

The Transmission of IP Datagrams over the SMDS Service

Status of this Memo

This memo defines a protocol for the transmission of IP and ARP packets over a Switched Multi-megabit Data Service Network configured as a logical IP subnetwork. This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This memo describes an initial use of IP and ARP in an SMDS service environment configured as a logical IP subnetwork, LIS (described below). The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of the SMDS Service in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

Acknowledgment

This memo draws heavily in both concept and text from [4], written by Jon Postel and Joyce K. Reynolds of ISI and [5], written by David Katz of Merit, Inc. The authors would also like to acknowledge the contributions of the IP Over SMDS Service working group of the Internet Engineering Task Force.

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 RECOMMENDED -- the 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.

Introduction

The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams and ARP requests and replies.

The characteristics of the SMDS Service and the SMDS Interface Protocol (SIP) are presented in [3], [6], and in [7]. Briefly, the SMDS Service is a connectionless, public, packet-switched data service. The operation and features of the SMDS Service are similar to those found in high-speed data networks such as LANs:

- o The SMDS Service provides a datagram packet transfer, where each data unit is handled and switched separately without the prior establishment of a network connection.
- o The SMDS Service exhibits high throughput and low delay, and provides the transparent transport and delivery of up to 9188 octets of user information in a single transmission.
- o No explicit flow control mechanisms are provided; instead, the rate of information transfer on the access paths is controlled both in the subscriber-to-network direction and in the network-to-subscriber direction through the use of an access class enforcement mechanism.
- o Both individually and group-addressed (multicast) packets can be transferred.
- o In addition to these LAN-like features, a set of addressing-related service features (source address validation, source and destination address screening) are provided to enable a subscriber or set of subscribers to create a logical private network, or closed user group, over the SMDS Service. The access control provided by the closed user group mechanism is supplied by the SMDS provider according to the specifications stated in [3].
- o SMDS addresses are 60 bits plus a 4 bit Address Type. The Address Type subfield occupies the 4 most significant bits of the destination and source address fields of the SIP Level 3 Protocol Data Unit (PDU). It contains the value 1100 to

indicate an individual address and the value 1110 for a 60-bit group address.

The SMDS Interface Protocol is based on the IEEE Standard 802.6, Distributed Queue Dual Bus (DQDB) Connectionless MAC protocol [8]. The SMDS service layer corresponds to the IEEE 802 MAC sublayer. The remainder of the Data Link Service is provided by the IEEE 802.2 Logical Link Control (LLC) service [9]. The resulting stack of services is illustrated in Figure 1:




Figure 1. Protocol stack for IP over SMDS Service

This memo describes an initial use of IP and ARP in an SMDS Service environment configured as a logical IP subnetwork (described below). It does not preclude subsequent treatment of SMDS Service in configurations other than logical IP subnetworks; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document does not address issues related to transparent data link layer interoperability.

Logical IP Subnetwork Configuration

This section describes the scenario for an SMDS Service that is configured with multiple logical IP subnetworks, LIS (described below). The scenario considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

In the LIS scenario, each separate administrative entity configures its hosts within a closed logical IP subnetwork. Each LIS operates and communicates independently of other LISs over the same network providing SMDS. Hosts connected to SMDS communicate directly to other hosts within the same LIS. Communication to hosts outside of an individual LIS is provided via an IP router. This router would simply be a station attached to the SMDS Service that has been configured to be a member of both logical IP subnetworks. This configuration results in a number of disjoint LISs operating over the

same network supporting the SMDS Service. It is recognized that with this configuration, hosts of differing IP networks would communicate via an intermediate router even though a direct path over the SMDS Service may be possible.

It is envisioned that the service will evolve to provide a more public interconnection, allowing machines directly connected to the SMDS Service to communicate without an intermediate router. However, the issues raised by such a large public interconnection, such as scalability of address resolution or propagation of routing updates, are beyond the scope of this paper. We anticipate that future RFCs will address these issues.

