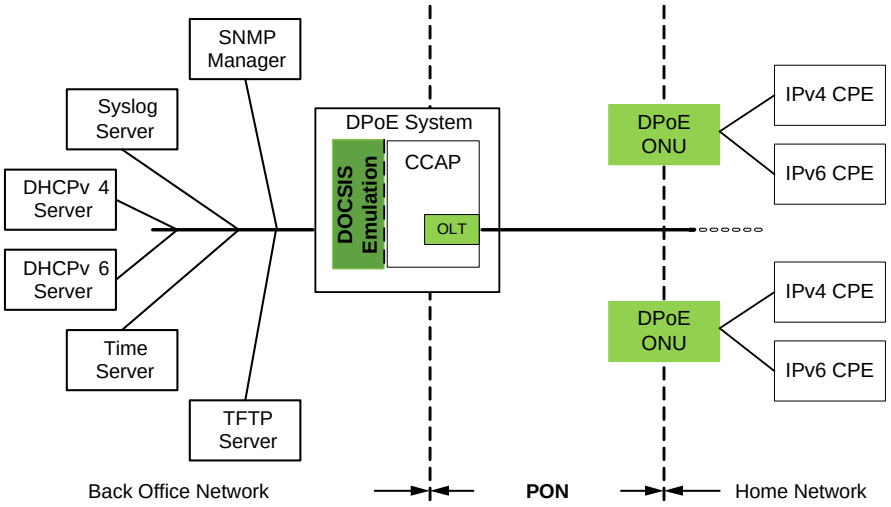

Figure 7–4 - DPoE Network using CCAP

7.6.2 DPoE Provisioning and Management

Interfaces and systems for managing DPoE devices may be provided through a DOCSIS Emulation module either running on the CCAP itself or running on an external server. Requirements for these interfaces are specified by the [DPoE OSSI] specification to provide Provisioning and Management support for the network. As shown

in Figure 7–5, requirements for the CMTS and CM MIBs are provided via the DPoE system by proxying for the OLT components on the CCAP line cards and the remote ONU devices.

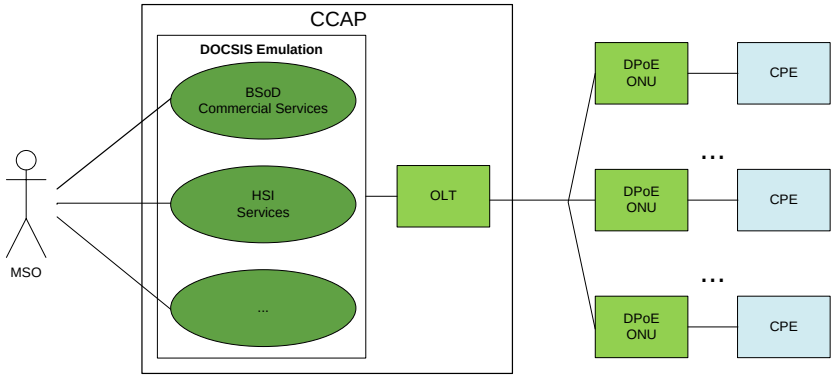

Figure 7–5 - Operator Interfaces to EPON Access Network

7.6.3 Provisioning and Management of OLT Devices

As CCAP is providing the functionality of the DPoE System, the OLT and associated PON interfaces are managed in a similar fashion to DOCSIS RF interfaces. Logical DOCSIS constructs, such as MAC domains, are implemented on top of the physical EPON interfaces. Relevant CMTS MIBs and DOCSIS CLI operations are mapped to their EPON equivalents to maintain the look-and-feel and, more importantly, the functionality expected of a DOCSIS RF interface.

7.7 Protocol Support

This section provides an overview of protocol support by the CCAP with respect to IP version, virtual private networks, routing, and multicast.

7.7.1 IP Versions

The CCAP supports IPv4 and IPv6 for both unicast and multicast traffic. The CCAP has the ability to forward traffic to both IPv4 and IPv6 devices.

