

PWE3 Working Group
Internet Draft
Expiration Date: April 2002

Jeremy Brayley
Gerald de Grace
John Shirron
Laurel Networks, Inc.

Rick Wilder
Masergy Communications

Nasser El-Aawar
Level 3 Communications, LLC.

Dimitri Stratton Vlachos
Mazu Networks, Inc.

Andrew G. Malis
Vinai Sirkay
Vivace Networks, Inc.

Tom Johnson
Litchfield Communications, Inc.

Jayakumar Jayakumar
Durai Chinniah

Chris Liljenstolpe
Cable & Wireless

Dan Tappan
Eric Rosen
Cisco Systems, Inc.

John RuteMiller
Marconi Networks

Laura Dominik
Qwest Communications, Inc.

Giles Heron
PacketExchange Ltd.

November 2001

PWE3: ATM service description
draft-brayley-pwe3-atm-service-01.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [section 10 of RFC 2026](#) [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/lid-abstracts.txt>.

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

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

Abstract

Generic requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3) have been described in [3]. This draft provides encapsulation formats and guidelines from transporting ATM services over a Packet Switched Network using IP, L2TP, or MPLS. It describes three cell relay services: VCC, VPC, and transparent port, as well as two methods for carrying AAL5 PDUs across a PSN.

Table of Contents

1	Conventions used in this document.....	2
2	Introduction.....	3
3	Terminology.....	4
4	General Requirements.....	5
5	ATM Service Encapsulation.....	5
5.1	ATM control Word.....	6
5.1.1	Setting the sequence number.....	7
5.1.2	Processing the sequence number.....	7
6	ATM Cell Relay Services.....	8
6.1	VCC Cell Relay Service.....	8
6.1.1	OAM Cell Support.....	8
6.2	VPC Cell Relay Service.....	9
6.2.1	OAM Cell Support.....	9
6.3	Transparent Port Cell Relay Service.....	10
6.3.1	OAM Cell Support.....	11
6.4	Cell Mode Encapsulation.....	11
7	Frame Based ATM Services.....	13
7.1	AAL5 Payload VCC Service.....	13
7.1.1	OAM Cell Support.....	15
7.2	AAL5 Transparent VCC Service.....	15
7.2.1	OAM Cell Support.....	17
7.2.2	Fragmentation.....	18
8	ILMI support.....	19
9	QoS considerations.....	19
10	Security Considerations.....	20
11	Intellectual Property Disclaimer.....	20
12	References.....	20
13	Authors' Addresses.....	21

[1](#) 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](#) [2].

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

[2](#) Introduction

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 benefit of this model to a service provider is threefold:

1. Leveraging of the existing systems and services to provide increased capacity from a packet switched core.
2. Preserving existing network operational processes and procedures used to maintain the legacy services.
3. 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.

This draft describes a method to carry ATM services over IP, L2TP and MPLS. It lists ATM specific requirements and provides encapsulation formats and semantics for connecting ATM edge networks through a core packet network using IP, L2TP or MPLS. The techniques described in this draft will allow ATM service providers to take advantage of new technologies in the core in order to provide ATM multi-services.

Figure 1, below displays the ATM services reference model. This model is adapted from [\[3\]](#).

