

Data-Over-Cable-Service-Interface Specifications Modular Headend Architecture

Edge QAM Video Stream Interface Specification

CM-SP-EQAM-VSI-I01-081107

ISSUED

Notice

This DOCSIS® specification is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document.

Neither CableLabs nor any member company is responsible to any party for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this document, or any document referenced herein. This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein.

© Copyright 2008 Cable Television Laboratories, Inc.
All rights reserved.

Document Status Sheet

Document Control Number: CM-SP-EQAM-VSI-I01-081107				
Document Title: Edge QAM Video Stream Interface Specification				
Revision History: I01 - Released 11/07/08				
Date: November 7, 2008				
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author/AVG Only	CL/Member	CL/Member/Vendor	Public

Key to Document Status Codes

- Work in Progress** An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.

- Draft** A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.

- Issued** A stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and for certification testing.

- Closed** A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

CableLabs[®], DOCSIS[®], EuroDOCSIS[™], eDOCSIS[™], M-CMTS[™], PacketCable[™], EuroPacketCable[™], PCMM[™], CableHome[®], CableOffice[™], OpenCable[™], OCAP[™], CableCARD[™], M-Card[™], DCAS[™], tru2way[™], and CablePC[™] are trademarks of Cable Television Laboratories, Inc.

Contents

1	SCOPE	1
1.1	INTRODUCTION AND PURPOSE	1
1.2	REQUIREMENTS AND CONVENTIONS.....	1
2	REFERENCES	2
2.1	NORMATIVE REFERENCES	2
2.2	INFORMATIVE REFERENCES	2
2.3	REFERENCE ACQUISITION.....	3
3	TERMS AND DEFINITIONS	4
4	ABBREVIATIONS AND ACRONYMS	5
5	TECHNICAL OVERVIEW	7
5.1	VIDEO EQAM REFERENCE ARCHITECTURE	7
5.1.1	<i>Video EQAM Block Diagram</i>	7
5.1.2	<i>Static vs. Dynamic Sessions</i>	8
5.1.3	<i>Input TS Processing</i>	9
5.1.4	<i>QAM Channel Processing</i>	10
5.1.5	<i>Redundancy</i>	10
5.1.6	<i>Summary of Session Types</i>	10
6	STREAM INPUT SPECIFICATIONS	12
6.1	STREAM INPUT ETHERNET INTERFACE SPECIFICATIONS	12
6.2	MPEG-2 TRANSPORT STREAM OVER IP INTERFACE SPECIFICATIONS	12
6.2.1	<i>MPEG-2 Input Transport Streams</i>	13
6.2.2	<i>IP Unicast and Multicast</i>	14
6.2.3	<i>Private Data Streams</i>	14
6.3	INPUT STREAM NETWORK DE-JITTERING SPECIFICATION	14
7	STREAM INPUT TO OUTPUT ROUTING	16
7.1	STATIC PORT MAP BASED ROUTING	16
7.2	DYNAMIC ERMI-2 SIGNALING BASED ROUTING	16
7.3	SIMULTANEOUS HANDLING OF DIFFERENT SESSION TYPES	17
7.4	MPTS INPUT ROUTING.....	18
7.5	INPUT TO OUTPUT LATENCY.....	18
8	QAM CHANNEL PROCESSING	19
8.1	MULTIPLEXING MODE	19
8.1.1	<i>General Multiplexing Requirements</i>	19
8.1.2	<i>Multiplexing with Input SPTS Streams</i>	23
8.1.3	<i>Multiplexing with Single Program Demux from Input MPTS</i>	23
8.1.4	<i>Multiplexing with Entire “Pre-conditioned” MPTS</i>	24
8.2	PASSTHROUGH MODE.....	24
9	ENCRYPTION AND ENCRYPTION INTERFACE	25
9.1	VIDEO PROGRAM ENCRYPTION	25
9.2	VIDEO PROGRAM SCRAMBLING.....	25
9.3	ENCRYPTION CONDITIONAL ACCESS CONTROL MESSAGES.....	25
10	STREAM OUTPUT SPECIFICATION	28

10.1 OUTPUT MPEG-2 TRANSPORT STREAM GENERAL REQUIREMENTS28

10.2 OUTPUT MPEG-2 TRANSPORT STREAM OPERATIONAL REQUIREMENTS.....28

 10.2.1 Output Stream Set up Time28

11 REDUNDANCY SPECIFICATIONS.....29

11.1 REDUNDANCY INTRODUCTION29

 11.1.1 Redundancy: Efficiency29

 11.1.2 Redundancy: Quality of Service Types29

 11.1.3 Redundancy: Failover transition.....30

 11.1.4 Redundancy: Input Interface basic rules.....30

 11.1.5 Redundancy: Management and Control.....30

11.2 REDUNDANCY GENERAL REQUIREMENTS30

 11.2.1 Input Interface Redundancy.....30

 11.2.2 Input TS Source Redundancy.....31

 11.2.3 Output TS Redundancy32

12 INPUT AND OUTPUT MONITORING33

12.1 INPUT TRANSPORT STREAM MONITORING.....33

12.2 INPUT TS PROGRAM STREAM BIT RATE MONITORING34

12.3 OUTPUT TS PROGRAM STREAM BANDWIDTH UTILIZATION MONITORING.....34

APPENDIX I ACKNOWLEDGEMENTS35

Figures

Figure 5–1 - Video EQAM Block Diagram 8

Figure 6–1 - MPEG-2 Transport Stream over UDP/IP 12

Tables

Table 5–1 - Summary of Dynamic Session Types 11

Table 5–2 - Summary of Static Session Types 11

Table 8–1 - Well-known PIDs and PID Range 22

1 SCOPE

1.1 Introduction and Purpose

This specification is a component of the Modular Headend Architecture. In particular, it defines the data plane requirements for receiving, processing, and transmitting MPEG transport streams in EQAMs compliant with the Video EQAM or Universal EQAM profiles described in the Architectural Overview of the Modular Headend Architecture [MHA TR].

It should be noted that this document refers to "MPEG video" support as a short-hand to designate support for the delivery of [ISO/IEC 13818-1] compliant MPEG transport streams that may carry audio material, video material, data, or a combination.

The EQAM supporting MPEG video has one or more IP network interfaces (e.g., Gigabit Ethernet) for receipt of IP-encapsulated MPEG transport stream packets that may arrive from a VOD or streaming server. The EQAM can be configured either statically (via [EQAM PMI]) or dynamically (via ERMI-2, as defined in [ERMI]) to process the incoming MPEG transport streams, generate a Multiple Program Transport Stream (MPTS), and then, from that MPTS, generate a QAM modulated RF waveform that is transmitted at a particular frequency on an RF output of the device.

1.2 Requirements and Conventions

Throughout this document, the words used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

2 REFERENCES

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references.

[DRFI]	Downstream Radio Frequency Interface, CM-SP-DRFI-I06-080215, February 15, 2008, Cable Television Laboratories, Inc.
[DTI]	DOCSIS Timing Interface, CM-SP-DTI-I04-061222, December 22, 2006, Cable Television Laboratories, Inc.
[EQAM PMI]	Edge QAM Provisioning and Management Interface Specification, CM-SP-PMI-D02-081107, November 7, 2008, Cable Television Laboratories, Inc.
[ERMI]	Edge Resource Management Interface Specification, CM-SP-ERMI-I03-081107, November 7, 2008, Cable Television Laboratories, Inc.
[ISO/IEC 13818-1]	ISO/IEC 13818-1:2000, Information technology, Generic coding of moving pictures, and associated audio information: Systems.
[ISO/IEC 8802-2]	ISO/IEC 8802-2:1998, Information technology, Telecommunications and information exchange between systems, Local and metropolitan area networks, Specific requirements, Part 2: Logical link control.
[ISO/IEC 8802-3]	ISO/IEC 8802-3:2000, 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.
[ITU-T J.83]	ITU-T Recommendation, J.83, (04/97) Digital multi-programme systems for television sound and data services for cable distribution.
[SCTE 52]	ANSI/SCTE 52 2008, Data Encryption Standard - Cipher Block Chaining Packet Encryption Specification.
[SCTE 54]	ANSI/SCTE 54 2006, Digital Video Service Multiplex and Transport System Standard for Cable Television.

2.2 Informative References

This specification uses the following informative references.

[CEP]	Content Encoding Profiles 2.0 Specification, MD-SP-VOD-CEP2.0-I02-070105, January 5, 2007, Cable Television Laboratories, Inc.
[E-BIF]	OpenCable Enhanced TV Binary Interchange Format 1.0 Specification, OC-SP-ETV-BIF1.0-I04-070921, September 21, 2007, Cable Television Laboratories, Inc.
[IANA]	Internet Assigned Numbers Authority, http://www.iana.org/assignments/multicasting-addresses .
[ISO/IEC 13818-2]	ISO/IEC 13818-2:2000 Information technology -- Generic coding of moving pictures and associated audio information: Video.

[ITU-T H.264]	ITU-T Recommendation, H.264 (05/03) Advanced video coding for generic audiovisual services.
[MHA TR]	Modular Headend Architecture Technical Report, CM-TR-MHA-V01-081107, November 7, 2008, Cable Television Laboratories, Inc.
[RFC 894]	IETF RFC 894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984.
[SCTE 18]	ANSI/SCTE 18 2007, Emergency Alert Messaging for Cable.
[SMPTE 421M]	SMPTE 421M-2006 Television, VC-1 Compressed Video Bitstream Format and Decoding Process.