7.7.2 VPN

The system will need to support VPN in order to support low cost layer 2 only VPN services provided over DOCSIS and expand offerings to high-touch managed L3 VPNs. When operating in routing mode, the CCAP is expected to support the following VPN-related RFCs:

- VPLS using BGP for auto-discovery and signaling, as specified in [RFC 4761]
- BGP/MPLS IP VPNs, as specified in [RFC 4364]
- The framework for L2VPNs, as specified in [RFC 4664]
- LDP (Label Distribution Protocol), as specified in [RFC 5036]
- Extensions to RSVP for LSP tunnels, as specified in [RFC 3209]
- Exclude routes extension to RSVP, as specified in [RFC 4874]

- Node behavior upon originating and receiving RSVP path error messages, as specified in [RFC 5711]

7.7.2.1 MPLS

The system is expected to support MPLS label switching services in order to integrate into existing and planned commercial services platforms. Major applications of MPLS are telecommunications traffic engineering and MPLS VPN. In the context of MPLS VPNs, the CCAP, operating in routing mode, is expected to support the following MPLS RFCs:

- Encapsulation MPLS in IP or Generic Routing Encapsulation (GRE), as specified in [RFC 4023]
- Encapsulation methods for transport of Ethernet over MPLS networks, as specified in [RFC 4448]
- Encoding of attributes for MPLS LSP establishment, as specified in [RFC 5420]

7.7.2.2 Multicast VPN (MVPN)

Multicast Virtual Private Network (MVPN) is a technology to deploy multicast service in an existing VPN or as part of a transport infrastructure. Multicast data is transmitted between private networks over a VPN infrastructure by encapsulating the original multicast packets.

When operating in routing mode, the CCAP is expected to support the following MVPN-related functionality:

- Intra-AS Multicast VPN (MVPN) membership discovery via BGP MCAST-VPN address family
- BGP C-multicast route exchange when the provider-edge – customer-edge (PE-CE) protocol is PIM-SM (SSM), PIM-SM (ASM), PIM-DM or IGMP
- IP/GRE based inclusive Provider tunnels (P-tunnels) signaled by PIM-SM (ASM)
- IP/GRE-based inclusive P-tunnels signaled by PIM-SM (SSM)
- MPLS inclusive P-tunnels signaled by RSVP-TE P2MP LSPs
- MPLS selective P-tunnels signaled by RSVP-TE P2MP LSPs

7.7.3 Routing

When operating in routing mode, the CCAP is expected to support the following routing protocols:

- RIPv2 on the access side as specified in [RFC 2453]
- IS-IS, as specified in [RFC 5303]
- Multi-Topology support for IS-IS, as specified in [RFC 5120]
- OSPFv2 as specified in [RFC 2328] and [RFC 5709]
- BGPv4, as specified in [RFC 4724]
- Multiprotocol Extensions for BGP-4, as specified in [RFC 4760]

7.7.4 Multicast

When operating in routing mode, the CCAP is expected to support Protocol Independent Multicast-Sparse Mode (PIM-SM) with SSM extensions, as specified in [RFC 4601]. In addition, the CCAP has the ability to join multiple IP multicast groups with PIM-SM. When the CCAP receives a multicast join request on an access interface for a specific multicast source and group, the CCAP uses the PIM-SM protocol to join that multicast flow, if needed.

For IPv4 on the access interfaces (DOCSIS, PON, etc.), the CCAP supports Internet Group Management Protocol version 3 (IGMPv3) with SSM extensions [RFC 3376] and also IGMPv2. For IPv6 on the access interfaces, the CCAP supports Multicast Listener Discovery version 2 protocol as specified in [RFC 3810]. The CCAP can allow support of IGMPv2, IGMPv3, and MLDv2 joins from the same access interface.

The CCAP is expected to support joining at least 4096 multicast groups.

7.8 Non-Routing Mode Description

A CCAP operating in non-routing mode is required to adhere to the bridging CMTS requirements of [MULPI], [SEC], and [OSSI]. The following sections describe the CCAP behavior when operating in non-routing mode. The ER complements the operation of a CCAP operating in non-routing mode, and is expected to perform routing and other protocols unsupported by that CCAP.

7.8.1 Unsupported Protocols

When operating in non-routing mode, the CCAP is not expected to support the following protocols:

- RIPv2 on the access side as specified in [RFC 2453]
  - Except for RIP Gleaning when SAV for static IP subscribers is implemented in non-routing mode
- IS-IS, as specified in [RFC 5303]
- Multi-Topology support for IS-IS, as specified in [RFC 5120]
- OSPFv2 as specified in [RFC 2328] and [RFC 5709]
- BGPv4, as specified in [RFC 4724]
- Multiprotocol Extensions for BGP-4, as specified in [RFC 4760]
- PIM-SM [RFC 4601]

7.8.2 Link and Path Redundancy

To ensure minimal disruption of service, the CCAP can be deployed in a manner that ensures high availability of services. Different methods provide similar functions to detect and remedy link and path failure to the CCAP, depending on the forwarding mode (routing or non-routing).

When multiple ER platforms are deployed (as shown earlier in Figure 5–1and Figure 5–2), it is expected that ER platforms will have methods to synchronize subscriber state information and multicast state information; this methodology is outside of the scope of this document.

