

Transport Working Group
Internet Draft
Expiration Date: November 2002

Matthew Bocci
Mustapha Aissaoui
John Fischer
Alcatel

Ghassem Koleyni
Nortel Networks

May 2002

Applicability Statement for AAL5 Transparent Frame Encapsulation over PSN

<[draft-bocci-pwe3-app-frame-over-psn-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

1. Abstract

This draft provides an applicability statement for the transparent AAL5 PDU frame mode encapsulation in [draft-fischer](#) [4].

Draft-fischer describes methods to carry ATM services over IP, L2TP or MPLS. The PSN (e.g. MPLS) is used to transport ATM layer services such as those defined by ITU-T as ATM transfer capabilities [2] and the ATM Forum as ATM service categories [3]. The basic requirement is to transparently transport ATM VCC service related information (e.g., traffic parameters, QoS, OAM, etc.) over the Pseudo Wire (PW), over the packet network. Transparent PDU frame mode enables bandwidth

efficiency gains to be realized for ATM VCCs that use AAL5, yet still provide full ATM transparency, including the correct sequencing of OAM cells on an AAL5 flow.

Table of contents

1.	Abstract.....	1
2.	Introduction.....	2
3.	Conventions used in this document.....	3
4.	Terminology.....	3
5.	Applicability Statement.....	4
5.1	Applicability.....	4
5.2	Implementation and deployment considerations.....	6
5.3	Limitations.....	6
6	ATM Service Encapsulation.....	7
6.1	Length and Sequence Number.....	7
6.1.1	Setting the length field.....	8
6.1.2	Processing the length field.....	9
6.1.3	Setting the sequence number.....	9
6.1.4	Processing the sequence number.....	9
7	ATM VCC Services.....	10
7.1	Transparent AAL5 PDU Frame Service.....	10
7.1.1	OAM Cell Support.....	12
7.1.2	Fragmentation.....	12
7.1.2.1	Procedures in the ATM-to-MPLS Direction.....	12
7.1.2.2	Procedures in the MPLS-to-ATM Direction.....	13
8	ILMI support.....	13
9	QoS considerations.....	13
10	ATM Pseudo-Wire over MPLS specific requirements.....	15
10.1	MPLS Transport Label.....	16
10.2	MPLS Pseudo Wire Label.....	16
11	ATM Pseudo-Wire over L2TP specific requirements.....	17
11.1	L2TP Session ID.....	18
11.2	Cookie.....	18
12	ATM Pseudo-Wire over IP specific requirements.....	18
12.1	C, K, and S bits.....	19
12.2	Protocol Type field.....	19
12.3	Key Field.....	19
12.4	GRE Sequence Number Field.....	19
13.	Security Considerations.....	20
14.	References.....	20
15.	Acknowledgement.....	21
16.	Author's Addresses.....	21

[2.](#) Introduction

This draft provides an applicability statement for the AAL5 transparent frame mode encapsulation in [draft-fischer](#) [4].

Draft-fischer describes methods to carry ATM services over an IP, L2TP or MPLS based PSN. The PSN is used to transport ATM layer services such as those defined by ITU-T as ATM transfer capabilities

[2] and by the ATM Forum as ATM service categories [3]. The basic requirement is to transparently transport the ATM VCC service related information (e.g., traffic parameters, QoS, OAM, etc.) over the Pseudo Wire (PW), over the packet network. The transparent AAL5 PDU mode is intended to be more efficient than the basic cell mode of [draft-fischer](#) [4], yet still provide full ATM transparency including the correct sequencing of OAM cells on an AAL5 flow.

Section 5 of this draft presents the applicability statement. Sections 6 to 13 are taken from [draft-fischer](#) [4]. They provide the details of the encapsulation methods as well as the OAM, ILMI, and QoS procedures for the basic cell mode.

3. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [16].

4. Terminology

Packet Switched Network - A Packet Switched Network (PSN) is a network using IP, MPLS or L2TP as the unit of switching.

Pseudo Wire Emulation Edge to Edge - Pseudo Wire Emulation Edge to Edge (PWE3) is a mechanism that emulates the essential attributes of a service (such as a T1 leased line or Frame Relay) over a PSN.

Customer Edge - A Customer Edge (CE) is a device where one end of an emulated service originates and terminates. The CE is not aware that it is using an emulated service rather than a "real" service.

Provider Edge - A Provider Edge (PE) is a device that provides PWE3 to a CE.

Pseudo Wire - A Pseudo Wire (PW) is a connection between two PEs carried over a PSN. The PE provides the adaptation between the CE and the PW.

Pseudo Wire PDU - A Pseudo Wire PDU is a PDU sent on the PW that contains all of the data and control information necessary to provide the desired service.

PSN Tunnel - A PSN Tunnel is a tunnel inside which multiple PWs can be nested so that they are transparent to core network devices.

Ingress - The point where the ATM service is encapsulated into a Pseudo Wire PDU (ATM to PSN direction.)

Egress - The point where the ATM service is de-capsulated from a Pseudo Wire PDU (PSN to ATM direction.)

CTD	Cell Transfer Delay
MTU	Maximum Transfer Unit
OAM	Operations, Administration, and Maintenance.
PVC	Permanent Virtual Connection. An ATM connection that is provisioning via a network management interface. The connection is not signalled.
VCC	Virtual Circuit Connection. An ATM connection that is switched based on the cell header's VCI.

