Network Working Group
Request for Comments: 4259
Category: Informational

M.-J. Montpetit
Motorola Connected Home Solutions
G. Fairhurst
University of Aberdeen
H. Clausen
TIC Systems
B. Collini-Nocker
H. Linder
University of Salzburg

November 2005

A Framework for Transmission of IP Datagrams over MPEG-2 Networks

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Abstract

This document describes an architecture for the transport of IP Datagrams over ISO MPEG-2 Transport Streams (TS). The MPEG-2 TS has been widely accepted not only for providing digital TV services but also as a subnetwork technology for building IP networks. Examples of systems using MPEG-2 include the Digital Video Broadcast (DVB) and Advanced Television Systems Committee (ATSC) Standards for Digital Television.

The document identifies the need for a set of Internet standards defining the interface between the MPEG-2 Transport Stream and an IP subnetwork. It suggests a new encapsulation method for IP datagrams and proposes protocols to perform IPv6/IPv4 address resolution, to associate IP packets with the properties of the Logical Channels provided by an MPEG-2 TS.

Table of Contents

1. Introduction

This document identifies requirements and an architecture for the transport of IP Datagrams over ISO MPEG-2 Transport Streams [ISO-MPEG]. The prime focus is the efficient and flexible delivery of IP services over those subnetworks that use the MPEG-2 Transport Stream (TS).

The architecture is designed to be compatible with services based on MPEG-2, for example the Digital Video Broadcast (DVB) architecture, the Advanced Television Systems Committee (ATSC) system [ATSC, ATSC-G], and other similar MPEG-2-based transmission systems. Such systems typically provide unidirectional (simplex) physical and link layer standards, and have been defined for a wide range of physical media (e.g., Terrestrial TV [ETSI-DVBT, ATSC-PSIP-TC], Satellite TV [ETSI-DVBS, ETSI-DVBS2, ATSC-S], Cable Transmission [ETSI-DVBC, ATSC-PSIP-TC, OPEN-CABLE], and data transmission over MPEG-2 [ETSI-MHP].

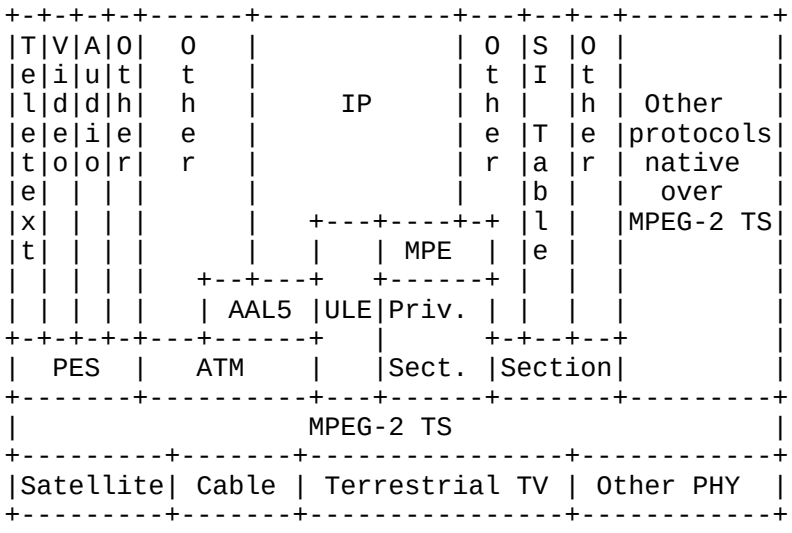

Figure 1: Overview of the MPEG-2 protocol stack

Although many MPEG-2 systems carry a mixture of data types, MPEG-2 components may be, and are, also used to build IP-only networks. Standard system components offer advantages of improved interoperability and larger deployment. However, some MPEG-2 networks do not implement all parts of a DVB / ATSC system, and may, for instance, support minimal, or no, signalling of Service Information (SI) tables.

1.1. Salient Features of the Architecture

The architecture defined in this document describes a set of protocols that support transmission of IP packets over the MPEG-2 TS. Key characteristics of these networks are that they may provide link-level broadcast capability, and that many supported applications require access to a very large number of subnetwork nodes.

Some, or all, of these protocols may also be applicable to other subnetworks, e.g., other MPEG-2 transmission networks, regenerative satellite links [ETSI-BSM], and some types of broadcast wireless links. The key goals of the architecture are to reduce complexity when using the system, while improving performance, increasing flexibility for IP services, and providing opportunities for better integration with IP services.

Since a majority of MPEG-2 transmission networks are bandwidth-limited, encapsulation protocols must therefore add minimal overhead to ensure good link efficiency while providing adequate network services. They also need to be simple to minimize processing, robust to errors and security threats, and extensible to a wide range of services.

In MPEG-2 systems, TS Logical Channels, are identified by their PID and provide multiplexing, addressing, and error reporting. The TS Logical Channel may also be used to provide Quality of Service (QoS). Mapping functions are required to relate TS Logical Channels to IP addresses, to map TS Logical Channels to IP-level QoS, and to associate IP flows with specific subnetwork capabilities. An important feature of the architecture is that these functions may be provided in a dynamic way, allowing transparent integration with other IP-layer protocols. Collectively, these will form an MPEG-2 TS Address Resolution (AR) protocol suite [IPDVB-AR].

2. Conventions Used in This Document

Adaptation Field: An optional variable-length extension field of the fixed-length TS Packet header, intended to convey clock references and timing and synchronization information as well as stuffing over an MPEG-2 Multiplex [ISO-MPEG].

ATSC: Advanced Television Systems Committee [ATSC]. A framework and a set of associated standards for the transmission of video, audio, and data using the ISO MPEG-2 standard [ISO-MPEG].

DSM-CC: Digital Storage Media Command and Control [ISO-DSMCC]. A format for transmission of data and control information defined by the ISO MPEG-2 standard that is carried in an MPEG-2 Private Section.

DVB: Digital Video Broadcast [ETSI-DVBC, ETSI-DVBRCS, ETSI-DVBS]. A framework and set of associated standards published by the European Telecommunications Standards Institute (ETSI) for the transmission of video, audio, and data, using the ISO MPEG-2 Standard [ISO-MPEG].

Encapsulator: A network device that receives PDUs and formats these into Payload Units (known here as SNDUs) for output as a stream of TS Packets.

Forward Direction: The dominant direction of data transfer over a network path. Data transfer in the forward direction is called "forward transfer". Packets travelling in the forward direction follow the forward path through the IP network.

MAC: Medium Access and Control. The link layer header of the Ethernet IEEE 802 standard of protocols, consisting of a 6B destination address, 6B source address, and 2B type field (see also NPA).

MPE: Multiprotocol Encapsulation [ETSI-DAT, ATSC-DAT, ATSC-DATG]. A scheme that encapsulates PDUs, forming a DSM-CC Table Section. Each Section is sent in a series of TS Packets using a single TS Logical Channel.

MPEG-2: A set of standards specified by the Motion Picture Experts Group (MPEG), and standardized by the International Standards Organisation (ISO) [ISO-MPEG].

NPA: Network Point of Attachment. Addresses primarily used for station (Receiver) identification within a local network (e.g., IEEE MAC address). An address may identify individual Receivers or groups of Receivers.

PAT: Program Association Table [ISO-MPEG]. An MPEG-2 PSI control table that associates program numbers with the PID value used to send the corresponding PMT. The PAT is sent using the well-known PID value of zero.

PDU: Protocol Data Unit. Examples of a PDU include Ethernet frames, IPv4 or IPv6 datagrams, and other network packets.

PES: Packetized Elementary Stream [ISO-MPEG]. A format of MPEG-2 TS packet payload usually used for video or audio information.

PID: Packet Identifier [ISO-MPEG]. A 13 bit field carried in the header of TS Packets. This is used to identify the TS Logical Channel to which a TS Packet belongs [ISO-MPEG]. The TS Packets forming the parts of a Table Section, PES, or other Payload Unit must

all carry the same PID value. The all 1s PID value indicates a Null TS Packet introduced to maintain a constant bit rate of a TS Multiplex. There is no required relationship between the PID values used for TS Logical Channels transmitted using different TS Multiplexes.

PMT: Program Map Table. An MPEG-2 PSI control table that associates the PID values used by the set of TS Logical Channels/Streams that comprise a program [ISO-MPEG]. The PID value which is used to send the PMT for a specific program is defined by an entry in the PAT.

PP: Payload Pointer [ISO-MPEG]. An optional one byte pointer that directly follows the TS Packet header. It contains the number of bytes between the end of the TS Packet header and the start of a Payload Unit. The presence of the Payload Pointer is indicated by the value of the PUSI bit in the TS Packet header. The Payload Pointer is present in DSM-CC and Table Sections; it is not present in TS Logical Channels that use the PES-format.

Private Section: A syntactic structure constructed in accordance with Table 2-30 of [ISO-MPEG]. The structure may be used to identify private information (i.e., not defined by [ISO-MPEG]) relating to one or more elementary streams, or a specific MPEG-2 program, or the entire TS. Other Standards bodies (e.g., ETSI, ATSC) have defined sets of table structures using the private_section structure. A Private Section is transmitted as a sequence of TS Packets using a TS Logical Channel. A TS Logical Channel may carry sections from more than one set of tables.

