

Classical IP and ARP over ATM

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This memo defines an initial application of classical IP and ARP in an Asynchronous Transfer Mode (ATM) network environment configured as a Logical IP Subnetwork (LIS) as described in Section 3. This memo does not preclude the subsequent development of ATM technology into areas other than a LIS; specifically, as single ATM networks grow to replace many ethernet local LAN segments and as these networks become globally connected, the application of IP and ARP will be treated differently. This memo considers only the application of ATM as a direct replacement for the "wires" and local LAN segments connecting IP end-stations ("members") and routers operating in the "classical" LAN-based paradigm. Issues raised by MAC level bridging and LAN emulation are beyond the scope of this paper.

This memo introduces general ATM technology and nomenclature. Readers are encouraged to review the ATM Forum and ITU-TS (formerly CCITT) references for more detailed information about ATM implementation agreements and standards.

Acknowledgments

This memo could not have come into being without the critical review from Jim Forster of Cisco Systems, Drew Perkins of FORE Systems, and Bryan Lyles, Steve Deering, and Berry Kercheval of XEROX PARC. The concepts and models presented in [1], written by Dave Piscitello and Joseph Lawrence, laid the structural groundwork for this work. ARP [3] written by Dave Plummer and Inverse ARP [12] written by Terry Bradley and Caralyn Brown are the foundation of ATMARP presented in this memo. This document could have not been completed without the expertise of the IP over ATM Working Group of the IETF and the ad hoc PVC committee at the Amsterdam IETF meeting.

1. Conventions

The following language conventions are used in the items of specification in this document:

- o MUST, SHALL, or MANDATORY -- the item is an absolute requirement of the specification.
- o SHOULD or RECOMMEND -- this item should generally be followed for all but exceptional circumstances.
- o MAY or OPTIONAL -- the item is truly optional and may be followed or ignored according to the needs of the implementor.

2. Introduction

The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams and ATM Address Resolution Protocol (ATMARP) requests and replies over ATM Adaptation Layer 5 (AAL5)[2,6].

Note: this memo defines only the operation of IP and address resolution over ATM, and is not meant to describe the operation of ATM networks. Any reference to virtual connections, permanent virtual connections, or switched virtual connections applies only to virtual channel connections used to support IP and address resolution over ATM, and thus are assumed to be using AAL5. This memo places no restrictions or requirements on virtual connections used for other purposes.

Initial deployment of ATM provides a LAN segment replacement for:

- 1) Local area networks (e.g., Ethernets, Token Rings and FDDI).
- 2) Local-area backbones between existing (non-ATM) LANs.
- 3) Dedicated circuits or frame relay PVCs between IP routers.

Note: In 1), local IP routers with one or more ATM interfaces will be able to connect islands of ATM networks. In 3), public or private ATM Wide Area networks will be used to connect IP routers, which in turn may or may not connect to local ATM networks. ATM WANs and LANs may be interconnected.

Private ATM networks (local or wide area) will use the private ATM address structure specified in the ATM Forum UNI specification [9]. This structure is modeled after the format of an OSI Network Service Access Point Address. A private ATM address uniquely identifies an

ATM endpoint. Public networks will use either the address structure specified in ITU-TS recommendation E.164 or the private network ATM address structure. An E.164 address uniquely identifies an interface to a public network.

The characteristics and features of ATM networks are different than those found in LANs:

- o ATM provides a Virtual Connection (VC) switched environment. VC setup may be done on either a Permanent Virtual Connection (PVC) or dynamic Switched Virtual Connection (SVC) basis. SVC call management signalling is performed via implementations of the Q.93B protocol [7,9].
- o Data to be passed by a VC is segmented into 53 octet quantities called cells (5 octets of ATM header and 48 octets of data).
- o The function of mapping user Protocol Data Units (PDUs) into the information field of the ATM cell and vice versa is performed in the ATM Adaptation Layer (AAL). When a VC is created a specific AAL type is associated with the VC. There are four different AAL types, which are referred to individually as "AAL1", "AAL2", "AAL3/4", and "AAL5". (Note: this memo concerns itself with the mapping of IP and ATMARP over AAL5 only. The other AAL types are mentioned for introductory purposes only.) The AAL type is known by the VC end points via the call setup mechanism and is not carried in the ATM cell header. For PVCs the AAL type is administratively configured at the end points when the Connection (circuit) is set up. For SVCs, the AAL type is communicated along the VC path via Q.93B as part of call setup establishment and the end points use the signaled information for configuration. ATM switches generally do not care about the AAL type of VCs. The AAL5 format specifies a packet format with a maximum size of (64K - 1) octets of user data. Cells for an AAL5 PDU are transmitted first to last, the last cell indicating the end of the PDU. ATM standards guarantee that on a given VC, cell ordering is preserved end-to-end. NOTE: AAL5 provides a non-assured data transfer service - it is up to higher-level protocols to provide retransmission.
- o ATM Forum signalling defines point-to-point and point-to-multipoint Connection setup [9]. Multipoint-to-multipoint VCs are not yet specified by ITU-TS or ATM Forum.
- o An ATM Forum ATM endpoint address is either encoded as an NSAP Address (NSAPA) or is an E.164 Public-UNI address [9]. In some cases, both an ATM endpoint address and an E.164 Public UNI address are needed by an ATMARP client to reach another host or