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

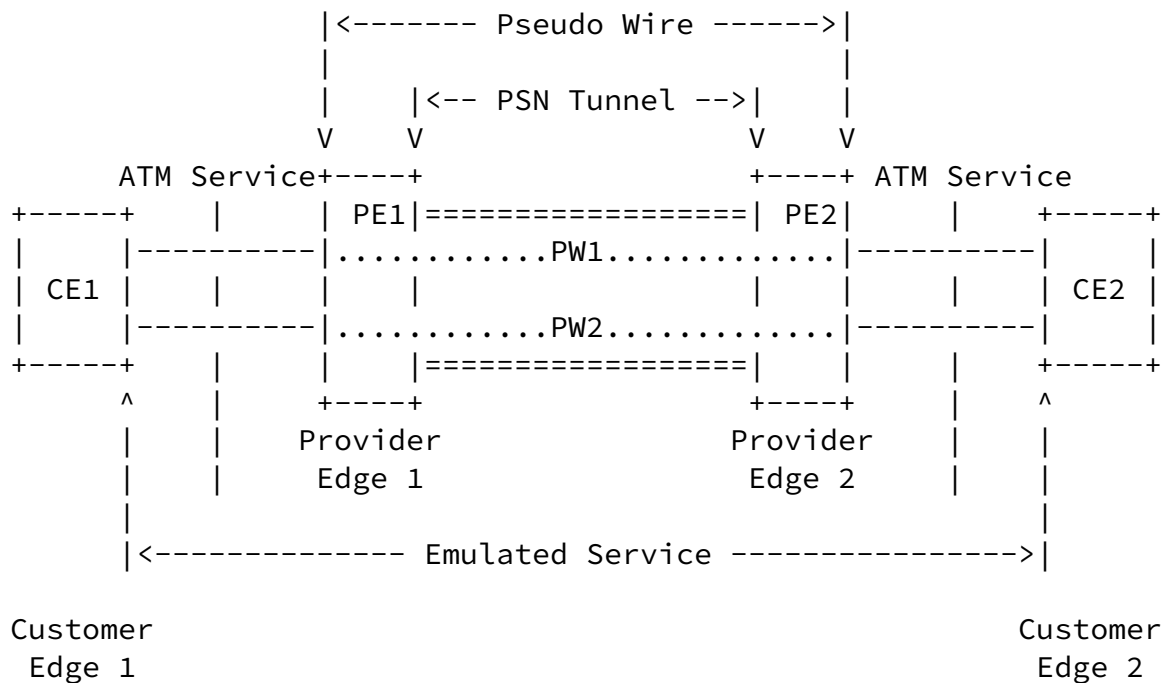


Figure 1: ATM Service Reference Model

3 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

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

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

[4](#) General Requirements

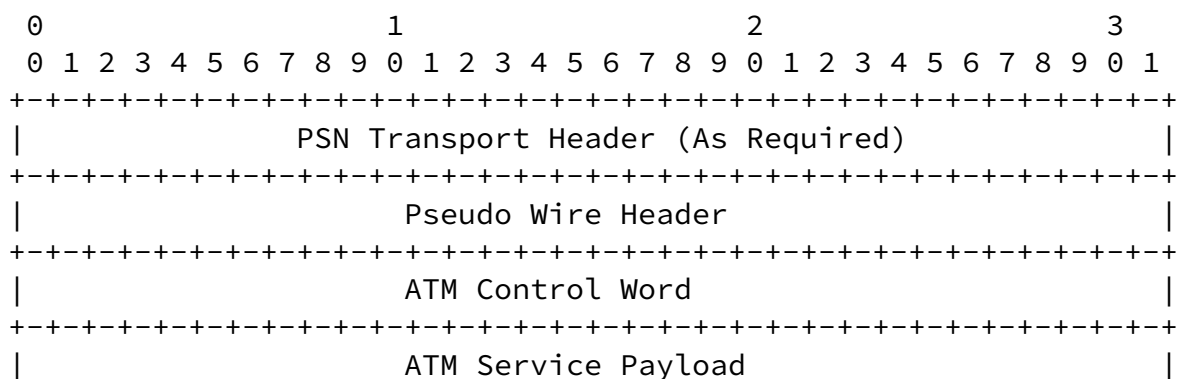
For transport of an ATM service across a PSN, the PSN SHOULD be able to:

1. Carry all AAL types transparently.
2. Carry multiple ATM connections (VPCs and/or VCCs).
3. Support ATM OAM applications.
4. Transport Cell Loss Priority (CLP) and Payload Type Indicator (PTI) information from the ATM cell header.
5. Provide a mechanism to detect mis-ordering of ATM cells over the PSN.
6. Support traffic contracts and the QoS commitments made to the ATM connections (through the use of existing IETF defined Diff-Serv techniques).

