

Network Working Group
Request for Comments: 4717
Category: Standards Track

L. Martini
J. Jayakumar
Cisco Systems, Inc.
M. Bocci
Alcatel
N. El-Aawar
Level 3 Communications, LLC
J. Brayley
ECI Telecom Inc.
G. Koleyni
Nortel Networks
December 2006

Encapsulation Methods for Transport of
Asynchronous Transfer Mode (ATM) over MPLS Networks

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2006).

Abstract

An Asynchronous Transfer Mode (ATM) Pseudowire (PW) is used to carry ATM cells over an MPLS network. This enables service providers to offer "emulated" ATM services over existing MPLS networks. This document specifies methods for the encapsulation of ATM cells within a pseudowire. It also specifies the procedures for using a PW to provide an ATM service.

Table of Contents

1. Introduction	3
2. Specification of Requirements	4
3. Applicability Statement	4
4. Terminology	4
5. General Encapsulation Method	6
5.1. The Control Word	6
5.1.1. The Generic Control Word	7
5.1.2. The Preferred Control Word	8
5.1.3. Setting the Sequence Number Field in the Control Word	9
5.2. MTU Requirements	9
5.3. MPLS Shim S Bit Value	10
5.4. MPLS Shim TTL Values	10
6. Encapsulation Mode Applicability	10
6.1. ATM N-to-One Cell Mode	11
6.2. ATM One-to-One Cell Encapsulation	13
6.3. AAL5 SDU Frame Encapsulation	13
6.4. AAL5 PDU Frame Encapsulation	14
7. ATM OAM Cell Support	15
7.1. VCC Case	15
7.2. VPC Case	16
7.3. SDU/PDU OAM Cell Emulation Mode	16
7.4. Defect Handling	17
8. ATM N-to-One Cell Mode	18
8.1. ATM N-to-One Service Encapsulation	19
9. ATM One-to-One Cell Mode	21
9.1. ATM One-to-One Service Encapsulation	21
9.2. Sequence Number	22
9.3. ATM VCC Cell Transport Service	22
9.4. ATM VPC Services	24
9.4.1. ATM VPC Cell Transport Services	25
10. ATM AAL5 CPCS-SDU Mode	26
10.1. Transparent AAL5 SDU Frame Encapsulation	27
11. AAL5 PDU Frame Mode	28
11.1. Transparent AAL5 PDU Frame Encapsulation	28
11.2. Fragmentation	30
11.2.1. Procedures in the ATM-to-PSN Direction	30
11.2.2. Procedures in the PSN-to-ATM Direction	31
12. Mapping of ATM and PSN Classes of Service	31
13. ILMI Support	32
14. ATM-Specific Interface Parameter Sub-TLVs	32
15. Congestion Control	32
16. Security Considerations	33
17. Normative References	34
18. Informative References	34
19. Significant Contributors	36

1. Introduction

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:

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

This document describes a method to carry ATM services over MPLS. It lists ATM-specific requirements and provides encapsulation formats and semantics for connecting ATM edge networks through a packet-switched network using MPLS.

Figure 1, below, displays the ATM services reference model. This model is adapted from [RFC3985].