router. Since the use of ATM endpoint addresses and E.164 public UNI addresses by ATMARP are analogous to the use of Ethernet addresses, the notion of "hardware address" is extended to encompass ATM addresses in the context of ATMARP, even though ATM addresses need not have hardware significance. ATM Forum NSAPAs use the same basic format as U.S. GOSIP NSAPAs [11]. Note: ATM Forum addresses should not be construed as being U.S. GOSIP NSAPAs. They are not, the administration is different, which fields get filled out are different, etc.

This memo describes the initial deployment of ATM within "classical" IP networks as a direct replacement for local area networks (ethernets) and for IP links which interconnect routers, either within or between administrative domains. The "classical" model here refers to the treatment of the ATM host adapter as a networking interface to the IP protocol stack operating in a LAN-based paradigm.

Characteristics of the classical model are:

- o The same maximum transmission unit (MTU) size is used for all VCs in a LIS [2]. (Refer to Section 5.)
- o Default LLC/SNAP encapsulation of IP packets.
- o End-to-end IP routing architecture stays the same.
- o IP addresses are resolved to ATM addresses by use of an ATMARP service within the LIS - ATMARPs stay within the LIS. From a client's perspective, the ATMARP architecture stays faithful to the basic ARP model presented in [3].
- o One IP subnet is used for many hosts and routers. Each VC directly connects two IP members within the same LIS.

Future memos will describe the operation of IP over ATM when ATM networks become globally deployed and interconnected.

The deployment of ATM into the Internet community is just beginning and will take many years to complete. During the early part of this period, we expect deployment to follow traditional IP subnet boundaries for the following reasons:

- o Administrators and managers of IP subnetworks will tend to initially follow the same models as they currently have deployed. The mindset of the community will change slowly over time as ATM increases its coverage and builds its credibility.

- o Policy administration practices rely on the security, access, routing, and filtering capability of IP Internet gateways: i.e., firewalls. ATM will not be allowed to "back-door" around these mechanisms until ATM provides better management capability than the existing services and practices.
- o Standards for global IP over ATM will take some time to complete and deploy.

This memo details the treatment of the classical model of IP and ATMARP over ATM. This memo does not preclude the subsequent treatment of ATM networks within the IP framework as ATM becomes globally deployed and interconnected; this will be the subject of future documents. This memo does not address issues related to transparent data link layer interoperability.

3. IP Subnetwork Configuration

In the LIS scenario, each separate administrative entity configures its hosts and routers within a closed logical IP subnetwork. Each LIS operates and communicates independently of other LISs on the same ATM network. Hosts connected to ATM communicate directly to other hosts within the same LIS. Communication to hosts outside of the local LIS is provided via an IP router. This router is an ATM Endpoint attached to the ATM network that is configured as a member of one or more LISs. This configuration may result in a number of disjoint LISs operating over the same ATM network. Hosts of differing IP subnets MUST communicate via an intermediate IP router even though it may be possible to open a direct VC between the two IP members over the ATM network.

The requirements for IP members (hosts, routers) operating in an ATM LIS configuration are:

- o All members have the same IP network/subnet number and address mask [8].
- o All members within a LIS are directly connected to the ATM network.
- o All members outside of the LIS are accessed via a router.
- o All members of a LIS MUST have a mechanism for resolving IP addresses to ATM addresses via ATMARP (based on [3]) and vice versa via InATMARP (based on [12]) when using SVCs. Refer to Section 6 "Address Resolution" in this memo.