[5](#) 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.

Figure 2 provides a general format for encapsulation of ATM cells (or frames) 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.

Brayley, et al.

Expires April 2002

[Page 5]

Internet Draft draft-brayley-pwe3-atm-service-01.txt November 2001

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

The ATM Control Word is inserted before the ATM service payload. It may contain a length and sequence number in addition to certain control bits needed to carry the service.

The ATM Service Payload is specific to the service being offered via the Pseudo Wire. It is defined in the following sections.


[5.1](#) ATM control Word

The ATM control word is part of the ATM specific header. It is not required for all services. The control word is designed to satisfy three requirements.

- Ability to detect out of order delivery of PDUs.
- Ability to detect padding added by certain link technologies.
- Control bits for ATM services.

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

The control word is defined as follows:



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

5.1.2 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.

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

- 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.

6 ATM Cell Relay Services

This section defines three types of ATM services that may be supported over the PSN: ATM VCC, ATM VPC, and ATM transparent port.

[6.1](#) VCC Cell Relay Service

The VCC cell relay service is characterized by the mapping of a single ATM VCC (VPI/VCI) to a Pseudo Wire. This service is fully transparent to the ATM Adaptation Layer.

The egress PE may choose to apply a different VCI other than the one that arrived at the ingress PE. The egress PE **MUST** choose the outgoing VCI based solely upon the Pseudo Wire header.

The VCC cell relay service is **REQUIRED**.

This service **MUST** use the cell mode encapsulation defined in [section 6.4](#).

[6.1.1](#) OAM Cell Support

When configured for a VCC cell relay service, both PEÆs **SHOULD** act as a VC switch in accordance with the OAM procedures defined in [\[4\]](#).

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 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.

[6.2](#) VPC Cell Relay Service

The VPC service is defined by mapping a single VPC (VPI) to a Pseudo Wire. As such it emulates as Virtual Path cross-connect across the PSN. All VCCs belonging to the VPC are carried transparently by the VPC service.

This service MUST use the cell mode encapsulation defined in [section 6.4](#).

The egress PE may choose to apply a different VPI other than the one that arrived at the ingress PE. The egress PE MUST choose the outgoing VPI based solely upon the Pseudo Wire header. As a VPC service, the egress PE MUST NOT change the VCI field.

The VPC cell relay service is REQUIRED.

[6.2.1](#) OAM Cell Support

When configured for a VPC cell relay service, both PEÆs SHOULD act as a VP cross-connect in accordance with the OAM procedures defined in [\[4\]](#).

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

- F4 AIS (segment and end-to-end)
- F4 RDI (segment and end-to-end)

- F4 loopback (segment and end-to-end)
- 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 ingress PE MUST be able to generate an F4 AIS upon reception of a lower layer defect (such as LOS).

The egress PE SHOULD be able to generate an F4 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 F4 AIS.

[6.3](#) Transparent Port Cell Relay Service

Transparent port encapsulation is used to emulate an ATM Port-to-Port connection over a PSN. This service is useful when one desires to connect two CEs without interfering at the VPC or VCC layer. The ingress PE SHOULD discard any idle/unassigned cells received from the ingress ATM port, and map all other received cells to a single Pseudo Wire.

The egress PE SHOULD not change the VPI, VCI, PTI, or CLP bits when it sends these cells on the egress ATM port. This service appears as a Layer 1 service (such as SONET/SDH) to CE devices. However the service provider benefits from increased transport efficiency due to statistical multiplexing.

The transparent port cell relay service may be used to migrate ATM services to a PSN without affecting current provisioning systems. Although this service specifies a mapping of an entire ATM port to a Pseudo Wire and PSN tunnel, the ingress PE should be able to inspect incoming cell headers to assign QoS characteristics on the PSN. For example, this allows a service provider to denote a group of VPC's or VCC's that receive a specific QoS treatment on the PSN, without requiring the service provider to use the VCC or VPC cell relay service.

This service MUST use the cell mode encapsulation defined in [section 6.4](#).