PSI: Program Specific Information [ISO-MPEG]. PSI is used to convey information about services carried in a TS Multiplex. It is carried in one of four specifically identified table section constructs [ISO-MPEG], see also SI Table.

PU: Payload Unit. A sequence of bytes sent using a TS. Examples of Payload Units include: an MPEG-2 Table Section or a ULE SNDU.

PUSI: Payload_Unit_Start_Indicator [ISO-MPEG]. A single bit flag carried in the TS Packet header. A PUSI value of zero indicates that the TS Packet does not carry the start of a new Payload Unit. A PUSI value of one indicates that the TS Packet does carry the start of a new Payload Unit. In ULE, a PUSI bit set to 1 also indicates the presence of a one byte Payload Pointer (PP).

Receiver: A piece of equipment that processes the signal from a TS Multiplex and performs filtering and forwarding of encapsulated PDUs to the network-layer service (or bridging module when operating at the link layer).

SI Table: Service Information Table [ISO-MPEG]. In this document, this term describes a table that is used to convey information about the services carried in a TS Multiplex, that has been defined by another standards body. A Table may consist of one or more Table Sections, however all sections of a particular SI Table must be carried over a single TS Logical Channel [ISO-MPEG].

SNDU: Sub-Network Data Unit. An encapsulated PDU sent as an MPEG-2 Payload Unit.

STB: Set-Top Box. A consumer equipment (Receiver) for reception of digital TV services.

Table Section: A Payload Unit carrying all or a part of an SI or PSI Table [ISO-MPEG].

TS: Transport Stream [ISO-MPEG], a method of transmission at the MPEG-2 level using TS Packets; it represents level 2 of the ISO/OSI reference model. See also TS Logical Channel and TS Multiplex.

TS Header: The 4-byte header of a TS Packet [ISO-MPEG].

TS Logical Channel: Transport Stream Logical Channel. In this document, this term identifies a channel at the MPEG-2 level [ISO-MPEG]. It exists at level 2 of the ISO/OSI reference model. All packets sent over a TS Logical Channel carry the same PID value (this value is unique within a specific TS Multiplex). According to MPEG-2, some TS Logical Channels are reserved for specific signalling. Other standards (e.g., ATSC, DVB) also reserve specific TS Logical Channels.

TS Multiplex: In this document, this term defines a set of MPEG-2 TS Logical Channels sent over a single lower layer connection. This may be a common physical link (i.e., a transmission at a specified symbol rate, FEC setting, and transmission frequency) or an encapsulation provided by another protocol layer (e.g., Ethernet, or RTP over IP). The same TS Logical Channel may be repeated over more than one TS Multiplex (possibly associated with a different PID value), for example to redistribute the same multicast content to two terrestrial TV transmission cells.

TS Packet: A fixed-length 188B unit of data sent over a TS Multiplex [ISO-MPEG]. Each TS Packet carries a 4B header, plus optional overhead including an Adaptation Field, encryption details and time stamp information to synchronize a set of related TS Logical Channels. It is also referred to as a TS_cell. Each TS Packet carries a PID value to associate it with a single TS Logical Channel.

ULE: Unidirectional Lightweight Encapsulation (ULE) [IPDVB-ULE]. A scheme that encapsulates PDUs, into SNDUs that are sent in a series of TS Packets using a single TS Logical Channel.

3. Architecture

The following sections introduce the components of the MPEG-2 Transmission Network and relate these to a networking framework.

3.1. MPEG-2 Transmission Networks

There are many possible topologies for MPEG-2 Transmission Networks. A number of example scenarios are briefly described below, and the following text relates specific functions to this set of scenarios.

A) Broadcast TV and Radio Delivery
The principal service in the Broadcast TV and Radio Delivery scenario is Digital TV and/or Radio and their associated data [MMUSIC-IMG, ETSI-IPDC]. Such networks typically contain two components: the contribution feed and the broadcast part. Contribution feeds provide communication from a typically small number of individual sites (usually at high quality) to the Hub of a broadcast network. The traffic carried on contribution feeds is typically encrypted, and is usually processed prior to being resent on the Broadcast part of the network. The Broadcast part uses a star topology centered on the Hub to reach a typically large number of down-stream Receivers. Although such networks may provide IP transmission, they do not necessarily provide access to the public Internet.

B) Broadcast Networks used as an ISP
Another scenario resembles that above, but includes the provision of IP services providing access to the public Internet. The IP traffic in this scenario is typically not related to the digital TV/Radio content, and the service may be operated by an independent operator such as unidirectional file delivery or bidirectional ISP access. The IP service must adhere to the full system specification used for the broadcast transmission, including allocation of PIDs and generation of appropriate MPEG-2 control information (e.g., DVB and ATSC SI tables).

C) Unidirectional Star IP Scenario
The Unidirectional Star IP Scenario utilizes a Hub station to provide a data network delivering a common bit stream to typically medium-sized groups of Receivers. MPEG-2 transmission technology provides the forward direction physical and link layers for this transmission; the return link (if required) is provided by other means. IP

services typically form the main proportion of the transmission traffic. Such networks do not necessarily implement the MPEG-2 control plane, i.e., PSI/SI tables.

D) Datacast Overlay
The Datacast Overlay scenario employs MPEG-2 physical and link layers to provide additional connectivity such as unidirectional multicast to supplement an existing IP-based Internet service. Examples of such a network includes IP Datacast to mobile wireless receivers [MMUSIC-IMG].

E) Point-to-Point Links
Point-to-Point connectivity may be provided using a pair of transmit and receive interfaces supporting the MPEG-2 physical and link layers. Typically, the transmission from a sender is received by only one or a small number of Receivers. Examples include the use of transmit/receive DVB-S terminals to provide satellite links between ISPs utilising BGP routing.

F) Two-Way IP Networks
Two-Way IP networks are typically satellite-based and star-based utilising a Hub station to deliver a common bit stream to medium-sized groups of receivers. A bidirectional service is provided over a common air-interface. The transmission technology in the forward direction at the physical and link layers is MPEG-2, which may also be used in the return direction. Such systems also usually include a control plane element to manage the (shared) return link capacity. A concrete example is the DVB-RCS system [ETSI-DVBRCS]. IP services typically form the main proportion of the transmission traffic.

Scenarios A-D employ unidirectional MPEG-2 Transmission Networks. For satellite-based networks, these typically have a star topology, with a central Hub providing service to large numbers of down-stream Receivers. Terrestrial networks may employ several transmission Hubs, each serving a particular coverage cell with a community of Receivers.

From an IP viewpoint, the service is typically either unidirectional multicast, or a bidirectional service in which some complementary link technology (e.g., modem, Local Multipoint Distribution Service (LMDS), General Packet Radio Service (GPRS)) is used to provide the return path from Receivers to the Internet. In this case, routing could be provided using UniDirectional Link Routing (UDLR) [RFC3077].

Note that only Scenarios A-B actually carry MPEG-2 video and audio (intended for reception by digital Set Top Boxes (STBs)) as the primary traffic. The other scenarios are IP-based data networks and need not necessarily implement the MPEG-2 control plane.

Scenarios E-F provide two-way connectivity using the MPEG-2 Transmission Network. Such networks provide direct support for bidirectional protocols above and below the IP layer.

The complete MPEG-2 transmission network may be managed by a transmission service operator. In such cases, the assignment of addresses and TS Logical Channels at Receivers are usually under the control of the service operator. Examples include a TV operator (Scenario A), or an ISP (Scenarios B-F). MPEG-2 transmission networks are also used for private networks. These typically involve a smaller number of Receivers and do not require the same level of centralized control. Examples include companies wishing to connect DVB-capable routers to form links within the Internet (Scenario B).

3.2. TS Logical Channels

An MPEG-2 Transport Multiplex offers a number of parallel channels, which are known here as TS Logical Channels. Each TS Logical Channel is uniquely identified by the Packet ID (PID) value that is carried in the header of each MPEG-2 TS Packet. The PID value is a 13 bit field; thus, the number of available channels ranges from 0 to 8191 decimal or 0x1FFF in hexadecimal, some of which are reserved for transmission of SI tables. Non-reserved TS Logical Channels may be used to carry audio [ISO-AUD], video [ISO-VID], IP packets [ISO-DSMCC, ETSI-DAT, ATSC-DAT], or other data [ISO-DSMCC, ETSI-DAT, ATSC-DAT]. The value 8191 decimal (0x1FFF) indicates a null packet that is used to maintain the physical bearer bit rate when there are no other MPEG-2 TS packets to be sent.

