Network Working Group Luca Martini
Internet Draft Eric C. Rosen
Expiration Date: October 2004 Cisco Systems, Inc.

Jeremy Brayley Matthew Bocci
Laurel Networks, Inc. Alcatel

Ghassem Koleyni
Nortel Networks.

April 2004

Encapsulation Methods for Transport of ATM Over IP and MPLS Networks

draft-ietf-pwe3-atm-encap-05.txt

Status of this Memo

Abstract

A framework for providing various Layer 1 and Layer 2 services over a Packet Switched Network has been described in [3]. This draft provides encapsulation formats and guidelines for transporting a variety of ATM services over a PSN.

Table of Contents

1. Specification of Requirements

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. Introduction

Packet Switched Networks (PSNs) have the potential to reduce the complexity of a service providers 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:

- Leveraging of the existing systems and services to provide increased capacity from a packet switched core.
- Preserving existing network operational processes and procedures used to maintain the legacy services.
- 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 document describes a method to carry ATM services over 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 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].

```
                  |<----- Pseudo Wire ---->|
                  |                        |
                  |  |<-- PSN Tunnel -->|  |
     ATM Service  V  V                  V  V  ATM Service
          |     +----+                  +----+     |
+----+    |     | PE1|==================| PE2|     |    +----+
|    |----------|............PW1.............|----------|    |
| CE1|    |     |    |                  |    |     |    |CE2 |
|    |----------|............PW2.............|----------|    |
+----+    |     |    |==================|    |     |    +----+
     ^          +----+                  +----+     |    ^
     |      Provider Edge 1         Provider Edge 2     |
     |                                                  |
     |<-------------- Emulated Service ---------------->|
Customer                                                Customer
Edge 1                                                  Edge 2
```

Figure 1: ATM Service Reference Model

QoS related issues are not discussed in this draft. This draft describes two methods of ATM cell encapsulation, One-to-one mode and N-to-one mode. This draft describes two methods of AAL5 encapsulation, PDU mode and SDU mode.

3. Terminology

One-to-one mode: The One-to-one mode specifies an encapsulation method which maps one ATM VCC (or one ATM VPC) to one Pseudo Wire.

N-to-one mode (N >= 1): The N-to-one mode specifies an encapsulation method which maps one or more ATM VCCs (or one or more ATM VPCs) to one Pseudo Wire.

Packet Switched Network - A Packet Switched Network (PSN) is an IP or MPLS network.

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 a service originates and/or terminates. The CE is not aware that it is using an emulated service rather than a native 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 PSN devices.

PSN Bound - The traffic direction where information from a CE is adapted to a PW, and PW-PDUs are sent into the PSN.

CE Bound - The traffic direction where PW-PDUs are received on a PW from the PSN, re-converted back in the emulated service, and sent out to a CE.

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 decapsulated from a Pseudo Wire PDU (PSN to ATM direction.)

CTD - Cell Transfer Delay

MTU - Maximum Transmission Unit

OAM - Operations And Maintenance.

PVC - Permanent Virtual Connection. An ATM connection that is provisioned via a network management interface. The connection is not signaled.

VCC Virtual Circuit Connection. An ATM connection that is switched based on the cell header's VCI.

VPC - Virtual Path Connection. An ATM connection that is switched based on the cell header's VPI.

4. General encapsulation method

This section describes the general encapsulation format for ATM over PSN pseudo wires.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Pseudo Wire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     ATM Control Word                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     ATM Service Payload                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: General format for ATM encapsulation over PSNs

The PSN Transport Header depends on the particular tunneling technology in use (L2TP or MPLS). This header is used to transport the encapsulated ATM information through the packet switched core.

The Pseudo Wire Header identifies a particular ATM service on a tunnel. In case of MPLS the Pseudo Wire Header is the MPLS label at the bottom of the MPLS label stack. In the Case of L2TP the Pseudo Wire Header is the L2TP header.

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.

4.1. MPLS Shim S Bit Value

The ingress LSR, PE1, MUST set the S bit of the PW label to a value of 1 to denote that the VC label is at the bottom of the stack.

4.2. MPLS Shim TTL Values

The setting of the TTL value in the PW label is application dependent, however in a strict point to point application the TTL SHOULD be appropriately set to 2.

4.3. The Control Word

There are four requirements that may need to be satisfied when transporting layer 2 protocols over an IP or MPLS backbone [8]:

-i. Sequentiality may need to be preserved.
-ii. 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.
-iii. Control bits carried in the header of the layer 2 frame may need to be transported.
-iv. To allow accurate packet inspection in an MPLS PSN, and/or to operate correctly over MPLS PSNs that have deployed equal-cost multiple-path load-balancing, a PW packet MUST NOT alias an IP packet.