The transparent port cell relay service is OPTIONAL.

[6.3.1](#) OAM Cell Support

This service is completely transparent to the F4 and F5 OAM layer. The PEs MUST pass all OAM and resource management cells.

If the ingress PE detects a physical layer defect (such as LOS) it SHOULD be able to notify the egress PE via a mechanism specific to the Pseudo Wire in use. When it receives such a notification, the egress PE SHOULD propagate the failure (such as sending a SONET Line AIS).

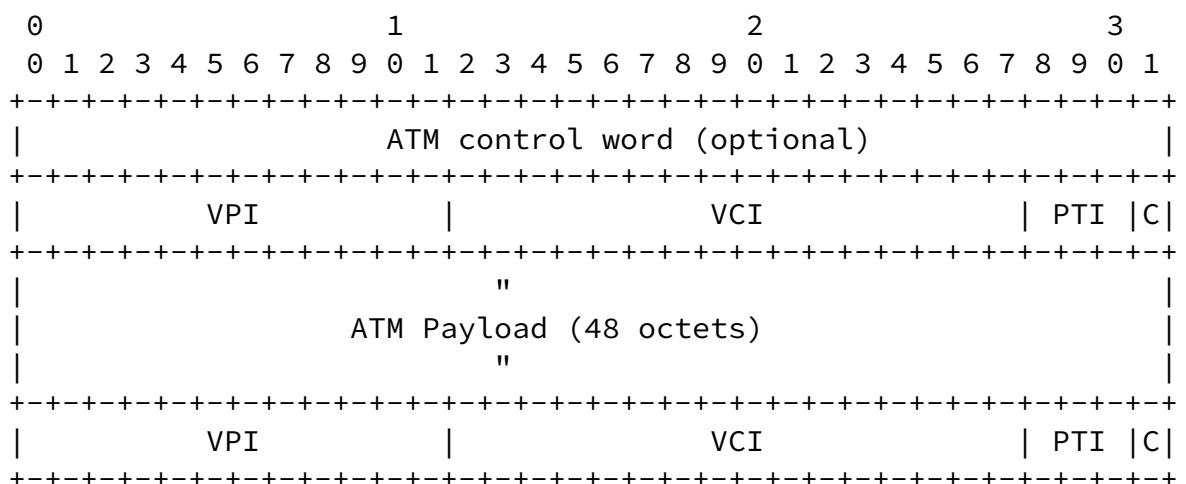
If the ingress PE cannot support the generation of OAM cells, it MAY notify the egress PE using a Pseudo Wire specific maintenance mechanism to be defined. For example, the ingress PE MAY withdraw the Pseudo Wire (VC label) associated with the service. Upon receiving such a notification, the egress PE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

6.4 Cell Mode Encapsulation

The encapsulation defined in this section is used for all cell relay services. The cell mode does not involve any higher layer AAL processing such as AAL5 segmentation and reassembly.

The cell mode MUST support the encapsulation of a single ATM cell into a Pseudo Wire PDU.

For increased transport efficiency, the ingress PE SHOULD be able to encapsulate multiple ATM cells into a Pseudo Wire PDU. The ingress and egress PE SHOULD agree to a maximum number of cells in a single Pseudo Wire PDU. This agreement may be accomplished via a Pseudo Wire specific signaling mechanism or via static configuration.






Figure 4: trunk mode encapsulation

The ATM control word is defined in [section 5.1](#). Its use is OPTIONAL. The control word is necessary only if sequencing is desired. If the ATM control word is used, then the Flag and Length fields should be set to 0 upon transmission and ignored upon receipt.

6.4.1 Review of header information

A review of the ATM cell header in the encapsulation mode is defined in this section. The review of the header is OPTIONAL. While information carried in the cell encapsulation is carried transparently through the packet based network, and does not require a SAR function, there is a need to review the header information of the traffic being transported. Inspection of the header information provides a mechanism to map characteristics of the transported information to the PSN. Each cell is inspected at the PE device and service requirements are mapped accordingly in the packet based network.