5. Applicability Statement

5.1 Applicability

The primary application supported by AAL5 PDU frame encapsulation over PSN is the transparent carriage of ATM layer services that use AAL5 to carry higher layer frames. The PDU frame mode takes advantage of the delineation of higher layer frames in the ATM layer to provide increased bandwidth efficiency compared with the basic cell encapsulation mode of [draft-fischer](#) [4]. The nature of the service, as defined by the ATM service category [3] or the ATM transfer capability [2] should be preserved. To provide this, the basic requirement of the ATM-PSN interworking function is to map the AAL5 PDU frames belonging to a VCC, together with any related OAM and protocol control information, into a PW.

Two network applications that utilize the PDU frame mode encapsulation described in [draft-fischer](#) [4] are:

a. The transport of multi-service ATM over a packet core network where AAL5 is used as the adaptation layer. Many service providers have multiple service networks and the Operational Support System capabilities needed to support these existing service offerings. Packet Switched Networks (PSNs) have the potential to reduce the complexity of a service provider's infrastructure by allowing virtually any existing digital service to be supported over a single networking infrastructure.

The benefits of this model to a service provider are threefold:

- i. Leveraging of the existing systems and services to provide increased capacity from a packet switched core.
- ii. Preserving existing network operational processes and procedures used to maintain the legacy services e.g. ATM OAM and ATM security.
- iii. Using the common packet switched network infrastructure to support both the core capacity requirements of existing services and the requirements of new services supported natively over the packet switched network.

b. L2 VPN service over a PSN infrastructure. In this case, VPN sites are connected through ATM VCCs, as in today's L2 VPNs. The transparent PDU frame mode encapsulation allows the VPN service provider to transparently extend this L2 connectivity over its PSN while achieving bandwidth efficiency gains over the basic cell mode and supporting ATM layer applications of the VPN customer, such as ATM security. The advantage is for the service provider to combine L2 and L3 services over the same PSN.

Figure 1 shows the reference model for carrying ATM services over a PSN. This model is adapted from [6].

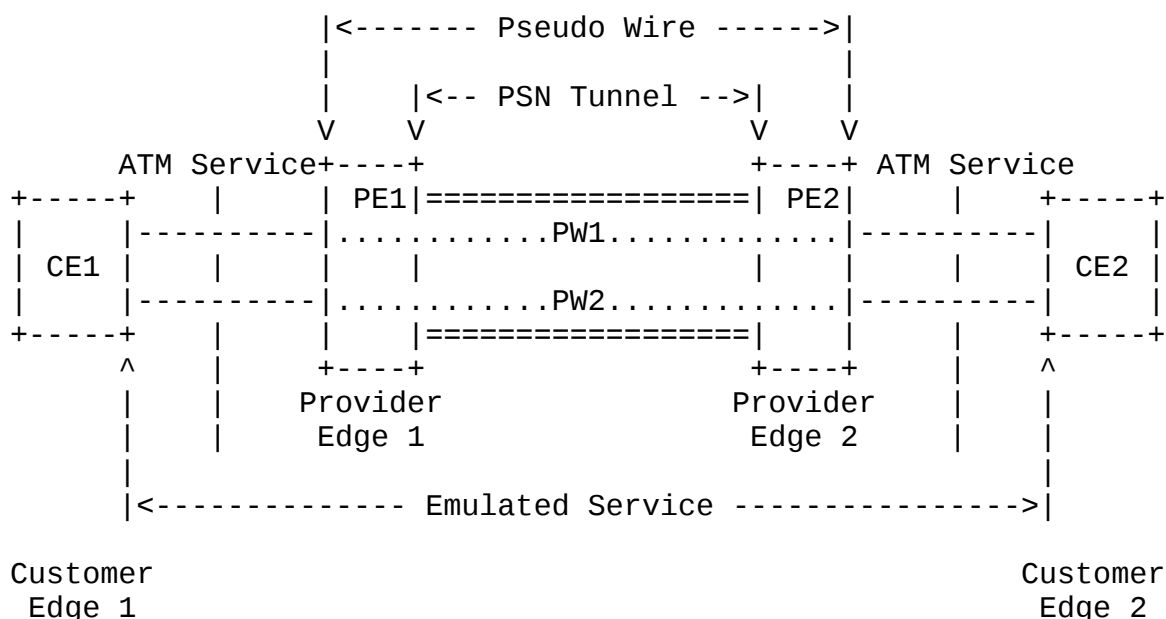


Figure 1: ATM-over-PSN Service Reference Model

An ATM VCC is carried over a PW. The PW corresponding to any VCC may be further tunneled in a transport PSN tunnel to achieve multiplexing gain and bandwidth efficiency.

When the QoS considerations in [Section 10](#) are respected, this ATM over PSN service encapsulation method provides end users with the same quality of service on any given VCC as per corresponding SLA, traffic contracts and the QoS commitments for that connection.

One important consideration to make when interworking is to allow OAM information to be treated as in the original network. The interworking function allows this transparency while performing AAL5 frame encapsulation. Fragmentation may be performed in order to maintain the position of the OAM cells with respect to the user cells.

Fragmentation may also be performed to maintain the size of the packet carrying the AAL5 PDU within the MTU of the link.

Cell Loss priority (CLP) field conveys the priority of the cell in the connection. The Explicit Forward Congestion Indicator (EFCI) field conveys the congestion state of ATM network. Information on both of these fields is obtained from the ATM cell header. CLP and EFCI fields are both part of the ATM service specific information header.

For ease of operation and to achieve transparency, the whole AAL5-PDU is encapsulated. In this case all necessary parameters such as CPCS-UU (CPCS User-to-User indicator), CPI (Common Part Indicator), Length (Length of the CPCS-SDU) and CRC (Cyclic Redundancy Check) are transported as part of the payload. Note that carrying of the full PDU also allows the simplification of the fragmentation operation since it is performed at cell boundaries and the CRC in the trailer of the AAL5 PDU can be used to check the integrity of the reassembled fragments.