2.3 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>
- Internet Engineering Task Force (IETF) Secretariat, 46000 Center Oak Plaza, Sterling, VA 20166, Phone +1-571-434-3500, Fax +1-571-434-3535, <http://www.ietf.org>
- International Telecommunication Union - Telecommunication Standardization Sector (ITU-T), <http://www.itu.int/itu-t/>
- International Organization for Standardization (ISO), <http://www.iso.org>
- Society of Cable Telecommunications Engineers (SCTE), <http://www.scte.org>
- The Society of Motion Picture and Television Engineers (SMPTE), 3 Barker Ave., White Plains, NY 10601, +1-914-761-1100, <http://www.smpte.org>

3 TERMS AND DEFINITIONS

This specification uses the following terms:

CA_descriptor	Conditional Access descriptor, as defined in ISO/IEC 13818-1. Each CA_descriptor provides the PID for EMMs for a particular conditional access system.
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.
ERMI-1	A registration interface between an ERM and an EQAM defined in [ERMI]. This interface is used to register and unregister EQAM resources (i.e., QAM channels) with an ERM.
ERMI-2	A control interface between an EQAM and an ERM defined in [ERMI]. This interface is used by an ERM to request QAM channel resource from an EQAM, and by an EQAM to deliver resources to an ERM.
MAC Domain	A grouping of layer two devices that can communicate with each other without using bridging or routing. In DOCSIS, a MAC domain is the group of CMs that are using upstream and downstream channels linked together through a MAC forwarding entity.
Service Group	An HFC service group (also known as a service group) is a portion of an HFC access network used to deliver a set of services to a population of cable modems that share a common spectrum of RF channels.
Jitter	Variability in the delay of a stream of incoming packets making up a flow, such as a voice communication.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

ARP	Address Resolution Protocol
ASM	Any Source Multicast
AVC	Advanced Video CODEC
CA	Conditional Access
CAS	Conditional Access System
CAT	Conditional Access Table
CBR	Constant Bit Rate
DSM-CC	Digital Storage Media Command and Control
DTS	Decoding Timestamp
DVB	Digital Video Broadcasting
EBIF	Enhanced TV Binary Interchange Format
ECM	Entitlement Control Message
ECMG	ECM Generator
EMM	Entitlement Management Message
EMMG	EMM Generators
EQAM	Edge QAM
ERM	Edge Resource Manager
ERMI	Edge Resource Management Interfaces
ES	Elementary Stream
EISS	Element used to create ETV Integrated Signaling Streams
ETV	Enhanced Television
HD	High Definition
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Management Protocol
IP	Internet Protocol
ISO	International Organization for Standardization
MPEG	Moving Picture Experts Group
MPTS	Multiple Program Transport Stream
PAT	Program Association Table
PCR	Program Clock Reference
PID	Packet Identifier

PMT	Program Map Table
PSI	Program Specific Information
PTS	Presentation Timestamp
QAM	Quadrature Amplitude Modulator
RF	Radio Frequency
RFC	Request for Comment
RTP	Real-time Transport Protocol.
SD	Standard Definition
SDV	Switched Digital Video
SPTS	Single Program Transport Stream
SSM	Source Specific Multicast
STB	Set-top Box
TS	Transport Stream
TSDT	Transport Stream Description Table
TSID	Transport Stream Identifier
UDP	User Datagram Protocol
VBR	Variable Bit Rate
VOD	Video on Demand

5 TECHNICAL OVERVIEW

The role of the EQAM in the video-on-demand and switched-digital-video architecture is to receive an IP unicast or multicast stream containing MPEG transport stream packets, and then produce that transport stream on one or more RF outputs for transmission over the hybrid fiber-coax cable plant.

5.1 Video EQAM Reference Architecture

5.1.1 Video EQAM Block Diagram

Figure 5–1 shows a conceptual block diagram of a video EQAM. The EQAM has a management and control interface through which configuration messages, resource management traffic, and network management information pass. It also has one or more input interfaces upon which the IP-encapsulated MPEG transport streams arrive. Those input interfaces feed input transport stream processing blocks that handle de-multiplexing the individual transport streams (based on UDP destination port number or IP multicast address information) from the input interface, providing de-jittering of those transport streams, and routing them to the appropriate QAM channel processing blocks.

Each QAM Channel Processing block handles the generation of an output MPTS for a single QAM channel. It could be in Passthrough mode, in which case it accepts a single MPEG Transport Stream from the input TS processing block, performs minimal processing, and forwards it to the QAM channel for modulation and transmission. Alternatively, it could be in Multiplexing mode, in which case it accepts a number of input transport streams (SPTS and/or MPTS), selects certain programs from those inputs, can perform program number remapping and PID remapping, generates PSI information, and multiplexes the selected programs into an output MPTS, which is then forwarded to the QAM channel for modulation and transmission.

Each physical RF Port on the output of the EQAM may produce multiple QAM channels (this example shows four), although that is not a required functionality in this specification.

The EQAM may be subdivided internally into multiple QAM Processing Blocks in which each input interface and output RF port is associated with a single QAM Processing Block. In this case, the input transport streams arriving at an input interface associated with one QAM Processing Block cannot be routed to RF Ports associated with a different QAM Processing Block. This might be the case for an EQAM implemented as a series of linecards in a chassis that provides a single management and control interface.

In the case where the EQAM can route any input transport stream to any QAM Channel, the EQAM could be considered to be implemented as a single QAM processing block.

This example shows three input interfaces and four RF Ports (each producing four QAM channels) per QAM Processing Block. This is simply meant as an illustrative example.

This block diagram is intended to provide insight into the functionality provided by the video EQAM; it is not intended to mandate a particular internal architecture or implementation.




Figure 5-1 - Video EQAM Block Diagram

5.1.2 Static vs. Dynamic Sessions

The term "session" in the context of an EQAM supporting VOD and SDV refers to the mapping of a single input transport stream (SPTS or MPTS), arriving encapsulated in UDP/IP, to a single output transport stream that is modulated on a particular QAM carrier.

The video EQAM supports two mechanisms for the establishment of sessions.

The first mechanism is invoked during initial boot-up and configuration, or during a maintenance window, and involves the EQAM being pre-configured to map particular input streams to particular output streams. This mechanism is called "static port mapping," "static session setup," or "table mode." In a static port mapping, the EQAM is provided with a particular UDP destination port number upon which to expect to receive a particular input SPTS, and an output QAM channel and MPEG Program Number that is to be the destination for the program contained in the input SPTS. When packets with the specified UDP destination port number are received on any input data interface within the QAM Processing Block associated with the specified QAM channel, the EQAM extracts the MPEG-TS packets and begins the processing steps that will result in those MPEG-TS packets being transmitted out the specified QAM channel, using the specified MPEG Program Number. If no packets are received with the designated UDP destination port number, the EQAM simply does not forward any MPEG packets for that program. In a static port mapped environment, each QAM channel on the EQAM might be configured with a large number of static port maps, only a handful of which might be in use at any given time.

The second mechanism is invoked at run-time, whenever a new SDV or VOD session needs to be transported by the EQAM. An external Edge Resource Manager (ERM) explicitly signals to the EQAM the details for the session, using a session setup message and, once the session is no longer needed, signals for the EQAM to tear down the session.

There are several different types of "dynamic sessions" that can be established on an EQAM to support a number of different applications.

In support of dynamic sessions, a QAM channel provides two modes: multiplexing mode and passthrough mode. In multiplexing mode, the QAM channel can simultaneously support several dynamic sessions (and potentially even a mix of static and dynamic sessions). In passthrough mode, the QAM channel only supports a single dynamic session. When there are no sessions active on a particular QAM channel, the channel is considered to be in an "idle" state, and only MPEG null packets and PAT are produced on the output. The first session to be created brings the QAM channel out of the idle state and selects which mode it will operate in.

When in multiplexing mode, dynamic multiplexing sessions can be created that map one or all of the programs from an input transport stream into an output multiplex. The output multiplex will generally contain multiple programs associated with multiple sessions. In many cases, each session will require the EQAM to remap the PIDs from the input program(s) to new PID values in order to prevent conflicts between the various sessions. For multiplexing sessions, only the MPEG packets corresponding to the elementary streams described in the PMT for each program will be multiplexed into the output transport stream. Any other MPEG packets arriving in the input transport stream will be discarded. At regular intervals, the EQAM will generate output PMTs and a PAT for the output multiplex, based on the PATs and PMTs that it receives on the input transport streams.

When in passthrough mode, the single dynamic passthrough session maps an entire input transport stream to an output QAM with minimal processing. For the passthrough session, each MPEG packet arriving on the input transport stream is expected to be produced on the output QAM channel in the order in which it was received.

An EQAM is required to support both the static session setup as well as the dynamic session setup. However, a compliant EQAM might not support both simultaneously.

5.1.3 Input TS Processing

The EQAM supports one or more input interfaces, each of which is capable of receiving multiple MPEG-2 transport streams encapsulated in UDP/IP/Ethernet.