- o All members of a LIS MUST have a mechanism for resolving VCs to IP addresses via InATMARP (based on [12]) when using PVCs. Refer to Section 6 "Address Resolution" in this memo.
- o All members within a LIS MUST be able to communicate via ATM with all other members in the same LIS; i.e., the virtual Connection topology underlying the intercommunication among the members is fully meshed.

The following list identifies a set of ATM specific parameters that MUST be implemented in each IP station connected to the ATM network:

- o ATM Hardware Address (atm\$ha). The ATM address of the individual IP station.
- o ATMARP Request Address (atm\$arp-req). atm\$arp-req is the ATM address of an individual ATMARP server located within the LIS. In an SVC environment, ATMARP requests are sent to this address for the resolution of target protocol addresses to target ATM addresses. That server MUST have authoritative responsibility for resolving ATMARP requests of all IP members within the LIS. Note: if the LIS is operating with PVCs only, then this parameter may be set to null and the IP station is not required to send ATMARP requests to the ATMARP server.

It is RECOMMENDED that routers providing LIS functionality over the ATM network also support the ability to interconnect multiple LISs. Routers that wish to provide interconnection of differing LISs MUST be able to support multiple sets of these parameters (one set for each connected LIS) and be able to associate each set of parameters to a specific IP network/ subnet number. In addition, it is RECOMMENDED that a router be able to provide this multiple LIS support with a single physical ATM interface that may have one or more individual ATM endpoint addresses. Note: this does not necessarily mean different End System Identifiers (ESIs) when NSAPAs are used. The last octet of an NSAPA is the NSAPA Selector (SEL) field which can be used to differentiate up to 256 different LISs for the same ESI. (Refer to Section 5.1.3.1, "Private Networks" in [9].)

4. Packet Format

Implementations MUST support IEEE 802.2 LLC/SNAP encapsulation as described in [2]. LLC/SNAP encapsulation is the default packet format for IP datagrams.

This memo recognizes that other encapsulation methods may be used however, in the absence of other knowledge or agreement, LLC/SNAP encapsulation is the default.

This memo recognizes the future deployment of end-to-end signalling within ATM that will allow negotiation of encapsulation method on a per-VC basis. Signalling negotiations are beyond the scope of this memo.

5. MTU Size

The default MTU size for IP members operating over the ATM network SHALL be 9180 octets. The LLC/SNAP header is 8 octets, therefore the default ATM AAL5 protocol data unit size is 9188 octets [2]. In classical IP subnets, values other than the default can be used if and only if all members in the LIS have been configured to use the non-default value.

This memo recognizes the future deployment of end-to-end signalling within ATM that will allow negotiation of MTU size on a per-VC basis. Signalling negotiations are beyond the scope of this document.

6. Address Resolution

Address resolution within an ATM logical IP subnet SHALL make use of the ATM Address Resolution Protocol (ATMARP) (based on [3]) and the Inverse ATM Address Resolution Protocol (InATMARP) (based on [12]) as defined in this memo. ATMARP is the same protocol as the ARP protocol presented in [3] with extensions needed to support ARP in a unicast server ATM environment. InATMARP is the same protocol as the original InARP protocol presented in [12] but applied to ATM networks. All IP stations MUST support these protocols as updated and extended in this memo. Use of these protocols differs depending on whether PVCs or SVCs are used.

6.1 Permanent Virtual Connections

An IP station MUST have a mechanism (eg. manual configuration) for determining what PVCs it has, and in particular which PVCs are being used with LLC/SNAP encapsulation. The details of the mechanism are beyond the scope of this memo.

All IP members supporting PVCs are required to use the Inverse ATM Address Resolution Protocol (InATMARP) (refer to [12]) on those VCs using LLC/SNAP encapsulation. In a strict PVC environment, the receiver SHALL infer the relevant VC from the VC on which the InATMARP request (InARP_REQUEST) or response (InARP_REPLY) was received. When the ATM source and/or target address is unknown, the corresponding ATM address length in the InATMARP packet MUST be set to zero (0) indicating a null length, otherwise the appropriate address field should be filled in and the corresponding length set appropriately. InATMARP packet format details are presented later in

this memo.