5.2 Implementation and deployment considerations

AAL5 transparent mode is only applicable to services that use AAL5 to carry higher layer frames over ATM VCCs.

5.3. Limitations

AAL5 frame encapsulation only supports point-to-point LSPs. Multi-point-to-point and point-to-multi-point are for further study (FFS).

To have bi-directional connectivity, as required in ATM, two LSPs should be configured, one for each direction (ATM-to-MPLS and MPLS-to-ATM) of the ATM connection.

Length of AAL5 frame may exceed the MTU of the PSN. This requires fragmentation, which may not be available to all nodes at the PW endpoint.

The maximum number of cells of an AAL5 PDU that may be reassembled before transport across the PSN may be limited by the cell transfers delay (CTD) and cell delay variation (CDV) objectives of the connection.

This mode does not preserve the value of the CLP bit for every ATM cell within an AAL5 PDU. Therefore, transparency of the CLP setting may be violated. Additionally, tagging of some cells may occur when tagging is not allowed by the conformance definition.

This mode does not preserve the EFCI state for every ATM cell within an AAL5 PDU. Therefore, transparency of the EFCI state may be violated.

6 ATM Service Encapsulation

This section describes the general encapsulation format for ATM over PSN pseudo wires, such as IP, L2TP, or MPLS. The specifics pertaining to each packet technology are covered in later sections. All encapsulation formats and procedures contained in the following sections are from [draft-fischer](#) [4].

Figure 2 provides a general format for encapsulation of ATM cells into packets.

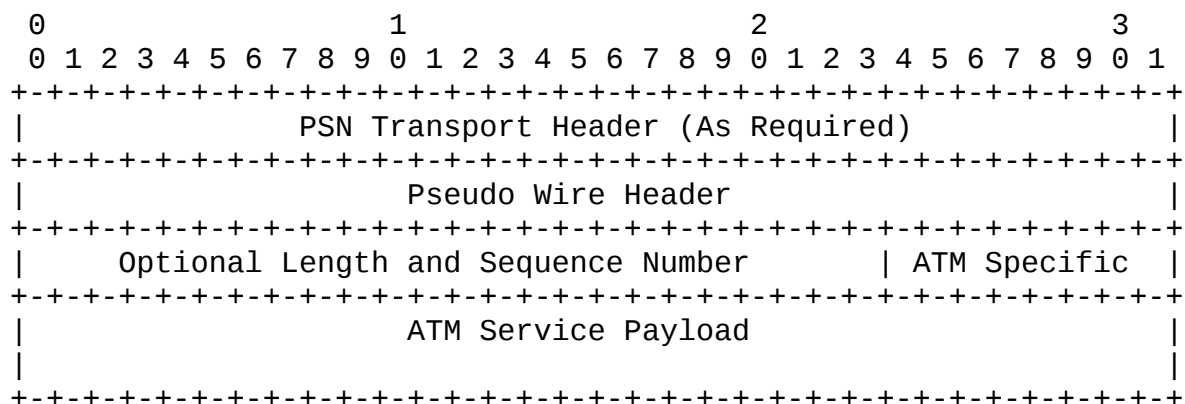


Figure 2: General format for ATM encapsulation over PSNs

The PSN Transport Header depends on the packet technology: IP, L2TP or MPLS. This header is used to transport the encapsulated ATM information through the packet switched core. This header is always present if the Pseudo Wire is MPLS.

The Pseudo Wire Header depends on the packet technology: IP, L2TP or MPLS. It identifies a particular ATM service within the PSN tunnel.

The Length and Sequence Number is inserted after the Pseudo Wire Header. This field is optional.

The ATM Specific Header is inserted before the ATM service payload. The ATM Specific Header contains control bits needed to carry the service. These are defined in the ATM service descriptions below. The length of ATM specific header may not always be one octet. It depends on the service type.

The ATM payload octet group is the payload of the service that is being encapsulated.

6.1 Length and Sequence Number

The length and sequence number are not required for all services. Length and sequence number are to satisfy these requirements.

- Order may need to be preserved.
- Small packets may need to be padded in order to be transmitted on a medium where the minimum transport unit is larger than the actual packet size.

The one-octet Length indicates length of the packet payload that includes, length of the length field, Sequence number length, the ATM specific header length and the payload length (i.e., Pseudo Wire PDU). The Length field is set to 0 by the ingress PE if not used and is ignored by the egress PE.

If the Pseudo Wire traverses a network link that requires a minimum frame size such as Ethernet as a practical example, with a minimum frame size of 64 octets, then such links will apply padding to the Pseudo Wire PDU to reach its minimum frame size. In this case the length field MUST be set to the PDU length. A mechanism is required for the egress PE to detect and remove such padding.

The Sequence Number is a 2-octet field that may be used to track packet order delivery. This field is set to 0 by the ingress PE if not used and is ignored by the egress PE. The sequence number space is a 16-bit, unsigned circular space. Processing of the sequence number field is OPTIONAL.

In all cases the egress PE MUST be aware of whether the ingress PE will send the length and sequence number over a specific Pseudo Wire. This may be achieved using static configuration or using Pseudo Wire specific signaling.

Length field is not required for cell mode

6.1.1 Setting the length field

The length field is required to enable the egress PE to remove any padding that may have resulted if the pseudo-wire traversed a network link that enforces a minimum frame size (e.g. Ethernet). Ethernet applies padding to frames that are smaller than 64 bytes, where this includes a minimum of 18 bytes for the Ethernet header and trailer.