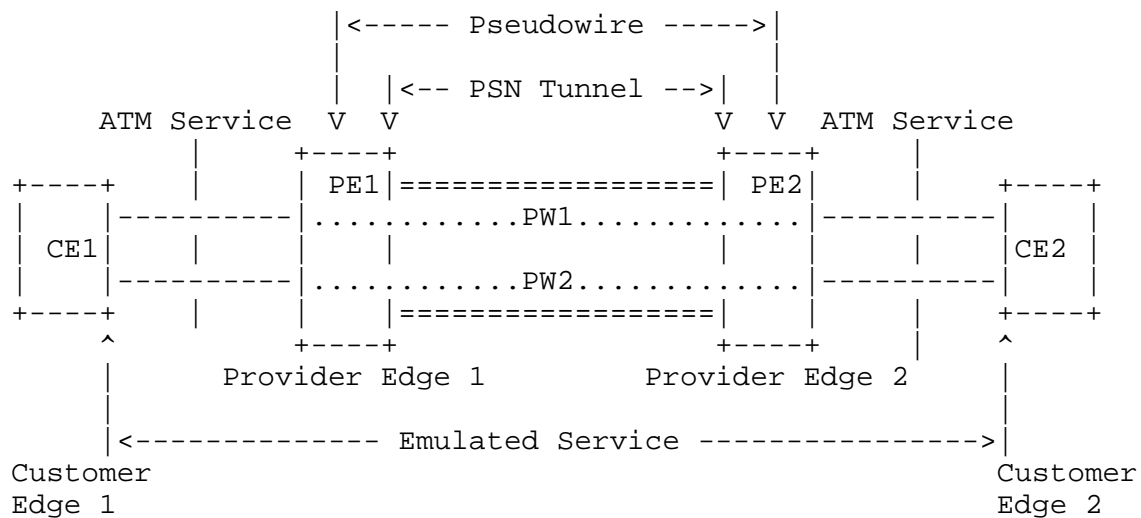


Figure 1: ATM Service Reference Model

2. 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 [RFC2119].

3. Applicability Statement

The ATM over PW service is not intended to perfectly emulate a traditional ATM service, but it can be used for applications that need an ATM transport service.

The following are notable differences between traditional ATM service and the protocol described in this document:

- ATM cell ordering can be preserved using the OPTIONAL sequence field in the control word; however, implementations are not required to support this feature. The use of this feature may impact other ATM quality of service (QoS) commitments.
- The QoS model for traditional ATM can be emulated. However, the detailed specification of ATM QoS emulation is outside the scope of this document. The emulation must be able to provide the required ATM QoS commitments for the end-user application.
- The ATM flow control mechanisms are transparent to the MPLS network and cannot reflect the status of the MPLS network.
- Control plane support for ATM SVCs, SVPs, SPVCs, and SPVPs is outside the scope of this document.

Note that the encapsulations described in this specification are identical to those described in [Y.1411] and [Y.1412].

4. Terminology

One-to-one mode: specifies an encapsulation method that maps one ATM Virtual Channel Connection (VCC) (or one ATM Virtual Path Connection (VPC)) to one pseudowire.

N-to-one mode ($N \geq 1$): specifies an encapsulation method that maps one or more ATM VCCs (or one or more ATM VPCs) to one pseudowire.

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

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

Customer Edge (CE): 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 (PE): a device that provides PWE3 to a CE.

Pseudowire (PW): a connection between two PEs carried over a PSN. The PE provides the adaptation between the CE and the PW.

Pseudowire PDU: a PDU sent on the PW that contains all of the data and control information necessary to provide the desired service.

PSN Tunnel: 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 pseudowire PDU (ATM to PSN direction).

Egress: the point where the ATM service is decapsulated from a pseudowire PDU (PSN to ATM direction).

CTD: Cell Transfer Delay.

MTU: Maximum Transmission Unit.

SDU: Service Data 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.

Additional terminology relevant to pseudowires and Layer 2 Virtual Private Networking (L2VPN) in general may be found in [RFC4026].

5. General Encapsulation Method

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

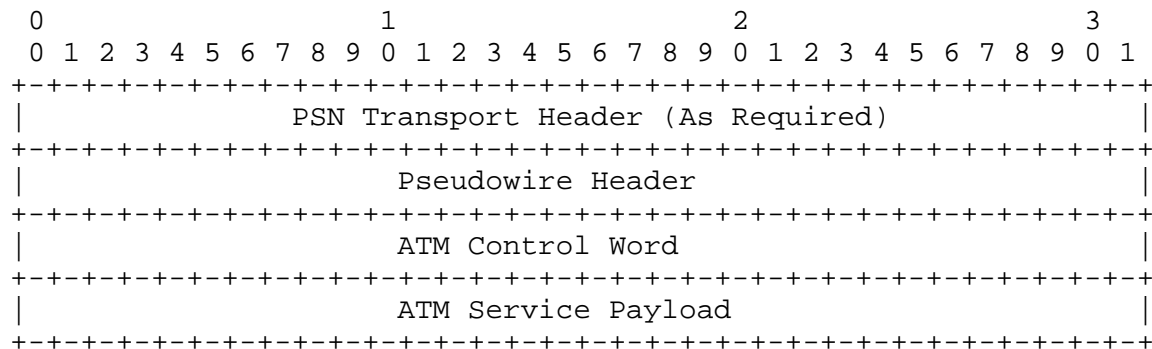


Figure 2: General format for ATM encapsulation over PSNs

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

The Pseudowire Header identifies a particular ATM service on a tunnel. In case of MPLS, the pseudowire header is one or more MPLS labels at the bottom of the MPLS label stack.