The PWE3 architecture document describes a generic control word and a preferred control word. This document makes use of both of these control words depending on the encapsulation mode. Both of these control words addresses all of the above requirements.

For some encapsulation modes, the control word is REQUIRED, and for others OPTIONAL. Where the control word is OPTIONAL implementations MUST support sending no control word, and MAY support sending a control word.

In all cases the egress router must be aware of whether the ingress router will send a control word over a specific pseudo wire. This may be achieved by configuration of the routers, or by signaling, for example as defined in [1].

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 control word must include a length field set to the PDU length. A mechanism is required for the egress PE to detect and remove such padding.

4.3.1. The Generic Control Word

This control word is used in the following encapsulation modes:

- ATM 1 to 1 Cell Mode

- AAL5 PDU Frame Mode

The PWE3 architecture document [8] provides the following structure for the generic control word:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0|          Specified by PW Encapsulation                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The detailed structure for the ATM 1 to 1 Cell Mode and for the AAL5 PDU Frame Mode is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |        Sequence Number        | ATM Specific  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

In the above diagram the first 4 bits MUST be set to 0 to indicate PW data. They MUST be ignored by the receiving PE.

The next four bits are reserved and MUST be set to 0 upon transmission and ignored upon reception.

The next 16 bits provide a sequence number that can be used to guarantee ordered packet delivery. The processing of the sequence number field is OPTIONAL.

The sequence number space is a 16 bit, unsigned circular space. The sequence number value 0 is used to indicate that the sequence number check alghorithm is not used.

The last 8 bits provide space for carrying ATM specific flags. These are defined in the protocol-specific details below.

There is no requirement for a length field for the One-to-one cell and PDU Frame modes because the PSN PDU is always greater than 64 bytes and so no padding is applied in Ethernet links in the PSN.

4.3.2. The Preferred Control Word

This control word is used in the following encapsulation modes:
- ATM N to 1 Cell Mode
- AAL5 SDU Frame Mode

It is defined as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Flags |Res|   Length  |     Sequence Number           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

In the above diagram the first 4 bits MUST be set to 0 to indicate PW data. They MUST be ignored by the receiving PE.

The next 4 bits provide space for carrying protocol specific flags. These are defined in the protocol-specific details below.

The next 6 bits provide a length field, which is used as follows: If the packet's length (defined as the length of the layer 2 payload plus the length of the control word) is less than 64 bytes, the length field MUST be set to the packet's length. Otherwise the length field MUST be set to zero. The value of the length field, if non-zero, can be used to remove any padding. When the packet reaches the service provider's egress router, it may be desirable to remove the padding before forwarding the packet. Note that the length field is not used in the N-to-1 mode , and MUST be set to 0.

The last 16 bits provide a sequence number that can be used to guarantee ordered packet delivery. The processing of the sequence number field is OPTIONAL.

The sequence number space is a 16 bit, unsigned circular space. The sequence number value 0 is used to indicate that the sequence number check alghorithm is not used.

4.3.3. Setting the sequence number field in the control word

This section applies to the sequence number field of both the Generic and Preferred Control Words.

For a given emulated VC, and a pair of routers PE1 and PE2, if PE1 supports packet sequencing then the following procedures should be used:

- the initial packet transmitted on the emulated VC MUST use sequence number 1

- subsequent packets MUST increment the sequence number by one for each packet

- when the transmit sequence number reaches the maximum 16 bit value (65535) the sequence number MUST wrap to 1

If the transmitting router PE1 does not support sequence number processing, then the sequence number field in the control word MUST be set to 0.

4.3.4. Sequence number field processing in the control word

This section applies to the sequence number field of both the Generic and Preferred Control Words.

If a router PE2 supports receive sequence number processing, then the following procedures should be used:

When an emulated VC is initially set up, the "expected sequence number" associated with it MUST be initialized to 1.

When a packet is received on that emulated VC, the sequence number should be processed as follows:

- if the sequence number on the packet is 0, then the packet passes the sequence number check.

- otherwise if the packet sequence number >= the expected sequence number and the packet sequence number - the expected sequence number < 32768, then the packet is in order.

- otherwise if the packet sequence number < the expected sequence number and the expected sequence number - the packet sequence number >= 32768, then the packet is in order.

- otherwise the packet is out of order.

If a packet is in order then, it can be delivered immediately. If the packet is in order, then the expected sequence number MUST be set using the algorithm:

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

Packets which are received out of order MAY be dropped or reordered at the discretion of the receiver.