The length field represents the size of the packet in bytes including the length, sequence number, ATM specific header and ATM service payload. If the size of the packet is larger than can be represented by the length field, the field MUST be set to 0. In addition, the length field MAY be set to 0 if the size of the packet prevents any padding from being applied.



The AAL5 SDU Frame service is the only service that can generate packets smaller than the Ethernet minimum and MUST set the length field accordingly. The length field MUST either be set to the size of the packet if the size is less than 46 bytes or to 0.

All cell transport services MUST always set the length field to 0 to indicate to the remote PE that no padding was applied.

[6.1.2](#) Processing the length field

Since length field is not used for cell mode, no processing is required.

[6.1.3](#) Setting the sequence number

The following procedures SHOULD be used by the ingress PE if sequencing is desired for a given ATM service:

- the initial PDU transmitted on the Pseudo Wire MUST use sequence number 1
- the PE MUST increment the sequence number by one for each subsequent PDU
- when the transmit sequence number reaches the maximum 16 bit value (65535) the sequence number MUST wrap to 1

If the ingress PE does not support sequence number processing, then the sequence number field in the control word MUST be set to 0.

[6.1.4](#) Processing the sequence number

If the egress PE supports receive sequence number processing, then the following procedures SHOULD be used:

When an ATM service is initially created, the "expected sequence number" associated with it MUST be initialized to 1.

When a PDU is received on the Pseudo Wire associated with the ATM service, the sequence number SHOULD be processed as follows:

- if the sequence number on the packet is 0, then the PDU passes the sequence number check
- otherwise if the PDU sequence number \geq the expected sequence number and the PDU sequence number - the expected sequence number < 32768 , then the PDU is in order.
- otherwise if the PDU sequence number $<$ the expected sequence number and the expected sequence number - the PDU sequence number ≥ 32768 , then the PDU is in order.
- otherwise the PDU is out of order.

If a PDU passes the sequence number check, or is in order then, it can be delivered immediately. If the PDU is in order, then the expected sequence number SHOULD be set using the algorithm:

```

expected_sequence_number := PDU_sequence_number + 1 mod 2**16
  if (expected_sequence_number = 0)
    then expected_sequence_number := 1;

```

Pseudo Wire PDUs that are received out of order MAY be dropped or reordered at the discretion of the egress PE.

If the egress PE does not support receive sequence number processing, then the sequence number field MAY be ignored.

7 ATM VCC Services

This section defines ATM AAL5 VCC services that may be supported over the PSN.

7.1 Transparent AAL5 PDU Frame Service

In this mode, the ingress PE encapsulates the entire CPCS-PDU including the PAD and trailer.

This mode MAY support fragmentation in order to maintain OAM cell sequencing.

Like the ATM AAL5 payload VCC service, the AAL5 transparent VCC service is intended to be more efficient than the VCC cell transport service. However, the AAL5 transparent VCC service carries the entire AAL5 CPCS-PDU, including the PAD and trailer. Note that the AAL5 CPCS-PDU is not processed _ i.e. an AAL5 frame with an invalid CRC or length field will be transported. One reason for this is that there may be a security agent that has scrambled the ATM cell payloads that form the AAL5 CPCS-PDU.

This service supports all OAM cell flows by using a fragmentation procedure that ensures that OAM cells are not repositioned in respect to AAL5 composite cells.

The AAL5 transparent VCC service is OPTIONAL.

[illegible]




Figure 3: AAL5 transparent service encapsulation

The first octet following the Pseudo Wire Header carries control information. The M, V, Res, and C bits are as defined earlier for VCC cell mode.

* Frg (Fragmentation) Bits

This field is used to support the fragmentation functionality described later in this section.

- 11 Single Segment Message (Beginning and End of Message)
- 10 Beginning of Message
- 00 Continuation of Message
- 01 End of Message

* E (EFCI) bit

This field is used to convey the EFCI state of the ATM cells. The EFCI state is indicated in the middle bit of each ATM cell's PTI field.

ATM-to-PSN direction (ingress): The EFCI field of the control byte is set to the EFCI state of the last cell of the AAL5 PDU or AAL5 fragment.

PSN-to-ATM direction (egress): The EFCI state of all constituent cells of the AAL5 PDU or AAL5 fragment is set to the value of the EFCI field in the control byte.

* C (CLP) bit

This field is used to convey the cell loss priority of the ATM cells.

ATM-to-PSN direction (ingress): The CLP field of the control byte is set to 1 if any of the constituent cells of the AAL5 PDU or AAL5 fragment has its CLP bit set to 1; otherwise this field is set to 0.

PSN-to-ATM direction (egress): The CLP bit of all constituent cells for an AAL5 PDU or AAL5 fragment is set to the value of the CLP field in the control byte.

The payload consists of the re-assembled AAL5 CPCS-PDU,

including the AAL5 padding and trailer or the AAL5 fragment.

[7.1.1](#) OAM Cell Support

When configured for the AAL5 transparent VCC service, both PE's SHOULD act as a VC switch, in accordance with the OAM procedures defined in [5].

The PEs SHOULD be able to pass the following OAM cells transparently:

- F5 AIS (segment and end-to-end)
- F5 RDI (segment and end-to-end)
- F5 loopback (segment and end-to-end)
- Resource Management
- Performance Management
- Continuity Check
- Security

The PEs SHALL use the single ATM VCC cell mode encapsulation ([Section 6.1](#), [draft-fischer\[4\]](#)) when passing an OAM cell.

The ingress PE SHOULD be able to generate an F5 AIS upon reception of a corresponding F4 AIS or lower layer defect (such as LOS).