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.

5.1. The Control Word

The Control Words defined in this section are based on the Generic PW MPLS Control Word as defined in [RFC4385]. They provide the ability to sequence individual frames on the PW, avoidance of equal-cost multiple-path load-balancing (ECMP) [RFC2992], and OAM mechanisms including VCCV [VCCV].

[RFC4385] states, "If a PW is sensitive to packet misordering and is being carried over an MPLS PSN that uses the contents of the MPLS payload to select the ECMP path, it MUST employ a mechanism which prevents packet misordering." This is necessary because ECMP implementations may examine the first nibble after the MPLS label stack to determine whether or not the labelled packet is IP. Thus, if the VPI of an ATM connection carried over the PW using N-to-one cell mode encapsulation, without a control word present, begins with 0x4 or 0x6, it could be mistaken for an IPv4 or IPv6 packet. This

could, depending on the configuration and topology of the MPLS network, lead to a situation where all packets for a given PW do not follow the same path. This may increase out-of-order frames on a given PW, or cause OAM packets to follow a different path than actual traffic (see section 4.4.3 on Frame Ordering).

The features that the control word provides may not be needed for a given ATM PW. For example, ECMP may not be present or active on a given MPLS network, strict frame sequencing may not be required, etc. If this is the case, and the control word is not REQUIRED by the encapsulation mode for other functions (such as length or the transport of ATM protocol specific information), the control word provides little value and is therefore OPTIONAL. Early ATM PW implementations have been deployed that do not include a control word or the ability to process one if present. To aid in backwards compatibility, future implementations MUST be able to send and receive frames without a control word present.

In all cases, the egress PE MUST be aware of whether the ingress PE will send a control word over a specific PW. This may be achieved by configuration of the PEs, or by signaling, as defined in [RFC4447].

If the pseudowire traverses a network link that requires a minimum frame size (Ethernet is a practical example), with a minimum frame size of 64 octets, then such links will apply padding to the pseudowire 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.

5.1.1. The Generic Control Word

This control word is used in the following encapsulation modes:


- ATM One-to-one Cell Mode
- AAL5 PDU Frame Mode

The PWE3 control word document [RFC4385] 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 One-to-one Cell Mode and for the AAL5 PDU Frame Mode is as follows:



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 algorithm 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; therefore, no padding is applied in Ethernet links in the PSN.

5.1.2. The Preferred Control Word

This control word is used in the following encapsulation modes:

- ATM N-to-one Cell Mode
- AAL5 SDU Frame Mode

It is defined as follows:



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-one 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 algorithm is not used.

5.1.3. 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 sequencing procedures defined in [RFC4385] **MUST** be used.

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

A simple extension of the processing algorithm in [RFC4385] **MAY** be used to detect lost packets.

If a PE router negotiated not to use receive sequence number processing, and it received a non-zero sequence number, then it **SHOULD** send a PW status message indicating a receive fault and disable the PW.

5.2. 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.

5.3. MPLS Shim S Bit Value

The ingress label switching router (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. For more information on setting the S Bit, see [RFC3032].

5.4. MPLS Shim TTL Values

The setting of the TTL value in the PW label is application dependent. In any case, [RFC3032] TTL processing procedure, including handling of expired TTLs, MUST be followed.

6. Encapsulation Mode Applicability

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

The N-to-one mode ($N \geq 1$) specifies an encapsulation method that maps one or more ATM VCCs (or one or more ATM VPCs) to one pseudowire. 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 pseudowire. 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 types.
- iii. AAL5: A 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 [TM4.0] or the ATM transfer capability [I.371], should be preserved.

6.1. ATM N-to-One 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 pseudowire. However, a service provider may wish to provision a single VCC to a pseudowire in order to satisfy QoS or restoration requirements.

The encapsulation also supports the binding of multiple VCCs/VPCs to a single pseudowire. 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 Layer 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 [TM4.0] or ATM transfer capability [I.371], 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 an L2 Frame Check Sequence (FCS) failure.
- viii. Cells can be concatenated from multiple VCCs or VPCs belonging to different service categories 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 Label Switched Paths (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.