In order to accommodate variability in the inter-packet interval (jitter) on the input interfaces, the EQAM provides a de-jittering function on the input transport streams. The de-jittering function ensures that, as long as the UDP/IP/Ethernet frames arrive within a configurable jitter tolerance window, the EQAM will be capable of forwarding the encapsulated MPEG packets via the output multiplex within the jitter bounds expected by the MPEG receiver (see [ISO/IEC 13818-1]). The EQAM performs de-jittering of the input transport streams at the session level. For example, if a session maps a single input program to an output multiplex, the EQAM ensures that the MPEG packets for that program are de-jittered appropriately. If a session maps an entire input MPTS to an output TS or output multiplex, the EQAM ensures that the entire MPTS is de-jittered, without altering the order of the MPEG packets within the MPTS.

The EQAM will route the incoming MPEG packets to the appropriate output QAM channels based on the sessions that are currently active on the EQAM.

5.1.4 QAM Channel Processing

As described above, each QAM channel that is not in the "idle" state is operating in one of two modes: multiplexing mode or passthrough mode, based on the type of session(s) that are active for the channel.

5.1.4.1 Multiplexing Mode

In multiplexing mode, the QAM channel generates an MPEG-2 compliant transport stream that contains programs from a number of different input transport streams. Sessions can either map a single program from an input transport stream (an SPTS, or a specific program de-multiplexed from an MPTS) to the output, in which case the program is mapped to a specific output program number, or they can map an entire input transport stream to the output, in which case the input programs are carried through without program number changes. A session can ask for PID remapping so that the EQAM can guarantee that there are no PID conflicts on the output, or it can ask for no PID remapping, in which case it is assumed that the incoming program(s) will not conflict with other programs in the output multiplex. In order to facilitate the use of "no PID remapping" in cases where it is needed, the EQAM supports operator-configured "Reserved PID Ranges" that the EQAM will not use as output PIDs for remapping. Thus, if the no PID remapping sessions utilize PIDs in the Reserved PID Ranges, they will at least be assured that they will not conflict with sessions for which the EQAM is performing PID remapping. In multiplexing mode, only those PIDs that are listed in the PMTs for the input programs will be included in the output multiplex.

The EQAM generates PAT and PMTs for the output multiplex at a configured rate. The output PAT will contain the configured TSID for the QAM channel as well as the list of programs contained in the multiplex. The output PMT for each program will carry all information from the corresponding input PMT with any PID and program number changes.

The EQAM will perform PCR correction on the output programs per [ISO/IEC 13818-1].

5.1.4.2 Passthrough Mode

In passthrough mode, the QAM channel supports a single session that routes an entire input transport stream to the QAM channel with minimal processing. The QAM channel performs MPEG null packet stuffing, when necessary, in order to create an output transport stream at the QAM channel information bit rate, and it performs PCR correction in order to provide a clock reference that complies with [ISO/IEC 13818-1]. Additionally, the QAM channel modifies the TSID in the PAT packets for the transport stream in order to advertise the TSID configured for the QAM channel.

5.1.5 Redundancy

For dynamic sessions delivered to the EQAM via IP multicast, a redundancy mechanism is defined. If an EQAM supports this redundancy mechanism, it can accept a list of two or more multicast options to receive the input transport stream. Each multicast option is described by a unique source address, group address, input interface tuple, and is provided in a rank-ordered list. The EQAM utilizes the highest ranked multicast option until it experiences a session loss timeout due to lack of data. Upon session loss, the EQAM selects the next-highest ranked multicast option.

5.1.6 Summary of Session Types

The following tables summarize the different session types that are defined as in scope of this specification. These tables are intended to be informative, and as such, do not on their own imply any requirements on the EQAM beyond those that are detailed in the normative sections of this specification. In the event that there is any discrepancy between what is presented in these tables and what is defined as normative requirements, the normative requirements take precedence.

Table 5-1 - Summary of Dynamic Session Types

QAM Channel Mode	TS Type	PID Remapping	Program Selection Signaling	Input IP Addressing	EQAM Requirement	Use Case
Passthrough	SPTS or MPTS	No-remapping	All Programs	Unicast or Multicast	Required	Pre-groomed, pre-muxed, pre-encrypted, ad-inserted streams
Multiplexing	SPTS	Remapping or No-remapping	Either Single Program or All Programs	Unicast or Multicast	Required	VOD, SDV, minicarousel, video rich navigation
Multiplexing	MPTS	No-remapping	All Programs	Unicast or Multicast	Required	Client DPI applications
Multiplexing	MPTS	Remapping	Either Single Program or All Programs	Unicast or Multicast	Recommended	Broadcast streams, pre-groomed streams, MPTS demux
Multiplexing	MPTS	No-remapping	Single Program	Unicast or Multicast	Optional	Grooming of pre-muxed streams

Table 5-2 - Summary of Static Session Types

QAM Channel Mode	TS Type	PID Remapping	Program Selection Signaling	Input IP Addressing	EQAM Requirement	Use Case
Multiplexing	SPTS	Remapping	Single Program	Unicast	Required	Legacy table-based VOD

6 STREAM INPUT SPECIFICATIONS

This section specifies the video EQAM stream input Ethernet interfaces, video stream over MPEG-2 transport stream/IP/UDP encapsulation, unicast, multicast, and network de-jittering specifications.

Any reference to IP in this specification is to be interpreted as IPv4. Support for IPv6 is not in scope for this version of the specification.

6.1 Stream Input Ethernet Interface Specifications

The EQAM MUST provide Ethernet interfaces for the data plane stream inputs. The Ethernet interfaces of the EQAM SHOULD be [ISO/IEC 8802-2] /IEEE 802.2 compliant. The Ethernet interfaces of the EQAM MUST be [ISO/IEC 8802-2]/IEEE 802.3 compliant.

The data plane Ethernet interfaces of the EQAM MUST support UDP/IP, ICMP, and ARP protocols. The EQAM SHOULD support "gratuitous ARP," in which ARP responses are sent periodically as announcements to avoid the possibility of EQAM IP/MAC addresses aging out of the ARP cache and/or Forwarding Database of upstream devices.

The EQAM MUST support data plane input streams that are encapsulated in the format of MPEG-2 Transport Stream over UDP/IP (MPEG2TS/UDP/IP), as specified in the next sub-section.

6.2 MPEG-2 Transport Stream over IP Interface Specifications

The transmission of MPEG-2 transport stream over IP is standardized by the IETF in RFC 2250 using RTP over UDP. Due to the embedded timestamps in the MPEG-2 transport stream, the RTP layer is not necessary in this application, and the much simpler MPEG-2 transport stream over UDP/IP (MPEG-2TS/UDP/IP) can be used for transmission of MPEG-2 transport streams over IP network transport.

The mapping of MPEG-2 transport packets directly on UDP/IP is shown in Figure 6–1 below.




Figure 6–1 - MPEG-2 Transport Stream over UDP/IP

The MPEG-2 transport stream is split into groups of an integer number of packets, usually seven packets, which form the payload of the UDP datagram. The UDP header is 8-bytes, containing a 16-bit source UDP port number, a 16-bit destination UDP port number, a 16-bit UDP length, and a 16-bit UDP checksum. In typical UDP applications, the UDP port identifies from and to which application the message is sent. In this application, the destination UDP port for each MPEG-2 transport stream is selected by the Resource Manager and utilized by the EQAM to differentiate input streams. No "well-known" or "registered" ports, per [IANA], have been defined for this application.

The UDP datagram then forms the payload of an IP packet. The IP header includes the 32-bit source IP address, 32-bit destination IP address, and 16-bit header checksum.

The IP packet is then encapsulated in an Ethernet frame, per [RFC 894]. The Ethernet frame contains, among other fields, a source and a destination MAC address, which enable Ethernet network equipment to forward the frame to the appropriate destination.

The EQAM MUST be able to receive 188-byte packet MPEG-2 transport stream input over UDP/IP with up to seven MPEG-2 packets per UDP payload. The transport stream can be Single Program Transport Stream (SPTS) or Multiple Program Transport Stream (MPTS); each UDP datagram could contain MPEG packets from a single transport stream, which can be either SPTS or MPTS.

6.2.1 MPEG-2 Input Transport Streams

The EQAM MUST support input MPEG-2 transport streams that may carry Standard Definition (SD) or High Definition (HD) video and other data streams.

Due to the nature of motion pictures and the compression algorithms, the encoded video program can be of Variable Bit Rate (VBR). When the video program is encapsulated into a MPEG-2 transport stream, it can be Constant Bit Rate (CBR) transport stream through null packet stuffing.

The EQAM MUST support CBR transport stream input. It is assumed that all transport streams are CBR, though video programs can be CBR or VBR.

There is no limitation on the input transport stream bit rate once it is within the bandwidth of the Ethernet interface.

The EQAM MUST be able to handle MPEG-2 Transport Stream, regardless of the content being carried; including video streams encoded with either MPEG-2, AVC/H.264, or VC-1 compression (either SD or HD), or data carousel.

Examples of data carousels are Mini Carousel for Switched Digital Video (SDV) services and ETV Data Carousel.

The SDV system transmits the Mini Carousel on all SDV QAM frequencies. The Mini Carousel is delivered to all edge SDV QAMs in an MPEG program format, complete with a PAT/PMT.

For Enhanced TV (ETV) application, the enhanced application data EBIF and signaling messages are encapsulated into the MPEG-2 transport stream through DSM-CC data carousel and delivered to the EQAM.