7.8.3 Non-Routing Mode Behavior Differences

7.8.3.1 Prefix Stability with IPv6

Topology changes, such as node splits from one CMTS to another, require special consideration when implementing a service that requires a stable IPv6 prefix. [MULPI] defines a mechanism, in the Prefix Stability at the CMTS section, to account for such changes through manipulation of the routing table within a CMTS. This mechanism does not apply to a CCAP operating in non-routing mode, as it does not participate in the operators' routing domain. With the non-routing mode forwarding architecture, only node splits from one ER to another will need special procedures to transfer stable IPv6 prefixes. The ER network design is highly individual to the operator, making the details of these procedures outside of the scope of this document.

7.8.3.2 Static IP Addressing

In addition to regular DHCP-based customers, the CCAP serves customers that are assigned static IP addresses. The CM serving such a customer obtains its IP address and configuration file by the normal DHCP+TFTP process. Normally, an operator-provided router exists behind the CM, which may get its CCAP-facing IP address through a DHCP transaction as well. However, the CPEs behind the CM/router will not invoke a DHCP transaction, and neither the CCAP nor the ER can rely on DHCP to learn their addresses. The CCAP operating in non-routing mode identifies CMs that serve static IP customers and learns these CPE's MAC/IP bindings from parameters in the CM's configuration file and by gleaning RIP messages coming through the CM. The ER learns addresses of static IP devices through RIP/ARP/ND.

In cases where an operator-provided router does not exist behind the CM, or when such a router does not use DHCP to obtain its CCAP-facing IP address, the ER may not be able to associate the CM with the static IP devices it

serves. In such cases, the ER should be provisioned with range(s) of IP addresses used for static IP customers; the CCAP operating in non-routing mode should have enforcement of SAV functions enabled.

7.8.3.3 Split-Horizon Forwarding

A CCAP operating in non-routing mode implements a slightly different forwarding model than a "regular" Layer 2 switch/bridge, even though it forwards traffic based on the source and destination MAC addresses. This forwarding model is known as a "residential bridge" or "split-horizon forwarding". The split-horizon forwarding model applies to two different types of ports: Access ports and Network ports.

All access ports belong to a "split-horizon group"; network ports belong to their own split-horizon group. Traffic cannot be forwarded by the bridge from one port of a split-horizon group to any port of the same split-horizon group. In other words, L2 communication is completely blocked between members of the same split-horizon group.

Thus, access ports can only forward traffic to network ports. Figure 7–6 shows the forwarding model for unicast traffic; flows indicated by green, solid arrows are allowed because they traverse the horizon, while the flows indicated by red, dashed arrows are denied communication since they are within the same split-horizon group.

It should be noted that throughout this section, the term "port" does not necessarily imply a physical port but rather a logical port. Thus, a single network port could encompass several physical ports in the case of LAG in the network direction or could represent a service group, a downstream bonding group (DBG), or a channel in the access direction (and thus, encompass only a fraction of a physical access/RF port).