The egress PE SHOULD be able to generate an F5 AIS based on a PSN failure (such as a PSN tunnel failure or LOS on the PSN port).

If the ingress PE cannot support the generation of OAM cells, it MAY notify the egress PE using a Pseudo Wire specific maintenance mechanism to be defined. For example, the ingress PE MAY withdraw the Pseudo Wire (VC label) associated with the service. Upon receiving such a notification, the egress PE SHOULD generate the appropriate F5 AIS.

[7.1.2](#) Fragmentation

The ingress PE may not always be able to reassemble a full AAL5 frame. This may be due to the AAL5 PDU exceeding the Pseudo Wire MTU or when OAM cells arrive during reassembly of the AAL5 PDU. In these cases, the AAL5 PDU shall be fragmented. In addition, fragmentation may be desirable to bound ATM cell delay.

If no fragmentation occurs, then the fragmentation bits are set to 11 (SSM, Single Segment Message).

When fragmentation occurs, the procedures described in the following subsections shall be followed.

[7.1.2.1](#) Procedures in the ATM-to-MPLS Direction

The following procedures shall apply while fragmenting AAL5 PDUs:

- Fragmentation shall always occur at cell boundaries within the

AAL5 PDU.

- For the first fragment, the FRG bits shall be set to 10 (BOM, Beginning Of Message).
- For the last fragment, the FRG bits shall be set to 01 (EOM, End Of Message).
- For all other fragments, the FRG bits shall be set to 00 (COM, Continuation Of Message).
- The E and C bits of the fragment shall be set as defined earlier in [section 7](#).

[7.1.2.2](#) Procedures in the MPLS-to-ATM Direction

The following procedures shall apply:

- The 3-bit PTI field of each ATM cell header is constructed as follows:
 - + The most significant bit is set to 0, indicating a user data cell.
 - + The middle bit is set to the E bit value of the fragment.
 - + The least significant bit is set to 1 for the last ATM cell of a fragment where the FRG bits are 01 (EOM) or 11(SSM); otherwise this bit is set to 0.
- The C bit of each ATM cell header is set to the value of the C bit of the control byte in Figure 5.

[8](#) ILMI support

Integrated Local Management Interface (ILMI) typically is used in ATM networks for neighbor resource availability detection, address registration, auto-configuration, and loss of connectivity detection. ILMI messages are sent as SNMP PDU's within ATM AAL5 cells.

A PE MAY provide an ATM ILMI to its attached CE. If the ingress PE receives an ILMI message indicating that the ATM service (VCC) is down, it MAY use a Pseudo Wire specific mechanism to notify the egress PE of the ATM service status. For example, a PE using an MPLS based Pseudo Wire may withdraw its advertised VC label.

When receiving such a notification, the egress PE MAY use ILMI to signal the ATM service status to its attached CE.

[9](#) QoS considerations

This section provides guidelines for the provision of QoS for the individual ATM PWs over the PSN. The objective is to provide the ability to support the traffic contracts and the QoS commitments made to the ATM connections [\[8\]](#). This section is informational and the provided guidelines SHOULD be treated as good practices and the mappings as examples only.

When ATM PW service is configured over a PSN, each ATM service category SHOULD be mapped to a compatible class of service in the PSN

network. A compatible class of service maintains the integrity of the service end to end. For example, the CBR service category SHOULD be mapped to a class of service with stringent loss and delay objectives. If the PSN implements the IP Diff-Serv framework to provide QoS classes, a class of service based on the EF PHB is a good candidate.

Furthermore, ATM service categories have support for multiple conformance definitions. A conformance definition specifies the conformance of cells of a connection at an interface, e.g., public UNI, in relation to the conformance algorithm and corresponding parameters specified in the connection traffic descriptor [15]. For example, the conformance definition specifies if the requested QoS parameters, e.g., CLR, apply to the aggregate CLP0+1 conforming cell flow or to the CLP0 conforming flow only. Thus, the conformance definition SHOULD be respected in the selected PSN class of service.

For example, a connection CLP1 cell flow in a VBR.3 conformance definition is treated as excess traffic in the ATM network and thus has no QoS guarantees associated with it. This flow SHOULD be provided a treatment no better than the treatment of the CLP0 cell flow in the PSN. This does not mean however that the PSN network should mirror the various conformance definitions of the ATM service categories.

In the remainder of this section, it is assumed that the PSN implements IP Diff-Serv framework to provide QoS.

All ATM traffic management functions specified in [3] are applicable for the ATM part of the ATM PW in the ingress PE and egress PE. In the ATM-to-PSN direction, each ATM connection MAY be policed and/or shaped to conform to its traffic descriptor in the ATM endpoint of the ATM PW in the PE. Whenever allowed by the conformance definition, a non-conforming CLP0 cell may be turned into a CLP1 cell and becomes conforming. Connection admission SHOULD be applied to make sure sufficient resources are available to carry the ATM PW over the transport LSP. The mapping of ATM service category and conformance definition to the Diff-Serv PHB determines the outgoing PHB. This is the PHB to be applied to the cells or packets of the ATM PW in the ingress PE and egress PE and inside the PSN. The PSN transport header of the outgoing PSN packet SHOULD be marked to identify the selected PHB. This consists of marking the DS field in the IP header in the case of IP PSN, or the EXP field in the transport shim header in the case of a MPLS PSN.

Figure 4 provides an example of mapping ATM service category and conformance definition to a Diff-Serv PHB.