The EQAM MUST detect the input stream bit rate and make sure it is within the provisioned bandwidth. If the input stream bit rate exceeds the assigned bandwidth, EQAM MUST report back the bandwidth violation to the ERM via ERMI-2 (as designated in [ERMI]). It is the ERM that decides how to handle this event, such as letting the higher priority stream pass through and dropping the low priority stream, enabling the trans-rating, or fixing the problem in some other way.

The EQAM MUST support input transport streams carrying MPEG-2 compliant PSI including PAT and PMT.

The EQAM MUST support input transport streams in which the PCR interval is less than 100 ms per MPEG-2 transport System specifications [ISO/IEC 13818-1].

The EQAM MUST support pre-encrypted input MPEG-2 transport streams.

6.2.2 IP Unicast and Multicast

The EQAM MUST support input MPEG-2 transport stream over UDP/IP unicast through its Ethernet interfaces.

For dynamic sessions, each unicast input stream is identified uniquely through its combined destination IP address and UDP port number. For static sessions, the destination UDP port number is the primary identifier; the destination IP address could be any of the IP addresses associated with the QAM Processing Block.

The EQAM MUST support input MPEG-2 transport stream over UDP/IP multicast through its Ethernet interfaces.

The EQAM MUST be able to join multiple multicast groups using Any Source Multicast (ASM).

The EQAM MUST be able to join multiple multicast groups using Source Specific Multicast (SSM).

The EQAM MUST support either or both PIM or IGMPv3. EQAMs with no IP routing capability MUST support IGMPv3.

The EQAM MAY be able to support multiple sessions that join the same multicast group (same source and group in the case of SSM) on different interfaces within the same input map. If the EQAM does not support this functionality, it MUST reject the session setup.

6.2.3 Private Data Streams

The EQAM MUST support private data that is carried in a PID, which is described in the PMT, along with other PIDs, to be included in the program stream. The program stream data rate is the summation of all the bit rates of all the PIDs that make up the program.

An example of the private data stream is the ETV EBIF data stream.

The EQAM is not required to support the introduction of private data streams (e.g., DVB-EIT) that are sent as separate MPEG2/UDP/IP input streams without PAT/PMT information.

6.3 Input Stream Network De-jittering Specification

In VOD or SDV session-based applications, the stream delivery delay and jitter have significant impacts on channel change latency, on-demand program request latency, and other time sensitive services.

Carrying MPEG-2 transport packets over a switched or routed IP/Ethernet network introduces jitter. The jitter is mainly due to the delay variation as the IP packets travel through Ethernet switches and IP routers. Each Ethernet switch might delay a video packet from a few milliseconds up to 50 milliseconds. The EQAM is required to correct and remove this kind of network-induced jitter.

The EQAM MUST provide de-jittering for each input MPEG-2 transport stream to handle packet delay variation up to at least 100 milliseconds. The EQAM SHOULD provide de-jittering for each input MPEG-2 transport stream to handle packet delay variation up to at least 200 milliseconds.

The EQAM de-jitter buffer MUST be configurable from five milliseconds to the maximum limit. The de-jitter buffer configuration is set for the EQAM through the [EQAM PMI] as a global (i.e., entire EQAM) configuration. The required de-jitter buffer depth is related to the data plane traffic network configuration.

This allows the de-jitter buffers to be tuned to provide the lowest possible latency based on the expected network jitter.

The EQAM MUST detect de-jitter buffer overflow for static sessions and detect de-jitter buffer underflow and overflow for dynamic sessions.

The EQAM MUST report de-jitter buffer underflow and overflow events to the network management system.

7 STREAM INPUT TO OUTPUT ROUTING

This section describes the routing from input to output of an EQAM for unicast, multicast, SPTS, and MPTS streams. There are primarily two mechanisms that an EQAM can use to determine where to route an input stream. One is static port map based routing and the second is dynamic ERMI-2 (as described in [ERMI]) signaling based routing. Both mechanisms are described in this section.

7.1 Static Port Map based Routing

For sessions that use the static port map mechanism, the EQAM uses a port map table to determine where to route the flow. The port map is a table that includes the input UDP destination port and output QAM channel and program number for each entry. This port map is configured on the EQAM via [EQAM PMI]. The EQAM reports this port map to the ERM via the UDP Map attribute in the ERMI-1 interface (as described in [ERMI]). The ERM uses this port map to select the appropriate destination UDP port that the VOD server must use for each flow. For sessions in this mode, the EQAM does not receive any ERMI-2 signaling.

For sessions that use static routing, the destination UDP port alone is used as a flow identifier. It implies that, irrespective of the destination IP address, the UDP port is mapped to a single output QAM channel and program number. It is the responsibility of the ERM to ensure that multiple flows using the same UDP port but different destination IP addresses are not forwarded to the EQAM.

For statically provisioned sessions, the EQAM MUST forward incoming unicast flows to the appropriate output QAM channel/program number according to the UDP port map that was advertised by the EQAM via the UDP Map attribute in the ERMI-1 interface.

A flow MUST be forwarded only to a single output QAM channel/program number.

This static mode is applicable only to unicast SPTS flows.

A static port mapped session is considered to be active when packets are received for that flow, and such a session is considered to have ended when no packets are received for the flow for a configured time interval. The time interval is used to determine whether a session has ended or is configured on the EQAM via the [EQAM PMI] interface.

Changes to static port maps for a QAM channel MUST be made only when no sessions (static or dynamic) are active on that QAM channel. The EQAM MAY reject any static port map changes that are made when active sessions are present on the QAM channel.

7.2 Dynamic ERMI-2 signaling based routing

For sessions using this mechanism, the EQAM receives signaling from the ERM via the ERMI-2 interface for each session setup and teardown. The signaling indicates to the EQAM where a particular flow must be routed. A unicast flow is identified by destination IP address and destination UDP port number. A multicast flow is identified by destination IP address and, optionally, the source IP address.

For Single Program Transport Streams (SPTS), the signaling includes the flow identifiers and the output QAM channel and program number. For Multiple Program Transport Streams (MPTS) in passthrough mode, the signaling includes the flow identifiers and the output QAM channel. For MPTS multiplexing mode, the signaling includes the flow identifiers, the output QAM channel, and the mapping between input and output program numbers. Both modes are described in greater detail in Section 8.

For dynamically provisioned sessions, the EQAM MUST forward incoming unicast flows to the appropriate output QAM channel/program number, as signaled by the ERM via the ERMI-2 interface. The ERMI-2 signaling includes the destination IP address and destination UDP port to identify the incoming unicast flow. A unicast flow MUST be forwarded only to a single output QAM channel/program number, as signaled via the ERMI-2 interface. If the EQAM receives multiple messages from the ERM to forward the same unicast flow to different output ports/program numbers, the EQAM MUST reject subsequent requests for the same flow.

The EQAM MUST support the delivery of video streams arriving in multicast datagrams to any or all output ports, based on signaling via the ERMI-2 interface. The ERMI-2 signaling includes the destination group address (and optionally, the source address) of the multicast stream to identify the incoming flow. For multicast streams, the EQAM does not use the UDP port to distinguish between flows. In the case of multicast, it is expected behavior to receive multiple messages from the ERM to forward the same multicast flow to different output ports/program numbers.

7.3 Simultaneous handling of different session types

This section describes the operation of the EQAM when it handles different session types on the same device. There are four possible modes of operation, all of which are described in detail:

- All static sessions on EQAM,
- All dynamic sessions on EQAM,
- Mix of static and dynamic sessions, but each QAM channel has only static or dynamic sessions, not a mix, or
- Mix of static and dynamic sessions within a QAM channel.

The EQAM MUST support the use of static port map based routing on all QAM channels on the device. The static port maps will be configured on the EQAM via the [EQAM PMI] interface. The EQAM advertises the static port map to the ERM via the ERMI-1 interface. The EQAM MAY reject any dynamic session signaling requests for any QAM channels when in this mode.

The EQAM MUST support the use of dynamic ERMI-2 signaling based routing on all QAM channels on the device.

The EQAM SHOULD support the use of static port map or dynamic ERMI-2 signaling based routing on a single device where each QAM channel would support all static or dynamic sessions. For output ports that are to use the static port map, the mapping is configured via the [EQAM PMI] interface. Output ports for which static port maps have not been configured are to be in dynamic mode. An EQAM that supports this capability MUST advertise configured static ports to the ERM via the ERMI-1 interface. The EQAM MAY reject any dynamic session signaling requests for QAM channels that have static port maps configured. Since static port maps do not include destination IP address, the EQAM MUST advertise dynamic port maps that do not conflict with static port maps of other QAM channels.

The EQAM MAY support the simultaneous use of both session types on the same QAM channel. The EQAM MUST advertise such capability via the Mixed Mode capability bit in the ERMI-1 interface. The static port maps (if any) are configured on the EQAM via the [EQAM PMI] interface. The EQAM MUST advertise any static port maps to the ERM via the ERMI-1 interface. In this mode, the algorithm used by the ERM to select a static or dynamic approach for each session is outside the scope of this specification. One possible use case is when the ERM always uses static sessions for VOD and dynamic sessions for SDV applications.

An EQAM that is capable of handling static and dynamic sessions on the same QAM channel MUST advertise non-conflicting UDP ports for the static ports and the dynamic UDP ports via the UDP Map attribute in the ERMI-1 interface.