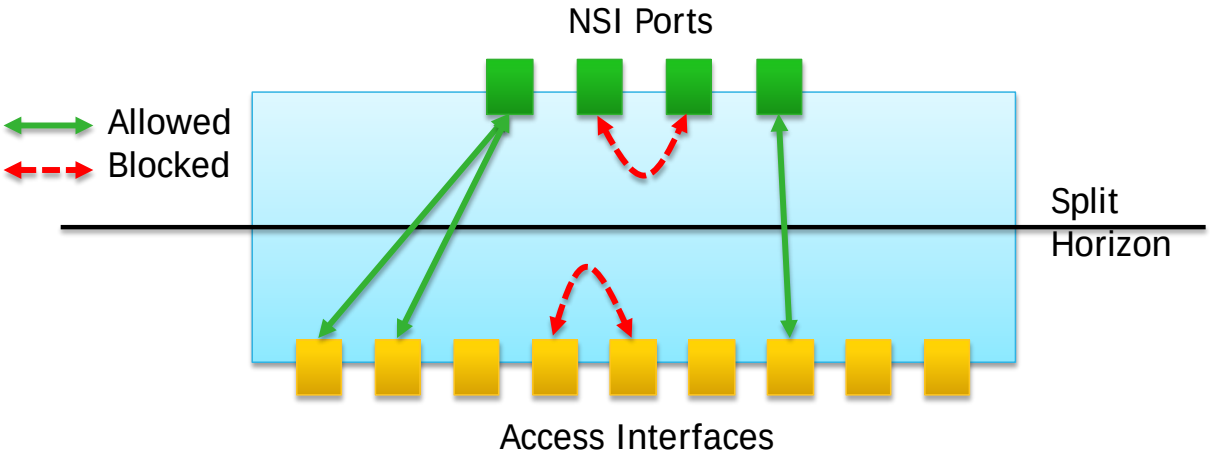

Figure 7–6 - Split Horizon Topology – Unicast Forwarding

The main purpose for split-horizon forwarding is to prevent any form of abuse by users connected to the same CCAP operating in non-routing mode that could otherwise forward traffic in an L2 fashion between them. This could potentially include non-IP protocols. By blocking any user-to-user L2 communication at the CCAP, all user-to-user communication is handled by the ER and is thus subject to the same policies and control regardless of where the two users might be located. The mechanisms needed to allow this user-to-user IP-level communication are explained in section 7.8.3.6.

7.8.3.4 Broadcast Traffic Behavior

Broadcast and multicast traffic are treated differently than unicast traffic by a CCAP in non-routing mode, but still follow the split-horizon rules outlined here.

As detailed in Section 7.8.3.6 (Proxy ARP/ND), the ER is not currently expected to send any DS broadcast. Nevertheless, it is envisioned that this may change, and that future CCAP features may require the ER to send broadcast in the downstream (DS). To support such future features, a CCAP operating in non-routing mode will

need to forward (and replicate) DS broadcast according to Virtual Local Area Network (VLAN) rules. For example, broadcast on a Business Services over DOCSIS (BSoD) VLAN will be forwarded to the downstream channel set (DCS) accessible by that BSoD CM; DS broadcast on a "general" VLAN will be forwarded to all MAC domains and default group service flows served by that VLAN. Accordingly, minimizing broadcast traffic (once a DS-broadcast feature is enabled) may require establishing a larger number of VLAN tags than just a single "general" one (e.g., VLAN per MAC domain). DS broadcast frames that are not VLAN tagged are assumed to be associated only with the CCAP host IP, and will not be forwarded to access ports.

In the upstream direction, broadcast is forwarded to the proper VLAN (on the proper port). In all cases, though, split-horizon rules still apply, thus no access-port-sourced (upstream) broadcast should ever reach the downstream of any access port. The non-routing mode implements no flooding due to an unknown destination MAC address. Again, no users shall see any other user's L2 frames.

Table 7–1 summarizes the allowed and denied communication flows.

Table 7–1 - Split horizon communication flows

Allowed and denied flows		Destination			
		Unicast		Broadcast	
		Access	Network	Access	Network
Source	Access	NO	YES	NO	YES
	Network	YES	NO	YES [2]	NO

All regular user traffic, sourced or destined from/to the Internet, is comprised exclusively of unicast traffic. Leaving aside multicast traffic (intended for IPTV-based video distribution), broadcast packets should still be investigated. Broadcast packets are used, in the case of residential subscribers, exclusively for L2/L3 bringup. Once IP-level connectivity has been established, be it IPv4 or IPv6, broadcast packets are not needed at all. Thus, two major cases will need to be investigated in terms of broadcast-packet use, namely DHCP (or DHCPv6) and ARP/ND.

7.8.3.5 DHCP behavior

We shall first focus on DHCPv4 specifics. During a host bring up, a CM, CPE or any other dynamic IP-enabled device will initiate a DHCP discover packet. This packet uses a broadcast destination MAC address (FF:FF:FF:FF:FF:FF) and is forwarded on the upstream of the CM. Once this packet reaches the CCAP operating in non-routing mode, it is sent to the control-plane of the CCAP. There it will be modified, option-82 information will be inserted, and some basic security checks will be performed. One of these checks is to verify that no option-82 information is already present in the packet, as this would signify a potential attack attempt by end-users. Some basic rate-limiting (ideally, on a per-CM or per-user basis) should also be performed. Additional security checks may also be performed, but they are out of the scope of this document.

The now-modified DHCP discover packet is forwarded to the appropriate NSI port and is received by the ER. The ER will process it, act on it as a DHCP relay (converting it into a unicast packet destined to the DHCP server) and eventually will process the reply (in the form of an "offer" or "ack" type of packet). In the great majority of cases, this reply will be sent to an unknown destination IP address, but the L2 address of the packet will be the MAC address of the client originally initiating the discover/request.

All DHCP replies (offers/acks) received on the network port coming from an ER will be intercepted and sent to the control-plane of the CCAP operating in non-routing mode. This is meant to allow the CCAP to statefully populate its IP/MAC tables based on the acknowledged addresses from the DHCP server, and to remove the option-82 information from the DHCP response.

[2] DHCP broadcast replies are intercepted by the control-plane to remove Option 82 (or its IPv6 equivalent), and are forwarded to the end-device only and do not use the DOCSIS broadcast channel.