The following is a list of the requirements for a LIS configuration:

- o All members have the same IP network/subnet number.
- o All stations within a LIS are accessed directly over SMDS.
- o All stations outside of the LIS are accessed via a router.
- o For each LIS a single SMDS group address has been configured that identifies all members of the LIS. Any packet transmitted with this address is delivered by SMDS Service to all members of the LIS.

The following list identifies a set of SMDS Service specific parameters that MUST be implemented in each IP station which would connect to the SMDS Service. The parameter values will be determined at SMDS subscription time and will be different for each LIS. Thus these parameters MUST be user configurable.

- o SMDS Hardware Address (smds\$ha). The SMDS Individual address of the IP station as determined at subscription time. Each host MUST be configured to accept datagrams destined for this address.
- o SMDS LIS Group Address(smds\$lis-ga). The SMDS Group address that has been configured at subscription time to identify the SMDS Subscriber Network Interfaces (SNI) of all members of the LIS connected to the SMDS Service. All members of the LIS MUST be prepared to accept datagrams addressed to smds\$lis-ga.
- o SMDS Arp Request Address (smds\$arp-req). The SMDS address (individual or group) to which arp requests are to be sent. In the initial LIS configuration this value is set to smds\$lis-ga. It is conceivable that in other configurations this value would be set to some address other than that of smds\$lis-ga (see

section on Address Resolution).

It is RECOMMENDED that routers providing LIS functionality over the SMDS service also support the ability to interconnect differing 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 with 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 SMDS interface that may have one or more individual SMDS addresses.

The following list identifies LIS specific parameters that MUST be configured in the network supporting the SMDS Service. For each LIS, the IP network administrator MUST request the configuration of these parameters at subscription time. The administrator of each LIS MUST update these parameters as each new station is added to the LIS.

- o SMDS LIS Group Address(sm_{ds}\$lis-ga). An SMDS Group address MUST be configured at subscription time to identify the SMDS Subscriber Network Interfaces (SNI) of all members of the LIS connected to the SMDS Service.
- o SMDS Address Screening Tables (Source and Destination). The use of SMDS screening tables is not necessary for the operation of IP over SMDS Service. If the SMDS screening tables are to be used, both source and destination tables for each SNI MUST be configured to allow, at minimum, both the direct communication between all hosts in the same LIS and the use of the SMDS LIS Group Address.

Packet Format

Service SHALL be encapsulated within the IEEE 802.2 LLC and IEEE 802.1A Sub-Network Access Protocol (SNAP) [10] Data Link layers and the 3-level SIP. The SNAP MUST be used with an Organizationally Unique Identifier Code indicating that the SNAP header contains the EtherType code as listed in Assigned Numbers [11] (see Figure 2). Note that values specified in this document follow Internet conventions: multi-byte fields are described in big-endian order and bits within bytes are described as most significant first [11].




Figure 2. Data Link Encapsulation

- o The value of HLPI in the SIP L3 Header is 1.
- o The total length of the LLC Header and the SNAP header is 8 octets.
- o The value of DSAP and SSAP in the LLC header is 170 (decimal), AA (Internet hexadecimal).
- o The Ctrl (Control) value in the LLC header is 3 (Indicates Type One Unnumbered Information).
- o The Org Code in the SNAP header is zero (000000 Internet hexadecimal).
- o The EtherType for IP is 2048 (decimal), 0800 (Internet hexadecimal). The EtherType for ARP is 2054 (decimal), 0806 (Internet hexadecimal).

IEEE 802.2 LLC Type One Unnumbered Information (UI) communication (which must be implemented by all conforming IEEE 802.2 stations) is used exclusively. The Higher Layer Protocol Id (HLPI) field in the SIP L3_PDU header MUST be set to the IEEE 802.6 assigned Protocol Id value for LLC (decimal 1) [8]. All frames MUST be transmitted in standard IEEE 802.2 LLC Type 1 Unnumbered Information format, with the DSAP and the SSAP fields of the IEEE 802.2 header set to the assigned global SAP value for SNAP (decimal 170) [10]. The 24-bit Org Code (Organizationally Unique Identifier Code) in the SNAP MUST be set to a value of zero, and the remaining 16 bits are set to the EtherType value from Assigned Numbers [11] (2048 for IP, 2054 for ARP).