An example, when implementing cell encapsulation mode for ATM traffic, is the review of the ATM header of each cell. It is through this examination that control mechanisms such as congestion management can be translated for transport in the PSN. This capability could also be used to support the mapping of ATM QoS to CoS.

Direct examination of the header provides a view of the CLP field and the payload type (PT) field on a per cell basis. The PT field provides the AAL indication, upstream congestion, and discrimination between data and OAM cells. Payload types carrying user information can also indicate whether congestion was

experienced by EFCI or whether the cell contains an indication to the AAL protocol.

A specific implementation is the mapping of the CLP bit into a MPLS based core network. In order to emulate drop precedence on a MPLS

tunnel, the CLP bit is associated with a pair of configurable MPLS EXP values. Cells with CLP = 0 are encapsulated into a MPLS packet with EXP = 000. Cells with CLP = 1 are encapsulated with EXP = 001. This information is carried in all levels of labels as necessary.

An additional requirement, for further study, is the mapping of the ATM EPD function into a corresponding PSN function.

7 Frame Based ATM Services

7.1 AAL5 Payload VCC Service

The AAL5 payload VCC service defines a mapping between the payload of an AAL5 VCC and a single Pseudo Wire. Therefore, it requires ATM segmentation and reassembly support on the PE. The AAL5 payload VCC service is designed to be more efficient than the VCC cell relay service.

The AAL5 payload VCC service is OPTIONAL.

Even the smallest TCP packet requires two ATM cells when sent over AAL5 on a native ATM device. It is desirable to avoid this padding on the Pseudo Wire. Therefore, once the ingress PE reassembles the AAL5 CPCS-PDU, the PE discards the PAD and CPCS-PDU trailer then inserts the resulting payload into a Pseudo Wire PDU.

The egress PE MUST regenerate the PAD and trailer before transmitting the AAL5 frame on the egress ATM port.

This service does allow the transport of OAM and RM cells, but does not attempt to maintain the relative order of these cells with respect to the cells that comprise the AAL5 CPCS-PDU. OAM cells that arrive during the reassembly of a single AAL5 CPCS-PDU are sent immediately on the Pseudo Wire, followed by the AAL5 payload. Therefore, the AAL5 payload VCC service may not be suitable for some ATM applications that require strict ordering of OAM cells (such as performance monitoring and security applications).

The AAL5 payload service encapsulation is shown below.

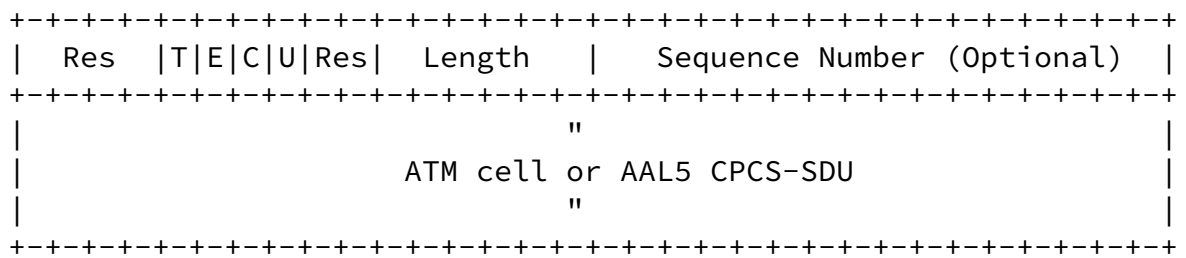


Figure 5: AAL5 payload service encapsulation

The AAL5 payload service encapsulation requires the ATM control word. The Flag bits are described below.

* Res (Reserved)

These bits are reserved and MUST be set to 0 upon transmission and ignored upon reception.

* T (Transport type) bit

Bit (T) of the control word indicates whether the packet contains an ATM cell or an AAL5 payload. If T = 1, the packet contains an ATM cell, encapsulated according to the VCC cell relay encapsulation of [section 6.4](#). If not set, the PDU contains an AAL5 payload. The ability to transport an ATM cell in the AAL5 mode is intended to provide a means of enabling OAM functionality over the AAL5 VCC.