6.2. 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 than 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, and x of N-to-one mode apply.

6.3. AAL5 SDU Frame Encapsulation

This OPTIONAL encapsulation supports the AAL5 model. This mode allows the transport of ATM AAL5 CSPS-SDUs traveling on a particular ATM PVC across the network to another ATM PVC. This encapsulation is used by a PW of type 0x0002 "ATM AAL5 SDU VCC transport" as allocated in [RFC4446].

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 [RFC2684]-encapsulated IP PDUs across a PSN.

The AAL5-SDU encapsulation requires Segmentation and Reassembly (SAR) 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 pseudowire 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 AAL5 PDU is scrambled using ATM security standards, a PE will not be able to extract the AAL5 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 [TM4.0].
- 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.

6.4. 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.

7. ATM OAM Cell Support

7.1. 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 [I.610].

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

- F5 Alarm Indication Signal (AIS) (segment and end-to-end)
- F5 Remote Defect Indicator (RDI) (segment and end-to-end)
- F5 loopback (segment and end-to-end)
- Resource Management
- Performance Management
- Continuity Check
- Security

However, if configured to be an administrative segment boundary, the PE SHOULD terminate and process F5 segment OAM cells.

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.

When the PE is operating in AAL5 CPCS-SDU transport mode if it does not support transport of ATM cells, the PE MUST discard incoming MPLS frames on an ATM PW that contain a PW label with the T bit set.

7.2. 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 [I.610].

The PEs SHOULD be able to process and pass the following OAM cells transparently according to [I.610]:

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

However, if configured to be an administrative segment boundary, the PE SHOULD terminate and process F4 segment OAM cells.

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.

7.3. SDU/PDU OAM Cell Emulation Mode

A PE operating in ATM SDU or PDU transport mode that does not support transport of OAM cells across a PW MAY provide OAM support on ATM PVCs using the following procedures:

- Loopback cells response

If an F5 end-to-end OAM cell is received from an ATM VC, by either PE that is transporting this ATM VC, with a loopback indication value of 1, and the PE has a label mapping for the ATM VC, then the PE MUST decrement the loopback indication value and loop back the cell on the ATM VC. Otherwise, the loopback cell MUST be discarded by the PE.

- AIS alarm

If an ingress PE, PE1, receives an AIS F4/F5 OAM cell, it MUST notify the remote PE of the failure. The remote PE, PE2, MUST in turn send F5 OAM AIS cells on the respective PVCs. Note that if the PE supports forwarding of OAM cells, then the received OAM AIS alarm cells MUST be forwarded along the PW as well.

- Interface failure

If the PE detects a physical interface failure, or the interface is administratively disabled, the PE MUST notify the remote PE for all VCs associated with the failure.

- PSN/PW failure detection

If the PE detects a failure in the PW, by receiving a label withdraw for a specific PW ID, or the targeted Label Distribution Protocol (LDP) session fails, or a PW status TLV notification is received, then a proper AIS F5 OAM cell MUST be generated for all the affected ATM PVCs. The AIS OAM alarm will be generated on the ATM output port of the PE that detected the failure.

7.4. 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




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 pseudowire-specific maintenance mechanism such as the PW status message defined in [RFC4447]. Alternatively, for example, the ingress PE MAY withdraw the pseudowire (PW label) 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 PW in one direction fails, then the complete bidirectional service is considered to have failed.

8. ATM N-to-One Cell Mode

The N-to-one mode ($N \geq 1$) described in this document 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.

8.1. ATM N-to-One Service Encapsulation

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




Figure 4: General format for ATM encapsulation over PSNs

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

The Pseudowire Header identifies a particular ATM service on a tunnel. Non-ATM services may also be carried on the PSN tunnel.

As shown above, in Figure 4, the ATM Control Word is inserted before the ATM service payload. It may contain a length field and a sequence number field in addition to certain control bits needed to carry the service.

The ATM Service Payload is specific to the service being offered via the pseudowire. 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 [FBATM] section 3.1.1, but without the trailer byte. The length of each frame, without the encapsulation headers, is a multiple of 52 bytes. 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 be signaled, for example using the methods described later in this document and in [RFC4447]. 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. These bits 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 change the CLP bit from zero to one in order 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:




Figure 5: Multiple Cell ATM Encapsulation

- * When multiple VCCs or VPCs are transported in one pseudowire, 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 Header Error Check (HEC) field) of the incoming cell.