A simple extension of the above processing algorithm can be used to detect lost packets.

If a router PE2 does not support receive sequence number processing, then the sequence number field MAY be ignored.

4.4. MTU Requirements

The network MUST be configured with an MTU that is sufficient to transport the largest encapsulation frames. If MPLS is used as the tunneling protocol, for example, this is likely to be 12 or more bytes greater than the largest frame size. Other tunneling protocols may have longer headers and require larger MTUs. If the ingress router determines that an encapsulated layer 2 PDU exceeds the MTU of the tunnel through which it must be sent, the PDU MUST be dropped. If an egress router receives an encapsulated layer 2 PDU whose payload length (i.e., the length of the PDU itself without any of the encapsulation headers), exceeds the MTU of the destination layer 2 interface, the PDU MUST be dropped.

5. Applicability

This Draft defines two methods for encapsulation of ATM cells, namely, One-to-one mode and N-to-one mode.

The N-to-one mode (N >= 1) specifies an encapsulation method that maps one or more ATM VCCs (or one or more ATM VPCs) to one Pseudo-Wire. This is the only REQUIRED mode. One format is used for both the VCC or VPC mapping to the tunnel. The 4-octet ATM header is unaltered in the encapsulation, thus the VPI/VCI is always present. Cells from one or more VCCs (or one or more VPCs) may be concatenated.

The One-to-one mode specifies an encapsulation method that maps one ATM VCC or one ATM VPC to one Pseudo-Wire. For VCCs, the VPI/VCI is not included. For VPCs, the VPI is not included. Cells from one VCC or one VPC may be concatenated. This mode is OPTIONAL.

Furthermore different OPTIONAL encapsulations are supported for ATM AAL5 transport: one for ATM AAL5 SDUs, and another for ATM AAL5 PDUs.

Three deployment models are supported by the encapsulations described in this document:

-i. Single ATM Connection: A PW carries the cells of only one ATM VCC or VPC. This supports both the transport of multiservice ATM and L2VPN service over a PSN for all AAL types.

-ii. Multiple ATM Connections: A PW carries the cells of multiple ATM VCCs and / or VPCs . This also supports both the transport of multiservice ATM and L2VPN service over a PSN for all AAL type.

-iii. AAL5: PW carries the AAL5 frames of only one ATM VCC. A large proportion of the data carried on ATM networks is frame based and therefore uses AAL5. The AAL5 mapping takes advantage of the delineation of higher layer frames in the ATM layer to provide increased bandwidth efficiency compared with the basic cell mapping. The nature of the service, as defined by the ATM service category [5] or the ATM transfer capability [6] should be preserved.

There are currently no OAM mechanisms defined for the PSN like those defined for ATM. Therefore the methods for the detection/consequent-actions of failures in the PSN are not specified. This also means that QoS/availability metrics cannot be specified for the PSN.

5.1. ATM N to 1 Cell Mode

This encapsulation supports both the Single and Multiple ATM Connection deployment models. This encapsulation is REQUIRED.

The encapsulation allows multiple VCCs/VPCs to be carried within a single pseudo wire. However, a service provider may wish to provision a single VCC to a pseudo wire in order to satisfy QoS or restoration requirements.

The encapsulation also supports the binding of multiple VCCs/VPCs to a single Pseudo Wire. This capability is useful in order to make more efficient use of the PW demultiplexing header space as well as to ease provisioning of the VCC/VPC services.

In the simplest case, this encapsulation can be used to transmit a single ATM cell per PSN PDU. However, in order to provide better PSN bandwidth efficiency, several ATM cells may optionally be encapsulated in a single PSN PDU. This process is called cell concatenation.

The encapsulation has the following attributes:

-i. Supports all ATM Adaptation Layers Types.

-ii. Non-terminating OAM/Admin cells are transported among the user cells in the same order as they are received. This requirement enables the use of various performance management and security applications.

-iii. In order to gain transport efficiency on the PSN, multiple cells may be encapsulated in a single PW PDU. This process is called cell concatenation. How many cells to insert or how long to wait for cell arrival before sending a PW PDU is an implementation decision. Cell concatenation adds latency and delay variation to a cell relay service.

-iv. The CLP bit from each cell may be mapped to a corresponding marking on the PW PDU. This allows the drop precedence to be preserved across the PSN.

-v. If the Single ATM connection deployment model is used, then it is simpler to provide an ATM layer service. The nature of the service, as defined by the ATM service category [5] or ATM transfer capability [6], should be preserved.

The limitations of the ATM N-to-one cell encapsulation are:

-vi. There is no currently defined method to translate the forward congestion indication (EFCI) to a corresponding function in the PSN. Nor is there a way to translate PSN congestion to the EFCI upon transmission by the egress PE.