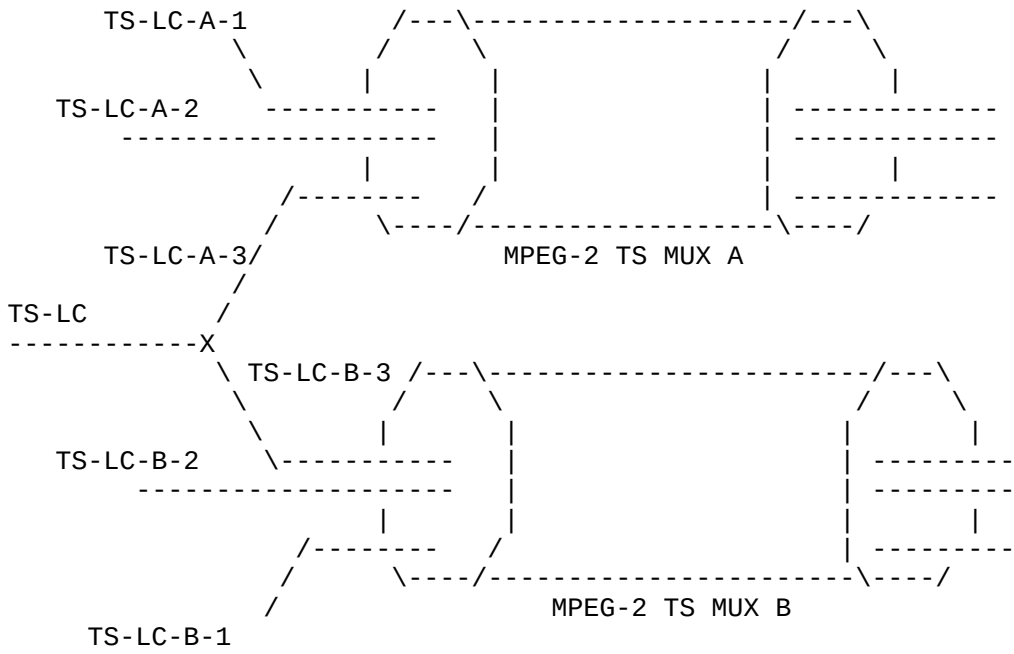

Figure 2: Example showing MPEG-2 TS Logical Channels carried Over 2 MPEG-2 TS Multiplexes.

TS Logical Channels are independently numbered on each MPEG-2 TS Multiplex (MUX). In most cases, the data sent over the TS Logical Channels will differ for different multiplexes. Figure 2 shows a set of TS Logical Channels sent using two MPEG-2 TS Multiplexes (A and B).

There are cases where the same data may be distributed over two or more multiplexes (e.g., some SI tables; multicast content that needs to be received by Receivers tuned to either MPEG-2 TS; unicast data where the Receiver may be in either/both of two potentially overlapping MPEG-2 transmission cells). In figure 2, each multiplex carries 3 MPEG-2 TS Logical Channels. These TS Logical Channels may differ (TS-LC-A-1, TS-LC-A-2, TS-LC-B-2, TS-LC-B-1), or may be common to both MPEG-2 TS Multiplexes (i.e., TS-LC-A-3 and TS-LC-B-3 carry identical content).

As can been seen, there are similarities between the way PIDs are used and the operation of virtual channels in ATM. However, unlike ATM, a PID defines a unidirectional broadcast channel and not a point-to-point link. Contrary to ATM, there is, as yet, no specified

standard interface for MPEG-2 connection setup, or for signaling mappings of IP flows to PIDs, or to set the Quality of Service, QoS, assigned to a TS Logical Channel.

3.3. Multiplexing and Re-Multiplexing

In a simple example, one or more TS Logical Channels are processed by an MPEG-2 multiplexor, resulting in a TS Multiplex. The TS Multiplex is forwarded over a physical bearer towards one or more Receivers (Figure 3).

In a more complex example, the same TS may be fed to multiple MPEG-2 multiplexors and these may, in turn, feed other MPEG-2 multiplexors (remultiplexing). Remultiplexing may occur in several places (and is common in Scenarios A and B of Section 3.1). One example is a satellite that provides on-board processing of the TS packets, multiplexing the TS Logical Channels received from one or more uplink physical bearers (TS Multiplex) to one (or more in the case of broadcast/multicast) down-link physical bearer (TS Multiplex). As part of the remultiplexing process, a remultiplexor may renumber the PID values associated with one or more TS Logical Channels to prevent clashes between input TS Logical Channels with the same PID carried on different input multiplexes. It may also modify and/or insert new SI data into the control plane.

In all cases, the final result is a "TS Multiplex" that is transmitted over the physical bearer towards the Receiver.

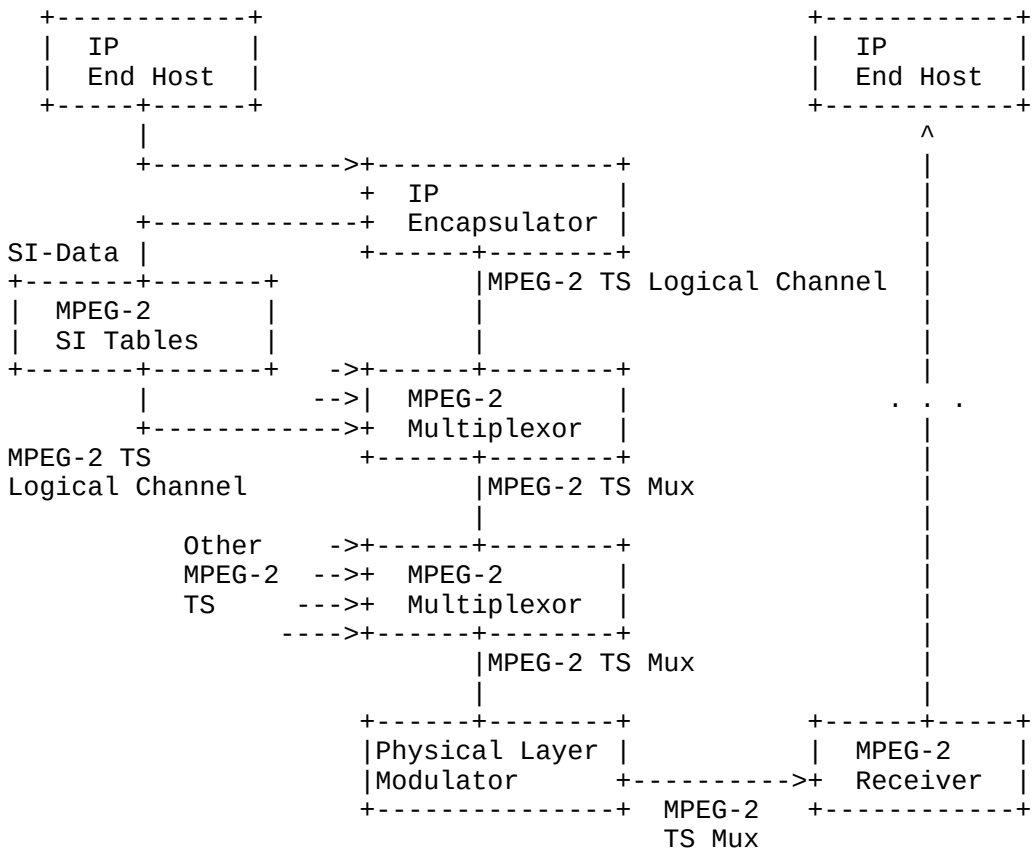

Figure 3: An example configuration for a unidirectional Service for IP transport over MPEG-2

3.4. IP Datagram Transmission

Packet data for transmission over an MPEG-2 Transport Multiplex is passed to an Encapsulator, sometimes known as a Gateway. This receives Protocol Data Units, PDUs, such as Ethernet frames or IP packets, and formats each into a Sub-Network Data Unit, SNDU, by adding an encapsulation header and trailer (see Section 4). The SNDUs are subsequently fragmented into a series of TS Packets.

To receive IP packets over an MPEG-2 TS Multiplex, a Receiver needs to identify the specific TS Multiplex (physical link) and also the TS Logical Channel (the PID value of a logical link). It is common for a number of MPEG-2 TS Logical Channels to carry SNDUs; therefore, a Receiver must filter (accept) IP packets sent with a number of PID values, and must independently reassemble each SNDU.

A Receiver that simultaneously receives from several TS Logical Channels must filter other unwanted TS Logical Channels by employing, for example, specific hardware support. Packets for one IP flow (i.e., a specific combination of IP source and destination addresses) must be sent using the same PID. It should not be assumed that all IP packets are carried on a single PID, as in some cable modem implementations, and multiple PIDs must be allowed in the architecture. Many current hardware filters limit the maximum number of active PIDs (e.g., 32), although if needed, future systems may reasonably be expected to support more.

In some cases, Receivers may need to select TS Logical Channels from a number of simultaneously active TS Multiplexes. To do this, they need multiple physical receive interfaces (e.g., radio frequency (RF) front-ends and demodulators). Some applications also envisage the concurrent reception of IP Packets over other media that may not necessarily use MPEG-2 transmission.

Bidirectional (duplex) transmission can be provided using an MPEG-2 Transmission Network by using one of a number of alternate return channel schemes [ETSI-RC]. Duplex IP paths may also be supported using non-MPEG-2 return links (e.g., in Scenarios B-D of section 3.1). One example of such an application is that of UniDirectional Link Routing, UDLR [RFC3077].

3.5. Motivation

The network layer protocols to be supported by this architecture include:

(i) IPv4 Unicast packets, destined for a single end host

(ii) IPv4 Broadcast packets, sent to all end systems in an IP network

(iii) IPv4 Multicast packets

(iv) IPv6 Unicast packets, destined for a single end host

(v) IPv6 Multicast packets

(vi) Packets with compressed IPv4 / IPv6 packet headers (e.g., [RFC2507, RFC3095])

(vii) Bridged Ethernet frames

(viii) Other network protocol packets (MPLS, potential new protocols)

The architecture will provide:

(i) Guidance on which MPEG-2 features are pre-requisites for the IP service, and identification of any optional fields that impact performance/correct operation.

