The ATM Forum
Technical Committee

ATM-MPLS Network Interworking
Version 2.0

af-aic-0178.001

August 2003

ATM-MPLS Network Interworking Version 2.0

af-aic-0178.001

Preface

In most team endeavors, the efforts of several individuals deserve special recognition. This specification was no exception. The ATM Forum gratefully thanks the following individuals and respective employers for their contributions to this document:

Jim Harford (Chairman)	AdvanceNet Systems
Matthew Bocci (Editor)	Alcatel
Khalid Ahmad	Nortel Networks
Mustapha Aissaoui	Alcatel
Ghassem Koleyni	Nortel Networks
Sohel Khan	Sprint
Adeel Siddiqui	Cable & Wireless
John Rutemiller	Marconi
Phuc Nguyen	Global One
Robert Dianda	Sprint
Eric Davalo	Maple Optical Systems

This specification uses three levels for indicating the degree of compliance necessary for specific functions, procedures, or coding. They are indicated by the use of key words as follows:

Requirement: "Shall" indicates a required function, procedure, or coding necessary for compliance. The word "shall" used in text indicates a conditional requirement when the operation described is dependent on whether or not an objective or option is chosen.

Objective: "Should" indicates an objective which is not required for compliance, but which is considered desirable.

Option: "May" indicates an optional operation without implying a desirability of one operation over another. That is, it identifies an operation that is allowed while still maintaining compliance.

Table of Contents

Table of Figures

1 Introduction

Many telecommunications carriers use ATM technology to deliver services (e.g. voice, leased line, frame relay, native ATM) at the edge of their networks. At the same time, there is a belief that it may be advantageous to employ MPLS technology within the network core. The use of ATM-MPLS network interworking allows a network operator to deploy MPLS in the core of the network while continuing to leverage ATM technology at the network edge.

2 Scope

Comprehensive network interworking includes user data, signalling, routing, and management (including security) aspects. The scope of this specification is to specify the architecture, certain interworking mechanisms, and certain procedures between ATM and MPLS networks. It particularly specifies:

- network interworking architecture,
- user plane requirements,
- encapsulation formats and semantics,
- traffic engineering and traffic management guidelines

for ATM-MPLS network interworking, where an intermediate MPLS network provides a transport service between two ATM networks.

This specification aims, where possible, to provide a transparent transport of ATM connections across the MPLS network. The MPLS network appears as one or more ATM links within an ATM network. This capability is applicable to both switched and permanent ATM connections and is subject to the ATM service and MPLS network limitations listed in Section 6.6.

This specification does not mandate or assume a particular link layer technology within the MPLS network.

At the time of publication, the one-to-one mode of ITU-T draft Recommendation Y.1411 [11] was aligned with the cell mode encapsulation of this specification. It is also anticipated that AAL5 PDU mode under discussion in ITU-T and the IETF would be aligned with the AAL5 PDU mode described in this specification.

All appendices to this specification are informative.

3 References

The following references contain provisions that, through reference in this text, constitute provisions of this specification. At the time of publication, the editions indicated were valid. All references are subject to revision, and parties to agreements based on this specification are encouraged to investigate the possibility of applying the most recent editions of the references indicated below.

[1] IETF: RFC 3032 (January 2001): MPLS Label Stack Encoding

[2] ITU-T: Recommendation I.610 (February 1999): B-ISDN operation and maintenance principles and functions

[3] ITU-T: Recommendation I.732 (March 1996): Functional characteristics of ATM equipment

[4] IETF: RFC 3031 (January 2001): Multiprotocol Label Switching Architecture

[5] ATM Forum: af-pnni-0055.000 (March 1996): Private Network-Network Interface Specification Version 1.0

[6] ATM Forum: af-sec-0100.002 (March 2001): ATM Security Specification Version 1.1

[7] ATM Forum: af-tm-0121.000 (March 1999): Traffic Management Specification Version 4.1

[8] ITU-T: Recommendation I.510 (March 1993): Definitions and General Principles for ISDN Interworking

[9] IETF: RFC 3270 (May 2002): Multi-Protocol Label Switching (MPLS) Support of Differentiated Services

[10] ATM Forum: af-cs-0197.000 (August 2003): ATM-MPLS Network Interworking Signalling Version 1.0

[11] ITU-T: Recommendation Y.1411 (2003): ATM-MPLS Network Interworking – Cell Mode User Plane Interworking

[12] IETF: RFC 2475 (December 1998): Architecture for Differentiated Services

[13] IETF: RFC 2991 (November 2000): Multipath Issues in Unicast and Multicast Next-Hop Selection

[14] IETF: RFC 2992 (November 2000): Analysis of an Equal-Cost Multi-Path Algorithm