Directly from [12]: "When the requesting station receives the InARP reply, it may complete the [ATM]ARP table entry and use the provided address information. Note: as with [ATM]ARP, information learned via In[ATM]ARP may be aged or invalidated under certain circumstances." It is the responsibility of each IP station supporting PVCs to re-validate [ATM]ARP table entries as part of the aging process. See Section 6.5 on "ATMARP Table Aging".

6.2 Switched Virtual Connections

SVCs require support for ATMARP in the non-broadcast, non-multicast environment that ATM networks currently provide. To meet this need a single ATMARP Server MUST be located within the LIS. This server MUST have authoritative responsibility for resolving the ATMARP requests of all IP members within the LIS.

The server itself does not actively establish connections. It depends on the clients in the LIS to initiate the ATMARP registration procedure. An individual client connects to the ATMARP server using a point-to-point VC. The server, upon the completion of an ATM call/connection of a new VC specifying LLC/SNAP encapsulation, will transmit an InATMARP request to determine the IP address of the client. The InATMARP reply from the client contains the information necessary for the ATMARP Server to build its ATMARP table cache. This information is used to generate replies to the ATMARP requests it receives.

The ATMARP Server mechanism requires that each client be administratively configured with the ATM address of the ATMARP Server atm\$arp-req as defined earlier in this memo. There is to be one and only one ATMARP Server operational per logical IP subnet. It is RECOMMENDED that the ATMARP Server also be an IP station. This station MUST be administratively configured to operate and recognize itself as the ATMARP Server for a LIS. The ATMARP Server MUST be configured with an IP address for each logical IP subnet it is serving to support InATMARP requests.

This memo recognizes that a single ATMARP Server is not as robust as multiple servers which synchronize their databases correctly. This document is defining the client-server interaction by using a simple, single server approach as a reference model, and does not prohibit more robust approaches which use the same client-server interface.

6.3 ATMARP Server Operational Requirements

The ATMARP server accepts ATM calls/connections from other ATM end points. At call setup and if the VC supports LLC/SNAP encapsulation, the ATMARP server will transmit to the originating ATM station an InATMARP request (InARP_REQUEST) for each logical IP subnet the server is configured to serve. After receiving an InATMARP reply (InARP_REPLY), the server will examine the IP address and the ATM address. The server will add (or update) the <ATM address, IP address> map entry and timestamp into its ATMARP table. If the InATMARP IP address duplicates a table entry IP address and the InATMARP ATM address does not match the table entry ATM address and there is an open VC associated with that table entry, the InATMARP information is discarded and no modifications to the table are made. ATMARP table entries persist until aged or invalidated. VC call tear down does not remove ATMARP table entries.

The ATMARP server, upon receiving an ATMARP request (ARP_REQUEST), will generate the corresponding ATMARP reply (ARP_REPLY) if it has an entry in its ATMARP table. Otherwise it will generate a negative ATMARP reply (ARP_NAK). The ARP_NAK response is an extension to the ARMARP protocol and is used to improve the robustness of the ATMARP server mechanism. With ARP_NAK, a client can determine the difference between a catastrophic server failure and an ATMARP table lookup failure. The ARP_NAK packet format is the same as the received ARP_REQUEST packet format with the operation code set to ARP_NAK, i.e., the ARP_REQUEST packet data is merely copied for transmission with the ARP_REQUEST operation code reset to ARP_NAK.

Updating the ATMARP table information timeout, the short form: when the server receives an ATMARP request over a VC, where the source IP and ATM address match the association already in the ATMARP table and the ATM address matches that associated with the VC, the server may update the timeout on the source ATMARP table entry: i.e., if the client is sending ATMARP requests to the server over the same VC that it used to register its ATMARP entry, the server should examine the ATMARP requests and note that the client is still "alive" by updating the timeout on the client's ATMARP table entry.

Adding robustness to the address resolution mechanism using ATMARP: when the server receives an ARP_REQUEST over a VC, it examines the source information. If there is no IP address associated with the VC over which the ATMARP request was received and if the source IP address is not associated with any other connection, then the server will add the <ATM address, IP address> entry and timestamp into its ATMARP table and associate the entry with this VC.