(ii) Standards to provide an efficient and flexible encapsulation scheme that may be easily implemented in an Encapsulator or Receiver. The payload encapsulation requires a type field for the SNDU to indicate the type of packet and a mechanism to signal which encapsulation is used on a certain PID.

(iii) Standards to associate a particular IP address with a Network Point of Attachment (NPA) that could or may not be a MAC Address. This process resembles the IPv4 Address Resolution Protocol, ARP, or IPv6 Neighbor Discovery, ND, protocol [IPDVB-AR]. In addition, the standard will be compatible with IPv6 autoconfiguration.

(iv) Standards to associate an MPEG-2 TS interface with one or more specific TS Logical Channels (PID, TS Multiplex). Bindings are required for both unicast transmission, and multicast reception. In the case of IPv4, this must also support network broadcast. To make the schemes robust to loss and state changes within the MPEG-2 transmission network, a soft-state approach may prove desirable.

(v) Standards to associate the capabilities of an MPEG-2 TS Logical Channel with IP flows. This includes mapping of QoS functions, such as IP QoS/DSCP and RSVP, to underlying MPEG-2 TS QoS, multi-homing and mobility. This capability could be associated by the AR standard proposed above.

(vi) Guidance on Security for IP transmission over MPEG-2. The framework must permit use of IPsec and clearly identify any security issues concerning the specified protocols. The security issues need to consider two cases: unidirectional transfer (in which communication is only from the sending IP end host to the receiving IP end host) and bidirectional transfer. Consideration should also be given to security of the TS Multiplex: the need for closed user groups and the use of MPEG-2 TS encryption.

(vii) Management of the IP transmission, including standardized SNMP MIBs and error reporting procedures. The need for and scope of this is to be determined.

The specified architecture and techniques should be suited to a range of systems employing the MPEG-2 TS, and may also suit other (sub)networks offering similar transfer capabilities.

The following section, 4, describes encapsulation issues. Sections 5 and 6 describe address resolution issues for unicast and multicast, respectively.

4. Encapsulation Protocol Requirements

This section identifies requirements and provides examples of mechanisms that may be used to perform the encapsulation of IPv4/v6 unicast and multicast packets over MPEG-2 Transmission Networks.

A network device, known as an Encapsulator receives PDUs (e.g., IP Packets or Ethernet frames) and formats these into Subnetwork Data Units, SNDUs. An encapsulation (or convergence) protocol transports each SNDU over the MPEG-2 TS service and provides the appropriate mechanisms to deliver the encapsulated PDU to the Receiver IP interface.

In forming an SNDU, the encapsulation protocol typically adds header fields that carry protocol control information, such as the length of SNDU, Receiver address, multiplexing information, payload type, sequence numbers, etc. The SNDU payload is typically followed by a trailer, which carries an Integrity Check (e.g., Cyclic Redundancy Check, CRC). Some protocols also add additional control information and/or padding to or after the trailer (figure 4).

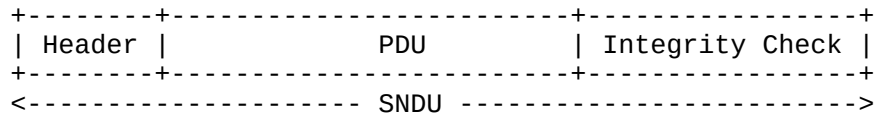

Figure 4: Encapsulation of a subnetwork PDU (e.g., IPv4 or IPv6 packet) to form an MPEG-2 Payload Unit.

Examples of existing encapsulation/convergence protocols include ATM AAL5 [ITU-AAL5] and MPEG-2 MPE [ETSI-DAT].

When required, an SNDU may be fragmented across a number of TS Packets (figure 5).

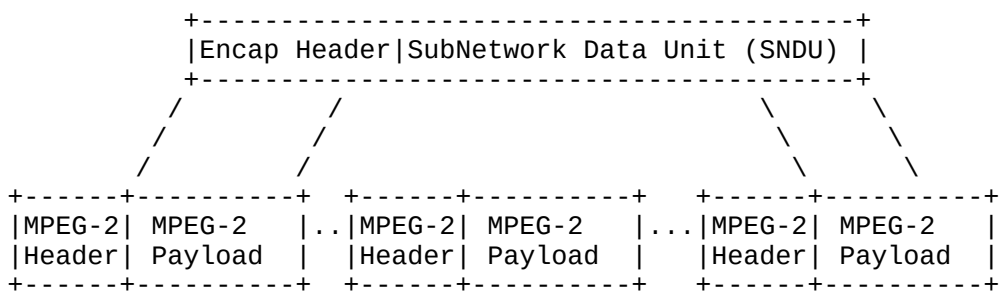

Figure 5: Encapsulation of a PDU (e.g., IP packet) into a Series of MPEG-2 TS Packets. Each TS Packet carries a header with a common Packet ID (PID) value denoting the MPEG-2 TS Logical Channel.

The DVB family of standards currently defines a mechanism for transporting an IP packet, or Ethernet frame using the Multi-Protocol Encapsulation (MPE) [ETSI-DAT]. An equivalent scheme is also supported in ATSC [ATSC-DAT, ATSC-DATG]. It allows transmission of IP packets or (by using LLC) Ethernet frames by encapsulation within a Table Section (with the format used by the control plane associated with the MPEG-2 transmission). The MPE specification includes a set of optional header components and requires decoding of the control headers. This processing is suboptimal for Internet traffic, since it incurs significant receiver processing overhead and some extra link overhead [CLC99].

The existing standards carry heritage from legacy implementations. These have reflected the limitations of technology at the time of their deployment (e.g., design decisions driven by audio/video considerations rather than IP networking requirements). IPv6, MPLS, and other network-layer protocols are not natively supported. Together, these justify the development of a new encapsulation that will be truly IP-centric. Carrying IP packets over a TS Logical Channel involves several convergence protocol functions. This section briefly describes these functions and highlights the requirements for a new encapsulation.

4.1. Payload Unit Delimitation

MPEG-2 indicates the start of a Payload Unit (PU) in a new TS Packet with a "payload_unit_start_indicator" (PUSI) [ISO-MPEG] carried in the 4B TS Packet header. The PUSI is a 1 bit flag that has normative meaning [ISO-MPEG] for TS Packets that carry PES Packets or PSI/SI data.

When the payload of a TS Packet contains PES data, a PUSI value of '1' indicates the TS Packet payload starts with the first byte of a PES Packet. A value of '0' indicates that no PES Packet starts in the TS Packet. If the PUSI is set to '1', then one, and only one, PES Packet starts in the TS Packet.

When the payload of the TS Packet contains PSI data, a PUSI value of '1' indicates the first byte of the TS Packet payload carries a Payload Pointer (PP) that indicates the position of the first byte of the Payload Unit (Table Section) being carried; if the TS Packet does not carry the first byte of a Table Section, the PUSI is set to '0', indicating that no Payload Pointer is present.

Using this PUSI bit, the start of the first Payload Unit in a TS Packet is exactly known by the Receiver, unless that TS Packet has been corrupted or lost in the transmission. In which case, the payload is discarded until the next TS Packet is received with a PUSI value of '1'.

The encapsulation should allow packing of more than one SNDU into the same TS Packet and should not limit the number of SNDUs that can be sent in a TS Packet. In addition, it should allow an IP Encapsulator to insert padding when there is an incomplete TS Packet payload. A mechanism needs to be identified to differentiate this padding from the case where another encapsulated SNDU follows.

A combination of the PUSI and a Length Indicator (see below) allows an efficient MPEG-2 convergence protocol to receive accurate delineation of packed SNDUs. The MPEG-2 standard [ISO-MPEG] does not specify how private data should use the PUSI bit.

4.2. Length Indicator

Most services using MPEG-2 include a length field in the Payload Unit header to allow the Receiver to identify the end of a Payload Unit (e.g., PES Packet, Section, or an SNDU).

When parts of more than two Payload Units are carried in the same TS Packet, only the start of the first is indicated by the Payload Pointer. Placement of a Length Indicator in the encapsulation header allows a Receiver to determine the number of bytes until the start of the next encapsulated SNDU. This placement also provides the opportunity for the Receiver to pre-allocate an appropriate-sized memory buffer to receive the reassembled SNDU.

A Length Indicator is required, and should be carried in the encapsulation header. This should support SNDUs of at least the MTU size offered by Ethernet (currently 1500 bytes). Although the IPv4

and IPv6 packet format permits an IP packet of size up to 64 KB, such packets are seldom seen on the current Internet. Since high speed links are often limited by the packet forwarding rate of routers, there has been a tendency for Internet core routers to support MTU values larger than 1500 bytes. A value of 16 KB is not uncommon in the core of the current Internet. This would seem a suitable maximum size for an MPEG-2 transmission network.

4.3. Next Level Protocol Type

Any IETF-defined encapsulation protocol should identify the payload type being transported (e.g., to differentiate IPv4, IPv6, etc). Most protocols use a type field to identify a specific process at the next higher layer that is the originator or the recipient of the payload (SNDU). This method is used by IPv4, IPv6, and also by the original Ethernet protocol (DIX). OSI uses the concept of a 'Selector' for this, (e.g., in the IEEE 802/ISO 8802 standards for CSMA/CD [LLC]; although in this case, a SNAP (subnetwork access protocol) header is also required for IP packets.

