

Multiprotocol Encapsulation over ATM Adaptation Layer 5

Status of this Memo

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 two encapsulations methods for carrying network interconnect traffic over ATM AAL5. The first method allows multiplexing of multiple protocols over a single ATM virtual circuit whereas the second method assumes that each protocol is carried over a separate ATM virtual circuit.

1. Introduction

Asynchronous Transfer Mode (ATM) based networks are of increasing interest for both local and wide area applications. This memo describes two different methods for carrying connectionless network interconnect traffic, routed and bridged Protocol Data Units (PDUs), over an ATM network. The first method allows multiplexing of multiple protocols over a single ATM virtual circuit. The protocol of a carried PDU is identified by prefixing the PDU by an IEEE 802.2 Logical Link Control (LLC) header. This method is in the following called "LLC Encapsulation" and a subset of it has been earlier defined for SMDS [1]. The second method does higher-layer protocol multiplexing implicitly by ATM Virtual Circuits (VCs). It is in the following called "VC Based Multiplexing".

ATM is a cell based transfer mode that requires variable length user information to be segmented and reassembled to/from short, fixed length cells. This memo doesn't specify a new Segmentation And Reassembly (SAR) method for bridged and routed PDUs. Instead, the PDUs are carried in the Payload field of Common Part Convergence Sublayer (CPCS) PDU of ATM Adaptation Layer type 5 (AAL5) [2].

Note that this memo only describes how routed and bridged PDUs are carried directly over the CPCS of AAL5, i.e., when the Service Specific Convergence Sublayer (SSCS) of AAL5 is empty. If Frame

Relay Service Specific Convergence Sublayer (FR-SSCS), as defined in I.36x.1 [3], is used over the CPCS of AAL5, then routed and bridged PDUs are carried using the NLPID multiplexing method described in RFC 1294 [4]. Appendix A (which is for information only) shows the format of the FR-SSCS-PDU as well as how IP and CLNP PDUs are encapsulated over FR-SSCS according to RFC 1294.


2. Selection of the Multiplexing Method

It is envisioned that VC Based Multiplexing will be dominant in environments where dynamic creation of large numbers of ATM VCs is fast and economical. These conditions are likely to first prevail in private ATM networks. LLC Encapsulation, on the other hand, may be desirable when it is not practical for one reason or another to have a separate VC for each carried protocol. This is the case, for example, if the ATM network only supports (semi) Permanent Virtual Circuits (PVCs) or if charging depends heavily on the number of simultaneous VCs.

When two ATM stations wish to exchange connectionless network interconnect traffic, selection of the multiplexing method is done either by manual configuration (in case of PVCs) or by B-ISDN signalling procedures (in case of Switched VCs). The details of B-ISDN signalling are still under study in CCITT [5]. It can, however, be assumed that B-ISDN signalling messages include a "Low layer compatibility" information element, which will allow negotiation of AAL5 and the carried (encapsulation) protocol.

3. AAL5 Frame Format

No matter which multiplexing method is selected, routed and bridged PDUs shall be encapsulated within the Payload field of AAL5 CPCS-PDU. The format of the AAL5 CPCS-PDU is given below:



The Payload field contains user information up to $2^{16} - 1$ octets.

The PAD field pads the CPCS-PDU to fit exactly into the ATM cells such that the last 48 octet cell payload created by the SAR sublayer will have the CPCS-PDU Trailer right justified in the cell.

The CPCS-UU (User-to-User indication) field is used to transparently transfer CPCS user to user information. The field has no function under the multiprotocol ATM encapsulation described in this memo and can be set to any value.

The CPI (Common Part Indicator) field aligns the CPCS-PDU trailer to 64 bits. Possible additional functions are for further study in CCITT. When only the 64 bit alignment function is used, this field shall be coded as 0x00.

The Length field indicates the length, in octets, of the Payload field. The maximum value for the Length field is 65535 octets. A Length field coded as 0x00 is used for the abort function.

The CRC field protects the entire CPCS-PDU except the CRC field itself.