* E (EFCI) Bit

The ingress PE device SHOULD set this bit to 1 if the EFCI bit of the final cell of those that transported the AAL5 payload is set to 1, or if the EFCI bit of the single ATM cell to be transported in the packet is set to 1. Otherwise this bit SHOULD be set to 0. The egress PE device SHOULD set the EFCI bit of all cells that transport the AAL5 payload to the value contained in this field.

* C (CLP) Bit

The ingress PE device SHOULD set this bit to 1 if the CLP bit of any of the ATM cells that transported the AAL5 payload are set to 1, or if the CLP bit of the single ATM cell that is to be transported in the packet is set to 1. Otherwise this bit SHOULD be set to 0. The egress PE device SHOULD set the CLP bit of all cells that transport the AAL5 CPCS-PDU to the value contained in this field.

* U (Command/Response) Bit

When FRF.8.1 Frame Relay / ATM PVC Service Interworking (see [5]) traffic is being transported, the CPCS-UU Least Significant Bit (LSB) of the AAL5 CPCS-PDU may contain the Frame Relay C/R bit.

The ingress PE device SHOULD copy this bit to the U bit of the control word. The egress PE device SHOULD copy the U bit to the CPCS-UU Least Significant Bit (LSB) of the AAL5 payload.

The Length and Sequence Number fields are described in [section 5.1](#).

[7.1.1](#) OAM Cell Support

Similar to the VCC cell relay service, both PEs SHOULD act as a VC switch in accordance with the OAM procedures defined in [4].

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
- Continuity Check

Because this service does not guarantee the original OAM cell position within the AAL5 composite cells, the following cell types are discarded by the ingress PE:

- Performance Management
- Security

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.

[7.2](#) AAL5 Transparent VCC Service

Like the ATM AAL5 payload VCC service, the AAL5 transparent VCC

service is intended to be more efficient than the VCC cell relay service. However, the AAL5 transparent VCC service carries the

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](http://www.ietf.org/internet-drafts/draft-brayley-pwe3-atm-service-01.txt) November 2001

entire AAL5 CPCS-PDU, including the PAD and trailer. 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.

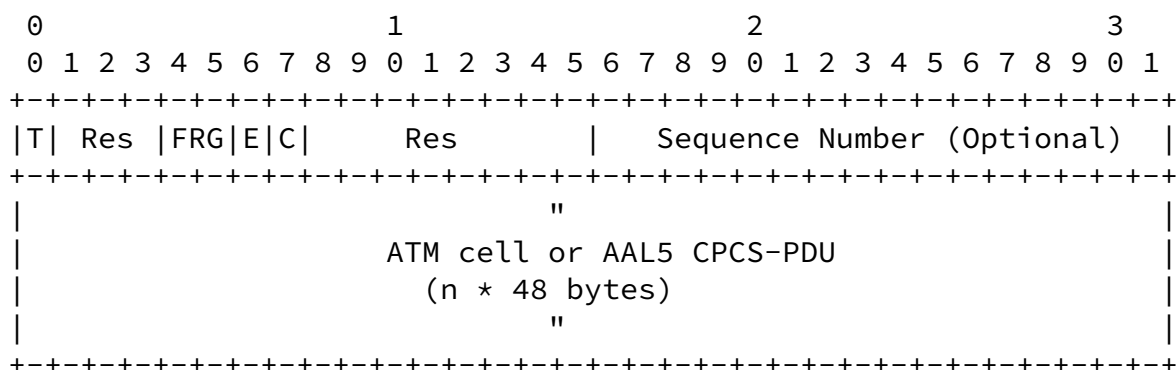


Figure 6: AAL5 transparent service encapsulation

The first octet following the Pseudo Wire Header carries control information.

* Res (Reserved)

These bits are reserved and MUST be set to 0 upon transmission and ignored upon reception.