A Next Level Protocol Type field is also required if compression (e.g., Robust Header Compression [RFC3095]) is supported. No compression method has currently been defined that is directly applicable to this architecture, however the ROHC framework defines a number of header compression techniques that may yield considerable improvement in throughput on links that have a limited capacity. Since many MPEG-2 Transmission Networks are wireless, the ROHC framework will be directly applicable for many applications. The benefit of ROHC is greatest for smaller SNDUs but does imply the need for additional processing at the Receiver to expand the received compressed packets. The selected type field should contain sufficient code points to support this technique.

It is thus a requirement to include a Next Level Protocol Type field in the encapsulation header. Such a field should specify values for at least IPv4, IPv6, and must allow for other values (e.g., MAC-level bridging).

4.4. L2 Subnet Addressing

In MPEG-2, the PID carried in the TS Packet header is used to identify individual services with the help of SI tables. This was primarily intended as a unidirectional (simplex) broadcast system. A TS Packet stream carries either tables or one PES Packet stream (i.e., compressed video or audio). Individual Receivers are not addressable at this level.

IPv4 and IPv6 allocate addresses to end hosts and intermediate systems (routers). Each system (or interface) is identified by a globally assigned address. ISO uses the concept of a hierarchically structured Network Service Access Point (NSAP) address to identify an end host user process in an Internet environment.

Within a local network, a completely different set of addresses for the Network Point of Attachment (NPA) is used; frequently these NPA addresses are referred to as Medium Access Control, MAC-level addresses. In the Internet they are also called hardware addresses. Whereas network layer addresses are used for routing, NPA addresses are primarily used for Receiver identification.

Receivers may use the NPA of a received SNDU to reject unwanted unicast packets within the (software) interface driver at the Receiver, but must also perform forwarding checks based on the IP address. IP multicast and broadcast may also filter using the NPA, but Receivers must also filter unwanted packets at the network layer based on source and destination IP addresses. This does not imply that each IP address must map to a unique NPA (more than one IP address may map to the same NPA). If a separate NPA address is not required, the IP address is sufficient for both functions.

If it is required to address an individual Receiver in an MPEG-2 transport system, this can be achieved either at the network level (IP address) or via a hardware-level NPA address (MAC-address). If both addresses are used, they must be mapped in either a static or a dynamic way (e.g., by an address resolution process). A similar requirement may also exist to identify the PID and TS multiplex on which services are carried.

Using an NPA address in an MPEG-2 TS may enhance security, in that a particular PDU may be targeted for a particular Receiver by specifying the corresponding Receiver NPA address. However, this is only a weak form of security, since the NPA filtering is often reconfigurable (frequently performed in software), and may be modified by a user to allow reception of specified (or all) packets, similar to promiscuous mode operation in Ethernet. If security is required, it should be applied at another place (e.g., link encryption, authentication of address resolution, IPsec, transport level security and/or application level security).

There are also cases where the use of an NPA is required (e.g., where a system operates as a router) and, if present, this should be carried in an encapsulation header where it may be used by Receivers as a pre-filter to discard unwanted SNDUs. The addresses allocated do not need to conform to the IEEE MAC address format. There are many cases where an NPA is not required, and network layer filtering

may be used. Therefore, a new encapsulation protocol should support an optional NPA.

4.5. Integrity Check

For the IP service, the probability of undetected packet error should be small (or negligible) [RFC3819]. Therefore, there is a need for a strong integrity check (e.g., Cyclic Redundancy Check or CRC) to verify correctness of a received PDU [RFC3819]. Such checks should be sufficient to detect incorrect operation of the encapsulator and Receiver (including reassembly errors following loss/corruption of TS Packets), in addition to protecting from loss and/or corruption by the transmission network (e.g., multiplexors and links).

Mechanisms exist in MPEG-2 Transmission Networks that may assist in detecting loss (e.g., the 4-bit continuity counter included in the MPEG-2 TS Packet header).

An encapsulation must provide a strong integrity check for each IP packet. The requirements for usage of a link CRC are provided in [RFC3819]. To ease hardware implementation, this check should be carried in a trailer following the SNDU. A CRC-32 is sufficient for operation with up to a 12 KB payload, and may still provide adequate protection for larger payloads.

4.6. Identification of Scope.

The MPE section header contains information that could be used by the Receiver to identify the scope of the (MAC) address carried as an NPA, and to prevent TS Packets intended for one scope from being received by another. Similar functionality may be achieved by ensuring that only IP packets that do not have overlapping scope are sent on the same TS Logical Channel. In some cases, this may imply the use of multiple TS Logical Channels.

4.7. Extension Headers

The evolution of the Internet service may require additional functions in the future. A flexible protocol should therefore provide a way to introduce new features when required, without having to provide additional out-of-band configuration.

IPv6 introduced the concept of extension headers that carry extra information necessary/desirable for certain subnetworks. The DOCSIS cable specification also allows a MAC header to carry extension headers to build operator-specific services. Thus, it is a requirement for the new encapsulation to allow extension headers.

4.8. Summary of Requirements for Encapsulation

The main requirements for an IP-centric encapsulation include:

- support of IPv4 and IPv6 packets
- support for Ethernet encapsulated packets
- flexibility to support other IP formats and protocols (e.g., ROHC, MPLS)
- easy implementation using either hardware or software processing
- low overhead/managed overhead
- a fully specified algorithm that allows a sender to pack multiple packets per SNDU and to easily locate packet fragments
- extensibility
- compatibility with legacy deployments
- ability to allow link encryption, when required
- capability to support a full network architecture including data, control, and management planes

5. Address Resolution Functions

Address Resolution (AR) provides a mechanism that associates layer 2 (L2) information with the IP address of a system [IPDVB-AR]. Many L2 technologies employ unicast AR at the sender: an IP system wishing to send an IP packet encapsulates it and places it into an L2 frame. It then identifies the appropriate L3 adjacency (e.g., next hop router, end host) and determines the appropriate L2 adjacency (e.g., MAC address in Ethernet) to which the frame should be sent so that the packet gets across the L2 link.

The L2 addresses discovered using AR are normally recorded in a data structure known as the arp/neighbor cache. The results of previous AR requests are usually cached. Further AR protocol exchanges may be required as communication proceeds to update or re-initialize the client cache state contents (i.e., purge/refresh the contents). For stability, and to allow network topology changes and client faults, the cache contents are normally "soft state"; that is, they are aged with respect to time and old entries are removed.

In some cases (e.g., ATM, X.25, MPEG-2 and many more), AR involves finding other information than the MAC address. This includes identifying other parameters required for L2 transmission, such as channel IDs (VCs in X.25, VCIs in ATM, or PIDs in MPEG-2 TS).

Address resolution has different purposes for unicast and multicast. Multicast address resolution is not required for many L2 networks, but is required where MPEG-2 transmission networks carry IP multicast packets using more than one TS Logical Channel.

5.1. Address Resolution for MPEG-2

There are three elements to the L2 information required to perform AR before an IP packet is sent over an MPEG-2 TS. These are:

(i) A Receiver ID (e.g., a 6B MAC/NPA address).
(ii) A PID or index to find a PID.
(iii) Tuner information (e.g., Transmit Frequency of the physical layer of a satellite/broadcast link

Elements (ii) and (iii) need to be de-referenced when the MPEG-2 Transmission Network includes (re)multiplexors that renumber the PID values of the TS Logical Channels that they process. In MPEG-2 [ISO-MPEG], this dereferencing is via indexes to the information (i.e., the Program Map Table, PMT, which is itself indexed via the Program Association Table, PAT). (Note that PIDs are not intended to be end-to-end identifiers.) However, although remultiplexing is common in broadcast TV networks (scenarios A and B), many private networks do not need to employ multiplexors that renumber PIDs (see Section 3.3).

The third element (iii) allows an AR client to resolve to a different MPEG TS Multiplex. This is used when there are several channels that may be used for communication (i.e., multiple outbound/inbound links). In a mesh system, this could be used to determine connectivity. This AR information is used in two ways at a Receiver:

(i) AR resolves an IP unicast or IPv4 broadcast address to the (MPEG TS Multiplex, PID, MAC/NPA address). This allows the Receiver to set L2 filters to let traffic pass to the IP layer. This is used for unicast, and IPv4 subnet broadcast.

(ii) AR resolves an IP multicast address to the (MPEG TS Multiplex, PID, MAC/NPA address), and allows the Receiver to set L2 filters enabling traffic to pass to the IP layer. A Receiver in an MPEG-2 TS Transmission Network needs to resolve the PID value and the tuning (if present) associated with a TS Logical Channel and (at least for unicast) the destination Receiver NPA address.

A star topology MPEG-2 TS transmission network is illustrated below, with two Receivers receiving a forward broadcast channel sent by a Hub. (A mesh system has some additional cases.) The forward broadcast channel consists of a "TS Multiplex" (a single physical

bearer) allowing communication with the terminals. These communicate using a set of return channels.