9. ATM One-to-One Cell Mode

The One-to-one mode described in this document 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 pseudowire.

9.1. ATM One-to-One Service Encapsulation

This section describes the general encapsulation format for ATM over pseudowires on an MPLS PSN. Figure 6 provides a general format for encapsulation of ATM cells into packets.




Figure 6: General format for One-to-one mode encapsulation over PSNs

The MPLS PSN Transport Header depends on how the MPLS network is configured. The Pseudowire Header identifies a particular ATM service within the PSN tunnel created by the PSN Transport Header.

This header is used to transport the encapsulated ATM information through the packet-switched core.

The generic control word is inserted after the Pseudowire Header. The presence of the control word is REQUIRED.

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.

9.2. Sequence Number

The sequence number is not required for all services.

Treatment of the sequence number is according to section 5.1.3.

9.3. ATM VCC Cell Transport Service

The VCC cell transport service is characterized by the mapping of a single ATM VCC (VPI/VCI) to a pseudowire. 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:




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 pseudowire PDU. The ingress and egress PE MUST agree to a maximum number of cells in a single pseudowire PDU. This agreement may be accomplished via a pseudowire-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.




Figure 8: Multiple ATM VCC Cell Encapsulation

9.4. ATM VPC Services

The VPC service is defined by mapping a single VPC (VPI) to a pseudowire. As such, it emulates a 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 pseudowire header. As a VPC service, the egress PE MUST NOT change the VCI field.

9.4.1. ATM VPC Cell Transport Services

The ATM VPC cell transport service is OPTIONAL.

This service MUST use the following cell mode encapsulation:

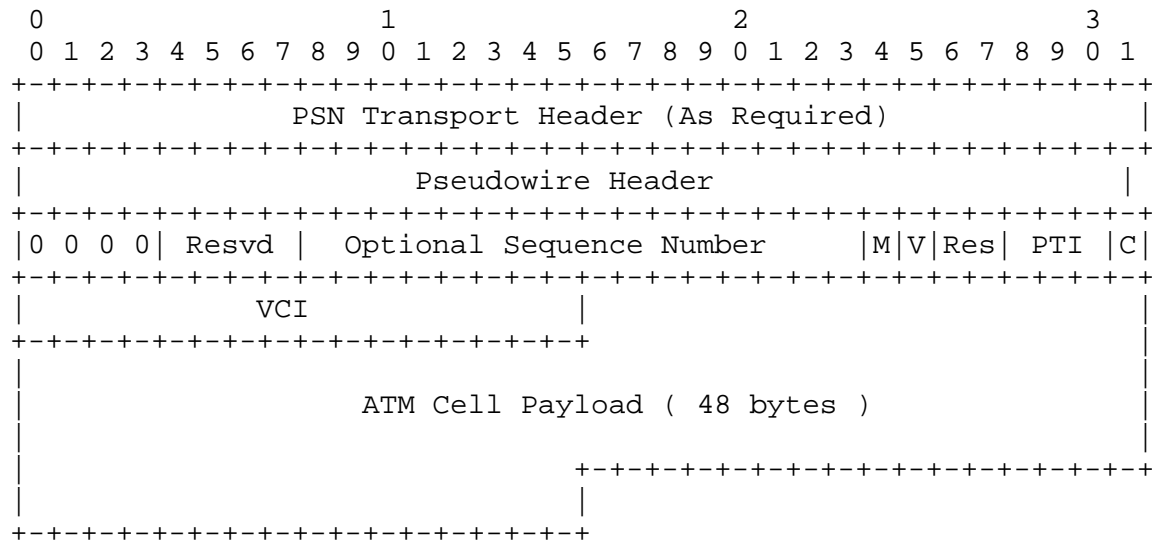


Figure 9: Single Cell VPC Encapsulation

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 pseudowire PDU. The ingress and egress PE MUST agree to a maximum number of cells in a single pseudowire PDU. This agreement may be accomplished via a pseudowire-specific signaling mechanism or via static configuration.

If the Egress PE supports cell concatenation, the ingress PE MUST only concatenate cells up to the "Maximum Number of concatenated ATM cells in a frame" interface parameter sub-TLV as received as part of the control protocol [RFC4447].