4 Acronyms and terminology

4.1 Acronyms

AAL ATM Adaptation Layer

AAL5 ATM Adaptation Layer Type 5

AF Assured Forwarding

ATM Asynchronous Transfer Mode

CAC Connection Admission Control

CDV Cell Delay Variation

CLP Cell Loss Priority

CPCS Common Part Convergence Sub-layer

CRC Cyclic Redundancy Code

CTD Cell Transfer Delay

DS Differentiated Services

EF Expedited Forwarding

EFCI Explicit Forward Congestion Indication

E-LSP EXP-Inferred Label Switched Path

EXP Experimental Bits

FEC Forwarding Equivalence Class

ILMI Integrated Link Management Interface

INE Interworking Network Element

ISH Interworking Specific Header

IWF Interworking Function

LDP Label Distribution Protocol

LER Label Edge Router

L-LSP Label-Inferred Label Switched Path

LSP Label Switched Path

LSR Label Switching Router

MPLS Multi-Protocol Label Switching

NMS Network Management System

OAM Operation Administration and Maintenance

PCI Protocol Control Information

PDU Protocol Data Unit

PHB Per-Hop Behavior

PM Performance Monitoring

PNNI Private Network-Network Interface

PSC Per-Hop Behavior Scheduling Class

PSN Packet Switched Network

PTI Payload Type Identifier

QoS Quality of Service

RCC Routing Control Channel

RM Resource Management

RSVP Resource Reservation Protocol

RSVP-TE Resource Reservation Protocol with Traffic Engineering

S-bit Stack bit

TLV Type Length Value

TTL Time To Live

UU User-to-User

VCC Virtual Channel Connection

VCI Virtual Channel Identifier

VPC Virtual Path Connection

VPI Virtual Path Identifier

4.2 Terminology

Interworking: The term interworking is used to express interactions between networks, between end systems, or between parts thereof, with the aim of providing a functional entity capable of supporting an end-to-end communication [8].

Network interworking: In network interworking, the PCI (Protocol Control Information) of the protocol used in two similar networks and the payload information are transferred, transparently, across an intermediate network by a pair of IWFs.

Interworking Function (IWF): An IWF includes the conversion between protocols and the mapping of one protocol to another. The functionality required between networks can be separated from the functionality, if any, required in end systems. The former functionality is considered to reside in an internetworking network element (INE). Additional details may be found in [8].

Interworking Network Element (INE): The INE is an entity where user plane, control plane, and management plane interworking functions (IWFs) may be implemented. The INE could be a standalone network element, part of the ATM switch, or part of a LSR located at the entrance to the MPLS network (LER).

Downstream INE: The INE receiving MPLS frames on an LSP.

Upstream INE: The INE sending MPLS frames on an LSP.

Cell Concatenation: The process of bundling a group of cells belonging to a VCC or a VPC into an MPLS frame. Note that this is not AAL5 reassembly.

Bundle: A set of one or more transport LSPs in each direction that together provide the appearance of a virtual ATM port interface to the ATM control protocols, including ATM signalling channel and PNNI routing control channel.

Per-Hop-Behaviour (PHB): The externally observable forwarding behavior applied at a DS-node (a node supporting differentiated services functions and behaviors) to a collection of packets with the same DS code points in a particular direction [12].

Per-platform label: A label in which the label value and context (i.e., FEC) are identical across all interfaces.

Label space: A set of unique labels assigned and distributed by the Downstream INE

5 Interworking Reference Model

Figure 1 shows the reference model for ATM-MPLS-ATM interworking, where an MPLS network interconnects two ATM networks. INEs perform network interworking between the MPLS network and the ATM networks, enabling end-to-end ATM services between users on different ATM networks to be carried across the MPLS network.

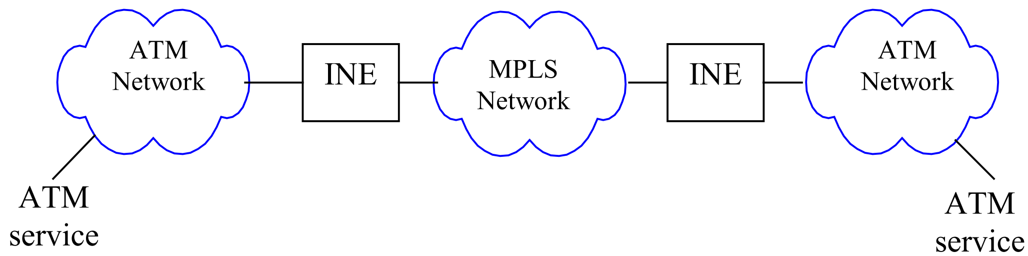

Figure 1: ATM-MPLS-ATM Interworking

6 ATM-MPLS-ATM User Plane Interworking Aspects

6.1 General Requirements

The general requirements for transparent transfer of ATM related information in the user plane are:

- The ability to multiplex multiple ATM connections (i.e. VPCs and/or VCCs) into an MPLS LSP.
- Support for the traffic contracts and the QoS commitments made to the ATM connections.
- The ability to transparently carry all AAL types.
- The ability to transparently carry all OAM cells, including the support for proper operation of OAM PM cells and OAM security cells.
- Transport of Resource Management (RM) cells.
- Transport of Cell Loss Priority (CLP) and Payload Type Indicator (PTI) information from the ATM cell header.
- The ability to encapsulate a single ATM cell within a single MPLS frame.
- The option to concatenate multiple cells into the same MPLS frame.
- The option to transport an AAL5 PDU within one or more MPLS frames for bandwidth efficiency.

6.2 Encapsulation Formats

Figure 2 shows the ATM-MPLS encapsulation format. It contains the Transport Label, the Interworking Label, the optional Reserved field, the optional Sequence Number field, and ATM-MPLS Interworking Specific Header and the Payload.

Bit

8 7 6 5 4 3 2 1

Transport label (4 octets)
Interworking label (4 octets)
Reserved (1 octet)
Sequence Number (2 octets)
ATM-MPLS interworking specific header(s) and payload(s)

Note: bit 8 is the most significant bit

Figure 2: ATM-MPLS Encapsulation Format

Figure 3 illustrates the relationship between the Transport Label and the Interworking Label.

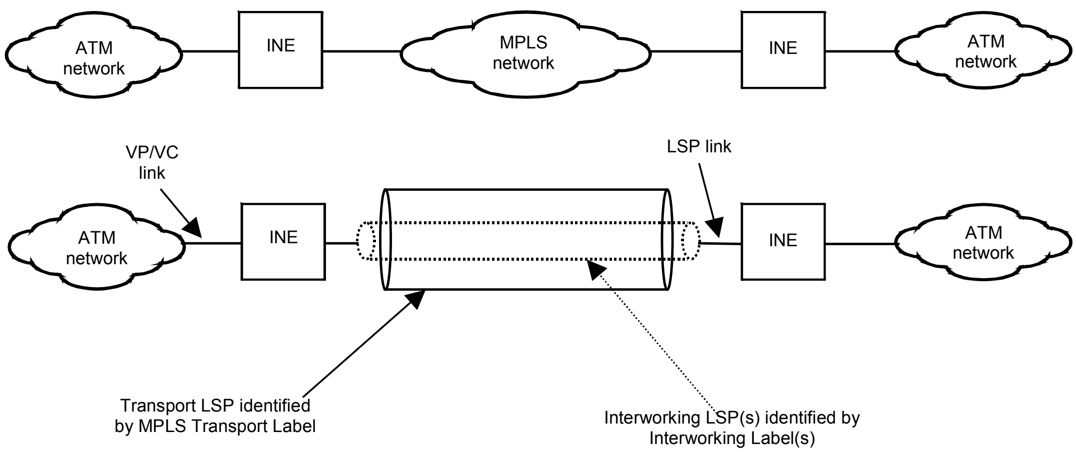

Figure 3: Relationship between transport label and the interworking label

6.2.1 Transport label

The 4-octet transport label identifies an LSP used to transport traffic between two ATM-MPLS interworking devices. Herein, this LSP is denoted a “Transport LSP”. The transport label is an MPLS label as defined by IETF RFCs 3031 and 3032. This label is visible to the core LSRs, which use it to forward the MPLS packet among the LSRs. The setting of the EXP and TTL fields of the transport label is outside the scope of this specification. The S bit is set to 0 for this label, indicating that this is not the bottom of the label stack.

Transport labels are assigned by label distribution protocols (e.g. LDP, RSVP-TE) or provisioned.

6.2.2 Interworking label

The 4-octet interworking label uniquely identifies an LSP, nested within the Transport LSP, that carries traffic associated with an ATM connection. Herein, this nested LSP is denoted an “Interworking LSP”. The interworking label is an MPLS label stack entry as defined by IETF RFCs 3031 [4] and 3032 [1]. One Transport LSP may carry more than one interworking LSP.

Interworking labels may be unique within an INE ("per-platform" as defined in RFC 3031). Alternatively, techniques may be used to create multiple label spaces within an INE, although interworking labels must be unique over at least a bundle between two peer INEs. However, the use of more than one label space within an INE imposes a number of restrictions on the MPLS features that may be used in the MPLS network. These features include LDP-based transport LSPs, shared-explicit style of RSVP-TE transport LSP establishment to support dynamic LSP modification, the ability to merge transport LSPs following LSP reroute, and penultimate hop popping. See Informative Appendix II for an example INE functional architecture that supports multiple label spaces.