The ERM must ensure that the output program numbers for different sessions do not conflict. If an EQAM receives a session request with an output program number that is in conflict with another session that is already active on the same QAM channel, the EQAM MUST reject the request. This applies to both static and dynamic sessions.

If statically configured and dynamically signaled sessions are supported by the EQAM on the same output QAM channel, then the EQAM MUST handle any combination of static and dynamic sessions within a single QAM channel. The EQAM can have all dynamically configured sessions, all statically configured sessions, or a mix of both types. Some of the typical use cases may include:

- a mix of VOD and SDV sessions, all being dynamically signaled,
- all VOD sessions being statically port mapped, or
- a mix of VOD sessions being statically port mapped and SDV sessions being dynamically signaled.

Packets received on flows other than those configured in the static port map or signaled via the ERMI-2 signaling interface MUST be dropped by the EQAM. Appropriate errors MAY be signaled via the [EQAM PMI] interface to indicate receipt of packets on unknown flows. A unicast flow is determined to be unknown if the UDP port doesn't belong to the static UDP map advertised via ERMI-1 interface and the UDP port has not been signaled via ERMI-2 interface at the time of session setup. An unknown multicast flow is one that has a multicast group address (and source address in case of source specific multicast) that has not been signaled via the ERMI-2 interface as part of session setup.

7.4 MPTS Input Routing

The EQAM routes MPTS streams to the output based on dynamic signaling received via the ERMI-2 interface. The stream may be unicast or multicast. All of the routing rules defined in Section 7.2 apply to MPTS streams also. Static port map based routing is not defined for MPTS streams.

7.5 Input to Output Latency

The EQAM must introduce minimal latency when routing the video flows to the output. Low latency is preferred since it can enhance the user experience, especially at session start and during trick modes. The latency introduced by the EQAM can be attributed to two primary causes. First is the jitter buffer latency. This latency is dependent on the jitter buffer configured for the EQAM. The second latency is the processing latency of the EQAM.

The [EQAM PMI] provides the ability to configure de-jitter buffering, although implementations are only required to support a maximum jitter of 100 msec (see Section 6.3). The EQAM should be configured to not delay streaming video by more than is required to accommodate expected delay variation conditions.

The processing latency of the EQAM will depend on a lot of factors, including multicast stream join latencies and PCR and PSI repetition rates. Hence, processing latency is not specified in this specification, but the intent is that EQAM products must minimize this latency to the extent possible.

8 QAM CHANNEL PROCESSING

This section describes the processing of the output MPTS of an EQAM with input SPTS and/or MPTS streams. There are primarily two processing modes of creating output MPTS for a QAM channel with input streams:

- Multiplexing mode, and
- Passthrough mode.

These two modes are mutually exclusive per QAM channel output stream and MUST be supported by the EQAM.

When a QAM channel is operational (i.e., the EQAM has completed provisioning and configuration and the QAM channel has an assigned TSID) with no sessions (either static or dynamic) assigned to it, it is in an "idle" state and can accept a dynamic session setup that selects the mode for the QAM channel. The EQAM MUST generate a PAT containing the configured TSID for each QAM channel that is in the idle state.

In multiplexing mode, the QAM channel can simultaneously support several dynamic sessions (and potentially even a mix of static and dynamic sessions). In passthrough mode, the QAM channel only supports a single dynamic session. The first session to be created brings the QAM channel out of the idle state and selects which mode it will operate in.

For a QAM channel in Multiplexing mode (either via active dynamic sessions or configured static sessions), the EQAM MUST reject any dynamic session setup requests that attempt to create Passthrough mode sessions for the QAM channel. For a QAM channel in Passthrough mode, the EQAM MUST reject any dynamic session setup requests that attempt to create Multiplexing mode sessions for the QAM channel. To change the mode of the QAM channel, all sessions of the current mode must first be torn down and the QAM channel possibly re-provisioned to remove any static sessions. This transitions the QAM channel to the idle state where the EQAM will accept a new dynamic session of either type.

8.1 Multiplexing Mode

The EQAM MUST support multiplexing mode.

Three types of sessions can be multiplexed into an output MPTS: 1) Input SPTS, 2) Single Program demux from input MPTS, and 3) Entire "pre-conditioned" MPTS.

The EQAM input transport streams can be SPTS CBR streams or MPTS CBR streams. The selected input streams, sometimes called groomed streams, are routed to the EQAM QAM channel processing engine to create an output MPTS CBR transport stream for a QAM channel through multiplexing. The QAM channel performs the digital QAM modulation and cable RF up-conversion.

8.1.1 General Multiplexing Requirements

This subsection describes the general output MPTS multiplexing specifications.

As indicated in Section 10.1, the EQAM output MPTS transport streams comply with all normative elements of [ISO/IEC 13818-1].

The EQAM MUST support multiplexing of MPEG-2 transport streams with a constant bit rate up to the QAM channel information bit rate.

The EQAM MUST simultaneously support multiplexing of constant bit rate standard definition and high definition video programs in one output MPTS.

The EQAM MUST NOT reorder MPEG ES packets within the context of a session. In the case of an SPTS input or a single program demuxed from an MPTS input, this implies that the ES packet ordering within the program is maintained. In the case of a pre-conditioned MPTS input, this implies that the ES packet ordering for the entire input multiplex is maintained.

The MPTS multiplexing SHOULD support the ability to filter specific PIDs from an input SPTS for static sessions, but the filtering configuration is left to proprietary means.

8.1.1.1 Output MPTS TSID

The EQAM multiplex output stream MUST contain the configured TSID for the corresponding QAM channel in the Program Association Table (PAT).

The TSID for output MPTS is configured via the EQAM configuration file, and is intended to be unique per head-end system; this unique TSID is needed for service group auto-discovery and validation.

8.1.1.2 MPTS Multiplexing with Specified Bit Rate

The multiplexing process of the EQAM generates a CBR MPTS stream for the QAM modulation channel that has a bit rate that matches the configured output channel QAM information bit rate (see Section 10.1). The information bit rate is defined by [ITU-T J.83], Annex A, B, or C, based on the QAM modulation mode and symbol rate.

Normally, the total bandwidth of the selected input programs for the multiplexing will be smaller than the bit rate of the output MPTS. In the multiplexing process, null packets MUST be used as stuffing packets to ensure the output MPTS is a constant bit rate stream at the configured output channel QAM information bit rate.

When the EQAM detects that the result of the multiplexing operation would exceed the bit rate of the output QAM channel, the EQAM will drop packets in an implementation specific manner to guarantee the constant bit rate of the output MPTS. As required in Section 12.3, these events are reported back to the ERM through ERMI-2. The ERM decides how to handle the event, e.g., by dropping a session according to session priority or taking another action.

8.1.1.3 MPTS Multiplexing with Timing Correction

The EQAM maintains timebase synchronization between audio and video elements and also data within a program in the multiplexing.

The EQAM MUST perform PCR correction with PCR time re-stamping per [ISO/IEC 13818-1]. The EQAM MAY also re-stamp the PTS and DTS time stamps of the individual elementary streams.

The output MPTS PCR tolerance MUST be within +/-500ns per [ISO/IEC 13818-1] MPEG-2 TS System Specifications.

8.1.1.4 PSI in MPTS Multiplexing

The PSI describes all information associated with programs within a TS. The PAT lists the programs by number contained in the TS and describes how to find the PMT for the individual programs, while the PMT identifies the individual elementary streams present per program.

The EQAM MUST process the input stream PSI tables and re-generate the output PSI tables based on the programs or streams in the output MPTS.

The EQAM MUST generate PAT per QAM channel and one PMT per program (a TS PAT section can declare a maximum of 42 programs).

The regenerated PSI MUST comply with [ISO/IEC 13818-1] and [SCTE 54] PSI constraints.

The output MPTS PSI table's repetition rates are configurable through [EQAM PMI] interface.

The output stream PAT/PMT repetition rate for VOD programs is intended to adhere to the CableLabs VOD Content Encoding Profiles Specifications [CEP]. It is recommended that the PAT repetition rate be eight per second and the PMT repetition rate be eight per second.

The EQAM MUST insert a minimum of four PAT per second per QAM channel.

The EQAM MUST insert a minimum of four PMT per second per program.

At any time the EQAM MUST guarantee that the output PMT contain all of the information from the most recently received input PMT (aside from the PID and Program Number values that are modified). In order to guarantee that all input PMT information is reproduced on the output, it is critical that the input PMT information is updated at a rate lower than the output repetition rate.

Whenever the contents of an output PSI table are modified from the previously transmitted instance of that table, the EQAM MUST update the version number field per [ISO/IEC 13818-1].

8.1.1.5 Program Numbers in MPTS Multiplexing

At a minimum, the EQAM MUST support an output MPTS with up to 20 simultaneous programs per QAM channel.

For a static port map based routing session, the program number of the output program in the MPTS is determined by the UDP port number of the input SPTS. For a dynamic session mapping a single input program to the output, the output program number for the session is dynamically passed to the EQAM in a session setup via ERMI-2. For a dynamic session mapping all input programs of an input TS to the output, the output program numbers for the session are learned via the input TS PAT.