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.




Figure 10: Multiple Cell VPC Encapsulation

10. ATM AAL5 CPCS-SDU Mode

The AAL5 payload VCC service defines a mapping between the payload of an AAL5 VCC and a single pseudowire. 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 pseudowire. Therefore, once the ingress PE reassembles the AAL5 CPCS-PDU, the PE discards the PAD and CPCS-PDU trailer, and then the ingress PE inserts the resulting payload into a pseudowire 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 it 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 pseudowire 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).

10.1. Transparent AAL5 SDU Frame Encapsulation

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

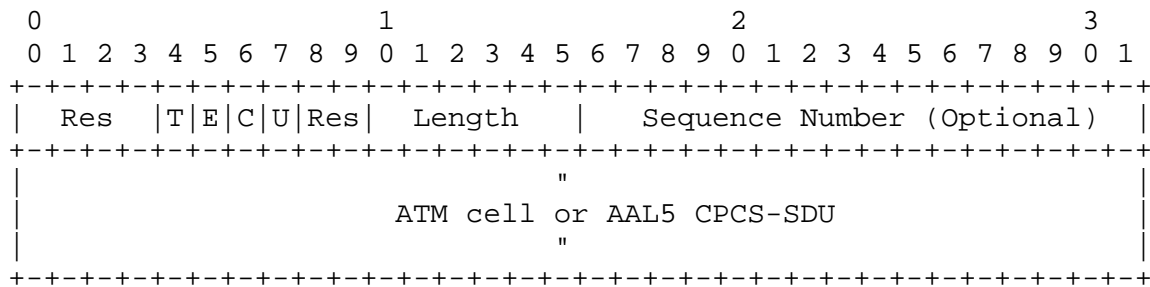


Figure 11: AAL5 CPCS-SDU Encapsulation

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 N-to-one cell relay encapsulation, Figure 4. 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 [RFC3916] 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.

11. AAL5 PDU Frame Mode

The AAL5 payload PDU service is OPTIONAL.

11.1. 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 procedures described in the "Fragmentation" section below, 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.




Figure 12: AAL5 transparent service encapsulation

The generic control word is inserted after the Pseudowire 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.

11.2. Fragmentation

The ingress PE may not always be able to reassemble a full AAL5 frame. This may be because the AAL5 PDU exceeds the pseudowire MTU or because 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.

11.2.1. 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 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).

11.2.2. 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.

12. Mapping of ATM and PSN Classes of Service

This section is provided for informational purposes, and for guidance only. This section should not be considered part of the standard proposed in this document.

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 Diffserv framework, a class of service based on the EF PHB is a good candidate.

Furthermore, ATM service categories have support for multiple conformance definitions [TM4.0]. 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 that 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 detail in sections 6 through 9 for each encapsulation mode.

13. ILMI Support

An MPLS edge PE MAY provide an ATM Integrated Local Management Interface (ILMI) to the ATM edge switch. If an ingress PE receives an ILMI message indicating that the ATM edge switch has deleted a VC, or if the physical interface goes down, it MUST send a PW status notification message for all PWs associated with the failure. When a PW label mapping is withdrawn, or PW status notification message is received, the egress PE MUST notify its client of this failure by deleting the VC using ILMI.

14. ATM-Specific Interface Parameter Sub-TLVs

The Interface parameter TLV is defined in [RFC4447], and the IANA registry with initial values for interface parameter sub-TLV types is defined in [RFC4446], but the ATM PW-specific interface parameter is specified as follows:

- 0x02 Maximum Number of concatenated ATM cells.

A 2-octet value specifying the maximum number of concatenated ATM cells that can be processed as a single PDU by the egress PE. An ingress PE transmitting concatenated cells on this PW can concatenate a number of cells up to the value of this parameter, but MUST NOT exceed it. This parameter is applicable only to PW types 3, 9, 0x0a, 0xc, [RFC4446], and 0xd and is REQUIRED for these PWC types. This parameter does not need to match in both directions of a specific PW.

15. Congestion Control

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion, with congestion characteristics depending on PSN type, network architecture, configuration, and loading. During congestion the PSN may exhibit packet loss that will impact the service carried by the ATM PW. In addition, since ATM PWs carry a variety of

services across the PSN, including but not restricted to TCP/IP, they may or may not behave in a TCP-friendly manner prescribed by [RFC2914]. In the presence of services that reduce transmission rate, ATM PWs may thus consume more than their fair share and in that case SHOULD be halted.