ATM Service Category	Conformance Definition	CLP Setting	Diff-Serv PHB	DSCP Marking
-------------------------	---------------------------	----------------	------------------	-----------------

CBR	CBR.1	0/1	EF	101110
rt-VBR	VBR.1	0/1	EF	101110
	VBR.2/VBR.3	0	AF11	001010
		1	AF12	001100
nrt-VBR	VBR.1	0/1	AF21	010010
	VBR.2/VBR.3	0	AF21	010010
		1	AF22	010100
ABR	ABR	0	AF31	011010
UBR+MDCR	UBR.1/UBR.2	0/1	AF31	011010
GFR	GFR.1/GFR.2	0	AF31	011010
		1	AF32	011100
UBR	UBR.1/UBR.2	0/1	DF	000000

Figure 4: Example of ATM Service Category to PHB Mapping

Note that an alternative mapping may not distinguish between the conformance definitions in a given ATM service category. In this case, the CLP0 and CLP1 flows of a ATM connection are provided with the same QoS in the PSN. As an example, all conformance definitions of the nrt-VBR service category MAY be mapped to the AF21 PHB in Figure 9.

When the PSN is MPLS based, the selected PHB for the ATM PW is conveyed in different ways depending if the transport LSP is an L-LSP or an E-LSP [15]. In the case of an L-LSP, the PHB Scheduling Class is signaled at the transport LSP establishment and is therefore inferred from the transport label value. The Drop Precedence of the individual PW packets is conveyed in the EXP field of the transport LSP shim header. In the case of an E-LSP, the PHB is conveyed in the EXP field of the transport LSP shim header. The actual values to be marked in the EXP field to reflect the example in Figure 9 is outside the scope of this document.

In the presence of congestion, the PE MAY mark the EFCI bit and MAY perform selective cell discard based on CLP setting, if allowed by the conformance definition, and in accordance with [15]. The method used to transfer the CLP and EFCI information of the individual cells into the ATM specific field of the PW packet is described in details in sections 7 and 8.

In the PSN-to-ATM direction, the ATM service category and conformance definition is obtained from the context accessed via the pseudo wire label of the ATM PW. The information needed to reconstruct the ATM header, including the setting of the CLP and EFCI fields, for the individual cells is contained in the ATM specific information field of the PW packet. The method used to determine the CLP and EFCI information of the individual cells from the ATM specific information field of the PW packet is described in details in sections 6 and 7.

[10](#) ATM Pseudo-Wire over MPLS specific requirements

Figure 5 provides a general format for interworking between ATM and MPLS. The ATM information is encapsulated inside two MPLS labels as defined in [9].

The Pseudo Wire Endpoint uses a unique MPLS label, the pseudo wire label, to identify each direction of an ATM connection. This label allows the PWE to access context information for the connection. As an example, the context may contain the information regarding connection type such as VCC or VPI/VCI value that need to be inserted into the ATM cell header in the MPLS-to-ATM direction.

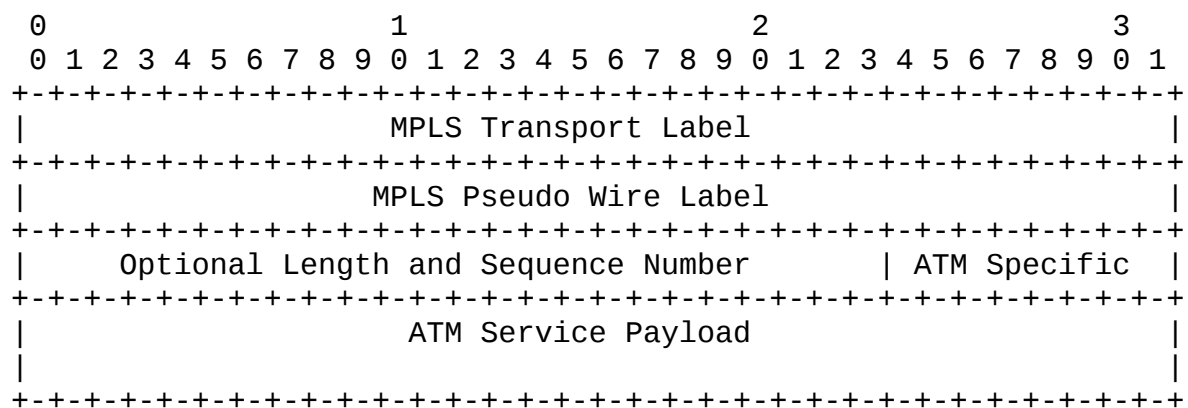


Figure 5: Format for ATM PW over a MPLS PSN

[10.1](#) MPLS Transport Label

The 4-octet MPLS transport label identifies an LSP used to transport traffic between two ATM-MPLS pseudo wire endpoints. This label is used to switch the transport LSP between core LSRs. Since an MPLS LSP is unidirectional, for the case of bi-directional traffic there will be two different pseudo wire LSPs, one for each direction of the connection. These may have different label values. Setting of the EXP and TTL is for further study. The S bit is set to 0 since this is not the last label in the MPLS label stack.

[10.2](#) MPLS Pseudo Wire Label

The 4-octet interworking label uniquely identifies one pseudo wire LSP, carried inside a MPLS transport LSP. The pseudo wire label has the structure of a standard MPLS shim header. More than one pseudo wire LSP may be supported by one MPLS transport LSP. The pseudo wire endpoint provides the association between the ATM connection or the ATM port and MPLS LSP by means of the 20-bit label field of the pseudo wire LSP. In this association, in the ATM-to-MPLS direction a mapping of the VCI/VPI of the ATM connection or the Port to the 20-bit label is performed, while in the MPLS-to-ATM direction the 20-bit label is mapped to a VPI/VCI of the ATM connection or to a Port. This association is signalled or provisioned between the two pseudo-wire endpoints.

