

Network Working Group
Request for Comments: 4947
Category: Informational

G. Fairhurst
University of Aberdeen
M.-J. Montpetit
Motorola Connected Home Solutions
July 2007

Address Resolution Mechanisms for 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

Copyright (C) The IETF Trust (2007).

Abstract

This document describes the process of binding/associating IPv4/IPv6 addresses with MPEG-2 Transport Streams (TS). This procedure is known as Address Resolution (AR) or Neighbor Discovery (ND). Such address resolution complements the higher-layer resource discovery tools that are used to advertise IP sessions.

In MPEG-2 Networks, an IP address must be associated with a Packet ID (PID) value and a specific Transmission Multiplex. This document reviews current methods appropriate to a range of technologies (such as DVB (Digital Video Broadcasting), ATSC (Advanced Television Systems Committee), DOCSIS (Data-Over-Cable Service Interface Specifications), and variants). It also describes the interaction with well-known protocols for address management including DHCP, ARP, and the ND protocol.

Table of Contents

1. Introduction	3
1.1. Bridging and Routing	4
2. Conventions Used in This Document	7
3. Address Resolution Requirements	10
3.1. Unicast Support	12
3.2. Multicast Support	12
4. MPEG-2 Address Resolution	14
4.1. Static Configuration	15
4.1.1. MPEG-2 Cable Networks	15
4.2. MPEG-2 Table-Based Address Resolution	16
4.2.1. IP/MAC Notification Table (INT) and Its Usage	17
4.2.2. Multicast Mapping Table (MMT) and Its Usage	18
4.2.3. Application Information Table (AIT) and Its Usage ..	18
4.2.4. Address Resolution in ATSC	19
4.2.5. Comparison of SI/PSI Table Approaches	19
4.3. IP-Based Address Resolution for TS Logical Channels	19
5. Mapping IP Addresses to MAC/NPA Addresses	21
5.1. Unidirectional Links Supporting Unidirectional Connectivity	22
5.2. Unidirectional Links with Bidirectional Connectivity	23
5.3. Bidirectional Links	25
5.4. AR Server	26
5.5. DHCP Tuning	27
5.6. IP Multicast AR	27
5.6.1. Multicast/Broadcast Addressing for UDLR	28
6. Link Layer Support	29
6.1. ULE without a Destination MAC/NPA Address (D=1)	30
6.2. ULE with a Destination MAC/NPA Address (D=0)	31
6.3. MPE without LLC/SNAP Encapsulation	31
6.4. MPE with LLC/SNAP Encapsulation	31
6.5. ULE with Bridging Header Extension (D=1)	32
6.6. ULE with Bridging Header Extension and NPA Address (D=0) ..	32
6.7. MPE with LLC/SNAP & Bridging	33
7. Conclusions	33
8. Security Considerations	34
9. Acknowledgments	35
10. References	35
10.1. Normative References	35
10.2. Informative References	36

1. Introduction

This document describes the process of binding/associating IPv4/IPv6 addresses with MPEG-2 Transport Streams (TS). This procedure is known as Address Resolution (AR), or Neighbor Discovery (ND). Such address resolution complements the higher layer resource discovery tools that are used to advertise IP sessions. The document reviews current methods appropriate to a range of technologies (DVB, ATSC, DOCSIS, and variants). It also describes the interaction with well-known protocols for address management including DHCP, ARP, and the ND protocol.

The MPEG-2 TS provides a time-division multiplexed (TDM) stream that may contain audio, video, and data information, including encapsulated IP Datagrams [RFC4259], defined in specification ISO/IEC 138181 [ISO-MPEG2]. Each Layer 2 (L2) frame, known as a TS Packet, contains a 4 byte header and a 184 byte payload. Each TS Packet is associated with a single TS Logical Channel, identified by a 13-bit Packet ID (PID) value that is carried in the MPEG-2 TS Packet header.

The MPEG-2 standard also defines a control plane that may be used to transmit control information to Receivers in the form of System Information (SI) Tables [ETSI-SI], [ETSI-SI1], or Program Specific Information (PSI) Tables.

To utilize the MPEG-2 TS as a L2 link supporting IP, a sender must associate an IP address with a particular Transmission Multiplex, and within the multiplex, identify the specific PID to be used. This document calls this mapping an AR function. In some AR schemes, the MPEG-2 TS address space is subdivided into logical contexts known as Platforms [ETSI-DAT]. Each Platform associates an IP service provider with a separate context that shares a common MPEG-2 TS (i.e., uses the same PID value).

MPEG-2 Receivers may use a Network Point of Attachment (NPA) [RFC4259] to uniquely identify a L2 node within an MPEG-2 transmission network. An example of an NPA is the IEEE Medium Access Control (MAC) address. Where such addresses are used, these must also be signalled by the AR procedure. Finally, address resolution could signal the format of the data being transmitted, for example, the encapsulation, with any L2 encryption method and any compression scheme [RFC4259].

The numbers of Receivers connected via a single MPEG-2 link may be much larger than found in other common LAN technologies (e.g., Ethernet). This has implications on design/configuration of the address resolution mechanisms. Current routing protocols and some multicast application protocols also do not scale to arbitrarily

large numbers of participants. Such networks do not by themselves introduce an appreciable subnetwork round trip delay, however many practical MPEG-2 transmission networks are built using links that may introduce a significant path delay (satellite links, use of dial-up modem return, cellular return, etc.). This higher delay may need to be accommodated by address resolution protocols that use this service.

1.1. Bridging and Routing

The following two figures illustrate the use of AR for a routed and a bridged subnetwork. Various other combinations of L2 and L3 forwarding may also be used over MPEG-2 links (including Receivers that are IP end hosts and end hosts directly connected to bridged LAN segments).




Figure 1: A routed MPEG-2 link

Figure 1 shows a routed MPEG-2 link feeding three downstream routers (R2-R4). AR takes place at the Encapsulator (R1) to identify each Receiver at Layer 2 within the IP subnetwork (R2, etc.).

When considering unicast communication from R1 to R2, several L2 addresses are involved:

1a is the L2 (sending) interface address of R1 on the MPEG-2 link.
 2b is the L2 (receiving) interface address of R2 on the MPEG-2 link.
 2a is the L2 (sending) interface address of R2 on the next hop link.

AR for the MPEG-2 link allows R1 to determine the L2 address (2b) corresponding to the next hop Receiver, router R2.

Figure 2 shows a bridged MPEG-2 link feeding three downstream bridges (B2-B4). AR takes place at the Encapsulator (B1) to identify each Receiver at L2 (B2-B4). AR also takes place across the IP subnetwork allowing the Feed router (R1) to identify the downstream Routers at Layer 2 (R2, etc.). The Encapsulator associates a destination MAC/NPA address with each bridged PDU sent on an MPEG-2 link. Two methods are defined by ULE (Unidirectional Lightweight Encapsulation) [RFC4326]:

The simplest method uses the L2 address of the transmitted frame. This is the MAC address corresponding to the destination within the L2 subnetwork (the next hop router, 2b of R2). This requires each Receiver (B2-B4) to associate the receiving MPEG-2 interface with the set of MAC addresses that exist on the L2 subnetworks that it feeds. Similar considerations apply when IP-based tunnels support L2 services (including the use of UDLR (Unidirectional Links) [RFC3077]).

It is also possible for a bridging Encapsulator (B1) to encapsulate a PDU with a link-specific header that also contains the MAC/NPA address associated with a Receiver L2 interface on the MPEG-2 link (Figure 2). In this case, the destination MAC/NPA address of the encapsulated frame is set to the Receiver MAC/NPA address (y), rather than the address of the final L2 destination. At a different level, an AR binding is also required for R1 to associate the destination L2 address 2b with R2. In a subnetwork using bridging, the systems R1 and R2 will normally use standard IETF-defined AR mechanisms (e.g., IPv4 Address Resolution Protocol (ARP) [RFC826] and the IPv6 Neighbor Discovery Protocol (ND) [RFC2461]) edge-to-edge across the IP subnetwork.




Figure 2: A bridged MPEG-2 link

Methods also exist to assign IP addresses to Receivers within a network (e.g., stateless autoconfiguration [RFC2461], DHCP [RFC2131], DHCPv6 [RFC3315], and stateless DHCPv6 [RFC3736]). Receivers may also participate in the remote configuration of the L3 IP addresses used in connected equipment (e.g., using DHCP-Relay [RFC3046]).