If the requesting device had set the broadcast bit, the DHCP response packet will be sent as an L2 broadcast from the ER to the CCAP operating in non-routing mode and then down the DOCSIS path. However, the CCAP will only forward this response to the DCS accessible by the originating device.

Figure 7–7 shows a simplified representation of the DHCP packet flow among all the elements involved.

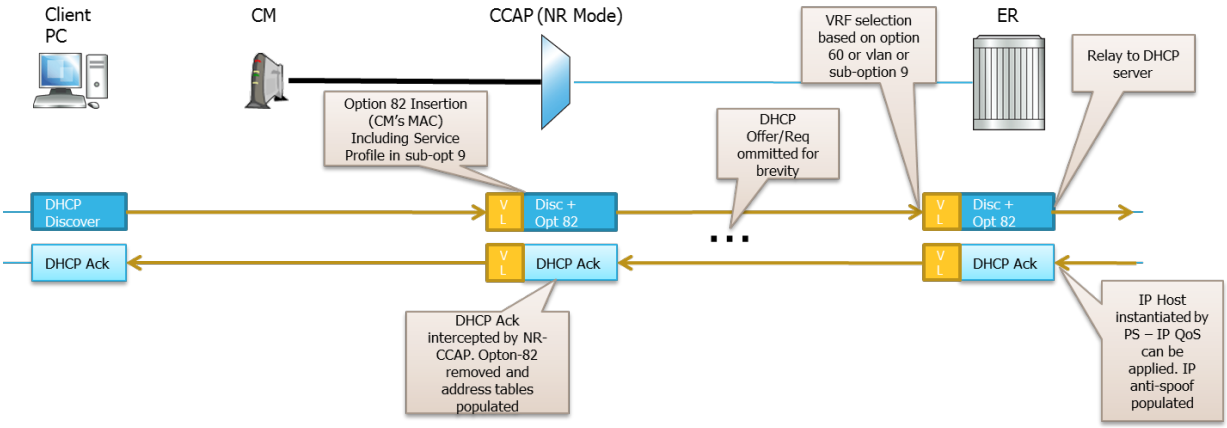

Figure 7–7 - DHCP Packet Flow

From an ER's perspective, all DHCP ack packets are inspected and a least-state table is built in the system. This allows the ER to keep an "authoritative" ARP table (since both MAC and IP are learned from the DHCP ack); the ER does not rely on ARP itself to populate the ARP table. This table is also the base for the ER's anti-spoof tables against which each subscriber's packets are matched to ensure that addresses (both L2 and L3) have not been stolen.

For IPv6, the behavior is similar, except that the DHCP agent in the CCAP operating in non-routing mode will behave according to [RFC 6221]. In this case, the DHCP message is encapsulated in a similar fashion as the typical DHCPv6 relay, but uses a multicast destination MAC address similar to that of a regular DHCP packet. This type of relay agent is called Lightweight DHCP Relay Agent (LDRA) and will add (in the upstream) the IPv6-equivalent options to option-82. The options needed are referenced in [MULPI] and include Interface-ID, CMTS capabilities option, Remote-ID, and CM-MAC address option. The packet handled via LDRA is then received by the ER which proceeds to relay it again as a regular DHCPv6 relay. On the return path, the LDRA will receive the DHCPv6 response packet from the upstream ER and will decapsulate it, remove the added options, and forward the packet down to the client on a single client-accessible DCS.

7.8.3.6 Proxy ARP/ND

The ER is required to implement proxy-ARP and proxy-ND. A CCAP operating in non-routing mode is not required to implement proxy-ARP and proxy-ND, but is required to forward upstream ARP broadcast and Neighbor Discovery multicast requests to the ER on the appropriate VLAN.

Furthermore, it should be noted that under normal operations, the ER will not generate downstream ARP requests for a customer, but instead would rely on the stateful population of its host tables through DHCP. This means the ER always has an authoritative ARP table. For residential services based on DHCP, this greatly reduces the need for any kind of ARP-level policing/rate limiting, as it completely eliminates the need for ER-initiated ARP requests.

In the case where certain subscribers have an entirely static configuration (that is, statically assigned IPs without the use of DHCP), special static hosts can be provisioned on the ER. These hosts can be populated in two manners: in the first, both L2 and L3 addresses are configured (very secure but operationally complex); in the second, only an L3 address is configured. In the case of L3-only hosts, the ER will not initiate an ARP request for the host itself, but will wait for an ARP request from the host towards the ER. Based on this information, it will populate its own ARP table and should keep the ARP entry valid as long as traffic is flowing from it.