6.4 ATMARP Client Operational Requirements

The ATMARP client is responsible for contacting the ATMARP server to register its own ATMARP information and to gain and refresh its own ATMARP entry/information about other IP members. This means, as noted above, that ATMARP clients MUST be configured with the ATM address of the ATMARP server. ATMARP clients MUST:

1. Initiate the VC connection to the ATMARP server for transmitting and receiving ATMARP and InATMARP packets.
2. Respond to ARP_REQUEST and InARP_REQUEST packets received on any VC appropriately. (Refer to Section 7, "Protocol Operation" in [12].)
3. Generate and transmit ARP_REQUEST packets to the ATMARP server and to process ARP_REPLY and ARP_NAK packets from the server appropriately. ARP_REPLY packets should be used to build/refresh its own client ATMARP table entries.
4. Generate and transmit InARP_REQUEST packets as needed and to process InARP_REPLY packets appropriately. InARP_REPLY packets should be used to build/refresh its own client ATMARP table entries. (Refer to Section 7, "Protocol Operation" in [12].)
5. Provide an ATMARP table aging function to remove its own old client ATMARP tables entries after a convenient period of time.

Note: if the client does not maintain an open VC to the server, the client MUST refresh its ATMARP information with the server at least once every 20 minutes. This is done by opening a VC to the server and exchanging the initial InATMARP packets.

6.5 ATMARP Table Aging

An ATMARP client or server MUST have knowledge of any open VCs it has (permanent or switched), their association with an ATMARP table entry, and in particular, which VCs support LLC/SNAP encapsulation.

Client ATMARP table entries are valid for a maximum time of 15 minutes.

Server ATMARP table entries are valid for a minimum time of 20 minutes.

Prior to aging an ATMARP table entry, an ATMARP server MUST generate an InARP_REQUEST on any open VC associated with that entry. If an InARP_REPLY is received, that table entry is updated and not deleted.

If there is no open VC associated with the table entry, the entry is deleted.

When an ATMARP table entry ages, an ATMARP client MUST invalidate the table entry. If there is no open VC associated with the invalidated entry, that entry is deleted. In the case of an invalidated entry and an open VC, the ATMARP client must revalidate the entry prior to transmitting any non address resolution traffic on that VC. In the case of a PVC, the client validates the entry by transmitting an InARP_REQUEST and updating the entry on receipt of an InARP_REPLY. In the case of an SVC, the client validates the entry by transmitting an ARP_REQUEST to the ATMARP Server and updating the entry on receipt of an ARP_REPLY. If a VC with an associated invalidated ATMARP table entry is closed, that table entry is removed.

6.6 ATMARP and InATMARP Packet Format

Internet addresses are assigned independently of ATM addresses. Each host implementation MUST know its own IP and ATM address(es) and MUST respond to address resolution requests appropriately. IP members MUST also use ATMARP and InATMARP to resolve IP addresses to ATM addresses when needed.

The ATMARP and InATMARP protocols use the same hardware type (ar\$hrd), protocol type (ar\$pro), and operation code (ar\$op) data formats as the ARP and InARP protocols [3,12]. The location of these fields within the ATMARP packet are in the same byte position as those in ARP and InARP packets. A unique hardware type value has been assigned for ATMARP. In addition, ATMARP makes use of an additional operation code for ARP_NAK. The remainder of the ATMARP/InATMARP packet format is different than the ARP/InARP packet format.

The ATMARP and InATMARP protocols have several fields that have the following format and values:

Data:


ar\$hrd	16 bits	Hardware type
ar\$pro	16 bits	Protocol type
ar\$shtl	8 bits	Type & length of source ATM number (q)
ar\$sstl	8 bits	Type & length of source ATM subaddress (r)
ar\$op	16 bits	Operation code (request, reply, or NAK)
ar\$spln	8 bits	Length of source protocol address (s)
ar\$thtl	8 bits	Type & length of target ATM number (x)
ar\$tstl	8 bits	Type & length of target ATM subaddress (y)
ar\$tpln	8 bits	Length of target protocol address (z)
ar\$sha	qoctets	source ATM number
ar\$ssa	roctets	source ATM subaddress

ar\$spa	soctets	source protocol address
ar\$tha	xoctets	target ATM number
ar\$tsa	yoctets	target ATM subaddress
ar\$tpa	zoctets	target protocol address

Where:

- ar\$hrd - assigned to ATM Forum address family and is 19 decimal (0x0013) [4].
- ar\$pro - see Assigned Numbers for protocol type number for the protocol using ATMARP. (IP is 0x0800).
- ar\$op - The operation type value (decimal):
 - ARP_REQUEST = 1
 - ARP_REPLY = 2
 - InARP_REQUEST = 8
 - InARP_REPLY = 9
 - ARP_NAK = 10
- ar\$spln - length in octets of the source protocol address. For IP ar\$spln is 4.
- ar\$tpln - length in octets of the target protocol address. For IP ar\$tpln is 4.
- ar\$sha - source ATM number (E.164 or ATM Forum NSAPA)
- ar\$ssa - source ATM subaddress (ATM Forum NSAPA)
- ar\$spa - source protocol address
- ar\$tha - target ATM number (E.164 or ATM Forum NSAPA)
- ar\$tsa - target ATM subaddress (ATM Forum NSAPA)
- ar\$tpa - target protocol address

The encoding of the 8-bit type and length value for ar\$sh1, ar\$st1, ar\$th1, and ar\$st1 is as follows:



Where:

- bit.8 (reserved) = 0 (for future use)
- bit.7 (type) = 0 ATM Forum NSAPA format
= 1 E.164 format
- bit.6-1 (length) = 6 bit unsigned octet length of address
(MSB = bit.6, LSB = bit.1)

ATM addresses in Q.93B (as defined by the ATM Forum UNI 3.0 signalling specification [9]) include a "Calling Party Number Information Element" and a "Calling Party Subaddress Information Element". These Information Elements (IEs) SHOULD map to ATMARP/InATMARP source ATM number and source ATM subaddress respectively. Furthermore, ATM Forum defines a "Called Party Number Information Element" and a "Called Party Subaddress Information Element". These IEs map to ATMARP/InATMARP target ATM number and target ATM subaddress respectively.

The ATM Forum defines three structures for the combined use of number and subaddress [9]:

	ATM Number	ATM Subaddress
	-----	-----
Structure 1	ATM Forum NSAPA	null
Structure 2	E.164	null
Structure 3	E.164	ATM Forum NSAPA

IP members MUST register their ATM endpoint address with their ATMARP server using the ATM address structure appropriate for their ATM network connection: i.e., LISs implemented over ATM LANs following ATM Forum UNI 3.0 should register using Structure 1; LISs implemented over an E.164 "public" ATM network should register using Structure 2. A LIS implemented over a combination of ATM LANs and public ATM networks may need to register using Structure 3. Implementations based on this memo MUST support all three ATM address structures.

ATMARP and InATMARP requests and replies for ATM address structures 1 and 2 MUST indicate a null ATM subaddress; i.e., ar\$st1.type = 1 and


ar\$stl.length = 0 and ar\$stl.type = 1 and ar\$stl.length = 0. When ar\$stl.length and ar\$stl.length = 0, the ar\$tsa and ar\$ssa fields are not present.

Note: the ATMARP packet format presented in this memo is general in nature in that the ATM number and ATM subaddress fields SHOULD map directly to the corresponding Q.93B fields used for ATM call/connection setup signalling messages. The IP over ATM Working Group expects ATM Forum NSAPA numbers (Structure 1) to predominate over E.164 numbers (Structure 2) as ATM endpoint identifiers within ATM LANs. The ATM Forum's VC Routing specification is not complete at this time and therefore its impact on the operational use of ATM Address Structure 3 is undefined. The ATM Forum will be defining this relationship in the future. It is for this reason that IP members need to support all three ATM address structures.

6.7 ATMARP/InATMARP Packet Encapsulation

ATMARP and InATMARP packets are to be encoded in AAL5 PDUs using LLC/SNAP encapsulation. The format of the AAL5 CPCS-SDU payload field for ATMARP/InATMARP PDUs is:

Payload Format for ATMARP/InATMARP PDUs:



The LLC value of 0xAA-AA-03 (3 octets) indicates the presence of a SNAP header.

The OUI value of 0x00-00-00 (3 octets) indicates that the following two-bytes is an ethertype.

The Etherbyte value of 0x08-06 (2 octets) indicates ARP [4].

The total size of the LLC/SNAP header is fixed at 8-octets. This aligns the start of the ATMARP packet on a 64-bit boundary relative to the start of the AAL5 CPCS-SDU.

The LLC/SNAP encapsulation for ATMARP/InATMARP presented here is consistent with the treatment of multiprotocol encapsulation of IP over ATM AAL5 as specified in [2] and in the format of ATMARP over IEEE 802 networks as specified in [5].