The remainder of this document describes current mechanisms and their use to associate an IP address with the corresponding TS Multiplex, PID value, the MAC/NPA address and/or Platform ID. A range of approaches is described, including Layer 2 mechanisms (using MPEG-2 SI tables), and protocols at the IP level (including ARP [RFC826] and ND [RFC2461]). Interactions and dependencies between these mechanisms and the encapsulation methods are described. The document does not propose or define a new protocol, but does provide guidance on issues that would need to be considered to supply IP-based address resolution.

2. Conventions Used in This Document

AIT: Application Information Table specified by the Multimedia Home Platform (MHP) specifications [ETSI-MHP]. This table may carry IPv4/IPv6 to MPEG-2 TS address resolution information.

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-MPEG2].

b: bit. For example, one byte consists of 8-bits.

B: Byte. Groups of bytes are represented in Internet byte order.

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

DVB: Digital Video Broadcasting [DVB]. 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.

DVB-RCS: Digital Video Broadcast Return Channel via Satellite. A bidirectional IPv4/IPv6 service employing low-cost Receivers [ETSI-RCS].

DVB-S: Digital Video Broadcast for Satellite [ETSI-DVBS].

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.

Feed Router: The router delivering the IP service over a Unidirectional Link.

INT: Internet/MAC Notification Table. A unidirectional address resolution mechanism using SI and/or PSI Tables.

L2: Layer 2, the link layer.

L3: Layer 3, the IP network layer.

MAC: Medium Access Control [IEEE-802.3]. A link layer protocol defined by the IEEE 802.3 standard (or by Ethernet v2).

MAC Address: A 6-byte link layer address of the format described by the Ethernet IEEE 802 standard (see also NPA).

MAC Header: The link layer header of the IEEE 802.3 standard [IEEE-802.3] or Ethernet v2. It consists of a 6-byte destination address, 6-byte source address, and 2 byte type field (see also NPA, LLC (Logical Link Control)).

MHP: Multimedia Home Platform. An integrated MPEG-2 multimedia Receiver, that may (in some cases) support IPv4/IPv6 services [ETSI-MHP].

MMT: Multicast Mapping Table (proprietary extension to DVB-RCS [ETSI-RCS] defining an AR table that maps IPv4 multicast addresses to PID values).

MPE: Multiprotocol Encapsulation [ETSI-DAT], [ATSC-A90]. A method that encapsulates PDUs, forming a DSM-CC Table Section. Each Section is sent in a series of TS Packets using a single Stream (TS Logical Channel).

MPEG-2: A set of standards specified by the Motion Picture Experts Group (MPEG), and standardized by the International Standards Organization (ISO/IEC 113818-1) [ISO-MPEG2], and ITU-T (in H.220).

NPA: Network Point of Attachment. A 6-byte destination address (resembling an IEEE MAC address) within the MPEG-2 transmission network that is used to identify individual Receivers or groups of Receivers [RFC4259].

PAT: Program Association Table. An MPEG-2 PSI control table. It associates each program with the PID value that is used to send the associated PMT (Program Map Table). The table is sent using the well-known PID value of 0x000, and is required for an MPEG-2 compliant Transport Stream.

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

PID: Packet Identifier [ISO-MPEG2]. A 13 bit field carried in the header of each TS Packet. This identifies the TS Logical Channel to which a TS Packet belongs [ISO-MPEG2]. The TS Packets that form the parts of a Table Section, or other Payload Unit must all carry the same PID value. A PID value of all ones 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-MPEG2]. The PID value used to send the PMT for a specific program is defined by an entry in the PAT.

Private Section: A syntactic structure constructed according to Table 2-30 of [ISO-MPEG2]. The structure may be used to identify private information (i.e., not defined by [ISO-MPEG2]) relating to one or more elementary streams, or a specific MPEG-2 program, or the entire Transport Stream. Other Standards bodies, e.g., ETSI and 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-MPEG2]. 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-MPEG2], see also SI Table.

Receiver: 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-MPEG2]. In this document, this term describes a table that is been defined by another standards body to convey information about the services carried in a TS Multiplex. 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-MPEG2].

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

Table Section: A Payload Unit carrying all or a part of an SI or PSI Table [ISO-MPEG2].

TS: Transport Stream [ISO-MPEG2], a method of transmission at the MPEG-2 level using TS Packets; it represents Layer 2 of the ISO/OSI reference model. See also TS Logical Channel and TS Multiplex.

TS Logical Channel: Transport Stream Logical Channel. In this document, this term identifies a channel at the MPEG-2 level [ISO-MPEG2]. This 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). The term "Stream" is defined in MPEG-2 [ISO-MPEG2]. This describes the

content carried by a specific TS Logical Channel (see ULE Stream). Some PID values are reserved (by MPEG-2) for specific signaling. Other standards (e.g., ATSC and DVB) also reserve specific PID values.

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) [RFC4259], 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-MPEG2]. Each TS Packet carries a 4B header.

UDL: Unidirectional link: A one-way transmission link. For example, and IP over DVB link using a broadcast satellite link.

ULE: Unidirectional Lightweight Encapsulation. A scheme that encapsulates PDUs, into SNDUs that are sent in a series of TS Packets using a single TS Logical Channel [RFC4326].

ULE Stream: An MPEG-2 TS Logical Channel that carries only ULE encapsulated PDUs. ULE Streams may be identified by definition of a `stream_type` in SI/PSI [RFC4326, ISO-MPEG2].

3. Address Resolution Requirements

The MPEG IP address resolution process is independent of the choice of encapsulation and needs to support a set of IP over MPEG-2 encapsulation formats, including Multi-Protocol Encapsulation (MPE) ([ETSI-DAT], [ATSC-A90]) and the IETF-defined Unidirectional Lightweight Encapsulation (ULE) [RFC4326].

The general IP over MPEG-2 AR requirements are summarized below:

- A scalable architecture that may support large numbers of systems within the MPEG-2 Network [RFC4259].
- A protocol version, to indicate the specific AR protocol in use and which may include the supported encapsulation method.
- A method (e.g., well-known L2/L3 address/addresses) to identify the AR Server sourcing the AR information.

- A method to represent IPv4/IPv6 AR information (including security mechanisms to authenticate the AR information to protect against address masquerading [RFC3756]).
- A method to install AR information associated with clients at the AR Server (registration).
- A method for transmission of AR information from an AR Server to clients that minimize the transmission cost (link-local multicast is preferable to subnet broadcast).
- Incremental update of the AR information held by clients.
- Procedures for purging clients of stale AR information.

An MPEG-2 transmission network may support multiple IP networks. If this is the case, it is important to recognize the scope within which an address is resolved to prevent packets from one addressed scope leaking into other scopes [RFC4259]. Examples of overlapping IP address assignments include:

- (i) Private unicast addresses (e.g., in IPv4, 10/8 prefix; 172.16/12 prefix; and 192.168/16 prefix). Packets with these addresses should be confined to one addressed area. IPv6 also defines link-local addresses that must not be forwarded beyond the link on which they were first sent.
- (ii) 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].
- (iii) Scoped multicast addresses [RFC2365] and [RFC2375]. Forwarding of these addresses is controlled by the scope associated with the address. The addresses are only valid within an addressed area (e.g., the 239/8 [RFC2365]).

Overlapping address assignments may also occur at L2, where the same MAC/NPA address is used to identify multiple Receivers [RFC4259]:

- (i) An MAC/NPA unicast address must be unique within the addressed area. The IEEE-assigned MAC addresses used in Ethernet LANs are globally unique. If the addresses are not globally unique, an address must only be re-used by Receivers in different addressed (scoped) areas.

- (ii) The MAC/NPA address broadcast address (a L2 address of all ones). Traffic with this address should be confined to one addressed area.
- (iii) IP and other protocols may view sets of L3 multicast addresses as link-local. This may produce unexpected results if frames with the corresponding multicast L2 addresses are distributed to systems in a different L3 network or multicast scope (Sections 3.2 and 5.6).

Reception of unicast packets destined for another addressed area will lead to an increase in the rate of received packets by systems connected via the network. Reception of the additional network traffic may contribute to processing load, but should not lead to unexpected protocol behaviour, providing that systems can be uniquely addressed at L2. It does however introduce a potential Denial of Service (DoS) opportunity. When the Receiver operates as an IP router, the receipt of such a packet can lead to unexpected protocol behaviour.

3.1. Unicast Support

Unicast address resolution is required at two levels.