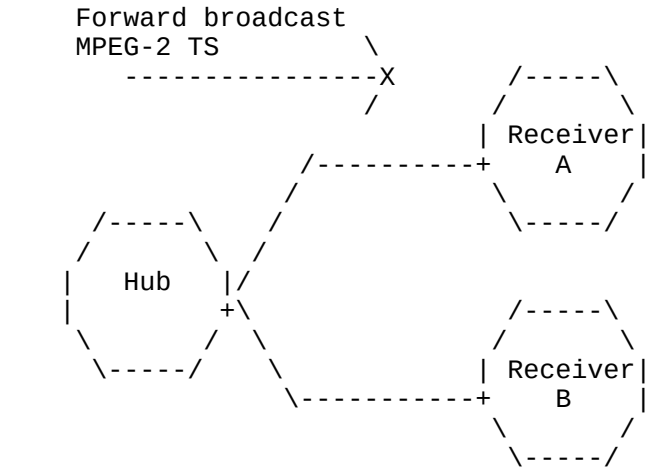

Figure 6: MPEG-2 Transmission Network with 2 Receivers

There are three possibilities for unicast AR:

(1) A system at a Receiver, A, needs to resolve an address of a system that is at the Hub;

(2) A system at a Receiver, A, needs to resolve an address that is at another Receiver, B;

(3) A host at the Hub needs to resolve an address that is at a Receiver. The sender (encapsulation gateway), uses AR to provide the MPEG TS Multiplex, PID, MAC/NPA address for sending unicast, IPv4 subnet broadcast and multicast packets.

If the Hub is an IP router, then case (1) and (2) are the same: The host at the Receiver does not know the difference. In these cases, the address to be determined is the L2 address of the device at the Hub to which the IP packet should be forwarded, which then relays the IP packet back to the forward (broadcast) MPEG-2 channel after AR (case 3).

If the Hub is an L2 bridge, then case 2 still has to relay the IP packet back to the outbound MPEG-2 channel. The AR protocol needs to resolve the specific Receiver L2 MAC address of B, but needs to send this on an L2 channel to the Hub. This requires Receivers to be informed of the L2 address of other Receivers.

An end host connected to the Hub needs to use the AR protocol to resolve the Receiver terminal MAC/NPA address. This requires the AR server at the Hub to be informed of the L2 addresses of other Receivers.

5.2. Scenarios for MPEG AR

An AR protocol may transmit AR information in three distinct ways:

(i) An MPEG-2 signalling table transmitted at the MPEG-2 level (e.g., within the control plane using a Table);

(ii) An MPEG-2 signalling table transmitted at the IP level (no implementations of this are known);

(iii) An address resolution protocol transported over IP (as in ND for IPv6)

There are three distinct cases in which AR may be used:

(i) Multiple TS-Muxes and the use of re-multiplexors, e.g., Digital Terrestrial, Satellite TV broadcast multiplexes. Many such systems employ remultiplexors that modify the PID values associated with TS Logical Channels as they pass through the MPEG-2 transmission network (as in Scenario A of Section 3.1).

(ii) Tuner configuration(s) that are fixed or controlled by some other process. In these systems, the PID value associated with a TS Logical Channel may be known by the Sender.

(iii) A service run over one TS Mux (i.e., uses only one PID, for example DOCSIS and some current DVB-RCS multicast systems). In these systems, the PID value of a TS Logical Channel may be known by the Sender.

5.2.1. Table-Based AR over MPEG-2

In current deployments of MPEG-2 networks, information about the set of MPEG-2 TS Logical Channels carried over a TS Multiplex is usually distributed via tables (service information, SI) sent using channels assigned a specific (well-known) set of PIDs. This was originally designed for audio/video distribution to STBs. This design requires access to the control plane by processing the SI table information (carried in MPEG-2 section format [ISO-DSMCC]). The scheme reflects the complexity of delivering and coordinating the various TS Logical Channels associated with a multimedia TV program.

One possible requirement to provide TS multiplex and PID information for IP services is to integrate additional information into the existing MPEG-2 tables, or to define additional tables specific to the IP service. The DVB INT and the A/92 Specification from ATSC [ATSC-A92] are examples of the realization of such a solution.

5.2.2. Table-Based AR over IP

AR information could be carried over a TS data channel (e.g., using an IP protocol similar to the Service Announcement Protocol, SAP). Implementing this independently of the SI tables would ease implementation, by allowing it to operate on systems where IP processing is performed in a software driver. It may also allow the technique to be more easily adapted to other similar delivery networks. It also is advantageous for networks that use the MPEG-2 TS, but do not necessarily support audio/video services and therefore do not need to provide interoperability with TV equipment (e.g., links used solely for connecting IP (sub)networks).

5.2.3. Query/Response AR over IP

A query/response protocol may be used at the IP level (similar to, or based on IPv6 Neighbor Advertisements of the ND protocol). The AR protocol may operate over an MPEG-2 TS Logical Channel using a previously agreed PID (e.g., configured, or communicated using a SI table). In this case, the AR could be performed by the target system itself (as in ARP and ND). This has good soft-state properties, and is very tolerant to failures. To find an address, a system sends a "query" to the network, and the "target" (or its proxy) replies.

5.3. Unicast Address Scoping

In some cases, an MPEG-2 Transmission Network may support multiple IP networks. When this is the case, it is important to recognize the context (scope) within which an address is resolved, to prevent packets from one addressed scope from leaking into other scopes.

An example of overlapping IP address assignments is the use of private unicast addresses (e.g., in IPv4, 10/8 prefix; 172.16/12 prefix; 192.168/16 prefix). These should be confined to the area to which they are addressed.

There is also a requirement for multicast address scoping (Section 6.2).

IP packets with these addresses must not be allowed to travel outside their intended scope, and may cause unexpected behaviour if allowed to do so. In addition, overlapping address assignments can arise when using level 2 NPA addresses:

(i) The NPA address must be unique within the TS Logical Channel. Universal IEEE MAC addresses used in Ethernet LANs are globally unique. If the NPA addresses are not globally unique, the same NPA address may be re-used by Receivers in different addressed areas.

(ii) The NPA broadcast address (all 1s MAC address). Traffic with this address should be confined to one addressed area.

Reception of unicast packets destined for another addressed area may lead to an increase in the rate of received packets by systems connected via the network. IP end hosts normally filter received unicast IP packets based on their assigned IP address. Reception of the additional network traffic may contribute to processing load but should not lead to unexpected protocol behaviour. However, it does introduce a potential Denial of Service (DoS) opportunity.

When the Receiver acts as an IP router, the receipt of such an IP packet may lead to unexpected protocol behaviour. This also provides a security vulnerability since arbitrary packets may be passed to the IP layer.

5.4. AR Authentication

In many AR designs, authentication has been overlooked because of the wired nature of most existing IP networks, which makes it easy to control hosts that are physically connected [RFC3819]. With wireless connections, this is changing: unauthorized hosts actually can claim L2 resources. The address resolution client (i.e., Receiver) may also need to verify the integrity and authenticity of the AR information that it receives. There are trust relationships both ways: clients need to know they have a valid server and that the resolution is valid. Servers should perform authorisation before they allow an L2 address to be used.

The MPEG-2 Transmission Network may also require access control to prevent unauthorized use of the TS Multiplex; however, this is an orthogonal issue to address resolution.

5.5. Requirements for Unicast AR over MPEG-2

The requirement for AR over MPEG-2 networks include:

(i) Use of a table-based approach to promote AR scaling. This requires definition of the frequency of update and volume of AR traffic generated.

(ii) Mechanisms to install AR information at the server (unsolicited registration).

(iii) Mechanisms to verify AR information held at the server (solicited responses). Appropriate timer values need to be defined.

(iv) An ability to purge client AR information (after IP network renumbering, etc.).

(v) Support of IP subnetwork scoping.

(vi) Appropriate security associations to authenticate the sender.

6. Multicast Support

This section addresses specific issues concerning IPv4 and IPv6 multicast [RFC1112] over MPEG-2 Transmission Networks. The primary goal of multicast support will be efficient filtering of IP multicast packets by the Receiver, and the mapping of IPv4 and IPv6 multicast addresses [RFC3171] to the associated PID value and TS Multiplex.

The design should permit a large number of active multicast groups, and should minimize the processing load at the Receiver when filtering and forwarding IP multicast packets. For example, schemes that may be easily implemented in hardware would be beneficial, since these may relieve drivers and operating systems from discarding unwanted multicast traffic [RFC3819].

Multicast mechanisms are used at more than one protocol level. The upstream router feeding the MPEG-2 Encapsulator may forward multicast traffic on the MPEG-2 TS Multiplex using a static or dynamic set of groups. When static forwarding is used, the set of IP multicast groups may also be configured or set using SNMP, Telnet, etc. A Receiver normally uses either an IP group management protocol (IGMP [RFC3376] for IPv4 or MLD [RFC2710][RFC3810] for IPv6) or a multicast routing protocol to establish tables that it uses to dynamically enable local forwarding of received groups. In a dynamic case, this

group membership information is fed back to the sender to enable it to start sending new groups and (if required) to update the tables that it produces for multicast AR.

Appropriate procedures need to identify the correct action when the same multicast group is available on more than one TS Logical Channel. This could arise when different end hosts act as senders to contribute IP packets with the same IP group destination address. The correct behaviour for SSM [RFC3569] addresses must also be considered. It may also arise when a sender duplicates the same IP group over several TS Logical Channels (or even different TS Multiplexes), and in this case a Receiver may potentially receive more than one copy of the same packet. At the IP level, the host/router may be unaware of this duplication.

6.1. Multicast AR Functions

The functions required for multicast AR may be summarized as:

(i) The Sender needs to know the L2 mapping of a multicast group.
(ii) The Receiver needs to know the L2 mapping of a multicast group.