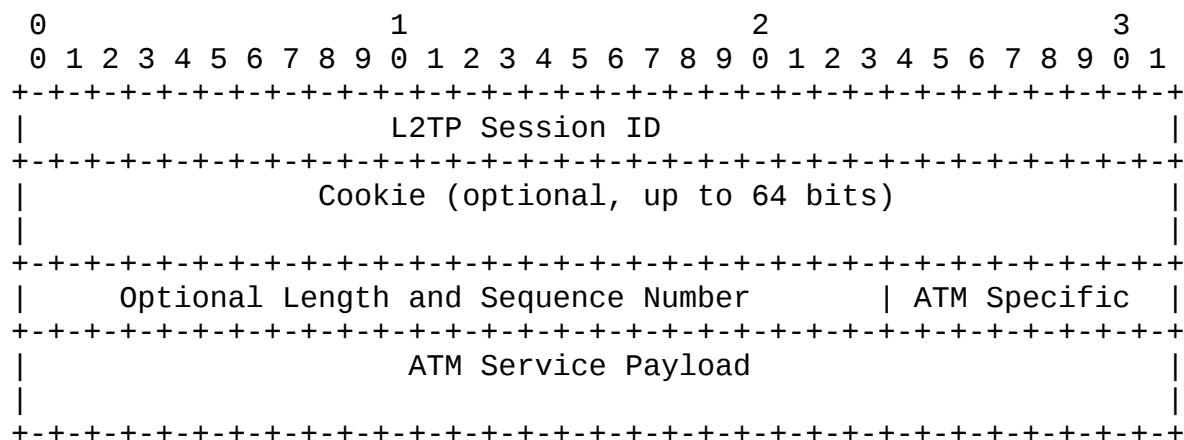
Since an MPLS LSP is unidirectional, for the case of bi-directional ATM VCCs there will be two different pseudo wire LSPs, one for each direction of the connection. These may have different label values. Setting of the EXP and TTL is for further study. The S bit is set to 1 since this is the last label in the bottom of the MPLS stack. The pseudo wire label is not visible to the LSRs within the MPLS core network.

11 ATM Pseudo-Wire over L2TP specific requirements

Figure 6 provides a general format for interworking between ATM and L2TP. This draft refers to L2TPv3, or L2TP base, as described in [11]. L2TP base extends the original L2TP [12] to operate directly over an IP PSN and to further generalize the control procedures to carry any tunneled Layer 2 protocol other than PPP. Any further reference to L2TP in this document assumes L2TPv3 or L2TP base as specified in [11].

The ATM information is encapsulated inside a L2TP tunnel packet. The L2TP tunnel is setup over an IPv4 PSN. The IPv4 protocol in the IPv4 header is set to 115 to indicate that the L2TP packet is encapsulated in an IPv4 packet [11]. Furthermore, L2TP can operate over IP or over UDP. The use of either mode is outside the scope of this document. The encapsulation format shown in Figure 11 represents the common headers for carrying L2TP packet over UDP and IP. If UDP is used, additional header information is required and is defined in [11].

The Pseudo Wire Endpoint uses a unique L2TP session ID to identify each direction of an ATM connection. This session ID is local to each PE and allows the PWE to identify each ATM PW in the L2TP tunnel in order to access context information for the ATM connection. As an example, the context may contain the information regarding connection type such as VCC VPI/VCI value that need to be inserted into the ATM cell header in the L2TP-to-ATM direction. Multiple PWs with locally unique Session IDs at each PE can be multiplexed into a L2TP tunnel.



11.1 L2TP Session ID


Each PE for the life of the session will give different Session IDs the same PW. Multiple PWs with locally unique Session IDs at each PE can be multiplexed into a L2TP tunnel. When the L2TP control connection is used for session establishment, Session IDs are selected and exchanged as Local Session ID Attribute Value Pairs (AVPs) during the creation of a PW. A session ID of zero is reserved for the carriage of L2TP control messages [11].

The optional Cookie field contains a variable length (maximum 64 bits), long word-aligned value used to check the association of a received packet with the PW identified by the Session ID. The Cookie MUST be configured with a random value utilizing all bits in the field [11]. The Cookie provides an additional level of guarantee, beyond the Session ID lookup, that a packet has been directed to the proper PW identified by the Session ID.

When the L2TP control connection is used for PW session establishment, random Cookie values are selected and exchanged as Assigned Cookie AVPs during the creation of a PW. The maximum size of the Cookie field is 64 bits.

Figure 7 provides a general format for carrying an ATM PW over an IP PSN. This is an alternative encapsulation to the one using L2TP, as described in [Section 11](#). The GRE encapsulation is used as specified in [\[13\]](#) and [\[14\]](#). The IPv4 protocol in the IPv4 header is set to 47 to indicate that the GRE packet is encapsulated in an IPv4 packet [\[13\]](#).

The ATM information is encapsulated inside a GRE/IP packet. The Pseudo Wire Endpoint uses a unique GRE Key to identify each direction of an ATM connection. As an example, the context may contain the information regarding connection type such as VCC VPI/VCI value that need to be inserted into the ATM cell header in the IP-to-ATM direction. Multiple PWs with unique GRE Keys can be multiplexed into a GRE/IP tunnel.