-vii. The ATM cell header checksum can detect a 2-bit error or detect and correct a single bit error in the cell header. Analogous functionality does not exist in most PSNs. A single bit error in a PW PDU will most likely cause the packet to be dropped due to a L2 FCS failure.

-viii. Cells can be concatenated from multiple VCCs or VPCs belonging to different service cathegories and qos requirements. In this case the PSN packet must receive treatment by the PSN to support the highest QoS of the ATM VCCs/VPCs carried.

-ix. Cell encapsulation only supports point-to-point LSPs. Multipoint-to-point and point-to-multi-point are for further study (FFS).

-x. The number of concatenated ATM cells is limited by the MTU size and the cell transfer delay (CTD) and cell delay variation (CDV) objectives of multiple ATM connections that are multiplexed into a single PW.

5.2. ATM One-to-One Cell Encapsulation

This OPTIONAL encapsulation supports the Single ATM Connection deployment model.

Like the N to one cell encapsulation mode, the One-to-one mode supports cell concatenation. The advantage of this encapsulation is that it utilizes less bandwidth that the N-to-one encapsulation, for a given number of concatenated cells. Since only one ATM VCC or VPC is carried on a PW, the VCI and/or VPI of the ATM VCC or VPC can be derived from the context of the PW using the PW label. These fields therefore do not need to be encapsulated for a VCC, and only the VCI needs to be encapsulated for a VPC. This encapsulation thus allows service providers to achieve a higher bandwidth efficiency on PSN links than the N-to-one encapsulation for a given number of concatenated cells.

The limitations vi,vii,ix,x of N to one mode apply.

5.3. AAL5 SDU Frame Encapsulation

This OPTIONAL encapsulation supports the AAL5 model.

The AAL5 SDU encapsulation is more efficient for small AAL5 SDUs than the VCC cell encapsulations. In turn it presents a more efficient alternative to the cell relay service when carrying RFC 2684 encapsulated IP PDUs across a PSN.

The AAL5-SDU encapsulation requires Segmentation and Reassembly on the PE-CE ATM interface. This SAR function is provided by common off-the-shelf hardware components. Once reassembled, the AAL5-SDU is carried via a Pseudo Wire to the egress PE. Herein lies another advantage of the AAL5-SDU encapsulation.

The limitations of the AAL5 SDU encapsulation are:

-i. If an ATM OAM cell is received at the ingress PE, it is sent before the cells of the surrounding AAL5 frame. Therefore, OAM cell reordering may occur, which may cause certain ATM OAM performance monitoring and ATM security applications to operate incorrectly.

-ii. If the ALL5 PDU is scrambled using ATM security standards, a PE will not be able to exctract the ALL5 SDU and therefore the whole PDU will be dropped.

-iii. The AAL5 PDU CRC is not transported across the PSN. The CRC must therefore be regenerated at the egress PE. Since the CRC has end-to-end significance in ATM security. This means that the AAL5 CRC may not be used to accurately check for errors on the end-to-end ATM VCC.

-iv. The 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.

-v. 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 [5].

-vi. 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.

5.4. AAL5 PDU Frame Encapsulation

This OPTIONAL encapsulation supports the AAL5 model.

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 main advantage of this AAL5 mode is that it is transparent to ATM OAM and ATM security applications.

One important consideration is to allow OAM information to be treated as in the original network. This encapsulation mode allows this transparency while performing AAL5 frame encapsulation. This mode supports fragmentation, which 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. Fragmentation provides a means for the PE to set the size of the PW packet to a different value than that of the original AAL5 PDU. This means that the PE has control on the delay and jitter provided to the ATM cells.

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 PDU.

Reassembly is not required at the egress PE for the PSN-to-ATM direction.

The limitations v and vi of the AAL5 SDU mode apply to this mode as well.

6. ATM OAM Cell Support

6.1. VCC Case

In general when configured for ATM VCC service, both PEs SHOULD act as a VC switch, in accordance with the OAM procedures defined in [7].

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

F4 OAM cells are inserted or extracted at the VP link termination. These OAM cells are not seen at the VC link termination and are therefore not sent across the PSN.

6.2. VPC Case

When configured for a VPC cell relay service, both PEs SHOULD act as a VP cross-connect in accordance with the OAM procedures defined in [7].

The PEs SHOULD be able to process and pass the following OAM cells transparently according to [7]:
- F4 AIS (segment and end-to-end)
- F4 RDI (segment and end-to-end)
- F4 loopback (segment and end-to-end)

F5 OAM are not inserted or extracted here. The PEs MUST 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 OAM cell MAY be encapsulated together with other user data cells if multiple cell encapsulation is used.