Whenever possible, ATM PWs should be run over traffic-engineered PSNs providing bandwidth allocation and admission control mechanisms. IntServ-enabled domains providing the Guaranteed Service (GS) or Diffserv-enabled domains using EF (expedited forwarding) are examples of traffic-engineered PSNs. Such PSNs will minimize loss and delay while providing some degree of isolation of the ATM PW's effects from neighboring streams.

It should be noted that when transporting ATM, Diffserv-enabled domains may use AF (Assured Forwarding) and/or DF (Default Forwarding) instead of EF, in order to place less burden on the network and gain additional statistical multiplexing advantage. In particular, Table 1 of Appendix "V" in [ATM-MPLS] contains a detailed mapping between ATM classes and Diffserv classes.

The PEs SHOULD monitor for congestion (by using explicit congestion notification, [VCCV], or by measuring packet loss) in order to ensure that the service using the ATM PW may be maintained. When a PE detects significant congestion while receiving the PW PDUs, the PE MAY use RM cells for ABR connections to notify the remote PE.

If the PW has been set up using the protocol defined in [RFC4447], then procedures specified in [RFC4447] for status notification can be used to disable packet transmission on the ingress PE from the egress PE. The PW may be restarted by manual intervention, or by automatic means after an appropriate waiting time.

16. Security Considerations

This document specifies only encapsulations, not the protocols used to carry the encapsulated packets across the PSN. Each such protocol may have its own set of security issues [RFC4447][RFC3985], 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 than a native ATM service.

17. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.

18. Informative References

- [FBATM] ATM Forum Specification af-fbatm-0151.000 (2000), "Frame Based ATM over SONET/SDH Transport (FAST)"
- [TM4.0] ATM Forum Specification af-tm-0121.000 (1999), "Traffic Management Specification Version 4.1"
- [I.371] ITU-T Recommendation I.371 (2000), "Traffic control and congestion control in B-ISDN".
- [I.610] ITU-T Recommendation I.610, (1999), "B-ISDN operation and maintenance principles and functions".
- [Y.1411] ITU-T Recommendation Y.1411 (2003), ATM-MPLS Network Interworking - Cell Mode user Plane Interworking
- [Y.1412] ITU-T Recommendation Y.1412 (2003), ATM-MPLS network interworking - Frame mode user plane interworking
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.

- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
- [VCCV] Nadeau, T., Pignataro, C., and R. Aggarwal, "Pseudowire Virtual Circuit Connectivity Verification (VCCV)", Work in Progress, June 2006.
- [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, November 2000.
- [ATM-MPLS] ATM Forum Specification af-aic-0178.001, "ATM-MPLS Network Interworking Version 2.0", August 2003.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.
- [RFC2684] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 2684, September 1999.

19. Significant Contributors

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

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

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

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

Steve Vogelsang
ECI Telecom
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: stephen.vogelsang@ecitele.com

Gerald de Grace
ECI Telecom
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: gerald.degrace@ecitele.com

John Shirron
ECI Telecom
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: john.shirron@ecitele.com

Andrew G. Malis
Verizon Communications
40 Sylvan Road
Waltham, MA
EMail: andrew.g.malis@verizon.com
Phone: 781-466-2362

Vinai Sirkay
Redback Networks
300 Holger Way
San Jose, CA 95134
EMail: vsirkay@redback.com

Chris Liljenstolpe
Alcatel
11600 Sallie Mae Dr.
9th Floor
Reston, VA 20193
EMail: chris.liljenstolpe@alcatel.com

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

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

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

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

John Rutenmiller
Marconi Networks
1000 Marconi Drive
Warrendale, PA 15086
EMail: John.Rutenmiller@marconi.com

Rick Wilder
Alcatel
45195 Business Court
Loudoun Gateway II Suite 300
M/S STERV-SMAE
Sterling, VA 20166
EMail: Rick.Wilder@alcatel.com

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

Authors' Addresses

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

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

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

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

Jeremy Brayley
ECI Telecom Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: jeremy.brayley@ecitele.com

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

Full Copyright Statement

Copyright (C) The IETF Trust (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST, AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