Since an MPLS LSP is unidirectional while ATM connections are bi-directional, there shall be two different interworking LSPs, one for each direction of the ATM connection. These may have different label values.

The IWF maintains context information associated with the ATM connection carried by an interworking LSP. This information is referenced by means of the 20-bit label field of the interworking label. The 20-bit label field (i.e., the label value) of the interworking label is assigned by the Downstream INE and is communicated to the Upstream INE. The context of the interworking label field implies:

- Connection type: VCC or VPC.
- VPI value to be inserted in the ATM cells in the MPLS to ATM direction.
- For VCC connection types, the VCI value to be inserted in the ATM cells in the MPLS to ATM direction.

This does not preclude the inclusion of other context information.

When it is intended that transport LSPs carry both interworking LSPs defined by this specification and LSPs for other MPLS applications, the downstream INE shall coordinate the allocation of the label space among MPLS applications. For example, this may be done by restricting the range of the label space allocated by each of these applications. Other coordination methods may also be used.

Procedures for the generation and parsing of the interworking label are as follows:

ATM-to-MPLS direction
In the case of a VPC, translation of the VPI to the 20 bit label field is performed. In the case of a VCC, the VPI and VCI are translated to the 20-bit label field. This association is signaled or provisioned between a pair of peer IWFs.

The S bit is set to 1 to indicate the bottom of the label stack.

The EXP and TTL fields are defined in IETF RFC 3032. The TTL field shall be set to a configurable value; the default is 32, which allows the use of PNNI tandem switching. The EXP bits may optionally be used as defined in IETF RFC 3270 [9].

MPLS-to-ATM direction
In the case of a VPC, translation of the 20 bit label field to the VPI is performed. In the case of a VCC, the 20-bit label field is translated to the VPI and VCI. This association is signaled or provisioned between a pair of peer IWFs. MPLS frames received with an invalid or unassigned interworking label are discarded.

6.2.3 Reserved Field

The reserved field is a 1-octet field that may optionally be present for interoperability with other ATM-MPLS Network Interworking specifications. When present, it shall be set to zero and ignored upon receipt. The presence of this field shall be signalled [10] or configured. If the sequence number field is present, this field shall also be present. This field allows for coexistence with the use of ECMP in the MPLS core network (see Section 6.6.2).

This field is part of the Common Interworking Indicators field of Y.1411 [11].

6.2.4 Sequence Number Field

The sequence number field is a 2-octet field that may be used to verify packet order delivery across the intermediate MPLS network. The presence and use of this field is optional. If present and used, this field contains a sequence number generated by the IWF in the ATM-to-MPLS direction. If present and unused, this field is set to all zeroes.

The sequence number space is a 16-bit, unsigned circular space. The presence of this field shall be signalled [10] or configured. The use of this field shall be configured at the IWF in the ATM-to-MPLS direction. If the reserved field specified in section 6.2.3 is present, this field shall also be present.

This field is part of the Common Interworking Indicators field of Y.1411 [11].

ATM-to-MPLS direction
If the sequence number field is used, then the following procedures apply in the ATM to MPLS direction:

- The sequence number shall be set to 1 for the first MPLS frame transmitted on the interworking LSP.
- For each subsequent MPLS frame, the sequence number shall be incremented by 1.
- If the result of incrementing is a value of 65535 for the current MPLS frame, the sequence number shall be reset to 1 for the next MPLS frame.

MPLS-to-ATM direction
If the sequence number field is present, then the following procedures apply in the MPLS to ATM direction:

- If Sequence_Number is 0, then the received packet shall be considered in order.
- If this is the first packet on the Interworking LSP, then Expected_Sequence_Number = 1.
- For each subsequent received packet, perform the following:
- If Sequence_Number >= Expected_Sequence_Number and Sequence_Number - Expected_Sequence_Number < 32768, then the received packet is in order.
- Else, if Sequence_Number < Expected_Sequence_Number and Expected_Sequence_Number – Sequence_Number >= 32768, then the received packet is in order.
- Else the received packet is out of order.
- If the received packet is in order, then Expected_Sequence_Number = Sequence_Number+1 mod 2^{16}
- If Expected_Sequence_Number = 0, then Expected_Sequence _Number=1

The treatment of packets that are received out of order is network specific. Correct ATM sequence integrity shall be preserved. Out of order packets may be dropped or re-ordered.

6.2.5 ATM-MPLS Interworking Specific Header (ISH)

This header is inserted before an ATM payload and identifies whether cells or frames are encapsulated, plus other protocol control information. Note that length of the ISH depends on whether the VCI is present. Furthermore, the format of the ISH is different for cell and for frame encapsulation, as specified in Section 6.3.