* T (Transport type) bit

Bit (T) of the control word indicates whether the packet contains an ATM cell. If T = 1, the packet contains an ATM cell, encapsulated according to the VCC cell relay encapsulation of [section 6.4](#). If T = 0, the PDU contains an AAL5 CPCS-PDU or CPCS-PDU fragment. The ability to transport a single ATM cell is intended to enable OAM functionality for the VCC.

* FRG (Fragmentation) Bits

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

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

* E (EFCI) bit

The ingress PE device SHOULD set this bit to 1 if the EFCI bit of the final cell of those that transported the AAL5 CPCS-PDU or CPCS-PDU fragment is set to 1, or if the EFCI bit of the single ATM cell to be transported in the Pseudo Wire PDU is set to 1. Otherwise this bit SHOULD be set to 0. The egress PE device SHOULD set the EFCI bit of all cells that transport the AAL5 CPCS-PDU or CPCS-PDU fragment to the value contained in this field.

* C (CLP) bit

The ingress PE device SHOULD set this bit to 1 if the CLP bit of any of the ATM cells that transported the AAL5 CPCS-PDU or CPCS-PDU fragment are set to 1, or if the CLP bit of the single ATM cell that is to be transported in the Pseudo Wire PDU is set to 1. Otherwise this bit SHOULD be set to 0. The egress PE device SHOULD set the CLP bit of all cells that transport the AAL5 CPCS-PDU or CPCS-PDU fragment to the value contained in this field.

The payload consists of a complete AAL5 CPCS-PDU, including the AAL5 padding and trailer or CPCS-PDU fragment.

[7.2.1](#) OAM Cell Support

When configured for the AAL5 transparent VCC service, both PEAs SHOULD act as a VC switch, in accordance with the OAM procedures defined in [\[4\]](#).

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 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).

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

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.

[7.2.2](#) 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 FRG bits are set to 11 (SSM, Single Segment Message).

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

[7.2.2.1](#) 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).

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

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.

Brayley, et al.

Expires April 2002

[Page 18]

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

[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 PDUs 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 or VPC) 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

The ingress PE should have the ability to maintain separation of ATM traffic classes (i.e. CBR, rt-VBR, nrt-VBR, ABR, and UBR) for each of the services transported across the PSN. The mechanism

used to maintain these traffic classes depends upon the PSN in use. For example, does the PSN support resource assignments per PSN tunnel? Can it support per PSN tunnel queuing?

The actual mechanisms to support the ATM traffic classes should be left up to the operator. This section offers some suggestions.

QoS assignment on the PSN requires close inspection of incoming cell headers. This includes mapping the VPI/VCI to a specific PSN traffic class and using the CLP bit to determine the PSN drop precedence. For example, when processing incoming cells for a CBR VCC service, the ingress PE may mark the outgoing Pseudo Wire PDUs with a particular DSCP or MPLS EXP. (Marking depends upon the PSN in use.) Downstream PSN devices should use this marking to map these PW PDU's to queuing and scheduling resources that emulate an ATM CBR service (i.e. low latency, guaranteed bandwidth).

If the PSN is MPLS based, the ingress PE may associate ATM services with E-LSPs or L-LSPs as defined in [9].

The PSN should also have the ability to maintain the ATM cell loss priority (CLP) of incoming cells. For example, in case of an MPLS based PSN, the ingress PE may mark both the PSN transport and Pseudo Wire labels with EXP = 010 or EXP = 011 depending upon the incoming cell's CLP value. (If the PW PDU contains multiple ATM cells the ingress PE should not mark the PW PDU to convey a single drop precedence.) For AAL5 services, the ingress PE should mark

the PW PDU using the same algorithm that determines the C (CLP) bit (i.e if any cell has CLP = 1 then the C bit should be set to 1.)

The following is an example of mapping ATM service classes and CLP to a Diff-Serv capable PSN.