At the lower level, the IP (or MAC) address needs to be associated with a specific TS Logical Channel (PID value) and the corresponding TS Multiplex (Section 4). Each Encapsulator within an MPEG-2 Network is associated with a set of unique TS Logical Channels (PID values) that it sources [ETSI-DAT, RFC4259]. Within a specific scope, the same unicast IP address may therefore be associated with more than one Stream, and each Stream contributes different content (e.g., when several different IP Encapsulators contribute IP flows destined to the same Receiver). MPEG-2 Networks may also replicate IP packets to send the same content (Simulcast) to different Receivers or via different TS Multiplexes. The configuration of the MPEG-2 Network must prevent a Receiver accepting duplicated copies of the same IP packet.

At the upper level, the AR procedure needs to associate an IP address with a specific MAC/NPA address (Section 5).

3.2. Multicast Support

Multicast is an important application for MPEG-2 transmission networks, since it exploits the advantages of native support for link broadcast. Multicast address resolution occurs at the network-level in associating a specific L2 address with an IP Group Destination Address (Section 5.6). In IPv4 and IPv6 over Ethernet, this

association is normally a direct mapping, and this is the default method also specified in both ULE [RFC4326] and MPE [ETSI-DAT].

Address resolution must also occur at the MPEG-2 level (Section 4). The goal of this multicast address resolution is to allow a Receiver to associate an IPv4 or IPv6 multicast address with a specific TS Logical Channel and the corresponding TS Multiplex [RFC4259]. This association needs to permit a large number of active multicast groups, and should minimize the processing load at the Receiver when filtering and forwarding IP multicast packets (e.g., by distributing the multicast traffic over a number of TS Logical Channels). Schemes that allow hardware filtering can be beneficial, since these may relieve the drivers and operating systems from discarding unwanted multicast traffic.

There are two specific functions required for address resolution in IP multicast over MPEG-2 Networks:

- (i) Mapping IP multicast groups to the underlying MPEG-2 TS Logical Channel (PID) and the MPEG-2 TS Multiplex at the Encapsulator.
- (ii) Provide signalling information to allow a Receiver to locate an IP multicast flow within an MPEG-2 TS Multiplex.

Methods are required to identify the scope of an address when an MPEG-2 Network supports several logical IP networks and carries groups within different multicast scopes [RFC4259].

Appropriate procedures need to specify the correct action when the same multicast group is available on separate TS Logical Channels. This could arise when different Encapsulators contribute IP packets with the same IP Group Destination Address in the ASM (Any-Source Multicast) address range. Another case arises when a Receiver could receive more than one copy of the same packet (e.g., when packets are replicated across different TS Logical Channels or even different TS Multiplexes, a method known as Simulcasting [ETSI-DAT]). At the IP level, the host/router may be unaware of this duplication and this needs to be detected by other means.

When the MPEG-2 Network is peered to the multicast-enabled Internet, an arbitrarily large number of IP multicast group destination addresses may be in use, and the set forwarded on the transmission network may be expected to vary significantly with time. Some uses of IP multicast employ a range of addresses to support a single application (e.g., ND [RFC2461], Layered Coding Transport (LCT) [RFC3451], and Wave and Equation Based Rate Control (WEBRC) [RFC3738]). The current set of active addresses may be determined dynamically via a multicast group membership protocol (e.g., Internet

Group Management Protocol (IGMP) [RFC3376] and Multicast Listener Discovery (MLD) [RFC3810]), via multicast routing (e.g., Protocol Independent Multicast (PIM) [RFC4601]) and/or other means (e.g., [RFC3819] and [RFC4605]), however each active address requires a binding by the AR method. Therefore, there are advantages in using a method that does not need to explicitly advertise an AR binding for each IP traffic flow, but is able to distribute traffic across a number of L2 TS Logical Channels (e.g., using a hash/mapping that resembles the mapping from IP addresses to MAC addresses [RFC1112, RFC2464]). Such methods can reduce the volume of AR information that needs to be distributed, and reduce the AR processing.

Section 5.6 describes the binding of IP multicast addresses to MAC/NPA addresses.

4. MPEG-2 Address Resolution

The first part of this section describes the role of MPEG-2 signalling to identify streams (TS Logical Channels [RFC4259]) within the L2 infrastructure.

At L2, the MPEG-2 Transport Stream [ISO-MPEG2] identifies the existence and format of a Stream, using a combination of two PSI tables: the Program Association Table (PAT) and entries in the program element loop of a Program Map Table (PMT). PMT Tables are sent infrequently and are typically small in size. The PAT is sent using the well-known PID value of 0X000. This table provides the correspondence between a program_number and a PID value. (The program_number is the numeric label associated with a program). Each program in the Table is associated with a specific PID value, used to identify a TS Logical Channel (i.e., a TS). The identified TS is used to send the PMT, which associates a set of PID values with the individual components of the program. This approach de-references the PID values when the MPEG-2 Network includes multiplexors or remultiplexors that renumber the PID values of the TS Logical Channels that they process.

In addition to signalling the Receiver with the PID value assigned to a Stream, PMT entries indicate the presence of Streams using ULE and MPE to the variety of devices that may operate in the MPEG-2 transmission network (multiplexors, remultiplexors, rate shapers, advertisement insertion equipment, etc.).

A multiplexor or remultiplexor may change the PID values associated with a Stream during the multiplexing process, the new value being reflected in an updated PMT. TS Packets that carry a PID value that is not associated with a PMT entry (an orphan PID), may, and usually will be dropped by ISO 13818-1 compliant L2 equipment, resulting in

the Stream not being forwarded across the transmission network. In networks that do not employ any intermediate devices (e.g., scenarios C,E,F of [RFC4259]), or where devices have other means to determine the set of PID values in use, the PMT table may still be sent (but is not required for this purpose).

Although the basic PMT information may be used to identify the existence of IP traffic, it does not associate a Stream with an IP prefix/address. The remainder of the section describes IP addresses resolution mechanisms relating to MPEG-2.

4.1. Static Configuration

The static mapping option, where IP addresses or flows are statically mapped to specific PIDs is the equivalent to signalling "out-of-band". The application programmer, installing engineer, or user receives the mapping via some outside means, not in the MPEG-2 TS. This is useful for testing, experimental networks, small subnetworks and closed domains.

A pre-defined set of IP addresses may be used within an MPEG-2 transmission network. Prior knowledge of the active set of addresses allows appropriate AR records to be constructed for each address, and to pre-assign the corresponding PID value (e.g., selected to optimize Receiver processing; to group related addresses to the same PID value; and/or to reflect a policy for usage of specific ranges of PID values). This presumes that the PID mappings are not modified during transmission (Section 4).

A single "well-known" PID is a specialization of this. This scheme is used by current DOCSIS cable modems [DOCSIS], where all IP traffic is placed into the specified TS stream. MAC filtering (and/or Section filtering in MPE) may be used to differentiate subnetworks.

4.1.1. MPEG-2 Cable Networks

Cable networks use a different transmission scheme for downstream (head-end to cable modem) and upstream (cable modem to head-end) transmissions.

IP/Ethernet packets are sent (on the downstream) to the cable modem(s) encapsulated in MPEG-2 TS Packets sent on a single well-known TS Logical Channel (PID). There is no use of in-band signalling tables. On the upstream, the common approach is to use Ethernet framing, rather than IP/Ethernet over MPEG-2, although other proprietary schemes also continue to be used.

Until the deployment of DOCSIS and EuroDOCSIS, most address resolution schemes for IP traffic in cable networks were proprietary, and did not usually employ a table-based address resolution method. Proprietary methods continue to be used in some cases where cable modems require interaction. In this case, equipment at the head-end may act as gateways between the cable modem and the Internet. These gateways receive L2 information and allocate an IP address.

DOCSIS uses DHCP for IP client configuration. The Cable Modem Terminal System (CMTS) provides a DHCP Server that allocates IP addresses to DOCSIS cable modems. The MPEG-2 transmission network provides a L2 bridged network to the cable modem (Section 1). This usually acts as a DHCP Relay for IP devices [RFC2131], [RFC3046], and [RFC3256]. Issues in deployment of IPv6 are described in [RFC4779].

4.2. MPEG-2 Table-Based Address Resolution

The information about the set of MPEG-2 Transport Streams carried over a TS Multiplex can be distributed via SI/PSI Tables. These tables are usually sent periodically (Section 4). This design requires access to and processing of the SI Table information by each Receiver [ETSI-SI], [ETSI-SI1]. This scheme reflects the complexity of delivering and coordinating the various Transport Streams associated with multimedia TV. A TS Multiplex may provide AR information for IP services by integrating additional information into the existing control tables or by transmitting additional SI Tables that are specific to the IP service.