6.3 Encapsulation Modes

Four modes of encapsulation are defined for ATM-MPLS-ATM network interworking:

- Single cell encapsulation
- Concatenated cell encapsulation, with no VCIP optimization
- Concatenated cell encapsulation, with VCIP optimization
- Frame (AAL5 PDU) encapsulation

Support for single cell encapsulation is mandatory. All other encapsulations are optional.

An objective of these encapsulation modes is to avoid the unnecessary transmission of the VPI and VCI fields, when such information can be derived from the context information referenced by the interworking label.

6.3.1 Cell Mode Encapsulation (Single and Concatenated Modes)

Note: This encapsulation mode is known in ITU-T Recommendation Y.1411 [11] as One-to-one mode.

In cell mode, one ATM cell or many concatenated ATM cells of an ATM VCC or VPC are encapsulated in a single MPLS frame. This specification mandates the support of one ATM cell per MPLS frame, which is the default case. Support for encapsulation of multiple ATM cells per MPLS frame is optional. When ATM cells are concatenated, they belong to the same ATM connection, either VCC or VPC. In the case of a VPC, the concatenated cells may belong to different VCCs within the VPC.

For VPCs, when concatenated cell mode encapsulation with no VCIP optimization is used, the VCI field shall be present for every cell within the MPLS frame.

For VPCs, when concatenated cell mode encapsulation with VCIP optimization is used, the VCI field shall only be present in the following cases:

- The cell is the first cell within the MPLS frame
- The previous cell within the MPLS frame belongs to a different VCC

For VCCs, the behavior of concatenated cell mode encapsulation is the same with or without VCIP optimization; the VCI field shall not be present within the MPLS frame.

The ISH headers and payloads for cell mode are shown in the following figures. Figure 4 shows the format for the single cell case, while Figure 5 shows the format for the concatenated cell case.

Bit

8	7	6	5	4	3	2	1
MODE	VCIP	RES		PTI			CLP
VCI (2 octets, if present)							
ATM PAYLOAD (48 octets)							

Note: bit 8 is the most significant bit

Figure 4: ISH and Payload for Cell Encapsulation of a Single Cell

Bit

8	7	6	5	4	3	2	1
MODE	VCIP	RES		PTI			CLP
VCI (2 octets, if present)							
ATM PAYLOAD (48 octets)							
MODE	VCIP	RES		PTI			CLP
VCI (2 octets, if present)							
ATM PAYLOAD (48 octets)							
MODE	VCIP	RES		PTI			CLP
VCI (2 octets, if present)							
ATM PAYLOAD (48 octets)							

Note: bit 8 is the most significant bit

Figure 5: ISHs and Payloads for Cell Encapsulation of Concatenated Cells

Description of the ATM-MPLS-ATM interworking specific header fields and the payloads:

MODE (bit 8): This field is set to 0 to indicate cell mode encapsulation.

VCIP (VCI Present, bit 7): This field is set to 1 when the VCI field is present, otherwise it is set to “0”.

RES (Reserved, bits 6-5): This field is set to 0 and ignored upon reception.

PTI (Payload Type Identifier, bits 4-2): This field carries the PTI coding of each encapsulated ATM cell.

CLP (Cell Loss Priority, bit 1): This field carries the CLP value of each encapsulated ATM cell.

VCI (2 octets): If present, this field carries the VCI value of the encapsulated ATM cell.

ATM PAYLOAD (48 octets): This field carries the payload of each ATM cell.

6.3.2 AAL5 PDU Mode Encapsulation

AAL5 PDU mode may be used to support AAL5 PDUs on an ATM VCC carried by an interworking LSP. ATM cells of an AAL5 PDU are re-assembled and the resulting partial (i.e. fragment) or complete AAL5 PDU is encapsulated in an MPLS frame.

In this mode, the entire AAL5 PDU (payload, padding, and whole AAL5 trailer, including UU, CPI, Length and CRC32) is encapsulated. The encapsulation for AAL5 PDU mode is shown in Figure 6.

Note: The use of this mode for AAL types other than AAL5 is not precluded. If this mode is used for any AAL type other than AAL5, maximum delay shall be bounded as described in section 6.6.1.1.

Note: This format is backwards compatible with the format defined as Frame mode in version 1 of this specification.

Bit

8	7	6	5	4	3	2	1
MODE	VCIP	RES			UU	EFCI	CLP
PAYLOAD							

Note: bit 8 is the most significant bit

Figure 6: ISH and Payload for AAL5 PDU Encapsulation

Description of ATM-MPLS interworking specific header fields and payload:

MODE (bit 8): This field is set to 1to indicate AAL5 PDU mode encapsulation.

VCIP (VCI Present, bit 7): This field is set to 0to indicate that no VCI field is present.