6.3. Defect Handling

Figure 3 illustrates four possible locations for defects on the PWE3 service:

- (a) On the ATM connection from CE to PE
- (b) On the ATM side of the PW
- (c) On the PSN side of the PE
- (d) In the PSN

```
                +----+                  +----+
+----+          | PE1|==================| PE2|          +----+
|    |---a------|b..c........PW1...d.........|----------|    |
| CE1|          |    |                  |    |          |CE2 |
|    |----------|............PW2.............|----------|    |
+----+          |    |==================|    |          +----+
     ^          +----+                  +----+          ^
     |      Provider Edge 1         Provider Edge 2     |
     |                                                  |
     |<-------------- Emulated Service ---------------->|
Customer                                                Customer
Edge 1                                                  Edge 2
```

Figure 3: Defect Locations

For failures at (a) or (b) in the VPC case the ingress PE MUST be able to generate an F4 AIS upon reception of a lower layer defect (such as LOS). In the VCC case, 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). These messages are sent across the PSN.

For failures at (c) or (d), in the VCC case 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). In the VPC case, 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 such as the PW status message defined in [1]. Alternatively, 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 (for VPC) or F5 AIS (for VCC).

If the ingress PE cannot support the generation of OAM cells, it MAY notify the egress PE using a Pseudo Wire specific maintenance mechanism such as the PW status message defined in [1]. Alternatively, 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.

If the PW in one direction fails, then the complete bidirectional service is considered to have failed.

7. ATM N-to-one Cell Mode

The N-to-one mode (N >= 1) described in this Draft allows a service provider to offer an ATM PVC or SVC based service across a network. The encapsulation allows multiple ATM VCCs or VPCs to be carried within a single PSN tunnel. A service provider may also use N-to-one mode to provision either one VCC or one VPC on a tunnel. This section defines the VCC and VPC cell relay services over a PSN and their applicability.

7.1. ATM N-to-one Service Encapsulation

This section describes the general encapsulation format for ATM over PSN pseudo wires.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Pseudo Wire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Flags |Res|   Length  |     Sequence Number           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     ATM Service Payload                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: General format for ATM encapsulation over PSNs

The PSN Transport Header depends on the particular tunneling technology in use (L2TP or MPLS). This header is used to transport the encapsulated ATM information through the packet switched core.

The Pseudo Wire Header identifies a particular ATM service on a tunnel. Non-ATM services may also be carried on 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.

In this encapsulation mode ATM cells are transported individually. The encapsulation of a single ATM cell is the only REQUIRED encapsulation for ATM. The encapsulation of more than one ATM cell in a PSN frame is OPTIONAL.

The ATM cell encapsulation consists of an OPTIONAL control word, and one or more ATM cells - each consisting of a 4 byte ATM cell header and the 48 byte ATM cell payload. This ATM cell header is defined as in the FAST encapsulation [4] section 3.1.1, but without the trailer byte. The length of each frame, without the encapsulation headers, is a multiple of 52 bytes long. The maximum number of ATM cells that can be fitted in a frame, in this fashion, is limited only by the network MTU and by the ability of the egress router to process them. The ingress router MUST NOT send more cells than the egress router is willing to receive. The number of cells that the egress router is willing to receive may either be configured in the ingress router or may be signaled, for example using the methods described in [1]. The number of cells encapsulated in a particular frame can be inferred by the frame length. The control word is OPTIONAL. If the control word is used then the flag, and length bits in the control word are not used, and MUST be set to 0 when transmitting, and MUST be ignored upon receipt.

The EFCI and CLP bits are carried across the network in the ATM cell header. The edge routers that implement this document MAY, when either adding or removing the encapsulation described herein, change the EFCI bit from zero to one in order to reflect congestion in the network that is known to the edge router, and the CLP bit from zero to one to reflect marking from edge policing of the ATM Sustained Cell Rate. The EFCI and CLP bits SHOULD NOT be changed from one to zero.

This diagram illustrates an encapsulation of two ATM cells:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                  Control word ( Optional )                    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |          VPI          |              VCI              | PTI |C|
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                  ATM Payload ( 48 bytes )                     |
    |                          "                                    |
    |                          "                                    |
    |                          "                                    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |          VPI          |              VCI              | PTI |C|
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                  ATM Payload ( 48 bytes )                     |
    |                          "                                    |
    |                          "                                    |
    |                          "                                    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Multiple Cell ATM Encapsulation

* When multiple VCCs or VPCs are transported in one pseudo-wire VPI/VCI values MUST be unique. When the multiple VCCs or VPCs, are from different a physical transmission path it may be necessary to assign unique VPI/VCI values to the ATM connections. If they are from the same physical transmission path, the VPI/VCI values are unique.