ATM traffic class	CLP	PSN marking

CBR	0	DSCP=000110 or EXP=110
CBR	1	DSCP=000111 or EXP=111
rt-VBR	0	DSCP=000100 or EXP=100
rt-VBR	1	DSCP=000101 or EXP=101
nrt-VBR	0	DSCP=000010 or EXP=010
nrt-VBR	1	DSCP=000011 or EXP=011
UBR	0	DSCP=000000 or EXP=000
UBR	1	DSCP=000001 or EXP=001

10 Security Considerations

This draft does not introduce any new security considerations to IP, L2TP or MPLS.

11 Intellectual Property Disclaimer

This document is being submitted for use in IETF standards discussions.

12 References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997
- [3] "Requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3)", [draft-ietf-pwe3-requirements-01.txt](#), Work in Progress
- [4] ITU-T Recommendation I.610 (February 1999): B-ISDN operation and maintenance principles and functions
- [5] "Frame Relay / ATM PVC Service Interworking Implementation Agreement (FRF.8.1)", Frame Relay Forum 2000.
- [6] "Frame Based ATM over SONET/SDH Transport (FAST)," ATM Forum 2000.
- [7] "Encapsulation Methods for Transport of Layer 2 Frames Over MPLS", [draft-martini-l2circuit-encap-mpls-03.txt](#), Work in Progress

- [8] "Requirements and framework for ATM network interworking over MPLS", [draft-koleyni-pwe3-atm-over-mpls-02.txt](#), Work in Progress
- [9] "MPLS Support of Differentiated Services", [draft-ietf-mpls-diff-ext-09.txt](#), Work in Progress

13 Authors' Addresses

Jeremy Brayley
Laurel Networks, Inc.
1300 Omega Drive
Pittsburgh, PA 15205
Email: jbrayley@laurelnetworks.com

Gerald de Grace
Laurel Networks, Inc.
1300 Omega Drive
Pittsburgh, PA 15205
Email: gdegrace@laurelnetworks.com

John Shirron
Laurel Networks, Inc.
1300 Omega Drive
Pittsburgh, PA 15205
Email: jshirron@laurelnetworks.com

Nasser El-Aawar
Level 3 Communications, LLC.
1025 Eldorado Blvd.
Broomfield, CO, 80021
Email: nna@level3.net

Andrew G. Malis
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Email: Andy.Malis@vivacenetworks.com

Vinai Sirkay
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Email: Vinai.Sirkay@vivacenetworks.com

Jayakumar Jayakumar
Cisco Systems, Inc.
225, E.Tasman, MS-SJ3/3,
San Jose, CA, 95134
Email: jjayakum@cisco.com

Durai Chinnaiah
Cisco Systems, Inc.
225, E.Tasman, MS-SJ3/3,
San Jose, CA, 95134
Email: dchinnai@cisco.com

Dan Tappan
Cisco Systems, Inc.
50 Apollo Drive
Chelmsford, MA, 01824
Email: tappan@cisco.com

Eric Rosen
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA, 01824
Email: erosen@cisco.com

Rick Wilder
Masergy Communications
2901 Telestar Ct.
Falls Church, VA 22042
Email: rwilder@masergy.com

Dimitri Stratton Vlachos
Mazu Networks, Inc.
125 Cambridgepark Drive
Cambridge, MA 02140
Email: d@mazunetworks.com

Thomas K. Johnson
Litchfield Communications, Inc.
76 Westbury Park Rd.
Watertown, CT 06795
Email: tom_johnson@litchfieldcomm.com

Chris Liljenstolpe
Cable & Wireless
11700 Plaza America Drive
Reston, VA 20190
Email: chris@cw.net

John Rutemiller
Marconi Networks
1000 Marconi Drive
Warrendale, PA 15086
Email: John.Rutemiller@marconi.com

Giles Heron
PacketExchange Ltd.
The Truman Brewery

Internet Draft [draft-brayley-pwe3-atm-service-01.txt](#) November 2001

91 Brick Lane
LONDON E1 6QL
United Kingdom
Tel.: +44 7880 506185
Email: giles@packetexchange.net

Laura Dominik
Qwest Communications, Inc.
600 Stinson Blvd.
Minneapolis, MN, 55413
Email: ldomini@qwest.com