RES (Reserved, bits 6-4): This field is set to 0and ignored upon reception. Note that bit 4 is used in version 1 of this specification to convey a field deprecated in this version.

UU (User-to-User, bit 3): This field carries the least significant bit of the PTI field of the last ATM cell encapsulated within the MPLS frame.

> **ATM-to-MPLS direction:** The UU field of the ISH is set to the CLP state of the last ATM cell to be included in the MPLS frame.
>
> **MPLS-to-ATM direction:** The least significant bit of the PTI field for the last ATM cell in the MPLS frame is set to the value of the UU bit of the ISH. The least significant PTI bit is set to 0 for all other cells in the MPLS frame is set to 0.

EFCI (Explicit Forward Congestion Indicator, bit 2): This field carries the middle bit of the PTI field of the last ATM cell encapsulated within the MPLS frame.

> **ATM-to-MPLS direction:** The EFCI field of the ISH is set to the EFCI state of the last ATM cell to be included in the MPLS frame.
>
> **MPLS-to-ATM direction:** The EFCI state of all constituent cells within the MPLS frame is set to the value of the EFCI field in the ISH.

CLP (Cell Loss Priority, bit 1): This field carries an aggregated CLP value of the encapsulated ATM cells.

ATM-to-MPLS direction: The CLP field of the ISH is set to 1 if any of the constituent ATM cells to be included in the MPLS frame has its CLP bit set to 1; otherwise this field is set to 0.

MPLS-to-ATM direction: The CLP bit of all constituent cells within the MPLS frame is set to the value of the CLP field in the ISH.

PAYLOAD: This field carries the re-assembled AAL5 PDU (including the AAL5 padding and trailer) or the AAL5 PDU fragment. The length of this field is a multiple of 48 octets.

6.3.2.1 Fragmentation

This section only applies to AAL5 PDU mode encapsulation.

In some cases, it is not possible to encapsulate a full AAL5 PDU in an MPLS frame. This occurs when the MPLS MTU is reached or if an OAM or RM cell arrives between the first cell and the last cell of an AAL5 PDU. In these cases, the AAL5 PDU shall be fragmented. In addition, fragmentation should be used to bound ATM cell delay. Fragments of an AAL5 PDU may be encapsulated and transmitted across the MPLS network before all of the cells constituting the AAL5 PDU are received at the upstream INE.

When fragmentation occurs, the procedures described in this section shall be followed.

Note: These procedures are backwards compatible with the Fragmentation procedures defined in version 1 of this specification.

Procedures in the ATM-to-MPLS direction

The following procedures shall apply:

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

ii. Setting of the UU, EFCI, and CLP fields in the ISH of the fragment shall be as per Section 6.3.2.

iii. If the arriving cell is an OAM or an RM cell, send the current MPLS frame and then send the OAM or RM cell using single cell mode as defined in section 6.3.1.

Procedures in the MPLS-to-ATM direction

The following procedures shall apply:

i. The 3-bit PTI field of each ATM cell header is constructed as follows:

   a) The most significant bit is set to 0, indicating a user data cell.
   b) The middle bit is set to the EFCI value of the ISH of the fragment.
   c) The least significant bit for the last ATM cell in the MPLS frame is set to the value of the UU bit of the ISH. The least significant PTI bit for all other cells in the MPLS frame is set to 0.

ii. The CLP bit of each ATM cell header is set to the value of the CLP field in the ISH.

iii. When a fragment is received, each constituent ATM cell is sent in correct order.

6.4 Configurable Parameters

For each ATM connection, the following parameters shall be signalled or provisioned:

- Interworking labels (one for each direction).
- Connection type (VCC or VPC).
- Encapsulation format (including the presence of the reserved and sequence number fields).
- Maximum number of cells per MPLS frame in concatenated cell mode and AAL5 PDU mode. This value is limited by the smallest link MTU in the transport LSP and QoS commitments of the connection.

6.5 ATM OAM and RM Cells

6.5.1 ATM-to-MPLS Direction

6.5.1.1 OAM Cells

Several types of OAM cells are defined in [2]. Applications, such as those identified in [6], utilize these OAM cells. These cells are categorized as:

- Fault management cells.
- Performance monitoring and reporting, both in forward and backward directions.
- User OAM cells (e.g. security OAM cells).

At the ATM layer, two types of OAM cell flows are identified: F4 (OAM flow on virtual path level) and F5 (OAM flow on virtual channel level). F4 and F5 OAM cells are either segment flows for communicating OAM related information within the boundary of the VPC or VCC, or end-to-end for information regarding end-to-end VPC or VCC operations. From an OAM perspective, the INE behaves as an ATM switch.