The data link encapsulation for IP packets is shown in Figure 3 and for ARP in Figure 4. All values shown are in Internet hexadecimal format.

SIP			LLC / SNAP					IP
SIP..	HLPI	...	DSAP	SSAP	Ctrl	Org Code	Ethertype	
SIP..	01	...	AA	AA	03	000000	0800	IP...

Figure 3. IP Data Link Encapsulation and Values

SIP			LLC / SNAP					ARP
SIP..	HLPI	...	DSAP	SSAP	Ctrl	Org Code	Ethertype	
SIP..	01	...	AA	AA	03	000000	0806	ARP...

Figure 4. ARP Data Link Encapsulation and Values

Address Resolution

The dynamic mapping of 32-bit Internet addresses to SMDS addresses SHALL be done via the dynamic discovery procedure of the Address Resolution Protocol (ARP) [2].

Internet addresses are assigned independent of SMDS addresses. Each host implementation MUST know its own Internet address and SMDS address and respond to Address Resolution requests appropriately. Hosts MUST also use ARP to map Internet addresses to SMDS addresses when needed.

The ARP protocol has several fields that parameterize its use in any specific context [2]. These fields are:

- ar\$hrd 16 - bits The Hardware Type Code
- ar\$pro 16 - bits The Protocol Type Code
- ar\$hln 8 - bits Octets in each hardware address
- ar\$pln 8 - bits Octets in each protocol address
- ar\$op 16 - bits Operation Code

- o The hardware type code assigned to SMDS addresses is 14 (decimal), 0E (Internet hexadecimal) [11].
- o The protocol type code for IP is 2048 (decimal), 0800 (Internet hexadecimal) [11].

- o The hardware address length for SMDS is 8.
- o The protocol address length for IP is 4.
- o The operation code is 1 for request and 2 for reply.

The SMDS hardware addresses in ARP packets (ar\$sha, ar\$tha) MUST be carried in SMDS native address format, with the most significant bit of the Address Type sub-field as the high order bit of the first octet. Although outside the scope of this document, it is RECOMMENDED that SMDS addresses be represented in this format in all higher layer Internet protocols (e.g., SNMP).

Traditionally, ARP requests are broadcast to all directly connected stations. For the SMDS Service, the ARP request packet is transmitted to the smds\$arp-req hardware address. In the LIS configuration, the smds\$arp-req address is set to smds\$lis-ga, (the SMDS group address that identifies all members of the LIS). It is conceivable that in a larger scale, public configuration, the smds\$arp-req address would be configured to the address of some ARP-server(s) instead of the group address that identifies the entire LIS.

IP Broadcast Address

There is no facility for complete hardware broadcast addressing over the SMDS Service. As discussed in the "LIS Configuration" section, an SMDS group address (smds\$lis-ga) SHALL be configured to include all stations in the same LIS. The broadcast Internet address (the address on that network with a host part of all binary ones) MUST be mapped to smds\$lis-ga (see also [12]).

IP Multicast Support

A method of supporting IP multicasting is specified in [13]. It would be desirable to fully utilize the SMDS group address capabilities to support IP multicasting. However, the method in [13] requires a Network Service Interface which provides multicast-like ability to provide dynamic access to the local network service interface operations:

- o JoinLocalGroup (group-address)
- o LeaveLocalGroup (group-address)

The SMDS group address ability does not currently support dynamic subscription and removal from group address lists. Therefore, it is RECOMMENDED that in the LIS configuration, if IP multicasting is to

be supported, the method of IP multicasting described for pure broadcast media, such as the Experimental Ethernet, be used. For this method, all Multicast IP addresses are mapped to the same SMDS address which the broadcast Internet address is mapped for a given LIS. Thus all Multicast IP addresses are mapped to smds\$lis-ga. Filtering of multicast packets MUST be performed in the destination host.