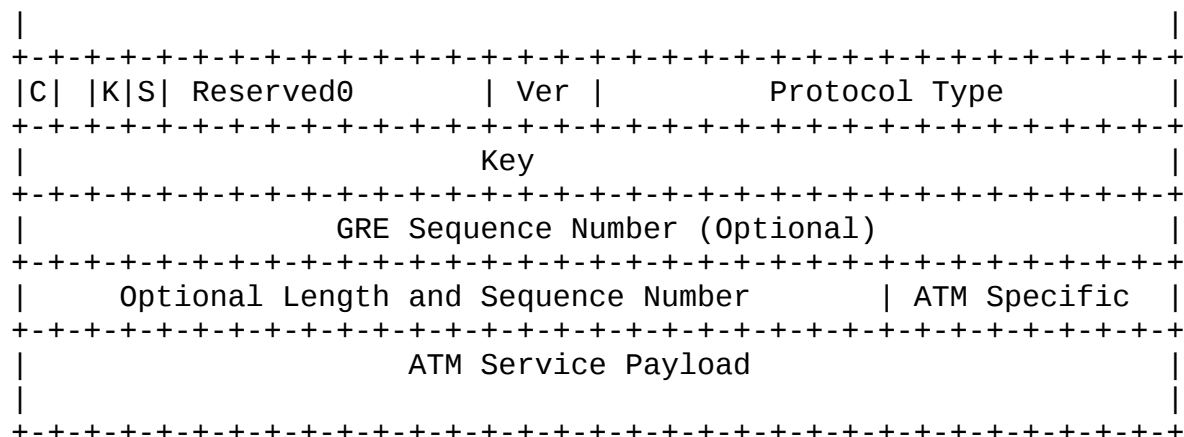


Figure 7: Format for ATM PW over an IP PSN

[12.1](#) C, K, and S bits

The Checksum field in the GRE header is not required for carrying ATM PW over IP PSN. Therefore the C bit is set to zero.

The Key field in the GRE header is always used (see [Section 12.3](#)). Therefore, the K bit is always set to 1.

If the GRE Sequence Number field is used, then the value of the K bit is set to 1. Otherwise, its value is set to zero.

[12.2](#) Protocol Type field

The Protocol Type field is set to a number to be allocated by IEEE for the purpose of encapsulating ATM PW over GRE.

[12.3](#) Key Field

The Key field contains a four-octet number that is inserted by the transmitting PE. The Key field is used by the receiving PE for identifying an individual ATM PW within a GRE/IP tunnel. Multiple PWs with unique GRE Keys can be multiplexed into a GRE/IP tunnel. The method by which the Key field value is negotiated between the transmitting PE and a receiving PE is further study.

[12.4](#) GRE Sequence Number Field

If the Sequence Number Present bit is set to 1, then it indicates that the Sequence Number field is present in the GRE header. Otherwise, the Sequence Number field is not present in the GRE header. The use of the Sequence Number field MUST comply to [\[14\]](#).

A specific ATM PWE network may choose to rely on the GRE protocol to track in-order delivery of ATM PW packets. Another way of tracking

in-order delivery is to use the PW Sequence number field as explained in [Section 5.1](#).

[13](#). Security Considerations

No additional security issues are introduced in this document. As ATM encapsulation to MPLS packet is related to MPLS. AAL5 frame encapsulation shares the security concerns associated with MPLS.

[14](#). References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] ITU-T Recommendation I.371 (2000), Traffic control and congestion control in B-ISDN.
- [3] ATM Forum Specification af-tm-0121.000 (1999), Traffic Management Specification Version 4.1.
- [4] IETF [draft-fischer-pwe3-atm-service-03.txt](#) (March 2002, work in progress), PWE3: ATM service description
- [5] ITU-T Recommendation I.610, (1999), B-ISDN operation and maintenance principles and functions.
- [6] IETF [draft-ietf-pwe3-requirements-02.txt](#) (November 2001, work in progress), Requirements for pseudo Wire Emulation Edge-to-Edge (PWE3).
- [7] IETF [draft-martini-l2circuit-encap-mpls-04.txt](#) Martini (November 2001, Work in Progress), Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks.
- [8] IETF [draft-koleyni-pwe3-atm-over-mpls-04.txt](#) (January 2002, Work in Progress), Requirements and framework for ATM network interworking over MPLS
- [9] IETF [RFC 3032](#) (2001), MPLS Label Stack Encoding.
- [10] ATM Forum af-aic-0178.000 (2001), ATM-MPLS Network Interworking.
- [11] IETF [draft-ietf-l2tpext-l2tp-base-01.txt](#) (October 2001, Work in Progress), Layer Two Tunneling Protocol (L2TP).
- [12] IETF [RFC 2661](#) (1999), Layer Two Tunneling Layer Two Tunneling Protocol (L2TP).
- [13] IETF [RFC 2784](#) (2000), Generic Routing Encapsulation (GRE).
- [14] IETF [RFC 2890](#) (2000), Key and Sequence Number Extensions to GRE.

- [15] IETF [draft-ietf-mpls-diff-ext-09.txt](#) (April 2001, Work in Progress), MPLS Support of Differentiated Services.
- [16] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[15](#). Acknowledgements

The authors like to acknowledge the contribution to this work by: TBD

[16](#). Author's Addresses

John Fischer
Alcatel
600 March Rd
Kanata, ON, Canada. K2K 2E6
Email: john.fischer@alcatel.com

Matthew Bocci
Alcatel
Grove House, Waltham Road Rd
White Waltham, Berks, UK. SL6 3TN
Email: matthew.bocci@alcatel.co.uk

Mustapha Aissaoui
Alcatel
600 March Rd
Kanata, ON, Canada. K2K 2E6
Email: mustapha.aissaoui@alcatel.com

Ghassem Koleyni
Nortel Networks
P O Box 3511, Station C Ottawa, Ontario,
K1Y 4H7 Canada
Email: ghassem@nortelnetworks.com

Full Copyright Statement

"Copyright (C) The Internet Society (May 2002). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into