OAM cells are always encapsulated with the cell mode encapsulation, regardless of the encapsulation format for user data. If the connection is in concatenated cell mode, then the OAM cell is included as one of the concatenated cells. For all other encapsulation modes, the OAM cell is sent in single cell mode.

For cell mode encapsulation of user data, OAM cells are encapsulated in the same manner as user data cells.

For AAL5 PDU mode encapsulation of user data, OAM cells that arrive during the reassembly of an AAL5 frame cause fragmentation procedures to be invoked. The partially reassembled AAL5 frame is sent as a fragment, immediately followed by the OAM cell. Reassembly of the AAL5 frame is then resumed. If an OAM cell arrives between AAL5 frames, then it is sent in cell mode encapsulation. This procedure ensures cell sequence integrity for user cells and OAM cells.

The general functional architecture of an ATM network element is provided in Figure 4-2/I.732 of ITU-T Recommendation I.732 [3]. This functional model is used below to describe the treatment of F4 and F5 OAM cells at the INE.

The INE performs switching at either the VP or the VC level. In the following sections, “AAL5 PDU mode encapsulation” refers to the encapsulation method for user data.

VP Switching - Cell Mode Encapsulation
F4 OAM cells are processed by the INE according to the procedures specified in [2] and are then sent across the LSP. F5 OAM are not inserted or extracted here and are therefore simply encapsulated and sent across the LSP.

VP Switching - AAL5 PDU mode Encapsulation
This case is not supported by this specification.

VC Switching - Cell Mode Encapsulation
F4 OAM cells are inserted or extracted at the VP link termination; such OAM cells are not seen at the VC link termination and are therefore not sent across the LSP. F5 OAM cells are inserted or extracted at the VC link termination or VC termination according to the procedures specified in [2] and are then sent across the LSP.

VC Switching - AAL5 PDU mode Encapsulation
This case is the same as “VC switching – Cell Mode Encapsulation”.

6.5.1.2 RM cells

RM cells use a PTI value of 110 [2] and are treated the same way as OAM cells with respect to the maintenance of cell ordering.

6.5.1.3 Cells with the PTI value Reserved for future VC functions

The PTI with binary value of 111 is reserved for future VC functions. Cells received with this value shall be treated the same as OAM cells with respect to the maintenance of cell ordering.

6.5.2 MPLS-to-ATM Direction

OAM and RM cells are received as single encapsulated cells. They are treated at the INE in accordance with procedures described in [2], [3] and [6].

6.6 Interworking Considerations

6.6.1 ATM

This section discusses restrictions on the applicability of this specification to the support of current ATM services.

6.6.1.1 General

An objective of this specification is to provide complete transparency for ATM services across the interworking function. Due to the nature of interworking across frame-based networks, several limitations are unavoidable.

The MPLS network may or may not maintain frame sequence ordering between the upstream INE and the downstream INE in steady state operation. In contrast, preservation of cell ordering is assumed as part of the ATM service. If the MPLS network does not maintain frame sequence ordering, then sequence numbers shall be used to detect the mis-ordering of MPLS frames. See section 6.2.4 for details.

The cell loss ratio may be adversely affected by transport across the MPLS network. An MPLS frame may be discarded due to a single bit error within the entire MPLS frame. By contrast, a separate HEC field protects ATM cell headers. Thus, ATM cells can be switched even if errors occur in the cell payload.

An INE has an implementation-specific contribution to the end-to-end maximum CTD and peak-to-peak CDV, when these parameters are specified for an ATM connection. The INE applies its contribution to the accumulated end-to-end values as specified in TM 4.1 [7]. Once the connection is established, the INE's contribution to the maximum CTD and peak-to-peak CDV shall not exceed the values previously advertised. Note that the INE contribution includes the time required for cell concatenation (cell encapsulation mode) or AAL5 PDU reassembly (AAL5 PDU encapsulation mode).

The delay commitment, if any, made to a connection, shall be taken into account when determining when to use fragmentation to bound delay at an INE.

At the time of publication, mechanisms for achieving QoS within an MPLS network were still evolving. This specification contains an informative appendix to assist network operators and equipment vendors in mapping ATM QoS to relevant MPLS mechanisms. However, careful network engineering will be required for the foreseeable future. This topic is for further study.

6.6.1.2 AAL5 PDU Mode Encapsulation

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 [7]. When tagging is not allowed by the conformance definition, AAL5 PDU mode shall not be used.

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.

6.6.2 MPLS

This section discusses restrictions on the applicability of this specification to current MPLS core networks.

There may be one or more label spaces for interworking labels within an INE. If multiple label spaces are created within an INE, they must be unique over at least a bundle between two INEs and subject to certain restrictions. If so, then the following MPLS functions may be affected:

- LSP merge following an LSP reroute may need to be disabled to prevent merging of bundles.