Examples of MPEG-2 Table usage that allows an MPEG-2 Receiver to identify the appropriate PID and the multiplex associated with a specific IP address include:

- (i) IP/MAC Notification Table (INT) in the DVB Data standard [ETSI-DAT]. This provides unidirectional address resolution of IPv4/IPv6 multicast addresses to an MPEG-2 TS.
- (ii) Application Information Table (AIT) in the Multimedia Home Platform (MHP) specifications [ETSI-MHP].
- (iii) Multicast Mapping Table (MMT) is an MPEG-2 Table employed by some DVB-RCS systems to provide unidirectional address resolution of IPv4 multicast addresses to an MPEG-2 TS.

The MMT and AIT are used for specific applications, whereas the INT [ETSI-DAT] is a more general DVB method that supports MAC, IPv4, and IPv6 AR when used in combination with the other MPEG-2 tables (Section 4).

4.2.1. IP/MAC Notification Table (INT) and Its Usage

The INT provides a set of descriptors to specify addressing in a DVB network. The use of this method is specified for Multiprotocol Encapsulation (MPE) [ETSI-DAT]. It provides a method for carrying information about the location of IP/L2 flows within a DVB network. A Platform_ID identifies the addressing scope for a set of IP/L2 streams and/or Receivers. A Platform may span several Transport Streams carried by one or multiple TS Multiplexes and represents a single IP network with a harmonized address space (scope). This allows for the coexistence of several independent IP/MAC address scopes within an MPEG-2 Network.

The INT allows both fully-specified IP addresses and prefix matching to reduce the size of the table (and hence enhance signalling efficiency). An IPv4/IPv6 "subnet mask" may be specified in full form or by using a slash notation (e.g., /127). IP multicast addresses can be specified with or without a source (address or range), although if a source address is specified, then only the slash notation may be used for prefixes.

In addition, for identification and security descriptors, the following descriptors are defined for address binding in INT tables:

- (i) target_MAC_address_descriptor: A descriptor to describe a single or set of MAC addresses (and their mask).
- (ii) target_MAC_address_range_descriptor: A descriptor that may be used to set filters.
- (iii) target_IP_address_descriptor: A descriptor describing a single or set of IPv4 unicast or multicast addresses (and their mask).
- (iv) target_IP_slash_descriptor: Allows definition and announcement of an IPv4 prefix.
- (v) target_IP_source_slash_descriptor: Uses source and destination addresses to target a single or set of systems.
- (vi) IP/MAC stream_location_descriptor: A descriptor that locates an IP/MAC stream in a DVB network.

The following descriptors provide corresponding functions for IPv6 addresses:

- target_IPv6_address_descriptor
- target_IPv6_slash_descriptor
- and target_IPv6_source_slash_descriptor

The ISP_access_mode_descriptor allows specification of a second address descriptor to access an ISP via an alternative non-DVB (possibly non-IP) network.

One key benefit is that the approach employs MPEG-2 signalling (Section 4) and is integrated with other signalling information. This allows the INT to operate in the presence of (re)multiplexors [RFC4259] and to refer to PID values that are carried in different TS Multiplexes. This makes it well-suited to a Broadcast TV Scenario [RFC4259].

The principal drawback is a need for an Encapsulator to introduce associated PSI/SI MPEG-2 control information. This control information needs to be processed at a Receiver. This requires access to information below the IP layer. The position of this processing within the protocol stack makes it hard to associate the results with IP Policy, management, and security functions. The use of centralized management prevents the implementation of a more dynamic scheme.

4.2.2. Multicast Mapping Table (MMT) and Its Usage

In DVB-RCS, unicast AR is seen as a part of a wider configuration and control function and does not employ a specific protocol.

A Multicast Mapping Table (MMT) may be carried in an MPEG-2 control table that associates a set of multicast addresses with the corresponding PID values [MMT]. This table allows a DVB-RCS Forward Link Subsystem (FLSS) to specify the mapping of IPv4 and IPv6 multicast addresses to PID values within a specific TS Multiplex. Receivers (DVB-RCS Return Channel Satellite Terminals (RCSTs)) may use this table to determine the PID values associated with an IP multicast flow that it requires to receive. The MMT is specified by the SatLabs Forum [MMT] and is not currently a part of the DVB-RCS specification.

4.2.3. Application Information Table (AIT) and Its Usage

The DVB Multimedia Home Platform (MHP) specification [ETSI-MHP] does not define a specific AR function. However, an Application Information Table (AIT) is defined that allows MHP Receivers to receive a variety of control information. The AIT uses an MPEG-2 signalling table, providing information about data broadcasts, the required activation state of applications carried by a broadcast stream, etc. This information allows a broadcaster to request that a Receiver change the activation state of an application, and to direct

applications to receive specific multicast packet flows (using IPv4 or IPv6 descriptors). In MHP, AR is not seen as a specific function, but as a part of a wider configuration and control function.

4.2.4. Address Resolution in ATSC

ATSC [ATSC-A54A] defines a system that allows transmission of IP packets within an MPEG-2 Network. An MPEG-2 Program (defined by the PMT) may contain one or more applications [ATSC-A90] that include IP multicast streams [ATSC-A92]. IP multicast data are signalled in the PMT using a stream_type indicator of value 0x0D. A MAC address list descriptor [SCTE-1] may also be included in the PMT.

The approach focuses on applications that serve the transmission network. A method is defined that uses MPEG-2 SI Tables to bind the IP multicast media streams and the corresponding Session Description Protocol (SDP) announcement streams to particular MPEG-2 Program Elements. Each application constitutes an independent network. The MPEG-2 Network boundaries establish the IP addressing scope.

4.2.5. Comparison of SI/PSI Table Approaches

The MPEG-2 methods based on SI/PSI meet the specified requirements of the groups that created them and each has their strength: the INT in terms of flexibility and extensibility, the MMT in its simplicity, and the AIT in its extensibility. However, they exhibit scalability constraints, represent technology specific solutions, and do not fully adopt IP-centric approaches that would enable easier use of the MPEG-2 bearer as a link technology within the wider Internet.

4.3. IP-Based Address Resolution for TS Logical Channels

As MPEG-2 Networks evolve to become multi-service networks, the use of IP protocols is becoming more prevalent. Most MPEG-2 Networks now use some IP protocols for operations and control and data delivery. Address resolution information could also be sent using IP transport. At the time of writing there is no standards-based IP-level AR protocol that supports the MPEG-2 TS.

There is an opportunity to define an IP-level method that could use an IP multicast protocol over a well-known IP multicast address to resolve an IP address to a TS Logical Channel (i.e., a Transport Stream). The advantages of using an IP-based address resolution include:

- (i) **Simplicity:**
The AR mechanism does not require interpretation of L2 tables; this is an advantage especially in the growing market share for home network and audio/video networked entities.
- (ii) **Uniformity:**
An IP-based protocol can provide a common method across different network scenarios for both IP to MAC address mappings and mapping to TS Logical Channels (PID value associated with a Stream).
- (iii) **Extensibility:**
IP-based AR mechanisms allow an independent evolution of the AR protocol. This includes dynamic methods to request address resolution and the ability to include other L2 information (e.g., encryption keys).
- (iv) **Integration:**
The information exchanged by IP-based AR protocols can easily be integrated as a part of the IP network layer, simplifying support for AAA, policy, Operations and Management (OAM), mobility, configuration control, etc., that combine AR with security.

The drawbacks of an IP-based method include:

- (i) It can not operate over an MPEG-2 Network that uses MPEG-2 remultiplexors [RFC4259] that modify the PID values associated with the TS Logical Channels during the multiplexing operation (Section 4). This makes the method unsuitable for use in deployed broadcast TV networks [RFC4259].
- (ii) IP-based methods can introduce concerns about the integrity of the information and authentication of the sender [RFC4259]. (These concerns are also applicable to MPEG-2 Table methods, but in this case the information is confined to the L2 network, or parts of the network where gateway devices isolate the MPEG-2 devices from the larger Internet creating virtual MPEG-2 private networks.) IP-based solutions should therefore implement security mechanisms that may be used to authenticate the sender and verify the integrity of the AR information as a part of a larger security framework.