* VPI

  The ingress router MUST copy the VPI field from the incoming cell into this field. For particular emulated VCs, the egress router MAY generate a new VPI and ignore the VPI contained in this field.

* VCI

  The ingress router MUST copy the VCI field from the incoming ATM cell header into this field. For particular emulated VCs, the egress router MAY generate a new VCI.

* PTI & CLP (C bit)

The PTI and CLP fields are the PTI and CLP fields of the incoming ATM cells. The cell headers of the cells within the packet are the ATM headers (without HEC) of the incoming cell.

8. ATM One-to-one Cell Mode

The One-to-one mode described in this Draft allows a service provider to offer an ATM PVC or SVC based service across a network. The encapsulation allows one ATM VCC or VPC to be carried within a single Pseudo-Wire.

8.1. ATM One-to-one 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 6 provides a general format for encapsulation of ATM cells into packets.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Pseudo Wire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |    Optional Sequence Number   | ATM Specific  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     ATM Service Payload                       |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: General format for One-to-one mode 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 generic control word is inserted after the Pseudo Wire Header. The presence of the control word is MANDATORY.

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.

8.2. Sequence Number

The sequence number is not required for all services.

Treatment of the sequence number is according to previous sections "Setting the sequence number", and "Processing the sequence number".

8.3. ATM VCC Cell Transport Service

The VCC cell transport 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 VCC single cell transport service is OPTIONAL. This service MUST use the following encapsulation format:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudo Wire Header                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |  Optional Sequence Number     |M|V|Res| PTI |C|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                    ATM Cell Payload ( 48 bytes )              |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: Single ATM VCC Cell Encapsulation

* M (transport mode) bit

  Bit (M) of the control byte indicates whether the packet contains an ATM cell or a frame payload. If set to 0, the packet contains an ATM cell. If set to 1, the PDU contains an AAL5 payload.

* V (VCI present) bit

Bit (V) of the control byte indicates whether the VCI field is present in the packet. If set to 1, the VCI field is present for the cell. If set to 0, no VCI field is present. In the case of a VCC, the VCI field is not required. For VPC, the VCI field is required and is transmitted with each cell.

* Reserved bits

The reserved bits should be set to 0 at the transmitter and ignored upon reception.

* PTI Bits

The 3-bit Payload Type Identifier (PTI) incorporates ATM Layer PTI coding of the cell. These bits are set to the value of the PTI of the encapsulated ATM cell.

* C (CLP) Bit

The Cell Loss Priority (CLP) field indicates CLP value of the encapsulated cell.

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.

When multiple cells are encapsulated in the same PSN packet, the ATM specific byte MUST be repeated for each cell. This means that 49 bytes are used to encapsulate each 53 byte ATM cell.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudo Wire Header                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |  Optional Sequence Number     |M|V|Res| PTI |C|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                   ATM Cell Payload ( 48 bytes )               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|M|V|Res| PTI |C|                                               |
```

```
   +-+-+-+-+-+-+-+-+                                                |
   |                   ATM Cell Payload ( 48 bytes )                |
   |                                                                |
   |               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |               |
   +-+-+-+-+-+-+-+-+
```

Figure 8: Multiple ATM VCC Cell Encapsulation

8.4. ATM VPC Services

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.

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.

8.4.1. ATM VPC Cell Transport Services

The ATM VPC cell transport service is OPTIONAL.

This service MUST use the following cell mode encapsulation:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudo Wire Header                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |  Optional Sequence Number     |M|V|Res| PTI |C|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             VCI               |                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
|                                                               |
|                   ATM Cell Payload ( 48 bytes )               |
|                                                               |
|                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 9: Single Cell VPC Encapsulation

The ATM control byte contains the same information as in the VCC encapsulation except for the VCI field.

* VCI Bits

  The 16-bit Virtual Circuit Identifier (VCI) incorporates ATM Layer VCI value of the cell.

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.

When multiple ATM cells are encapsulated in the same PSN packet, the ATM specific byte MUST be repeated for each cell. This means that 51 bytes are used to encapsulate each 53 byte ATM cell.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               PSN Transport Header (As Required)              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudo Wire Header                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0| Resvd |  Optional Sequence Number     |M|V|Res| PTI |C|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             VCI               |                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
|                                                               |
|                   ATM Cell Payload (48 bytes)                 |
|                                                               |
|                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               |M|V|Res| PTI |C|       VCI     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   VCI         |                                               |
+-+-+-+-+-+-+-+-+                                               |
|                   ATM Cell Payload (48 bytes)                 |
|                                                               |
|               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               |
+-+-+-+-+-+-+-+-+
```

Figure 10: Multiple Cell VPC Encapsulation

9. ATM AAL5 CPCS-SDU Mode

The AAL5 payload VCC service defines a mapping between the payload of an AAL5 VCC and a single Pseudo Wire. The AAL5 payload VCC service requires ATM segmentation and reassembly support on the PE.