In the Internet, multicast AR is normally a mapping function rather than a one-to-one association using a protocol. In Ethernet, the sender maps an IP address to an L2 MAC address, and the Receiver uses the same mapping to determine the L2 address to set an L2 hardware/software filter entry.

A typical sequence of actions for the dynamic case is:

L3) Populate the IP L3 membership tables at the Receiver.

L3) Receivers send/forward IP L3 membership tables to the Hub

L3) Dynamic/static forwarding at hub/upstream router of IP L3 groups

L2) Populate the IP AR tables at the encapsulator gateway (i.e., Map IP L3 mcast groups to L2 PIDs)

L2) Distribute the AR information to Receivers

L2) Set Receiver L2 multicast filters for IP groups in the membership table.

To be flexible, AR must associate a TS Logical Channel (PID) not only with a group address, but possibly also a QoS class and other appropriate MPEG-2 TS attributes. Explicit per group AR to individual L2 addresses is to be avoided.

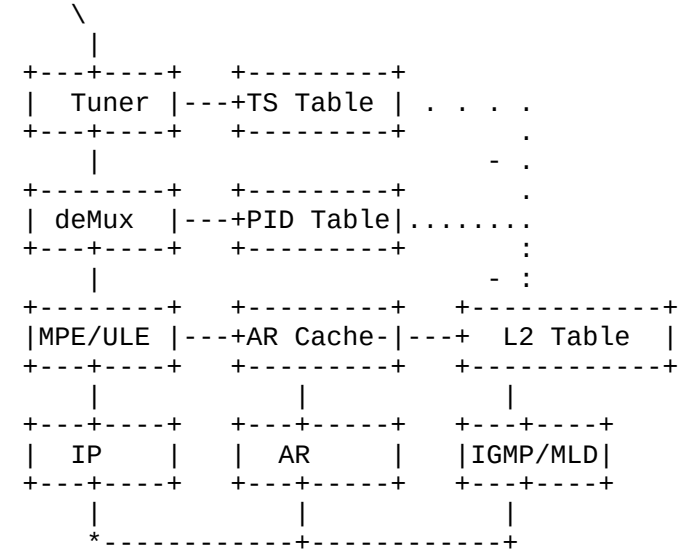

Figure 7: Receiver Processing Architecture

6.2. Multicast Address Scoping

As in unicast, it is important to recognize the context (scope) within which a multicast IP address is resolved, to prevent packets from one addressed scope leaking into other scopes.

Examples of overlapping IP multicast address assignments include:

(i) Local scope multicast addresses. These are only valid within the local area (examples for IPv4 include: 224.0.0/24; 224.0.1/24). Similar cases exist for some IPv6 multicast addresses [RFC2375].

(ii) Scoped multicast addresses [RFC2365] [RFC 2375]. Forwarding of these addresses is controlled by the scope associated with the address. The addresses are only valid with an addressed area (e.g. the 239/8 [RFC2365]).

(iii) Other non-IP protocols may also view sets of MAC multicast addresses as link-local, and may produce unexpected results if distributed across several private networks.

IP packets with these addresses must not be allowed to travel outside their intended scope (see Section 5.3). Performing multicast AR at the IP level can enable providers to offer independently scoped addresses and would need to use multiple Multicast AR servers, one per multicast domain.

6.3. Requirements for Multicast over MPEG-2

The requirements for supporting multicast include, but are not restricted to:

(i) Encapsulating multicast packets for transmission using an MPEG-2 TS.

(ii) Mapping IP multicast groups to the underlying MPEG-2 TS Logical Channel (PID) and the MPEG-2 TS Multiplex.

(iii) Providing AR information to allow a Receiver to locate an IP multicast flow within an MPEG-2 TS Multiplex.

(iv) Error Reporting.

7. Summary

This document describes the architecture for a set of protocols to perform efficient and flexible support for IP network services over networks built upon the MPEG-2 Transport Stream (TS). It also describes existing approaches. The focus is on IP networking, the mechanisms that are used, and their applicability to supporting IP unicast and multicast services.

The requirements for a new encapsulation of IPv4 and IPv6 packets is described, outlining the limitations of current methods and the need for a streamlined IP-centric approach.

The architecture also describes MPEG-2 Address Resolution (AR) procedures to allow dynamic configuration of the sender and Receiver using an MPEG-2 transmission link/network. These support IPv4 and IPv6 services in both the unicast and multicast modes. Resolution protocols will support IP packet transmission using both the Multiprotocol Encapsulation (MPE), which is currently widely deployed, and also any IETF-defined encapsulation (e.g., ULE [IPDVB-ULE]).

8. Security Considerations

When the MPEG-2 transmission network is not using a wireline network, the normal security issues relating to the use of wireless links for transport of Internet traffic should be considered [RFC3819].

End-to-end security (including confidentiality, authentication, integrity and access control) is closely associated with the end user assets that are protected. This close association cannot be ensured when providing security mechanisms only within a subnetwork (e.g., an MPEG-2 Transmission Network). Several security mechanisms that can be used end-to-end have already been deployed in the general Internet and are enjoying increasing use. Important examples include:

- Transport Layer Security (TLS), which is primarily used to protect web commerce;

- Pretty Good Privacy (PGP) and S/MIME, primarily used to protect and authenticate email and software distributions;

- Secure Shell (SSH), used to secure remote access and file transfer;

- IPsec, a general purpose encryption and authentication mechanism above IP that can be used by any IP application.

However, subnetwork security is also important [RFC3819] and should be encouraged, on the principle that it is better than the default situation, which all too often is no security at all. Users of especially vulnerable subnets (such as radio/broadcast networks and/or shared media Internet access) often have control over, at most, one endpoint - usually a client - and therefore cannot enforce the use of end-to-end mechanisms.

A related role for subnetwork security is to protect users against traffic analysis, i.e., identifying the communicating parties (by IP or MAC address) and determining their communication patterns, even when their actual contents are protected by strong end-to-end security mechanisms. (This is important for networks such as broadcast/radio, where eavesdropping is easy.)

Encryption performed at the Transport Stream (encrypting the payload of all TS-Packets with the same PID) encrypts/scrambles all parts of the SNDU, including the layer 2 MAC/NPA address. Encryption at the section level in MPE may also optionally encrypt the layer 2 MAC/NPA address in addition to the PDU data [ETSI-DAT]. In both cases, encryption of the MAC/NPA address requires a Receiver to decrypt all encrypted data, before it can then filter the PDUs with the set of

MAC/NPA addresses that it wishes to receive. This method also has the drawback that all Receivers must share a common encryption key. Encryption of the MPE MAC address is therefore not permitted in some systems (e.g., [ETSI-DVBRCS]).

Where it is possible for an attacker to inject traffic into the subnetwork control plane, subnetwork security can additionally protect the subnetwork assets. This threat must specifically be considered for the protocols used for subnetwork control functions (e.g., address resolution, management, configuration). Possible threats include theft of service and denial of service; shared media subnets tend to be especially vulnerable to such attacks [RFC3819].

Appropriate security functions must therefore be provided for IPDVB control protocols [RFC3819], particularly when the control functions are provided above the IP-layer using IP-based protocols. Internet level security mechanisms (e.g., IPsec) can mitigate such threats.

In general, End-to-End security is recommended for users of any communication path, especially when it includes a wireless/radio or broadcast link, where a range of security techniques already exist. Specification of security mechanisms at the application layer, or within the MPEG-2 transmission network, are the concerns of organisations beyond the IETF. The complexity of any such security mechanisms should be considered carefully so that it will not unduly impact IP operations.

8.1. Link Encryption

Link level encryption of IP traffic is commonly used in broadcast/radio links to supplement End-to-End security (e.g., provided by TLS, SSH, Open PGP, S/MIME, IPsec). The encryption and key exchange methods vary significantly, depending on the intended application. For example, DVB-S/DVB-RCS operated by Access Network Operators may wish to provide their customers (Internet Service Providers, ISP) with security services. Common security services are: terminal authentication and data confidentiality (for unicast and multicast) between an encapsulation gateway and Receivers. A common objective is to provide the same level of privacy as terrestrial links. An ISP may also wish to provide end-to-end security services to the end-users (based on well-known mechanisms such as IPsec).

Therefore, it is important to understand that both security solutions (Access Network Operators to ISP and ISP to end-users) may coexist.

MPE supports optional link encryption [ETSI-DAT]. A pair of bits within the MPE protocol header indicate whether encryption (scrambling) is used. For encrypted PDUs, the header bits indicate which of a pair of previously selected encryption keys is to be used.

It is recommended that any new encapsulation defined by the IETF allows Transport Stream encryption and also supports optional link level encryption/authentication of the SNDU payload. In ULE [IPDVB-ULE], this may be provided in a flexible way using Extension Headers. This requires definition of a mandatory header extension, but has the advantage that it decouples specification of the security functions from the encapsulation functions. This method also supports encryption of the NPA/MAC addresses.

9. IANA Considerations

A set of protocols that meet these requirements will require the IANA to make assignments. This document in itself, however, does not require any IANA involvement.

10. Acknowledgements

The authors wish to thank Isabel Amonou, Torsten Jaekel, Pierre Loyer, Luoma Juha-Pekka, and Rod Walsh for their detailed inputs. We also wish to acknowledge the input provided by the members of the IETF ipdvb WG.