An IP-level method could use an IP multicast protocol running an AR Server (see also Section 5.4) over a well-known (or discovered) IP multicast address. To satisfy the requirement for scalability to networks with a large number of systems (Section 1), a single packet needs to transport multiple AR records and define the intended scope

for each address. Methods that employ prefix matching are desirable (e.g., where a range of source/destination addresses are matched to a single entry). It can also be beneficial to use methods that permit a range of IP addresses to be mapped to a set of TS Logical Channels (e.g., a hashing technique similar to the mapping of IP Group Destination Addresses to Ethernet MAC addresses [RFC1112] [RFC2464]).

5. Mapping IP Addresses to MAC/NPA Addresses

This section reviews IETF protocols that may be used to assign and manage the mapping of IP addresses to/from MAC/NPA addresses over MPEG-2 Networks.

An IP Encapsulator requires AR information to select an appropriate MAC/NPA address in the SNDU header [RFC4259] (Section 6). The information to complete this header may be taken directly from a neighbor/ARP cache, or may require the Encapsulator to retrieve the information using an AR protocol. The way in which this information is collected will depend upon whether the Encapsulator functions as a Router (at L3) or a Bridge (at L2) (Section 1.1).

Two IETF-defined protocols for mapping IP addresses to MAC/NPA addresses are the Address Resolution Protocol, ARP [RFC826], and the Neighbor Discovery protocol, ND [RFC2461], respectively for IPv4 and IPv6. Both protocols are normally used in a bidirectional mode, although both also permit unsolicited transmission of mappings. The IPv6 mapping defined in [RFC2464] can result in a large number of active MAC multicast addresses (e.g., one for each end host).

ARP requires support for L2 broadcast packets. A large number of Receivers can lead to a proportional increase in ARP traffic, a concern for bandwidth-limited networks. Transmission delay can also impact protocol performance.

ARP also has a number of security vulnerabilities. ARP spoofing is where a system can be fooled by a rogue device that sends a fictitious ARP RESPONSE that includes the IP address of a legitimate network system and the MAC of a rogue system. This causes legitimate systems on the network to update their ARP tables with the false mapping and then send future packets to the rogue system instead of the legitimate system. Using this method, a rogue system can see (and modify) packets sent through the network.

Secure ARP (SARP) uses a secure tunnel (e.g., between each client and a server at a wireless access point or router) [RFC4346]. The router ignores any ARP RESPONSEs not associated with clients using the secure tunnels. Therefore, only legitimate ARP RESPONSEs are used

for updating ARP tables. SARP requires the installation of software at each client. It suffers from the same scalability issues as the standard ARP.

The ND protocol uses a set of IP multicast addresses. In large networks, many multicast addresses are used, but each client typically only listens to a restricted set of group destination addresses and little traffic is usually sent in each group. Therefore, Layer 2 AR for MPEG-2 Networks must support this in a scalable manner.

A large number of ND messages may cause a large demand for performing asymmetric operations. The base ND protocol limits the rate at which multicast responses to solicitations can be sent. Configurations may need to be tuned when operating with large numbers of Receivers.

The default parameters specified in the ND protocol [RFC2461] can introduce interoperability problems (e.g., a failure to resolve when the link RTT (round-trip time) exceed 3 seconds) and performance degradation (duplicate ND messages with a link RTT > 1 second) when used in networks where the link RTT is significantly larger than experienced by Ethernet LANs. Tuning of the protocol parameters (e.g., RTR_SOLICITATION_INTERVAL) is therefore recommended when using network links with appreciable delay (Section 6.3.2 of [RFC2461]).

ND has similar security vulnerabilities to ARP. The Secure Neighbor Discovery (SEND) [RFC3971] was developed to address known security vulnerabilities in ND [RFC3756]. It can also reduce the AR traffic compared to ND. In addition, SEND does not require the configuration of per-host keys and can coexist with the use of both SEND and insecure ND on the same link.

The ND Protocol is also used by IPv6 systems to perform other functions beyond address resolution, including Router Solicitation / Advertisement, Duplicate Address Detection (DAD), Neighbor Unreachability Detection (NUD), and Redirect. These functions are useful for hosts, even when address resolution is not required.

5.1. Unidirectional Links Supporting Unidirectional Connectivity

MPEG-2 Networks may provide a Unidirectional Broadcast Link (UDL), with no return path. Such links may be used for unicast applications that do not require a return path (e.g., based on UDP), but commonly are used for IP multicast content distribution.




Figure 3: Unidirectional connectivity

The ARP and ND protocols require bidirectional L2/L3 connectivity. They do not provide an appropriate method to resolve the remote (destination) address in a unidirectional environment.

Unidirectional links therefore require a separate out-of-band configuration method to establish the appropriate AR information at the Encapsulator and Receivers. ULE [RFC4326] defines a mode in which the MAC/NPA address is omitted from the SNDU. In some scenarios, this may relieve an Encapsulator of the need for L2 AR.

5.2. Unidirectional Links with Bidirectional Connectivity

Bidirectional connectivity may be realized using a unidirectional link in combination with another network path. Common combinations are a Feed link using MPEG-2 satellite transmission and a return link using terrestrial network infrastructure. This topology is often known as a Hybrid network and has asymmetric network routing.




Figure 4: Bidirectional connectivity

The Unidirectional Link Routing (UDLR) [RFC3077] protocol may be used to overcome issues associated with asymmetric routing. The Dynamic Tunnel Configuration Protocol (DTCP) enables automatic configuration of the return path. UDLR hides the unidirectional routing from the IP and upper layer protocols by providing a L2 tunnelling mechanism that emulates a bidirectional broadcast link at L2. A network using UDLR has a topology where a Feed Router and all Receivers form a logical Local Area Network. Encapsulating L2 frames allows them to be sent through an Internet Path (i.e., bridging).

Since many unidirectional links employ wireless technology for the forward (Feed) link, there may be an appreciable cost associated with forwarding traffic on the Feed link. Therefore, it is often desirable to prevent forwarding unnecessary traffic (e.g., for multicast this implies control of which groups are forwarded). The implications of forwarding in the return direction must also be considered (e.g., asymmetric capacity and loss [RFC3449]). This suggests a need to minimize the volume and frequency of control messages.

Three different AR cases may be identified (each considers sending an IP packet to a next-hop IP address that is not currently cached by the sender):

- (i) A Feed Router needs a Receiver MAC/NPA address.

This occurs when a Feed Router sends an IP packet using the Feed UDL to a Receiver whose MAC/NPA address is unknown. In IPv4, the Feed Router sends an ARP REQUEST with the IP address of the Receiver. The Receiver that recognizes its IP address replies with an ARP RESPONSE to the MAC/NPA address of the Feed Router (e.g., using a UDLR tunnel). The Feed Router may then address IP packets to the unicast MAC/NPA address associated with the Receiver. The ULE encapsulation format also permits packets to be sent without specifying a MAC/NPA address, where this is desirable (Section 6.1 and 6.5).

- (ii) A Receiver needs the Feed Router MAC/NPA address.

This occurs when a Receiver sends an IP packet to a Feed Router whose MAC/NPA address is unknown. In IPv4, the Receiver sends an ARP REQUEST with the IP address of the Feed Router (e.g., using a UDLR tunnel). The Feed Router replies with an ARP RESPONSE using the Feed UDL. The Receiver may then address IP packets to the MAC/NPA address of the recipient.

(iii) A Receiver needs another Receiver MAC/NPA address.

This occurs when a Receiver sends an IP packet to another Receiver whose MAC/NPA address is unknown. In IPv4, the Receiver sends an ARP REQUEST with the IP address of the remote Receiver (e.g., using a UDLR tunnel to the Feed Router). The request is forwarded over the Feed UDL. The target Receiver replies with an ARP RESPONSE (e.g., using a UDLR tunnel). The Feed Router forwards the response on the UDL. The Receiver may then address IP packets to the MAC/NPA address of the recipient.

These 3 cases allow any system connected to the UDL to obtain the MAC/NPA address of any other system. Similar exchanges may be performed using the ND protocol for IPv6.

A long round trip delay (via the UDL and UDLR tunnel) impacts the performance of the reactive address resolution procedures provided by ARP, ND, and SEND. In contrast to Ethernet, during the interval when resolution is taking place, many IP packets may be received that are addressed to the AR Target address. The ARP specification allows an interface to discard these packets while awaiting the response to the resolution request. An appropriately sized buffer would however prevent this loss.