The AAL5 payload CPCS-SDU 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. All OAM cells, regardless of their type, that arrive during the reassembly of a single AAL5 CPCS-PDU are sent immediately on the Pseudo Wire using N-to-one cell encapsulation, followed by the AAL5 payload. Therefore, the AAL5 payload VCC service will not be suitable for ATM applications that require strict ordering of OAM cells (such as performance monitoring and security applications).

9.1. Transparent AAL5 SDU Frame Encapsulation

The AAL5 CPCS-SDU is prepended by the following header:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Res  |T|E|C|U|Res|  Length   |   Sequence Number (Optional)  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               "                               |
|                     ATM cell or AAL5 CPCS-SDU                 |
|                               "                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 11: AAL5 CPCS-SDU 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 admin cell or an AAL5 payload. If T = 1, the packet contains an ATM admin cell, encapsulated according to the VCC cell relay encapsulation, figure 8. If not set, the PDU contains an AAL5 payload. The ability to transport an ATM cell in the AAL5 SDU mode is intended to provide a means of enabling administrative functionality over the AAL5 VCC (though it does not endeavor to preserve user-cell and admin-cell arrival/transport ordering).

* E (EFCI) Bit

The ingress router, PE1, SHOULD set this bit to 1 if the EFCI bit of the final cell of those that transported the AAL5 CPCS-SDU 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 router, PE2, SHOULD set the EFCI bit of all cells that transport the AAL5 CPCS-SDU to the value contained in this field.

* C (CLP) Bit

The ingress router, PE1, SHOULD set this bit to 1 if the CLP bit of any of the ATM cells that transported the AAL5 CPCS-SDU is set to 1, or if the CLP 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 router, PE2, SHOULD set the CLP bit of all cells that transport the AAL5 CPCS-SDU to the value contained in this field.

* U (Command / Response Field) Bit

When FRF.8.1 Frame Relay / ATM PVC Service Interworking [3] 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 router, PE1, SHOULD copy this bit to the U bit of the control word. The egress router, PE2, SHOULD copy the U bit to the CPCS-UU Least Significant Bit (LSB) of the AAL5 CPCS PDU.

10. AAL5 PDU frame mode

The AAL5 payload PDU service is OPTIONAL.

10.1. Transparent AAL5 PDU Frame Encapsulation

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.

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |               PSN Transport Header (As Required)              |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                      Pseudo Wire Header                       |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |0 0 0 0| Resvd |   Optional Sequence Number    |M|V| Res |U|E|C|
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                              "                                |
 |                        AAL5 CPCS-PDU                          |
 |                      (n * 48 bytes)                           |
 |                              "                                |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 12: AAL5 transparent service encapsulation

The generic control word is inserted after the Pseudo Wire Header. The presence of the control word is MANDATORY.

The M, V, Res, and C bits are as defined earlier for VCC One-to-one cell mode.

* U Bit

  This field indicates whether this frame contains the last cell of an AAL5 PDU and represents the value of the ATM User-to-User bit for the last ATM cell of the PSN frame. Note: The ATM User-to-

User bit is the least significant bit of the PTI field in the ATM header. This field is used to support the fragmentation functionality described later in this section.

* 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.

10.2. 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.

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

10.2.1. Procedures in the ATM-to-PSN Direction

The following procedures shall apply while fragmenting AAL5 PDUs:

- Fragmentation shall always occur at cell boundaries within the AAL5 PDU.

- Set the UU bit to the value of the ATM User-to-User bit in the cell header of the most recently received ATM cell.
- The E and C bits of the fragment shall be set as defined earlier in section 9.
- If the arriving cell is an OAM or an RM cell, send the current PSN frame and then send the OAM or RM cell using One-to-one single cell encapsulation (VCC).

10.2.2. Procedures in the PSN-to-ATM Direction

The following procedures shall apply:

- The 3-bit PTI field of each ATM cell header is constructed as follows:
  - -i. The most significant bit is set to 0, indicating a user data cell.
  - -ii. The middle bit is set to the E bit value of the fragment.
  - -iii. The least significant bit for the last ATM cell in the PSN frame is set to the value of the UU bit of Figure 12.
  - -iv. The least significant PTI bit is set to 0 for all other cells in the PSN frame.
- The CLP bit of each ATM cell header is set to the value of the C bit of the control byte in Figure 12.
- When a fragment is received, each constituent ATM cell is sent in correct order.