The EQAM will reject a session SETUP that would result in a program number conflict with an existing session (either dynamically created or statically configured). The EQAM MUST verify that there are no program number conflicts present in the output MPTS, or potentially present, due to statically configured sessions. If the EQAM discovers an output program number conflict after session setup, the latest added session that causes the program number conflict MUST NOT be included in the PAT. The conflicting session SHOULD NOT be routed to the output. This could occur, for example, in the case of a pre-conditioned input MPTS where program number remapping is disabled and the EQAM would not be made aware of the program numbers for the MPTS during session setup. The EQAM MUST report the program number conflict event back to the ERM via an ERMI-2 ANNOUNCE (described in [ERMI]).

8.1.1.6 Unique PIDs in MPTS Multiplexing

All the component streams in a program are identified by their PIDs.

The EQAM MUST support generating an output MPTS with up to eight PIDs per program. The EQAM SHOULD support generating an output MPTS with up to sixteen PIDs per program.

The EQAM MUST comply with the PID assignments per [SCTE 54] for its output MPTS multiplexing. The following table defines the well-known PIDs and PID assignment range.

Table 8–1 - Well-known PIDs and PID Range

PID Range	Description
0x0000	Well-known PID for PAT
0x0001	Well-known PID for CAT
0x0002	Well-known PID for TSDT
0x0003-0x000F	Reserved
0x0010-0x001F	Used for DVB
0x0020-0x002F	Used for ARIB
0x0030-0x1FEF	May be used for elementary streams, private data, or PMT
0x1FFB	PSIP and EA data
0x1FFE	Well-known PID for DOCSIS
0x1FFF	Well-known PID for null packet

In multiplexing mode, the EQAM MUST guarantee PID uniqueness within the output MPTS.

The EQAM MUST perform PID re-mapping, when needed, to guarantee that there is no PID conflict in the output MPTS.

The well-known PIDs, as listed in Table 8–1, MUST NOT be used by the EQAM for output PID re-mapped values.

The configuration interface [EQAM PMI] allows the definition of one or more Reserved PID Ranges. The EQAM reports the Reserved PID Ranges to the ERM via ERMI-1 (described in [ERMI]). The EQAM MUST NOT use PIDs in the Reserved PID Ranges for output PID re-mapping.

Some input programs may require no PID re-mapping, as instructed by ERMI-2 (described in [ERMI]) session setup messages; for such programs, the EQAM MUST pass the PMT PID and all PIDs described in the PMT to the output without re-mapping.

An example of a program requiring no PID remapping is the SDV Mini Carousel program with:

- MPEG Program Map Table PID set to 0x1FED (8173), or
- MPEG Mini Carousel PID set to 0x1FEE (8174).

This program would be signaled as no PID-remapping in the session setup so that the PIDs are carried through to the output unchanged, and these PIDs would be reserved via the configuration interface, so that PID remapping will not cause conflicts with this program.

When multiplexing cannot avoid PID conflict in the output MPTS, the latest added session that causes the PID conflict MUST NOT be routed to the output TS. The EQAM reports this condition back to the ERM via ERMI-2.

When multiplexing an input program into an output TS (using either static or dynamic session setups), the EQAM MUST route all PIDs described in the program number's input PMT to the multiplexed output TS.

Some input TSs may contain PIDs that are not referenced in the PAT or in any PMT. Such a PID is called a ghost PID. The EQAM MUST NOT route ghost PIDs to the multiplexed output TS.

To assist in the prevention of incorrect program content being displayed by legacy STBs that do not check for PMT version number changes, the EQAM MUST implement a round-robin PID cycling algorithm for selecting output PIDs for remapping.

8.1.1.7 Pre-encrypted Stream Multiplexing

The EQAM MUST support multiplexing of pre-encrypted input streams, as defined in Section 9.

8.1.1.8 Private Data Stream Multiplexing

There are three cases in which the EQAM is required to handle private data: 1) private descriptor data sent in the PMT itself, 2) private data sent as separate packets on the PMT PID, and 3) private data sent on a separate PID where the PID is indicated in the PMT.

Case 1: The EQAM MUST include, in the output PMT, all private descriptor data received in the input PMT.

Case 2: Except in the case of “Privacy Mode” support, the EQAM MUST forward all non-PMT packets received on the input PMT PID to the output PMT PID unmodified.

Case 3: The EQAM MUST support PID re-mapping and multiplexing of all input elementary streams signaled in the input PMT.

Support for encapsulation and delivery of additional private data received via an out-of-band mechanism is outside the scope of this specification.

Examples of the private data stream include:

- ETV Data Carousel MPEG-2 TS with ETV application EBIF Elementary Stream (ES), ETV signaling, and trigger EISS ES,
- SDV Mini Carousel Message Stream, and
- [SCTE 18] Emergency Alert System (EAS) stream.

The EQAM MUST support input private data streams on the PMT PID (both within the PMT itself and as separate packets) with bit rates in the range of 50 kbps to 2 Mbps.

8.1.2 Multiplexing with Input SPTS Streams

For VOD sessions, the input program to the EQAM is unicast SPTS, typically CBR. For SDV sessions, the input program to the EQAM is multicast SPTS, typically rate clamped in the staging process. From the resource management and multiplexing point of view, the rate clamped SPTS can be treated as CBR SPTS.

The EQAM MUST support multiplexing of input CBR MPEG-2 SPTS streams into an output MPEG-2 MPTS. The input SPTS may carry audio, video, and data, with compression of MPEG-2 [ISO/IEC 13818-2], AVC/H.264 [ITU-T H.264], or VC-1 [SMPTE 421M] typically used for the video programs.

8.1.3 Multiplexing with Single Program Demux from Input MPTS

An individual program from an input MPTS can be de-multiplexed and then re-multiplexed with other input programs into an output MPTS for QAM channel.

The EQAM SHOULD support de-multiplexing of an individual program from an input MPTS and re-multiplexing of the single program demux with other programs to create an output MPTS for a QAM channel. The EQAM MAY support such re-multiplexing of a single program demux with no PID remapping.

The other programs in the output multiplex can include programs from SPTS inputs, other single programs de-multiplexed from MPTS inputs, as well as “pre-conditioned” MPTS inputs.

The individual program in the MPTS input can be de-multiplexed and mapped to a different output program.

8.1.4 Multiplexing with Entire “Pre-conditioned” MPTS

In some applications there are pre-conditioned input MPTS streams that need to be multiplexed with additional programs or with other pre-conditioned MPTSs to generate a single output MPTS for QAM modulation. Two such applications are: 1) a pre-conditioned "ad set" MPTS for client DPI applications, or 2) a statistically multiplexed MPTS (CBR MPTS containing VBR programs).

The EQAM MUST support multiplexing pre-conditioned MPTS inputs with additional programs to generate the output MPTS for a QAM channel. The other programs in the output multiplex can be SPTS inputs, single programs demuxed from an MPTS input, or other pre-conditioned MPTS inputs.

In certain cases (e.g., the pre-conditioned ad set MPTS) the PIDs and program numbers are communicated to the client STB via an out-of-band mechanism, and so cannot be changed by the EQAM. In these cases, the entire MPTS transport is set up as a single session via ERMI-2 with the "no PID remapping" option selected and the EQAM MUST NOT alter the program numbers or PIDs. To avoid PID conflicts with other sessions for which PID remapping is enabled, the EQAM is configured via [EQAM PMI] to reserve a set of PIDs to which it will not remap. Configuration of a reserved program number space on the ERM is necessary to avoid program number conflicts, but this is out of scope of this specification.

In other cases (e.g., the statistically multiplexed MPTS), the PIDs and program numbers can be remapped. In these cases, each program in the MPTS is set up as a separate session via ERMI-2, but the sessions are linked together via a "StatMux Group."

8.2 Passthrough Mode

For some applications and systems, a pre-multiplexed MPTS has been created with a bit rate at or below the output QAM channel information bit rate and delivered to the EQAM for direct QAM modulation without re-multiplexing. Such MPEG-2 MPTS processing and delivery is called Passthrough mode.

The EQAM MUST support Passthrough mode.

The input MPTS SHOULD be checked for packet loss and that the bit rate does not exceed the QAM channel information bit rate. In order to control the TSID space, the EQAM MUST modify the passthrough MPTS PAT by inserting the configured TSID for the QAM channel. The EQAM MUST perform PCR correction on the passthrough video/audio streams.

Passthrough mode for a QAM channel is signaled to the EQAM by the ERM via a session setup message on the ERMI-2 interface.

In passthrough mode, the EQAM MUST pass all PIDs from the input TS through to the output TS, regardless of whether those PIDs are listed in the PAT or PMTs. The EQAM passes all program PIDs and any ghost PIDs that reside in the input TS on to the output TS. Passthrough mode implies a one input TS to one output TS relationship where no program number or PID remapping is applied. No input TS program number handle is used to reference input programs in passthrough mode, for all input programs are simply passed through to the output TS.

The EQAM SHOULD verify that there are no program number conflicts present in the output MPTS. If a conflict is detected, the EQAM SHOULD report the program number conflict event back to the ERM via ERMI-2 (described in [ERMI]).

9 ENCRYPTION AND ENCRYPTION INTERFACE

9.1 Video Program Encryption

The EQAM MAY support video program encryption. Examples of video program encryption methods are the proprietary Cisco PowerKEY, the proprietary Motorola MediaCipher/Privacy Mode encryption, or the DVB SimulCrypt.