Traditionally, address resolution requests are broadcast to all directly connected IP members within a LIS. It is conceivable in the future that larger scaled ATM networks may handle ATMARP requests to destinations outside the originating LIS, perhaps even globally; issues raised by ATMARP'ing outside the LIS or by a global ATMARP mechanism are beyond the scope of this memo.

7. IP Broadcast Address

ATM does not support broadcast addressing, therefore there are no mappings available from IP broadcast addresses to ATM broadcast services. Note: this lack of mapping does not restrict members from transmitting or receiving IP datagrams specifying any of the four standard IP broadcast address forms as described in [8]. Members, upon receiving an IP broadcast or IP subnet broadcast for their LIS, MUST process the packet as if addressed to that station.

8. IP Multicast Address

ATM does not support multicast address services, therefore there are no mappings available from IP multicast addresses to ATM multicast services. Current IP multicast implementations (i.e., MBONE and IP tunneling, see [10]) will continue to operate over ATM based logical IP subnets if operated in the WAN configuration.

This memo recognizes the future development of ATM multicast service addressing by the ATM Forum. When available and widely implemented, the roll-over from the current IP multicast architecture to this new ATM architecture will be straightforward.

9. Security

Not all of the security issues relating to IP over ATM are clearly understood at this time, due to the fluid state of ATM specifications, newness of the technology, and other factors.

It is believed that ATM and IP facilities for authenticated call management, authenticated end-to-end communications, and data encryption will be needed in globally connected ATM networks. Such future security facilities and their use by IP networks are beyond the scope of this memo.

There are known security issues relating to host impersonation via the address resolution protocols used in the Internet [13]. No special security mechanisms have been added to the address resolution mechanism defined here for use with networks using IP over ATM.

10. Open Issues

- o Interim Local Management Interface (ILMI) services will not be generally implemented initially by some providers and vendors and will not be used to obtain the ATM address network prefix from the network [9]. Meta-signalling does provide some of this functionality and in the future we need to document the options.
- o Well known ATM address(es) for ATMARP servers? It would be very handy if a mechanism were available for determining the "well known" ATM address(es) for the client's ATMARP server in the LIS.
- o There are many VC management issues which have not yet been addressed by this specification and which await the unwary implementor. For example, one problem that has not yet been resolved is how two IP members decide which of duplicate VCs can be released without causing VC thrashing. If two IP stations simultaneously established VCs to each other, it is tempting to allow only one of these VCs to be established, or to release one of these VCs immediately after it is established. If both IP stations simultaneously decide to release opposite VCs, a thrashing effect can be created where VCs are repeatedly established and immediately released. For the time being, the safest strategy is to allow duplicate VCs to be established and simply age them like any other VCs.

References

- [1] Piscitello, D., and J. Lawrence, "IP and ARP over the SMDS Service", RFC 1209, Bell Communications Research, March 1991.
- [2] Heinanen, J., "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 1483, Telecom Finland, July 1993.
- [3] Plummer, D., "An Ethernet Address Resolution Protocol - or - Converting Network Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, MIT, November 1982.
- [4] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1340, USC/Information Sciences Institute, July 1992.

- [5] Postel, J., and J. Reynolds, "A Standard for the Transmission of IP Datagrams over IEEE 802 Networks", STD 43, RFC 1042, USC/Information Sciences Institute, February 1988.
- [6] CCITT, "Draft Recommendation I.363", CCITT Study Group XVIII, Geneva, 19-29 January 1993.
- [7] CCITT, "Draft text for Q.93B", CCITT Study Group XI, 23 September - 2 October 1992.
- [8] Braden, R., "Requirements for Internet Hosts -- Communication Layers", STD 3, RFC 1122, USC/Information Sciences Institute, October 1989.
- [9] ATM Forum, "ATM User-Network Interface Specification Version 3.0.", ATM Forum, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, June 1993.
- [10] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, Stanford University, August 1989.
- [11] Colella, R., and Gardner, E., and R. Callon, "Guidelines for OSI NSAP Allocation in the Internet", RFC 1237, NIST, Mitre, DEC, July 1991.
- [12] Bradely, T., and C. Brown, "Inverse Address Resolution Protocol", RFC 1293, Wellfleet Communications, Inc., January 1992.
- [13] Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, Vol. 19, Issue 2, pp. 32-48, 1989.

Security Considerations

Security issues are discussed in Section 9.

Author's Address

Mark Laubach
Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto, CA 94304

Phone: 415-857-3513
Fax: 415-857-8526
EMail: laubach@hpl.hp.com