11. Mapping of ATM and PSN Classes of Service

This section is informational.

When ATM PW service is configured over a PSN, the ATM service category of a connection 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, a class of service based on the EF PHB is a good candidate.

Furthermore, ATM service categories have support for multiple conformance definitions [5]. Some are CLP blind, e.g., CBR, meaning that the QoS objectives apply to the aggregate CLP0+1 conforming cell flow. Some are CLP significant, e.g., VBR.3, meaning that the QoS objectives apply to the CLP0 conforming cell flow only.

When the PSN is MPLS based, a mapping between the CLP bit and the EXP field can be performed to provide visibility of the cell loss priority in the MPLS network. The actual value to be marked in the EXP field depends on the ATM service category, the ATM conformance definition, and the type of tunnel LSP used (E-LSP or L-LSP). The details of this mapping are outside the scope of this document. Operators have the flexibility to design a specific mapping which satisfies their own requirements.

In both the ATM-to-PSN and PSN-to-ATM directions, the method used to transfer the CLP and EFCI information of the individual cells into the ATM specific field, or flags, of the PW packet is described in details in sections 6 through 9 for each encapsulation mode.

12. Security Considerations

This document specifies only encapsulations, and not the protocols used to carry the encapsulated packets across the PSN. Each such protocol may have its own set of security issues, but those issues are not affected by the encapsulations specified herein. Note that the security of the transported ATM service will only be as good as the security of the PSN. This level of security might be less rigorous then a native ATM service.

13. Intellectual Property Disclaimer

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

14. References

[1] "Transport of Layer 2 Frames Over MPLS", draft-ietf-pwe3-control-protocol-06.txt. (work in progress)

[2] "MPLS Label Stack Encoding", E. Rosen, Y. Rekhter, D. Tappan, G. Fedorkow, D. Farinacci, T. Li, A. Conta. RFC3032

[3] "Requirements for Peudo Wire Emulation Edge-to-Edge (PWE3", draft-ietf-pwe3-requirements-08.txt. (work in Progress)

[4] ATM Forum Specification fb-fbatm-0151.000 (2000) ,Frame Based ATM over SONET/SDH Transport (FAST)

[5] ATM Forum Specification af-tm-0121.000 (1999), Traffic Management Specification Version 4.1.

[6] ITU-T Recommendation I.371 (2000), Traffic control and congestion control in B-ISDN.

[7] ITU-T Recommendation I.610, (1999), B-ISDN operation and maintenance principles and functions.

[8] "PWE3 Architecture", draft-ietf-pwe3-arch-07.txt. (work in progress)

15. Author Information

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO, 80112
e-mail: lmartini@cisco.com

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

Giles Heron
Tellabs
Abbey Place
24-28 Easton Street
High Wycombe
Bucks
HP11 1NT
UK
e-mail: giles.heron@tellabs.com

Dimitri Stratton Vlachos
Mazu Networks, Inc.
125 Cambridgepark Drive
Cambridge, MA 02140

e-mail: d@mazunetworks.com

Dan Tappan
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
e-mail: tappan@cisco.com

Jayakumar Jayakumar,
Cisco Systems Inc.
170, W.Tasman,
San Jose , CA, 95134
e-mail: jjayakum@cisco.com

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
E-mail: erosen@cisco.com

Steve Vogelsang
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
e-mail: sjv@laurelnetworks.com

Jeremy Brayley
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
e-mail: jbrayley@laurelnetworks.com

Gerald de Grace
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205

e-mail: gdegrace@laurelnetworks.com

John Shirron
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
e-mail: jshirron@laurelnetworks.com

Andrew G. Malis
Tellabs
90 Rio Robles Dr.
San Jose, CA 95134
e-mail: Andy.Malis@tellabs.com

Vinai Sirkay
Reliance Infocomm
Dhirubai Ambani Knowledge City
Navi Mumbai 400 709
India
e-mail: vinai@sirkay.com

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

Kireeti Kompella
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089
e-mail: kireeti@juniper.net

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

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

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

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

Tom Walsh
Lucent Technologies
1 Robbins Road
Westford, MA 01886 USA
e-mail: tdwalsh@lucent.com

John Rutemiller
Marconi Networks
1000 Marconi Drive
Warrendale, PA 15086
e-mail: John.Rutemiller@marconi.com

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

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