9.2 Video Program Scrambling

The EQAM SHOULD support scrambling of MPEG programs using [SCTE 52], DVB Common Scrambling Algorithm, or other Single/Double/Triple-DES and AES (128bit) per ATIS IIF-DSA.

9.3 Encryption Conditional Access Control Messages

The following terms define the different modes of encryption that an EQAM MAY support:

- Clear - no encryption is performed.
- Pre-Encryption - content is encrypted prior to delivery to the EQAM.
- Non-Session - real time encryption that does not require per-session encryptor provisioning (covers fixed key and tier based encryption).
- Session - real time encryption based on a per session encryptor provisioning model.

In the case of Clear mode video stream, content is delivered to the EQAM in the 'clear' and processed as normal by the EQAM without any encryption processing applied.

Non-Session and Session modes receive the same 'clear' video stream content, apply encryption to the content, and add CAS related information that may be required to enable decryption of the content. Non-session mode can apply the same encryption to multiple video streams, whereas session mode applies unique encryption per MPEG program.

In the case of Pre-encryption mode, the video stream content is delivered already encrypted to the EQAM, along with any required CAS traffic.

A number of different CAS architectures exist, such as Cisco PowerKEY, Motorola MediaCipher, DVB SimulCrypt, and Motorola Privacy Mode. Some of these CAS systems send control information to support the encryption in the same output transport stream on which the video stream is transmitted. Other systems transmit some of this control information on out-of-band transport channels.

Different scrambling algorithms exist, depending on the CAS system deployed, including Single/Double/Triple-DES, DVB CSA, and AES (128bit) per ATIS IIF-DSA.

Entitlement Management Messages (EMMs) and Entitlement Control Messages (ECMs) provide entitlement information and carry scrambling control words, respectively, to STB devices. EMMs (created by EMM Generators, or EMMGs) can transmit multicast entitlement information to update large numbers of STBs concurrently or unicast entitlement information for individual STBs (primarily for VOD). ECMs (created by ECM Generators, or ECMGs) convey the keys (using proprietary encryption schemes) to STBs to decrypt the encrypted video stream content. The STBs need information in the EMMs to enable the ECM content to be accessed.

The EMMs are processed by the STBs and provide sufficient information (along with any smartcard or internal STB credentials) to enable the processing of ECMs. ECMs are associated with one or more video streams, depending on session or tier based encryption. ECMs are typically sent multiple times per second and MAY also be updated frequently in order to transmit updated scrambling control words. When EMMs are transmitted on the same transport stream as the video stream, they are discovered by the STB through the use of the Conditional Access Table (CAT), which contains a CA_descriptor and a corresponding EMM PID for each Conditional Access system present in the transport stream.

ECMs are associated with program content through the use of CA_descriptors added to the Program Map Table (PMT) that accompanies every video stream. All such video streams are encrypted using the same control word. The ECM is associated with a video stream by referencing the ECM in a CA_descriptor added to the first loop of the PMT table. Per MPEG standards, it is also legal to place the same CA_descriptor in all the individual elementary stream descriptor loops.

Depending on the CAS architecture, such as DVB SimulCrypt, it is also possible to have multiple ECMs associated with a single video stream, and each ECM associated to a different vendor's CAS system. In this setup, the video stream(s) are scrambled using a control word that is supplied to different ECMGs. Each ECMG creates a unique ECM message that is sent with video stream content. The associated PMT is updated to include unique CA_descriptors for each CA system. An STB receives the correct ECM messages by extracting the ECM PID associated with the CA System it needs.

The EQAM MUST support clear mode.

The EQAM MUST support pre-encrypted content.

The EQAM MAY support session encryption.

The EQAM MAY support non-session encryption.

When dealing with pre-encrypted content, the EQAM MUST process incoming PMTs and identify any ECM PIDs required by decoding any CA_descriptors that may be present in the PMT. The ECMs PIDs MUST be passed along with the content. If the EQAM performs PID remapping, the ECM PID must be remapped and the corresponding CA_descriptor MUST be updated accordingly.

The EQAM MUST process any CAT tables that may be present in the input stream to identify EMM PIDs that MUST be carried with the video stream content.

The EQAM MUST update the CAT table if multiple input streams contain different EMM CA Systems. Each CA System MUST have a unique CA_descriptor.

Any ECM messages MUST be carried with the associated video content through the EQAM, meaning they MUST also pass through any de-jitter buffers to ensure the ECMs remain in sync with the video stream.

When dealing with Session or Non-Session encryption:

The EQAM MUST insert ECMs into the output transport stream along with the video stream content.

The EQAM MUST insert the correct CA_descriptor in the PMT for the particular video streams.

The EQAM MAY insert a CA_descriptor into the PMT first loop for program/service level scrambling.

The EQAM MAY insert a CA_descriptor into the PMT ES loop for elementary stream scrambling.

If the CAS system requires it, the EQAM MUST create a CAT table and include CA_descriptors for any EMM streams it MUST transmit.

If the CAS system requires it, the EQAM MUST receive an EMM stream and transmit it as MPEG2-TS packets using the EMM PID signaled in the CAT.

The EQAM MUST support the required ECM interface to transmit CA related information and receive ECM messages in response.

The EQAM MUST transmit the received ECM messages at the signaled rate.

The EQAM MUST perform PID remapping of any EMMs and ECMs destined for a single transport stream to ensure PID uniqueness on the output (unless signaled otherwise).

10 STREAM OUTPUT SPECIFICATION

This section defines the EQAM output MPEG-2 transport stream general requirements, bit rate, and operational requirements.

10.1 Output MPEG-2 Transport Stream General Requirements

The EQAM MUST generate an output stream as an MPEG2 MPTS over QAM modulated RF carrier.

The MPTS MUST be CBR at a bit rate that matches the QAM channel information bit rate.

The EQAM MUST direct an MPTS output stream to the appropriate QAM channel.

The EQAM output MPTS MUST comply with [ISO/IEC 13818-1] MPEG-2 system specifications. EQAMs intended for North American operation MUST additionally comply with [SCTE 54].

Output stream MUST contain the configured TSID for the corresponding QAM channel in the Program Association Table (PAT). The TSID is used for service group auto discovery and validation by the STB, and as an identifier for the QAM channel in the headend architecture. To ensure proper operation, each QAM channel is to be configured with a unique TSID.

10.2 Output MPEG-2 Transport Stream Operational Requirements

10.2.1 Output Stream Set up Time

Output Stream activation execution Time MUST support the EQAM input provisioning messages flow. Per SPTS, there is no delay on Output Stream Activation Command to preserve system integrity.

11 REDUNDANCY SPECIFICATIONS

EQAM products are in constant evolution, integrating more and more applications and functionalities, as well as increasing QAM channel density. EQAM products should support some form of redundancy to ensure high service availability to the end-user.

As such, redundancy mechanisms should be considered for the following system components:

- Reference Clocks/DTI,
- Internal processing,
- Management and Control,
- Ethernet input ports,
- Transport stream, and
- RF output ports.

The first three elements (reference clocks/DTI, internal processing, management and control) are systems components that can be momentarily non-functional without disturbing video stream flow.

The last three elements (Ethernet input ports, Transport stream, RF output ports) are key system components that are critical in video stream flow, and therefore need to be specified in the standard.

11.1 Redundancy Introduction

11.1.1 Redundancy: Efficiency

The EQAM redundancy efficiency is dependent on several system performance capabilities:

- Failure detection delay - Time required to detect failure.
- Failure recovery time - Time to fix the failure.
- Hardware capability - Availability of duplicate hardware components used for failover of critical system components.

11.1.2 Redundancy: Quality of Service Types

Two types of redundancy scenarios can be recognized:

- Movement of an affected session from a failed QAM channel to a working QAM channel at a different channel frequency through the use of signaling, and
- Substitution of a failed QAM channel with a working QAM channel where the session is maintained at the same channel frequency.

These two scenarios should be subject to the same general requirements.

The following defines different EQAM redundancy types that can be used to define different quality of service. Ideally, redundancy support should be completely user transparent.

- Type I - User transparent, with no loss of service.
- Type II - Not user transparent, with no loss of service.
- Type III - Not user transparent, with momentary loss of service.

The EQAM supporting redundancy must at least support Type III.

Definition: “Loss of service” means loss of set-top box (STB) synchronization.

11.1.3 Redundancy: Failover transition

It is recommended during system failure and upon system restoration to avoid exceeding the allowed disturbance on the RF output port in order to minimize the end-user video experience. Set-top box should stay synchronized at all layers.

11.1.4 Redundancy: Input Interface basic rules

The EQAM uses standard input interfaces, such as Gigabit Ethernet or 10Gigabit Ethernet interfaces. The input interfaces support a variety of traffic (management and control, data, and different types of MPEG TS packets). Different redundancy approaches are possible, but as a basic rule, it is recommended that the input interface be backed up without ERM intervention.

The redundant input interface can be active on request only (hot-warm) or can be active at all times (hot-hot).

11.1.5 Redundancy: Management and Control

The Management and Control plane (such as SNMP or ERMI interface) can use the data interfaces instead of separate external lower speed management and control interfaces.

Under that scenario, the same redundancy capabilities explained in Section 11.1.4 are applicable.

11.2 Redundancy General Requirements

The EQAM MAY support redundancy.