7.8.3.7 Multicast Traffic Behavior

The DOCSIS portion of multicast forwarding is implemented by the CCAP in an identical fashion regardless of the operation mode (routing or non-routing). The NSI segment, however, is slightly different, since a CCAP in non-routing mode receives multicast traffic on a separate, multicast-only VLAN. Based on the learned joins received from each CM, a CCAP operating in non-routing mode will replicate the specific multicast groups received on the multicast VLAN and will forward them to the appropriate group service flows. The ER completes the multicast functionality by performing IGMP/MLDv2-to-PIM translation.

Multicast control traffic between the CCAP in non-routing mode and the ER is intended to minimize the capacity utilization of the link between them, and should function as follows:

- Individual CMs send IGMP joins and leaves to the CCAP operating in non-routing mode.
- These messages are received by the CCAP, which implements IGMP-proxy.
- The CCAP, upon receiving a customer-sourced join for an un-received channel, issues its own IGMP join on the multicast-specific VLAN. Subsequent customer joins for that channel will not generate a new join towards the ER.
- When the last watching customer sends an IGMP leave – and thus no more clients requested that particular group – the CCAP issues an IGMP leave towards the ER.

The CCAP IGMP proxy mechanism should also replicate the ER-initiated IGMP GMQ (General Membership Query) to which the clients should respond in order to fulfill the liveliness requirements of the protocol. The CCAP implements MLDv2 proxy in a similar fashion.

7.8.3.8 Encapsulation

When the CCAP is operating in non-routing mode, the L3 forwarding functionality is moved from the CCAP to the ER. A simple encapsulation method is required to enable connectivity between the CCAP and ER through other L2 and/or L3 network devices, while keeping the CCAP data segregated from other traffic that may exist in the network. Currently, VLAN tagging is the preferred method for facilitating this segregation, although it is envisioned that MPLS may replace VLAN in the future.

VLAN configuration is static in both the ER and the CCAP and configured to facilitate point-to-point tunnels. Accordingly, the VLAN IDs used for this purpose are unique in the local network and are not shared between separate CCAP platforms (with a possible exception of the VLAN for DOCSIS multicast). The static VLAN configuration consists of:

- A single unique VLAN tag for all DOCSIS multicast services.
- A single unique VLAN tag for traffic destined to devices behind each business CM (BSoD). If two or more business CMs, connected to the same CCAP, are serving multiple locations of the same BSoD VPN service, each CM's traffic will still be assigned a unique VLAN tag. Note that traffic to/from the CM entity itself (as an IP host) is not part of the BSoD service and will be tagged with the "general" VLAN tag as below.
- At least one unique VLAN tag for general traffic to all other devices and services served by the CCAP. Multiple such VLAN tags may be configured if a greater degree of segregation is desired (e.g., a unique VLAN tag per MAC domain).

These VLAN tags are required only for the link between the CCAP and the ER. Accordingly, each of these two entities adds the tags to Ethernet frames destined to the other entity and removes the tag from frames arriving from the other entity before forwarding these frames. The frames tagged over the CCAP-ER link may or may not be already VLAN tagged, and that original tagging will not be affected by the CCAP-ER link.

Note that traffic destined to the video EQAM function of the CCAP is directed to the CCAP as an IP host, and as such, this traffic is not required to be VLAN tagged. Similarly, management traffic directed to/from the CCAP as an IP host is not required to be VLAN tagged.

7.8.3.9 Subnet Query and Control Point Discovery Support

The back office systems (policy server, call management server, delivery function) need to know which CCAP can carry out an operation pertaining to a specific end device. Thus, these systems need to know which subnets are handled by each CCAP. One current, commonly deployed method is a static configuration of a back office element with the IP addresses of all CCAP devices. The configuration does not include the subnets that each CCAP handles (which may occasionally change); instead, the back office device periodically queries each CCAP for its subnet scope(s), thus constructing and maintaining an association database of the CCAP control points and the end point devices they control. Another method makes use of PacketCable 2.0 Control Point Discovery (CPD) as an alternative to static configuration of the back office element. A CPD request message is sent by the back office device to the end-point IP address, but is intercepted by the control point (CCAP). The CCAP then generates a response to the requestor, identifying itself as the control point for that IP address. If requested, that response includes the subnet associated with that end-point IP address.

In the case of a CCAP operating in a non-routing mode, the CCAP does not require configuration of the subnets that it serves to perform data path forwarding. Nevertheless, it is desirable that this subnet configuration will be provided to the CCAP to serve the functionality described above. If such subnet configuration was not provided, the CCAP should attempt to reconstruct the subnets it serves from information gathered through DHCP and RIP gleaning. In its response to subnet queries, a CCAP in non-routing mode only uses reconstructed subnet configuration if explicit subnet static configuration is not available.