11. References

11.1. Normative References

[ISO-MPEG] ISO/IEC DIS 13818-1:2000, "Information Technology; Generic Coding of Moving Pictures and Associated Audio Information Systems", International Standards Organisation (ISO).

[ETSI-DAT] EN 301 192, "Digital Video Broadcasting (DVB); DVB Specifications for Data Broadcasting", European Telecommunications Standards Institute (ETSI).

11.2. Informative References

[ATSC] A/53C, "ATSC Digital Television Standard", Advanced Television Systems Committee (ATSC), Doc. A/53C, 2004.

[ATSC-DAT] A/90, "ATSC Data Broadcast Standard", Advanced Television Systems Committee (ATSC), Doc. A/090, 2000.

[ATSC-DATG] A/91, "Recommended Practice: Implementation Guidelines for the ATSC Data Broadcast Standard", Advanced Television Systems Committee (ATSC), Doc. A/91, 2001.

[ATSC-A92] A/92, "Delivery of IP Multicast Sessions over ATSC Data Broadcast", Advanced Television Systems Committee (ATSC), Doc. A/92, 2002.

[ATSC-G] A/54A, "Guide to the use of the ATSC Digital Television Standard", Advanced Television Systems Committee (ATSC), Doc. A/54A, 2003.

[ATSC-PSIP-TC] A/65B, "Program and System Information Protocol for Terrestrial Broadcast and Cable", Advanced Television Systems Committee (ATSC), Doc. A/65B, 2003.

[ATSC-S] A/80, "Modulation and Coding Requirements for Digital TV (DTV) Applications over Satellite", Advanced Television Systems Committee (ATSC), Doc. A/80, 1999.

[CLC99] Clausen, H., Linder, H., and Collini-Nocker, B., "Internet over Broadcast Satellites", IEEE Commun. Mag. 1999, pp.146-151.

[ETSI-BSM] TS 102 292, "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia Services and Architectures; Functional Architecture for IP Interworking with BSM networks", European Telecommunications Standards Institute (ETSI).

[ETSI-DVBC] EN 300 800, "Digital Video Broadcasting (DVB); DVB interaction channel for Cable TV distribution systems (CATV)", European Telecommunications Standards Institute (ETSI).

[ETSI-DVBRCS] EN 301 790, "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems", European Telecommunications Standards Institute (ETSI).

[ETSI-DVBS] EN 301 421,"Digital Video Broadcasting (DVB); Modulation and Coding for DBS satellite systems at 11/12 GHz", European Telecommunications Standards Institute (ETSI).

[ETSI-DVBS2] EN 302 207, "Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services,News Gathering and Other Broadband Satellite Applications", European Telecommunications Standards Institute (ETSI).

[ETSI-DVBT] EN 300 744, "Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for digital terrestrial television (DVB-T)", European Telecommunications Standards Institute (ETSI).

[ETSI-IPDC] "IP Datacast Specification", DVB Interim Specification CNMS 1026 v1.0.0,(Work in Progress), April 2004.

[ETSI-MHP] TS 101 812, "Digital Video Broadcasting (DVB); Multimedia Home Platform (MHP) Specification", v1.2.1, European Telecommunications Standards Institute (ETSI), June 2002.

[ETSI-RC] ETS 300 802, "Digital Video Broadcasting (DVB); Network-independent protocols for DVB interactive services", European Telecommunications Standards Institute (ETSI).

[ETSI-SI] EN 300 468, "Digital Video Broadcasting (DVB); Specification for Service Information (SI) in DVB systems", European Telecommunications Standards Institute (ETSI).

[IPDVB-ULE] Fairhurst, G. and B. Collini-Nocker, "Unidirectional Lightweight Encapsulation (ULE) for transmission of IP datagrams over an MPEG-2 Transport Stream", Work in Progress, June 2005.

[IPDVB-AR] Fairhurst, G. and M-J. Montpetit, "Address Resolution for IP datagrams over MPEG-2 networks", Work in Progress, 2005.

[ISO-AUD] ISO/IEC 13818-3:1995, "Information technology; Generic coding of moving pictures and associated audio information; Part 3: Audio", International Standards Organisation (ISO).

[ISO-DSMCC] ISO/IEC IS 13818-6, "Information technology; Generic coding of moving pictures and associated audio information; Part 6: Extensions for DSM-CC", International Standards Organisation (ISO).

[ISO-VID] ISO/IEC DIS 13818-2:1998, "Information technology; Generic coding of moving pictures and associated audio information; Video", International Standards Organisation (ISO).

[ITU-AAL5] ITU-T I.363.5, "B-ISDN ATM Adaptation Layer Specification Type AAL5", International Standards Organisation (ISO), 1996.

[LLC] ISO/IEC 8802.2, "Information technology; Telecommunications and information exchange between systems; Local and metropolitan area networks; Specific requirements; Part 2: Logical Link Control", International Standards Organisation (ISO), 1998.

[MMUSIC-IMG] Nomura, Y., Walsh, R., Luoma, J-P., Ott, J., and H. Schulzrinne, "Requirements for Internet Media Guides", Work in Progress, June 2004.

[OPEN-CABLE] "Open Cable Application Platform Specification; OCAP 2.0 Profile", OC-SP-OCAP2.0-I01-020419, Cable Labs, April 2002.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.

[RFC2375] Hinden, R. and S. Deering, "IPv6 Multicast Address Assignments", RFC 2375, July 1998.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[RFC2507] Degermark, M., Nordgren, B., and S. Pink, "IP Header Compression", RFC 2507, February 1999.

[RFC3077] Duros, E., Dabbous, W., Izumiyama, H., Fujii, N., and Y. Zhang, "A Link-Layer Tunneling Mechanism for Unidirectional Links", RFC 3077, March 2001.

[RFC3095] Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed", RFC 3095, July 2001.

[RFC3171] Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 3171, August 2001.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3569] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC3819] Karn, P., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, July 2004 .

Appendix A: MPEG-2 Encapsulation Mechanisms

Transmitting packet data over an MPEG-2 transmission network requires that individual PDUs (e.g., IPv4, IPv6 packets, or bridged Ethernet Frames) are encapsulated using a convergence protocol. The following encapsulations are currently standardized for MPEG-2 transmission networks:

(i) Multi-Protocol Encapsulation (MPE).

The MPE specification of DVB [ETSI-DAT] uses private Sections for the transport of IP packets and uses encapsulation that is similar to the IEEE LAN/MAN standards [LLC]. Data packets are encapsulated in datagram sections that are compliant with the DSMCC section format for private data. Some Receivers may exploit section processing hardware to perform a first-level filtering of the packets that arrive at the Receiver.

This encapsulation makes use of a MAC-level Network Point of Attachment address. The address format conforms to the ISO/IEEE standards for LAN/MAN [LLC]. The 48-bit MAC address field contains the MAC address of the destination; it is distributed over six 8-bit fields, labelled MAC_address_1 to MAC_address_6. The MAC_address_1 field contains the most significant byte of the MAC address, while MAC_address_6 contains the least significant byte. How many of these bytes are significant is optional and defined by the value of the broadcast descriptor table [ETSI-DAT] sent separately over another MPEG-2 TS within the TS multiplex.

MPE is currently a widely deployed scheme. Due to Investments in existing systems, usage is likely to continue in current and future MPEG-2 Transmission Networks. ATSC provides a scheme similar to MPE [ATSC-DAT] with some small differences.

(ii) Data Piping.

The Data Piping profile [ETSI-DAT] is a minimum overhead, simple and flexible profile that makes no assumptions concerning the format of the data being sent. In this profile, the Receiver is intended to provide PID filtering, packet reassembly according to [ETSI-SI], error detection, and optional Conditional Access (link encryption).

The specification allows the user data stream to be unstructured or organized into packets. The specific structure is transparent to the Receiver. It may conform to any protocol, e.g., IP, Ethernet, NFS, FDDI, MPEG-2 PES, etc.

(iii) Data Streaming.

The data broadcast specification profile [ETSI-DAT] for PES tunnels (Data Streaming) supports unicast and multicast data services that require a stream-oriented delivery of data packets. This encapsulation maps an IP packet into a single PES Packet payload.

Two different types of PES headers can be selected via the stream_id values [ISO-MPEG]. The private_stream_2 value permits the use of the short PES header with limited overhead, while the private_stream_1 value makes available the scrambling control and the timing and clock reference features of the PES layer.

Authors' Addresses

Marie J. Montpetit
Motorola Connected Home Solutions
45 Hayden Avenue 4th Floor
Lexington MA 02130
USA

EMail: mmontpetit@motorola.com

Godred Fairhurst
Department of Engineering
University of Aberdeen
Aberdeen, AB24 3UE
UK

EMail: gorry@erg.abdn.ac.uk
Web: http://www.erg.abdn.ac.uk/users/gorry

Horst D. Clausen
TIC Systems
Lawrence, Kansas

EMail: h.d.clausen@ieee.org

Bernhard Collini-Nocker
Department of Scientific Computing
University of Salzburg
Jakob Haringer Str. 2
5020 Salzburg
Austria

EMail: bnocker@cosy.sbg.ac.at
Web: http://www.network-research.org

Hilmar Linder
Department of Scientific Computing
University of Salzburg
Jakob Haringer Str. 2
5020 Salzburg
Austria

EMail: hlinder@cosy.sbg.ac.at
Web: http://www.network-research.org

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.