If the EQAM offers redundancy support, it MUST comply with the Redundancy requirements in this section. The EQAM MAY support additional redundancy mechanisms above and beyond what is specified in this section.

The maximum EQAM recovery time SHOULD be less than one second.

11.2.1 Input Interface Redundancy

The EQAM MUST support 1: N Ethernet input interface redundancy ($N \geq 1$).

The EQAM MUST detect input interface failure and manage failover from preferred interface to backup interface without ERM intervention.

The input interface failure parameters (such as no traffic detection) MUST be programmable.

The EQAM MUST report failure and redirection via the ERM Interface and the management interface.

11.2.2 Input TS Source Redundancy

Transport Stream (TS) is defined as SPTS or MPTS transport stream.

The EQAM MUST support 1:N multicast group redundancy ($N \geq 1$).

The EQAM MUST be able to select between the TS options using the appropriate address information. The EQAM MUST join the multicast group associated with the selected TS option. The multicast group/source address/input interface combinations are provided by the ERM via ERMI-2.

The EQAM MUST support at least two TS options.

The EQAM MUST be capable of receiving multiple TS options on the same input interface.

For "hot-hot" redundancy, the EQAM MUST declare TS failure when the Hot-Hot Multicast Session Loss timeout expires for that TS. The "hot-hot" redundancy implies simultaneous JOINS on any combination of multicast group/source/physical interface. For "hot-warm" redundancy, the EQAM MUST declare TS failure when the Hot-Warm Multicast Session Loss timeout expires for the Primary TS. The "hot-warm" redundancy implies one JOIN at a time for one multicast group/source/physical interface combination, holding the backup(s) multicast group/source/physical interface in reserve for fail-over. Upon TS failure, the EQAM MUST send an ANNOUNCE message indicating failover to redundant source. If the EQAM has detected that no TS options are available for the session, the EQAM MUST leave any multicast groups for the session and send an ANNOUNCE message indicating server resource unavailable.

We would thus have the following sequence of events:

1. Monitor traffic on the selected TS:
 - a. If there is traffic, set all of the TS options as "not tried yet";
 - b. If the appropriate Multicast Session Loss timeout is triggered, set the selected TS as "tried" and go to step 2;
2. The EQAM sends an ANNOUNCE message to the ERM to signal failover to the redundant source; go to step 3;
3. If there are any TS options in the list that are "not tried yet," switch* to the highest ranked of those options, set it as "tried," and go to step 1; otherwise go to step 4;
4. If all TS options have been marked as "tried," the EQAM sends an ANNOUNCE message to the ERM to signal server resource unavailable, leaves all multicast groups, and waits for the ERM to teardown the session.

*In the case of "hot-warm" redundancy, the EQAM will leave the failed group and join the multicast group for the newly selected TS option.

If the EQAM supports monitoring multiple TS options simultaneously, it MAY bypass the step of selecting and waiting for a Multicast Session Loss timeout on a TS option that it has already detected as failed. If the EQAM supports monitoring multiple TS options simultaneously, it MAY automatically revert to selecting a higher ranked TS option whenever one becomes available.

For both "hot-hot" and "hot-warm" there is no restriction on different combinations of multicast group/source/physical interface.

The EQAM MUST detect TS failure on the currently selected TS and switch to an alternate TS without ERM intervention.

The EQAM MUST support the “hot-warm” redundancy scheme. The EQAM SHOULD support the “hot-hot” redundancy scheme.

11.2.3 Output TS Redundancy

The EQAM MUST detect RF output signal failure.

The EQAM MUST isolate the RF output without ERM intervention if an RF output signal failure is detected.

The EQAM SHOULD activate a backup of the RF output without ERM intervention if an RF output signal failure is detected.

The EQAM SHOULD report the RF output signal failure and backup action to the management plane.

12 INPUT AND OUTPUT MONITORING

12.1 Input Transport Stream Monitoring

The EQAM MUST detect and report GigE link down via the [EQAM PMI] interface. The EQAM MUST also send an update of the Input Map via ERMI-1 (in [ERMI]) to the ERM to indicate the GigE link failure.

PID conflicts may arise for streams that use the “no remap” option. These may be detected only after session setup has completed and the actual stream is received by the EQAM. Such PID conflicts MUST be reported by the EQAM via the ANNOUNCE message in ERMI-2.

For dynamically configured sessions, the EQAM MUST also detect and report input session loss. A unicast session is determined to be lost when no packets are received for that flow for the Unicast Session Loss Time Interval. A default time interval of 2 seconds MUST be supported by the EQAM. The time interval SHOULD be configurable between 0.5 seconds and 6 seconds. For unicast flows, a session is identified by destination IP address and destination UDP port. A unicast session loss is reported to the ERM using an ANNOUNCE message via the ERMI-2 (in [ERMI]) interface. The unicast session loss timer is started when the SETUP response is sent, and is restarted whenever a packet is received for the session.

A multicast session is determined to be lost when no packets are received for that flow for the Multicast Session Loss Time Interval. The time interval SHOULD be configurable between 30 msec and 6 seconds. For multicast flows, a session is identified by source and destination IP address. Even when redundant sources may be available for multicast flows, the EQAM MUST report input session loss via the [EQAM PMI] interface. The EQAM MUST report to the ERM, via the ANNOUNCE message in ERMI-2 interface, that it is unable to join multicast source only after it has tried all available secondary multicast sources. The multicast session loss timer is started when the multicast group is joined, and is restarted whenever a packet is received for the multicast group.

For dynamic sessions, the EQAM MUST, via the [EQAM PMI] interface, provide indication of buffer underflows on an input transport stream basis.

The EQAM MAY detect and report packets received on flows other than those configured in the static port map or signaled via the ERMI-2 signaling interface. Appropriate errors are signaled via the [EQAM PMI] interface to indicate receipt of packets on unknown flows. A unicast flow is determined to be unknown if the UDP port doesn't belong to the static UDP map advertised via ERMI-1 interface and the UDP port has not been signaled via ERMI-2 interface at the time of session setup. An unknown multicast flow is one in which the multicast group address (and source address in case of source specific multicast) has not been signaled via ERMI-2 interface as part of session setup.

The EQAM SHOULD report via the [EQAM PMI] interface if it is unable to find a PAT or PMT in the input transport stream within five seconds of the stream being active. In the case of dynamic sessions, the EQAM MAY request the tearing down of such a session via ERMI-2.

The EQAM MUST make available, via [EQAM PMI], the program number, input and output PID value, and stream type for each incoming program on each input interface.

The EQAM SHOULD monitor the incoming stream's PCR interval and maintain a count of the number of times a PCR interval exceeds 100msec.

12.2 Input TS Program Stream Bit Rate Monitoring

For dynamic sessions, the EQAM MUST monitor the input bitrate of each transport stream and report to the ERM via the Bandwidth Exceeded Limit ANNOUNCE code in the ERMI-2 interface if the bitrate exceeds the configured rate. The threshold for error reporting and the mechanism by which the bitrate is monitored are left to the vendor.

The EQAM SHOULD measure the average and maximum bitrate for input transport streams and report the measured rates via the [EQAM PMI] interface. The mechanism by which the average and maximum bitrate are measured is left to the vendor.

The EQAM MUST report an error via [EQAM PMI] if the private data stream bitrate exceeds the bitrate signaled via ERMI-2.

12.3 Output TS Program Stream Bandwidth Utilization Monitoring

The EQAM MUST monitor the output TS bandwidth utilization on a per QAM channel basis. The EQAM SHOULD report, via [EQAM PMI], when output bandwidth required by the active sessions exceeds the total output bandwidth of the QAM channel.

If configured with a bandwidth threshold, the EQAM SHOULD report instances when output TS bandwidth utilization exceeds the configured threshold.

Appendix I Acknowledgements

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals who contributed to the development of this specification:

Contributor	Company Affiliation
Ian Wheelock	ARRIS
Tom Williams	ARRIS
Doug Jones	BigBand Networks
Victor Hou	Broadcom
Neil Buchen	Cisco
Sangeeta Ramakrishnan	Cisco
Charlie Bergren	CableLabs
Deepak Kharbanda	CableLabs
Greg White	CableLabs
Weidong Mao	Comcast
Walt Michel	Comcast
Ash Dittakavi	Cox
Muhammad Raza	Cox
Joon-Young Jung	ETRI
Woongshik You	ETRI
Noam Adika	Harmonic Video Systems
Daniel Cohen	Harmonic Video Systems
Robert Smith	LiquidStream
Ian Macfarlane	Pace Networks
Xuduan Lin	RGB Networks
Kirk Erichsen	Time Warner Cable
Chuck Hasek	Time Warner Cable
Scott Davis	Vecima Networks

We would particularly like to thank the following individuals for their significant contributions:

Xuduan Lin, for his work in drafting the Stream Input and QAM Channel Processing sections and for acting as the lead editor for the specification, Sangeeta Ramakrishnan for her work in drafting the Input to Output Routing and the Input and Output Monitoring sections, Ian Wheelock for his work in drafting the Encryption and Encryption Interface section, Robert Smith for his work in drafting the Stream Output and the Redundancy sections, and Neil Buchen, Daniel Cohen, and Victor Hou for their extensive review and comments.

We would like to thank Weidong Mao and Chuck Hasek for their guidance during the development of this specification.

Finally, we would like to thank Greg White for his work in leading the development of this specification.