In case (iii), the time to complete address resolution may be reduced by the use of an AR Server at the Feed (Section 5.4).

Using DHCP requires prior establishment of the L2 connectivity to a DHCP Server. The delay in establishing return connectivity in UDLR networks that use DHCP, may make it beneficial to increase the frequency of the DTCP HELLO message. Further information about tuning DHCP is provided in Section 5.5.

5.3. Bidirectional Links

Bidirectional IP networks can be and are constructed by a combination of two MPEG-2 transmission links. One link is usually a broadcast link that feeds a set of remote Receivers. Links are also provided from Receivers so that the combined link functions as a full duplex interface. Examples of this use include two-way DVB-S satellite links and the DVB-RCS system.

5.4. AR Server

An AR Server can be used to distribute AR information to Receivers in an MPEG-2 Network. In some topologies, this may significantly reduce the time taken for Receivers to discover AR information.

The AR Server can operate as a proxy responding on behalf of Receivers to received AR requests. When an IPv4 AR request is received (e.g., Receiver ARP REQUEST), an AR Server responds by (proxy) sending an AR response, providing the appropriate IP to MAC/NPA binding (mapping the IP address to the L2 address).

Information may also be sent unsolicited by the AR Server using multicast/broadcast to update the ARP/neighbor cache at the Receivers without the need for explicit requests. The unsolicited method can improve scaling in large networks. Scaling could be further improved by distributing a single broadcast/multicast AR message that binds multiple IP and MAC/NPA addresses. This reduces the network capacity consumed and simplifies client/server processing in networks with large numbers of clients.

An AR Server can be implemented using IETF-defined Protocols by configuring the subnetwork so that AR Requests from Receivers are intercepted rather than forwarded to the Feed/broadcast link. The intercepted messages are sent to an AR Server. The AR Server maintains a set of MAC/NPA address bindings. These may be configured or may be learned by monitoring ARP messages sent by Receivers. Currently defined IETF protocols only allow one binding per message (i.e., there is no optimization to conserve L2 bandwidth).

Equivalent methods could provide IPv6 AR. Procedures for intercepting ND messages are defined in [RFC4389]. To perform an AR Server function, the AR information must also be cached. A caching AR proxy stores the system state within a middle-box device. This resembles a classic man-in-the-middle security attack; interactions with SEND are described in [SP-ND].

Methods are needed to purge stale AR data from the cache. The consistency of the cache must also be considered when the Receiver bindings can change (e.g., IP mobility, network topology changes, or intermittent Receiver connectivity). In these cases, the use of old (stale) information can result in IP packets being directed to an inappropriate L2 address, with consequent packet loss.

Current IETF-defined methods provide bindings of IP addresses to MAC/NPA, but do not allow the bindings to other L2 information pertinent to MPEG-2 Networks, requiring the use of other methods for

this function (Section 4). AR Servers can also be implemented using non-IETF AR protocols to provide the AR information required by Receivers.

5.5. DHCP Tuning

DHCP [RFC2131] and DHCPv6 [RFC3315] may be used over MPEG-2 Networks with bidirectional connectivity. DHCP consists of two components: a protocol for delivering system-specific configuration parameters from a DHCP Server to a DHCP Client (e.g., default router, DNS server) and a mechanism for the allocation of network addresses to systems.

The configuration of DHCP Servers and DHCP Clients should take into account the local link round trip delay (possibly including the additional delay from bridging, e.g., using UDLR). A large number of clients can make it desirable to tune the DHCP lease duration and the size of the address pool. Appropriate timer values should also be selected: the DHCP messages retransmission timeout, and the maximum delay that a DHCP Server waits before deciding that the absence of an ICMP echo response indicates that the relevant address is free.

DHCP Clients may retransmit DHCP messages if they do not receive a response. Some client implementations specify a timeout for the DHCPDISCOVER message that is small (e.g., suited to Ethernet delay, rather than appropriate to an MPEG-2 Network) providing insufficient time for a DHCP Server to respond to a DHCPDISCOVER retransmission before expiry of the check on the lease availability (by an ICMP Echo Request), resulting in potential address conflict. This value may need to be tuned for MPEG-2 Networks.

5.6. IP Multicast AR

Section 3.2 describes the multicast address resolution requirements. This section describes L3 address bindings when the destination network-layer address is an IP multicast Group Destination Address.

In MPE [ETSI-DAT], a mapping is specified for the MAC Address based on the IP multicast address for IPv4 [RFC1112] and IPv6 [RFC2464]. (A variant of DVB (DVB-H) uses a modified MAC header [ETSI-DAT]).

In ULE [RFC4326], the L2 NPA address is optional, and is not necessarily required when the Receiver is able to perform efficient L3 multicast address filtering. When present, a mapping is defined based on the IP multicast address for IPv4 [RFC1112] and IPv6 [RFC2464].

The L2 group addressing method specified in [RFC1112] and [RFC2464] can result in more than one IP destination address being mapped to the same L2 address. In Source-Specific Multicast, SSM [RFC3569], multicast groups are identified by the combination of the IP source and IP destination addresses. Therefore, senders may independently select an IP group destination address that could map to the same L2 address if forwarded onto the same L2 link. The resulting addressing overlap at L2 can increase the volume of traffic forwarded to L3, where it then needs to be filtered.

These considerations are the same as for Ethernet LANs, and may not be of concern to Receivers that can perform efficient L3 filtering. Section 3 noted that an MPEG-2 Network may need to support multiple addressing scopes at the network and link layers. Separation of the different groups into different Transport Streams is one remedy (with signalling of IP to PID value mappings). Another approach is to employ alternate MAC/NPA mappings to those defined in [RFC1112] and [RFC2464], but such mappings need to be consistently bound at the Encapsulator and Receiver, using AR procedures in a scalable manner.

5.6.1. Multicast/Broadcast Addressing for UDLR

UDLR is a Layer 2 solution, in which a Receiver may send multicast/broadcast frames that are subsequently forwarded natively by a Feed Router (using the topology in Figure 2), and are finally received at the Feed interface of the originating Receiver. This multicast forwarding does not include the normal L3 Reverse Path Forwarding (RPF) check or L2 spanning tree checks, the processing of the IP Time To Live (TTL) field or the filtering of administratively scoped multicast addresses. This raises a need to carefully consider multicast support. To avoid forwarding loops, RFC 3077 notes that a Receiver needs to be configured with appropriate filter rules to ensure that it discards packets that originate from an attached network and are later received over the Feed link.

When the encapsulation includes an MAC/NPA source address, re-broadcast packets may be filtered at the link layer using a filter that discards L2 addresses that are local to the Receiver. In some circumstances, systems can send packets with an unknown (all-zero) MAC source address (e.g., IGMP Proxy Queriers [RFC4605]), where the source at L2 can not be determined at the Receiver. These packets need to be silently discarded, which may prevent running the associated services on the Receiver.

Some encapsulation formats also do not include an MAC/NPA source address (Table 1). Multicast packets may therefore alternatively be discarded at the IP layer if their IP source address matches a local IP address (or address range). Systems can send packets with an

all-zero IP source address (e.g., BOOTP (bootstrap protocol) [RFC951], DHCP [RFC2131] and ND [RFC2461]), where the source at L3 can not be determined at the Receiver these packets need to be silently discarded. This may prevent the associated services at a Receiver, e.g., participation in IPv6 Duplicate Address Detection or running a DHCP server.

6. Link Layer Support

This section considers link layer (L2) support for address resolution in MPEG-2 Networks. It considers two issues: The code-point used at L2 and the efficiency of encapsulation for transmission required to support the AR method. The table below summarizes the options for both MPE ([ETSI-DAT], [ATSC-A90]) and ULE [RFC4326] encapsulations.

[RFC4840] describes issues and concerns that may arise when a link can support multiple encapsulations. In particular, it identifies problems that arise when end hosts that belong to the same IP network employ different incompatible encapsulation methods. An Encapsulator must therefore use only one method (e.g., ULE or MPE) to support a single IP network (i.e., set of IPv4 systems sharing the same subnet broadcast address or same IPv6 prefix). All Receivers in an IP network must receive all IP packets that use a broadcast (directed to all systems in the IP network) or a local-scope multicast address (Section 3). Packets with these addresses are used by many IP-based protocols including service discovery, IP AR, and routing protocols. Systems that fail to receive these packets can suffer connectivity failure or incorrect behaviour (e.g., they may be unable to participate in IP-based discovery, configuration, routing, and announcement protocols). Consistent delivery can be ensured by transmitting link-local multicast or broadcast packets using the same Stream that is used for unicast packets directed to this network. A Receiver could simultaneously use more than one L2 AR mechanism. This presents a potential conflict when the Receiver receives two different bindings for the same identifier. When multiple systems advertise AR bindings for the same identifiers (e.g., Encapsulators), they must ensure that the advertised information is consistent. Conflicts may also arise when L2 protocols duplicate the functions of IP-based AR mechanisms.