4. LLC Encapsulation

LLC Encapsulation is needed when several protocols are carried over the same VC. In order to allow the receiver to properly process the incoming AAL5 CPCS-PDU, the Payload Field must contain information

necessary to identify the protocol of the routed or bridged PDU. In LLC Encapsulation this information is encoded in an LLC header placed in front of the carried PDU.

Although this memo only deals with protocols that operate over LLC Type 1 (unacknowledged connectionless mode) service, the same encapsulation principle applies also to protocols operating over LLC Type 2 (connection-mode) service. In the latter case the format and/or contents of the LLC header would differ from what is shown below.

4.1. LLC Encapsulation for Routed Protocols

In LLC Encapsulation the protocol of the routed PDU is identified by prefixing the PDU by an IEEE 802.2 LLC header, which is possibly followed by an IEEE 802.1a SubNetwork Attachment Point (SNAP) header. In LLC Type 1 operation, the LLC header consists of three one octet fields:

```

+-----+-----+-----+
| DSAP | SSAP | Ctrl |
+-----+-----+-----+

```

In LLC Encapsulation for routed protocols, the Control field has always value 0x03 specifying Unnumbered Information Command PDU.

The LLC header value 0xFE-FE-03 identifies that a routed ISO PDU (see [6] and Appendix B) follows. The Control field value 0x03 specifies Unnumbered Information Command PDU. For routed ISO PDUs the format of the AAL5 CPCS-PDU Payload field shall thus be as follows:

```

      Payload Format for Routed ISO PDUs
+-----+
|          LLC  0xFE-FE-03          |
+-----+
|          .          |
|          ISO PDU    |
|          (up to 2^16 - 4 octets) |
|          .          |
+-----+

```

The routed ISO protocol is identified by a one octet NLPID field that is part of Protocol Data. NLPID values are administered by ISO and CCITT. They are defined in ISO/IEC TR 9577 [6] and some of the currently defined ones are listed in Appendix C.

An NLPID value of 0x00 is defined in ISO/IEC TR 9577 as the Null Network Layer or Inactive Set. Since it has no significance within

the context of this encapsulation scheme, a NLPID value of 0x00 is invalid under the ATM encapsulation.

It would also be possible to use the above encapsulation for IP, since, although not an ISO protocol, IP has an NLPID value 0xCC defined for it. This format must not be used. Instead, IP is encapsulated like all other routed non-ISO protocols by identifying it in the SNAP header that immediately follows the LLC header.

The presence of a SNAP header is indicated by the LLC header value 0xAA-AA-03. A SNAP header is of the form

```

+-----+-----+-----+-----+-----+
|           OUI           |     PID     |
+-----+-----+-----+-----+

```

The three-octet Organizationally Unique Identifier (OUI) identifies an organization which administers the meaning of the following two octet Protocol Identifier (PID). Together they identify a distinct routed or bridged protocol. The OUI value 0x00-00-00 specifies that the following PID is an EtherType.


The format of the AAL5 CPCS-PDU Payload field for routed non-ISO PDUs shall thus be as follows:

```

Payload Format for Routed non-ISO PDUs
+-----+
|           LLC  0xAA-AA-03           |
+-----+
|           OUI  0x00-00-00           |
+-----+
|           EtherType (2 octets)       |
+-----+
|           .                           |
|           Non-ISO PDU                 |
|           (up to 2^16 - 9 octets)     |
|           .                           |
+-----+

```

In the particular case of an Internet IP PDU, the Ethertype value is 0x08-00:



This is compatible with RFC 1042 [7]. Any changes in the header format specified in RFC 1042 should be followed by this memo.

4.2. LLC Encapsulation for Bridged Protocols

In LLC Encapsulation bridged PDUs are encapsulated by identifying the type of the bridged media in the SNAP header. As with routed non-ISO protocols, the presence of the SNAP header is indicated by the LLC header value 0xAA-AA-03. With bridged protocols the OUI value in the SNAP header is the 802.1 organization code 0x00-80-C2 and the actual type of the bridged media is specified by the two octet PID. Additionally, the PID indicates whether the original Frame Check Sequence (FCS) is preserved within the bridged PDU. The media type (PID) values that can be used in ATM encapsulation are listed in Appendix B.