Trailer Formats

Some versions of Unix 4.x BSD use a different encapsulation method in order to get better network performance with the VAX virtual memory architecture. Trailers SHALL not be used over the SMDS Service.

Byte Order

As described in Appendix B of the Internet Protocol specification [1], the IP datagram is transmitted over the SMDS Service as a series of 8-bit bytes. The byte order of the IP datagram shall be mapped directly onto the native SMDS byte order.

MAC Sublayer Details

Packet Size

The SMDS Service defines a maximum service data unit size of 9188 information octets. This leaves 9180 octets for user data after the LLC/SNAP header is taken into account. Therefore, the MTU for IP stations operating over the network supporting the SMDS Service SHALL be 9180 octets.

There is no minimum packet size restriction defined for the SMDS Service.

Other MAC Sublayer Issues

The SMDS Service requires that the publicly administered 60-bit address plus 4-bit type field format SHALL be used in both source and destination address fields of the SIP L3_PDU [3].

IEEE 802.2 Details

While not necessary for supporting IP and ARP, all implementations MUST support IEEE 802.2 standard Class I service in order to be compliant with IEEE 802.2. Some of the functions are not related directly to the support of the SNAP SAP (e.g., responding to XID and TEST commands directed to the null or global SAP addresses), but are part of a general LLC implementation. Both [4] and [5] describe the

minimum functionality necessary for a conformant station.
Implementors should also consult IEEE Std. 802.2 [14] for details.

REFERENCES

1. Postel, J., "Internet Protocol", RFC 791, USC/Information Sciences Institute, September 1981.
2. Plummer, D., "An Ethernet Address Resolution Protocol - or - Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", RFC 826, MIT, November 1982.
3. "Generic Systems Requirements in support of Switched Multi-megabit Data Service", Technical Advisory TA-TSY-000772, Bellcore Technical Advisory, Issue 3, October 1989.
4. Postel, J., and J. Reynolds, "A Standard for the Transmission of IP Datagrams over IEEE 802 Networks", RFC 1042, USC/Information Sciences Institute, February 1988.
5. Katz, D., "A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks", RFC 1188, Merit/NSFNET, October 1990.
6. Dix, F., Kelly, M., and R. Klessig, "Access to a Public Switched Multi-Megabit Data Service Offering", ACM SIGCOMM CCR, July 1990.
7. Hemrick, C. and L. Lang, "Introduction to Switched Multi-megabit Data Service (SMDS), an Early Broadband Service", publication pending in the Proceedings of the XIII International Switching Symposium (ISS 90), May 27, 1990 - June 1, 1990.
8. Institute of Electrical & Electronic Engineers, Inc. IEEE Standard 802.6, "Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN) Standard", December 1990.
9. IEEE, "IEEE Standards for Local Area Networks: Logical Link Control", IEEE, New York, New York, 1985.
10. IEEE, "Draft Standard P802.1A--Overview and Architecture", 1989.
11. Reynolds, J., and J. Postel, "Assigned Numbers", RFC 1060, USC/Information Sciences Institute, March 1990.
12. Braden, R., and J. Postel, "Requirements for Internet Gateways", RFC 1009, USC/Information Sciences Institute, June 1987.

13. Deering, S., "Host Extensions for IP Multicasting", RFC 1112, Stanford University, August 1989.
14. IEEE, "ANSI/IEEE Std 802.2-1985, ISO Draft International Standard 8802/2", IEEE, New York, New York, 1985.

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Dave Piscitello
Bell Communications Research
331 Newman Springs Road
Red Bank, NJ 07701

Phone: (908) 758-2286

EMail: dave@sabre.bellcore.com

Joseph Lawrence
Bell Communications Research
331 Newman Springs Road
Red Bank, NJ 07701

Phone: (908) 758-4146

EMail: jcl@sabre.bellcore.com