8 CCAP IMPLEMENTATIONS

The following paragraphs describe implementations that support current headend use cases. It is expected that the CCAP could evolve to support extended frequency ranges and other access network technologies, such as EoC, RFoG, and PON over Coax.

8.1 CCAP Interface Options

8.1.1 Hybrid-Fiber Coax Interfaces

The CCAP is expected to implement upstream and downstream RF interfaces on separate downstream and upstream physical interface cards (PICs). This separation of PICs allows upstream and downstream capacity to be changed independently. The CCAP could be implemented with a combined PIC (a card that has both upstream and downstream interfaces), although this is not expected to be a typical configuration.

8.1.1.1 Downstream RF Interfaces

The CCAP is expected to support a downstream RF interface ratio of one downstream RF port per downstream service group. A large CCAP is expected to support a minimum of 40 – 60 downstream ports, and a small CCAP is expected to support a minimum of 16 – 20 downstream ports. Downstream physical interface cards (DPICs) are expected to support a minimum of 8 downstream RF ports per card.

Each downstream RF port supports the following edge-to-edge frequency ranges for North American devices:

- 54 – 1002 MHz or
- 108 – 1002 MHz

The North American channel width is specified as 6 MHz.

For European devices, each downstream RF port supports an edge-to-edge frequency range of 86 – 1006 MHz, with a channel width of 8 MHz.

In addition, a single downstream RF port is expected to be capable of supporting up to 158 QAMs of any type. Each port will typically support 32 – 64 narrowcast QAMs and up to 96 broadcast QAMs.

The preferred implementation of a downstream RF interface is in the form of an F-connector, but ganged 75 Ohm MCX (Micro-Coaxial) interfaces in a universal cable holder (UCH) could also be implemented. All RF interfaces are expected to be located at the rear of the chassis.

8.1.1.2 Upstream RF Interfaces

The CCAP is expected to support an upstream RF interface ratio of one upstream RF port per upstream service group. A large CCAP is expected to support at least 80 – 120 upstream ports (with 120 being the preferred minimum), and a small chassis is expected to support 32 – 40 upstream ports. Upstream physical interface cards (UPICs) are expected to support at least 16 upstream RF ports.

Each upstream RF port supports the following edge-to-edge frequency ranges for North American devices:

- 5 – 42 MHz or
- 5 – 85 MHz

For European devices, each upstream RF port supports an edge-to-edge frequency range of 5 – 65 MHz, as specified in Annex B of [PHY].

Both European and North American devices support channel widths of 1.6, 3.2, and 6.4 MHz. They are also expected to support between 4 – 6 DOCSIS RF upstream channels per RF port.

The preferred implementation of an upstream RF interface is in the form of ganged 75 Ohm MCX interfaces in a UCH, but an F-connector can also be implemented. All RF interfaces are expected to be located at the rear of the chassis.

8.1.2 Ethernet Passive Optical Network (EPON) Interfaces

Fiber access represents one method of offering higher bandwidth to subscribers. Increasing bandwidth may be required to meet the increasing data demands of business customers. EPON minimizes investment in the access infrastructure while delivering more bandwidth and greater service flexibility. It is the least costly method of constructing fiber to the subscriber and it has the service control and flexibility to offer any business or consumer service.

8.1.2.1 Split Ratios and Customers Served

The EPON split ratio for CCAP is expected to be 128:1. Each CCAP chassis is expected to support at least 16 active EPON interfaces, and 24 or more is preferable. Two RF service groups should be served by each EPON interface. This suggests a density of approximately 128 business customers per 1000 residential customers.

In EPON networks distance and subscriber density are inversely correlated. For example, the CCAP is expected to support 128 business customers on a PON at a distance up to 5km, up to 64 customers at 10km, up to 32 customers at 20km, 16 customers at 30km, 8 customers at 40km, and up to 4 customers at a distance of 50km.

8.1.2.2 EPON Redundancy

The CCAP chassis can optionally support N+1 EPON redundancy. Other redundant configurations are possible as well. High availability and reliability are critical to enterprise business services and will be an important option when diversely routed fiber is available to the subscriber.

8.1.2.3 EPON Connectors

EPON line cards should support SFP-type connectors. This form factor allows for maximum flexibility for wavelength and port selection. SC/PV is also an option.