In ULE, the bridging format may be used in combination with the normal mode to address packets to a Receiver (all ULE Receivers are required to implement both methods). Frames carrying IP packets using the ULE Bridging mode, that have a destination address corresponding to the MAC address of the Receiver and have an IP address corresponding to a Receiver interface, will be delivered to the IP stack of the Receiver. All bridged IP multicast and broadcast frames will also be copied to the IP stack of the Receiver.

Receivers must filter (discard) frames that are received with a source address that matches an address of the Receiver itself [802.1D]. It must also prevent forwarding frames already sent on a connected network. For each network interface, it must therefore filter received frames where the frame source address matches a unicast destination address associated with a different network interface [802.1D].

L2 Encapsulation	PDU	L2 Frame Header Fields		
	overhead [bytes]	src mac	dst mac	type
6.1 ULE without dst MAC address	8	-	-	x
6.2 ULE with dst MAC address	14	-	x	x
6.3 MPE without LLC/SNAP	16	-	x	-
6.4 MPE with LLC/SNAP	24	-	x	x
6.5 ULE with Bridging extension	22	x	x	x
6.6 ULE with Bridging & NPA	28	x	x	x
6.7 MPE with LLC/SNAP&Bridging	38	x	x	x

Table 1: L2 Support and Overhead (x =supported, - =not supported)

The remainder of the section describes IETF-specified AR methods for use with these encapsulation formats. Most of these methods rely on bidirectional communications (see Sections 5.1, 5.2, and 5.3 for a discussion of this).

6.1. ULE without a Destination MAC/NPA Address (D=1)

The ULE encapsulation supports a mode (D=1) where the MAC/NPA address is not present in the encapsulated frame. This mode may be used with both IPv4 and IPv6. When used, the Receiver is expected to perform L3 filtering of packets based on their IP destination address [RFC4326]. This requires careful consideration of the network topology when a Receiver is an IP router, or delivers data to an IP router (a simple case where this is permitted arises in the connection of stub networks at a Receiver that have no connectivity to other networks). Since there is no MAC/NPA address in the SNDU, ARP and the ND protocol are not required for AR.

IPv6 systems can automatically configure their IPv6 network address based upon a local MAC address [RFC2462]. To use auto-configuration, the IP driver at the Receiver may need to access the MAC/NPA address of the receiving interface, even though this value is not being used to filter received SNDUs.

Even when not used for AR, the ND protocol may still be required to support DAD, and other IPv6 network-layer functions. This protocol uses a block of IPv6 multicast addresses, which need to be carried by the L2 network. However, since this encapsulation format does not provide a MAC source address, there are topologies (e.g., Section 5.6.1) where a system can not differentiate DAD packets that were originally sent by itself and were re-broadcast, from those that may have been sent by another system with the same L3 address. Therefore, DAD can not be used with this encapsulation format in topologies where this L2 forwarding may occur.

6.2. ULE with a Destination MAC/NPA Address (D=0)

The IPv4 Address Resolution Protocol (ARP) [RFC826] is identified by an IEEE EtherType and may be used over ULE [RFC4326]. Although no MAC source address is present in the ULE SNDU, the ARP protocol still communicates the source MAC (hardware) address in the ARP record payload of any query messages that it generates.

The IPv6 ND protocol is supported. The protocol uses a block of IPv6 multicast addresses, which need to be carried by the L2 network. The protocol uses a block of IPv6 multicast addresses, which need to be carried by the L2 network. However, since this encapsulation format does not provide a MAC source address, there are topologies (e.g., Section 5.6.1) where a system can not differentiate DAD packets that were originally sent by itself and were re-broadcast, from those that may have been sent by another system with the same L3 address. Therefore, DAD can not be used with this encapsulation format in topologies where this L2 forwarding may occur.

6.3. MPE without LLC/SNAP Encapsulation

This is the default (and sometimes only) mode specified by most MPE Encapsulators. MPE does not provide an EtherType field and therefore can not support the Address Resolution Protocol (ARP) [RFC826].

IPv6 is not supported in this encapsulation format, and therefore it is not appropriate to consider the ND protocol.

6.4. MPE with LLC/SNAP Encapsulation

The LLC/SNAP (Sub-Network Access Protocol) format of MPE provides an EtherType field and therefore may support ARP [RFC826]. There is no specification to define how this is performed. No MAC source address is present in the SNDU, although the protocol communicates the source MAC address in the ARP record payload of any query messages that it generates.

The IPv6 ND protocol is supported using The LLC/SNAP format of MPE. This requires specific multicast addresses to be carried by the L2 network. The IPv6 ND protocol is supported. The protocol uses a block of IPv6 multicast addresses, which need to be carried by the L2 network. However, since this encapsulation format does not provide a MAC source address, there are topologies (e.g., Section 5.6.1) where a system can not differentiate DAD packets that were originally sent by itself and were re-broadcast, from those that may have been sent by another system with the same L3 address. Therefore, DAD can not be used with this encapsulation format in topologies where this L2 forwarding may occur.

6.5. ULE with Bridging Header Extension (D=1)

The ULE encapsulation supports a bridging extension header that supplies both a source and destination MAC address. This can be used without an NPA address (D=1). When no other Extension Headers precede this Extension, the MAC destination address has the same position in the ULE SNDU as that used for an NPA destination address. The Receiver may optionally be configured so that the MAC destination address value is identical to a Receiver NPA address.

At the Encapsulator, the ULE MAC/NPA destination address is determined by a L2 forwarding decision. Received frames may be forwarded or may be addressed to the Receiver itself. As in other L2 LANs, the Receiver may choose to filter received frames based on a configured MAC destination address filter. ARP and ND messages may be carried within a PDU that is bridged by this encapsulation format. Where the topology may result in subsequent reception of re-broadcast copies of multicast frames that were originally sent by a Receiver (e.g., Section 5.6.1), the system must discard frames that are received with a source address that it used in frames sent from the same interface [802.1D]. This prevents duplication on the bridged network (e.g., this would otherwise invoke DAD).

6.6. ULE with Bridging Header Extension and NPA Address (D=0)

The combination of an NPA address (D=0) and a bridging extension header are allowed in ULE. This SNDU format supplies both a source and destination MAC address and a NPA destination address (i.e., Receiver MAC/NPA address).

At the Encapsulator, the value of the ULE MAC/NPA destination address is determined by a L2 forwarding decision. At the Receiver, frames may be forwarded or may be addressed to the Receiver itself. As in other L2 LANs, the Receiver may choose to filter received frames based on a configured MAC destination address filter. ARP and ND messages may be carried within a PDU that is bridged by this

encapsulation format. Where the topology may result in the subsequent reception of re-broadcast copies of multicast frames, that were originally sent by a Receiver (e.g., Section 5.6.1), the system must discard frames that are received with a source address that it used in frames sent from the same interface [802.1D]. This prevents duplication on the bridged network (e.g., this would otherwise invoke DAD).

6.7. MPE with LLC/SNAP & Bridging

The LLC/SNAP format MPE frames may optionally support an IEEE bridging header [LLC]. This header supplies both a source and destination MAC address, at the expense of larger encapsulation overhead. The format defines two MAC destination addresses, one associated with the MPE SNDU (i.e., Receiver MAC address) and one with the bridged MAC frame (i.e., the MAC address of the intended recipient in the remote LAN).

At the Encapsulator, the MPE MAC destination address is determined by a L2 forwarding decision. There is currently no formal description of the Receiver processing for this encapsulation format. A Receiver may forward frames or they may be addressed to the Receiver itself. As in other L2 LANs, the Receiver may choose to filter received frames based on a configured MAC destination address filter. ARP and ND messages may be carried within a PDU that is bridged by this encapsulation format. The MPE MAC destination address is determined by a L2 forwarding decision. Where the topology may result in a subsequent reception of re-broadcast copies of multicast frames, that were originally sent by a Receiver (e.g., Section 5.6.1), the system must discard frames that are received with a source address that it used in frames sent from the same interface [802.1D]. This prevents duplication on the bridged network (e.g., this would otherwise invoke DAD).