The AAL5 CPCS-PDU Payload field carrying a bridged PDU shall, therefore, have one of the following formats. Padding is added after the PID field if necessary in order to align the user information field of the bridged PDU at a four octet boundary.

Payload Format for Bridged Ethernet/802.3 PDUs

```

+-----+
|          LLC  0xAA-AA-03          |
+-----+
|          OUI  0x00-80-C2          |
+-----+
|          PID  0x00-01 or 0x00-07  |
+-----+
|          PAD  0x00-00              |
+-----+
|          MAC destination address   |
+-----+
|          (remainder of MAC frame)  |
+-----+
|          LAN FCS (if PID is 0x00-01) |
+-----+


```

Payload Format for Bridged 802.4 PDUs


```


+-----+
|          LLC  0xAA-AA-03          |
+-----+
|          OUI  0x00-80-C2          |
+-----+
|          PID  0x00-02 or 0x00-08  |
+-----+
|          PAD  0x00-00-00          |
+-----+
|          Frame Control (1 octet)   |
+-----+
|          MAC destination address   |
+-----+
|          (remainder of MAC frame)  |
+-----+
|          LAN FCS (if PID is 0x00-02) |
+-----+

```



Note that the 802.5 Access Control (AC) field has no significance outside the local 802.5 subnetwork. It can thus be regarded as the last octet of the three octet PAD field, which can be set to any value (XX).






Note that in bridged 802.6 PDUs, there is only one choice for the PID value, since the presence of a CRC-32 is indicated by the CIB bit in the header of the MAC frame.

The Common Protocol Data Unit (PDU) Header and Trailer are conveyed to allow pipelining at the egress bridge to an 802.6 subnetwork. Specifically, the Common PDU Header contains the BAsize field, which contains the length of the PDU. If this field is not available to the egress 802.6 bridge, then that bridge cannot begin to transmit the segmented PDU until it has received the entire PDU, calculated the length, and inserted the length into the BAsize field. If the field is available, the egress 802.6 bridge can extract the length from the BAsize field of the Common PDU Header, insert it into the corresponding field of the first segment, and immediately transmit the segment onto the 802.6 subnetwork. Thus, the bridge can begin transmitting the 802.6 PDU before it has received the complete PDU.

Note that the Common PDU Header and Trailer of the encapsulated frame should not be simply copied to the outgoing 802.6 subnetwork because the encapsulated BETag value may conflict with the previous BETag value transmitted by that bridge.

An ingress 802.6 bridge can abort an AAL5 CPCS-PDU by setting its Length field to zero. If the egress bridge has already begun transmitting segments of the PDU to an 802.6 subnetwork and then

notices that the AAL5 CPCS-PDU has been aborted, it may immediately generate an EOM cell that causes the 802.6 PDU to be rejected at the receiving bridge. Such an EOM cell could, for example, contain an invalid value in the Length field of the Common PDU Trailer.




5. VC Based Multiplexing

In VC Based Multiplexing, the carried network interconnect protocol is identified implicitly by the VC connecting the two ATM stations, i.e. each protocol must be carried over a separate VC. There is therefore no need to include explicit multiplexing information in the Payload of the AAL5 CPCS-PDU. This results in minimal bandwidth and processing overhead.

As indicated above, the carried protocol can be either manually configured or negotiated dynamically during call establishment using signalling procedures. The signalling details will be defined later in other RFCs when the relevant standards have become available.



5.1. VC Based Multiplexing of Routed Protocols

PDU's of routed protocols shall be carried as such in the Payload of the AAL5 CPCS-PDU. The format of the AAL5 CPCS-PDU Payload field thus becomes:





5.2. VC Based Multiplexing of Bridged Protocols

PDUs of bridged protocols shall be carried in the Payload of the AAL5 CPCS-PDU exactly as described in section 4.2 except that only the fields after the PID field are included. The AAL5 CPCS-PDU Payload field carrying a bridged PDU shall, therefore, have one of the following formats.



Note that the 802.5 Access Control (AC) field has no significance outside the local 802.5 subnetwork. It can thus be regarded as the last octet of the three octet PAD field, which in case of 802.5 can be set to any value (XX).