8.1.3 Network-Side Interface (NSI)

The CCAP receives any data to be transmitted downstream (e.g., Internet content, IP video, PacketCable voice data, and DSG data) through the Network-Side Interface (NSI), which would consist of at least 160 Gbps of data on one or more physical interfaces in order to support a fully loaded large chassis (at least 80 Gbps for a small chassis). The CCAP will be expected to implement multiple, redundant NSIs on the Switch Routing line card. These interfaces should be deployed with standards-based pluggable optics. While initially deploying 10 GbE interfaces, CCAP deployments are expected to migrate to 100 GbE interfaces as the technology matures. 10 GbE interfaces should comply with [IEEE 802.3ae] and 100 GbE interfaces should comply with [IEEE 802.3ba]. Each NSI port should support untagged IEEE 802.3 Ethernet encapsulation. Individual Network-to-Network Interfaces (NNIs) may be provided for HSI business services using a 10G EPON interface.

8.2 CCAP Chassis Sizing

The CCAP chassis may be deployed in a large chassis, designed to support a minimum of 40 downstream RF ports. The CCAP could also be implemented in a smaller chassis, supporting at least 16 downstream RF ports. The smaller implementation will provide great value to smaller hub sites in use by MSOs.

8.3 Duplicate MAC Addresses Handling

Various CPE devices, including wireless routers, provide users the capability of modifying their MAC address. Unfortunately, it is very common that users use such "simple" MAC addresses as 1234-5678-9ABC or 1111-1111-1111, greatly increasing the probability that two CPE devices under the same CCAP will have the same MAC address. Other CPE devices, in particular those of lower cost non-conforming brands, tend to re-use MACs from a small pool or just "borrow" other vendor's OUI. Regardless of the origin though, this present an issue. Normally, a DHCP server will assign the same IP address to DHCP requests from duplicate MAC CPEs, resulting in both Layer 2 and Layer 3 address conflict. This is an erroneous case, which in some cases indicates a hacking attempt. A cable operator may wish the CCAP to prevent the connection of a device showing a duplicate MAC address to the network (including a DHCP transaction), and flag the occurrence.

Alternatively, a DOCSIS aware DHCP server can be configured to accept requests from duplicate MACs and grant different IP addresses based on Option 82 differentiation (Relay Agent Remote ID = CM MAC address). A routing CCAP can handle such a scenario since forwarding decisions to/from these CPEs will be based on their IP addresses. However, a CCAP operating in non-routing mode must make forwarding decisions based on the MAC address, and will not be able to differentiate between the duplicate MAC CPEs when forwarding in the DS direction.

Appendix I Acknowledgments

We wish to thank the following contributors for their efforts toward creating versions 1 and 2 of this report:

Tom Cloonan	ARRIS
Tal Laufer	BigBand Networks
Andrew Chagnon	Broadcom Corporation
Victor Hou	Broadcom Corporation
Matthew Schmitt	CableLabs
John Bevilacqua	Comcast Corporation
Saif Rahman	Comcast Corporation
Jorge Salinger	Comcast Corporation
Joe Solomon	Comcast Corporation
Jeff Finkelstein	Cox Communications
Mannix O'Connor	Hitachi Communication Technologies
Bill Welch	Juniper Networks
Jean-Francis Kisovec	LiquidXstream Systems
Charles Corbalis	RGB Networks
Kirk Erichsen	Time Warner Cable
Mike Kelsen	Time Warner Cable

In addition, we thank the following contributors for their efforts toward creating version 3 of this report:

Rex Coldren	Alcatel-Lucent
Diego Garcia Del Rio	Alcatel-Lucent
Marty Glapa	Alcatel-Lucent
Jeff Howe	ARRIS
Tal Laufer	ARRIS
Rei Brockett	Aurora Networks
Victor Hou	Broadcom Corporation
Volker Leisse	Cable Europe Labs
Jennifer Andreoli-Fang	CableLabs
Matthew Schmitt	CableLabs
Alon Bernstein	Cisco Systems
Niclas Comstedt	Cisco Systems
Pawel Sowinski	Cisco Systems
John Bevilacqua	Comcast Corporation
Jason Combs	Comcast Corporation
Steven Fager	Comcast Corporation

Peter Hutnick	Comcast Corporation
Walt Michel	Comcast Corporation
Saif Rahman	Comcast Corporation
Jorge Salinger	Comcast Corporation
Steve Sigman	Comcast Corporation
Joe Solomon	Comcast Corporation
Robert Smith	CommScope
Adi Bonen	Harmonic
Alexander Fertelmeister	Harmonic
Amir Leventer	Harmonic
Adam Levy	Harmonic
Asaf Matatyaou	Harmonic
Colby Barth	Juniper Networks
Bill Welch	Juniper Networks
John Ulm	Motorola
Nasir Ansari	Rogers Cable
Niem Dang	Time Warner Cable
Kirk Erichsen	Time Warner Cable