7. Conclusions

This document describes addressing and address resolution issues for IP protocols over MPEG-2 transmission networks using both wired and wireless technologies. A number of specific IETF protocols are discussed along with their expected behaviour over MPEG-2 transmission networks. Recommendations for their usage are provided.

There is no single common approach used in all MPEG-2 Networks. A static binding may be configured for IP addresses and PIDs (as in some cable networks). In broadcast networks, this information is normally provided by the Encapsulator/Multiplexor and carried in signalling tables (e.g., AIT in MHP, the IP Notification Table, INT,

of DVB and the DVB-RCS Multicast Mapping Table, and MMT). This document has reviewed the status of these current address resolution mechanisms in MPEG-2 transmission networks and defined their usage.

The document also considers a unified IP-based method for AR that could be independent of the physical layer, but does not define a new protocol. It examines the design criteria for a method, with recommendations to ensure scalability and improve support for the IP protocol stack.

8. Security Considerations

The normal security issues relating to the use of wireless links for transmission of Internet traffic should be considered.

L2 signalling in MPEG-2 transmission networks is currently provided by (periodic) broadcasting of information in the control plane using PSI/SI tables (Section 4). A loss or modification of the SI information may result in an inability to identify the TS Logical Channel (PID) that is used for a service. This will prevent reception of the intended IP packet stream.

There are known security issues relating to the use of unsecured address resolution [RFC3756]. Readers are also referred to the known security issues when mapping IP addresses to MAC/NPA addresses using ARP [RFC826] and ND [RFC2461]. It is recommended that AR protocols support authentication of the source of AR messages and the integrity of the AR information, this avoids known security vulnerabilities resulting from insertion of unauthorized AR messages within a L2 infrastructure. For IPv6, the SEND protocol [RFC3971] may be used in place of ND. This defines security mechanisms that can protect AR.

AR protocols can also be protected by the use of L2 security methods (e.g., Encryption of the ULE SNDU [IPDVB-SEC]). When these methods are used, the security of ARP and ND can be comparable to that of a private LAN: A Receiver will only accept ARP or ND transmissions from the set of peer senders that share a common group encryption and common group authentication key provided by the L2 key management.

AR Servers (Section 5.4) are susceptible to the same kind of security issues as end hosts using unsecured AR. These issues include hijacking traffic and denial-of-service within the subnet. Malicious nodes within the subnet can take advantage of this property, and hijack traffic. In addition, an AR Server is essentially a legitimate man-in-the-middle, which implies that there is a need to distinguish such proxies from unwanted man-in-the-middle attackers. This document does not introduce any new mechanisms for the

protection of these AR functions (e.g., authenticating servers, or defining AR Servers that interoperate with the SEND protocol [SP-ND]).

9. Acknowledgments

The authors wish to thank the IPDVB WG members for their inputs and in particular, Rod Walsh, Jun Takei, and Michael Mercurio. The authors also acknowledge the support of the European Space Agency. Martin Stiernerling contributed descriptions of scenarios, configuration, and provided extensive proof reading. Hidetaka Izumiyama contributed on UDLR and IPv6 issues. A number of issues discussed in the UDLR working group have also provided valuable inputs to this document (summarized in "Experiments with RFC 3077", July 2003).

10. References

10.1. Normative References

- [ETSI-DAT] EN 301 192, "Specifications for Data Broadcasting", v1.3.1, European Telecommunications Standards Institute (ETSI), May 2003.
- [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-SI] EN 300 468, "Digital Video Broadcasting (DVB); Specification for Service Information (SI) in DVB systems", v1.7.1, European Telecommunications Standards Institute (ETSI), December 2005.
- [ISO-MPEG2] ISO/IEC IS 13818-1, "Information technology -- Generic coding of moving pictures and associated audio information -- Part 1: Systems", International Standards Organization (ISO), 2000.
- [RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, November 1982.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC3077] Duros, E., Dabbous, W., Izumiyama, H., Fujii, N., and Y. Zhang, "A Link-Layer Tunneling Mechanism for Unidirectional Links", RFC 3077, March 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.
- [RFC4326] Fairhurst, G. and B. Collini-Nocker, "Unidirectional Lightweight Encapsulation (ULE) for Transmission of IP Datagrams over an MPEG-2 Transport Stream (TS)", RFC 4326, December 2005.

10.2. Informative References

- [802.1D] IEEE 802.1D, "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges", IEEE, 2004.
- [802.3] IEEE 802.3, "Local and metropolitan area networks-Specific requirements Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", IEEE Computer Society, (also ISO/IEC 8802-3), 2002.
- [ATSC] A/53C, "ATSC Digital Television Standard", Advanced Television Systems Committee (ATSC), Doc. A/53C, 2004.
- [ATSC-A54A] A/54A, "Guide to the use of the ATSC Digital Television Standard", Advanced Television Systems Committee (ATSC), Doc. A/54A, 2003.
- [ATSC-A90] A/90, "ATSC Data Broadcast Standard", Advanced Television Systems Committee (ATSC), Doc. A/90, 2000.

- [ATSC-A92] A/92, "Delivery of IP Multicast Sessions over ATSC Data Broadcast", Advanced Television Systems Committee (ATSC), Doc. A/92, 2002.
- [DOCSIS] "Data-Over-Cable Service Interface Specifications, DOCSIS 2.0, Radio Frequency Interface Specification", CableLabs, document CM-SP-RFiv2.0-I10-051209, 2005.
- [DVB] Digital Video Broadcasting (DVB) Project.
<http://www.dvb.org>.
- [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-RCS] EN 301 790, "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution Systems", European Telecommunications Standards Institute (ETSI).
- [ETSI-SI1] TR 101 162, "Digital Video Broadcasting (DVB); Allocation of Service Information (SI) codes for DVB systems", European Telecommunications Standards Institute (ETSI).
- [IPDVB-SEC] H. Cruickshank, S. Iyengar, L. Duquerroy, P. Pillai, "Security requirements for the Unidirectional Lightweight Encapsulation (ULE) protocol", Work in Progress, May 2007.
- [ISO-DSMCC] ISO/IEC IS 13818-6, "Information technology -- Generic coding of moving pictures and associated audio information -- Part 6: Extensions for DSM-CC is a full software implementation", International Standards Organization (ISO), 2002.
- [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 Organization (ISO), 1998.
- [MMT] "SatLabs System Recommendations, Part 1, General Specifications", Version 2.0, SatLabs Forum, 2006.
http://satlabs.org/pdf/SatLabs_System_Recommendations_v2.0_general.pdf.

- [RFC951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, September 1985.
- [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.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.
- [RFC3256] Jones, D. and R. Woundy, "The DOCSIS (Data-Over-Cable Service Interface Specifications) Device Class DHCP (Dynamic Host Configuration Protocol) Relay Agent Information Sub-option", RFC 3256, April 2002.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G., and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", BCP 69, RFC 3449, December 2002.
- [RFC3451] Luby, M., Gemmell, J., Vicisano, L., Rizzo, L., Handley, M., and J. Crowcroft, "Layered Coding Transport (LCT) Building Block", RFC 3451, December 2002.
- [RFC3569] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.
- [RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [RFC3738] Luby, M. and V. Goyal, "Wave and Equation Based Rate Control (WEBRC) Building Block", RFC 3738, April 2004.
- [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.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC4259] Weis, B., "The Use of RSA/SHA-1 Signatures within Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4359, January 2006.
- [RFC4346] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", RFC 4346, April 2006.
- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, April 2006.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [RFC4779] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks", RFC 4779, January 2007.
- [RFC4840] Aboba, B., Davies, E., and D. Thaler, "Multiple Encapsulation Methods Considered Harmful", RFC 4840, April 2007.
- [SCTE-1] "IP Multicast for Digital MPEG Networks", SCTE DVS 311r6, March 2002.
- [SP-ND] Daley, G., "Securing Proxy Neighbour Discovery Problem Statement", Work in Progress, February 2005.

Authors' Addresses

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

EMail: gorry@erg.abdn.ac.uk
URL: <http://www.erg.abdn.ac.uk/users/gorry>

Marie-Jose Montpetit
Motorola Connected Home Solutions
Advanced Technology
55 Hayden Avenue, 3rd Floor
Lexington, Massachusetts 02421
USA

EMail: mmontpetit@motorola.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