In case of Ethernet, 802.3, 802.4, 802.5, and FDDI PDUs the presense or absence of the trailing LAN FCS shall be identified implicitly by the VC, since the PID field is not included. PDUs with the LAN FCS and PDUs without the LAN FCS are thus considered to belong to different protocols even if the bridged media type would be the same.

6. Bridging in an ATM Network

An ATM interface acting as a bridge must be able to flood, forward, and filter bridged PDUs.

Flooding is performed by sending the PDU to all possible appropriate destinations. In the ATM environment this means sending the PDU through each relevant VC. This may be accomplished by explicitly copying it to each VC or by using a multicast VC.

To forward a PDU, a bridge must be able to associate a destination MAC address with a VC. It is unreasonable and perhaps impossible to require bridges to statically configure an association of every

possible destination MAC address with a VC. Therefore, ATM bridges

must provide enough information to allow an ATM interface to dynamically learn about foreign destinations beyond the set of ATM stations.

To accomplish dynamic learning, a bridged PDU shall conform to the encapsulation described within section 4. In this way, the receiving ATM interface will know to look into the bridged PDU and learn the association between foreign destination and an ATM station.

7. For Further Study

Due to incomplete standardization of ATM multicasting, addressing, and signalling mechanisms, details related to the negotiation of the multiplexing method as well as address resolution had to be left for further RFCs.

Acknowledgements


This document has evolved from RFCs [1] and [4] from which much of the material has been adopted. Thanks to their authors T. Bradley, C. Brown, A. Malis, D. Piscitello, and C. Lawrence. In addition, the expertise of the ATM working group of the IETF has been invaluable in completing the document. Special thanks Brian Carpenter of CERN, Rao Cherukuri of IBM, Dan Grossman of Motorola, Joel Halpern of Network Systems, Bob Hinden of Sun Microsystems, and Gary Kessler of MAN Technology Corporation for their detailed contributions.

Security Considerations

Security issues are not addressed in this memo.


References

- [1] Piscitello, D. and Lawrence, C., "The Transmission of IP Datagrams over the SMDS Service". RFC 1209, Bell Communications Research, March 1991.
- [2] CCITT, "Draft Recommendation I.363". CCITT Study Group XVIII, Geneva, 19 - 29 January, 1993.
- [3] CCITT, "Draft Recommendation I.36x.1". CCITT Study Group XVIII, Geneva, 19-29 January, 1993.
- [4] Bradley, T., Brown, C., and Malis, A., "Multiprotocol Interconnect over Frame Relay". RFC 1294, Wellfleet Communications, Inc. and BBN Communications, January 1992.



Note that according to RFC 1294 the Q.922 Address field shall be either 2 or 4 octets, i.e., a 3 octet Address field is not supported.

In the particular case of a CLNP PDU, the NLPID is 0x81 and the FR-SSCS-PDU has the following format:



Note that in case of ISO protocols the NLPID field forms the first octet of the PDU itself and shall thus not be repeated.

The above encapsulation applies only to those routed protocols that have a unique NLPID assigned. For other routed protocols (and for bridged protocols), it is necessary to provide another mechanism for easy protocol identification. This can be achieved by using an NLPID value 0x80 to indicate that an IEEE 802.1a SubNetwork Attachment Point (SNAP) header follows.

See RFC 1294 for more details related to multiprotocol encapsulation over FRCS.

Appendix B. List of Locally Assigned values of OUI 00-80-C2

with preserved FCS	w/o preserved FCS	Media
-----	-----	-----
0x00-01	0x00-07	802.3/Ethernet
0x00-02	0x00-08	802.4
0x00-03	0x00-09	802.5
0x00-04	0x00-0A	FDDI
0x00-05	0x00-0B	802.6
	0x00-0D	Fragments
	0x00-0E	BPDUs

Appendix C. Partial List of NLPIDs

0x00	Null Network Layer or Inactive Set (not used with ATM)
0x80	SNAP
0x81	ISO CLNP
0x82	ISO ESIS
0x83	ISO ISIS
0xCC	Internet IP

Author's Address

Juha Heinanen
 Telecom Finland
 PO Box 228
 SF-33101 Tampere
 Finland

Phone: +358 49 500 958

Email: Juha.Heinanen@datanet.tele.fi