- Penultimate Hop Popping (PHP) may need to be disabled to prevent loss of the transport label which may be needed to associate interworking LSPs with bundles.

The absence of the reserved field (Section 6.2.3) may adversely affect the ability to deploy equal-cost multi-path (ECMP) (RFC 2991 [13] and RFC 2992 [14]) load balancing in the MPLS network. Without this field, core nodes in the MPLS network will not be able to distinguish between MPLS frames encapsulating IP packets and MPLS frames conforming to this specification. Since IP packets are carried in MPLS label stacks without any protocol identifier, historic values of the IP version number are used to distinguish between IP and non-IP MPLS payloads. If the reserved field is present and its first four bits are set to 0, then MPLS frames conforming to this specification will not be mistaken assumed to be carrying native IP packets.

7 ATM-MPLS-ATM Control Plane Interworking Aspects

It is the intention of this specification that af-cs-0197.000 [10] should be considered a companion specification, although other control mechanisms can be used. When af-cs-0197.000 [10] is used, this section applies.

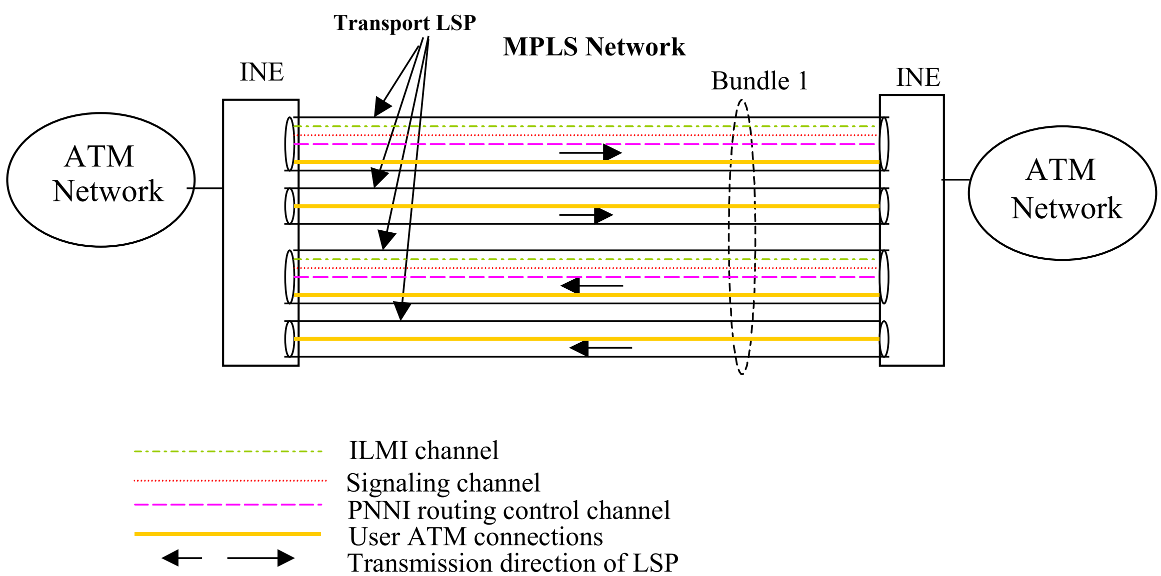

Figure 7: Connections to Support Control Plane

From a PNNI perspective, a “bundle” of transport LSPs is considered to be an abstraction of a PNNI physical link established between two PNNI nodes. For the ATM-MPLS-ATM network interworking architecture, the role of PNNI routing protocols and ILMI is the same as in an ATM PNNI network. The role of PNNI signalling is to establish interworking LSPs between INEs during ATM VCC or VPC establishment and to perform related signalling functions defined in the PNNI specification.

A bundle is a set of one or more transport LSPs in each direction that together provide the appearance of a virtual ATM port interface to the ATM control protocols. The ATM signalling channel within a bundle controls all of the ATM user connections in that bundle. Within each bundle, any transport LSP that carries ATM control channels shall support the QoS requirements of those control channels.

There are several models to group transport LSPs between a given pair of peer INEs into a bundle:

i. One bi-directional transport LSP pair. This is the degenerate case of a bundle. There is a one-to-one relationship between bi-directional transport LSP pairs and virtual ATM port interfaces; no additional abstraction of a PNNI physical link is needed. However, this will result in multiple PNNI/AINI links between peer INEs where there is more than one transport LSP pair. This is the only required model.

ii. All transport LSPs. This is for further study.

iii. A subset of transport LSPs. This is for further study.

Each bundle between a given pair of peer INEs shall be uniquely identified by provisioning. Identification of bundles is for further study.

Figure 7 shows various interworking LSPs (user ATM connections, signalling channel, routing control channel, and ILMI channel) aggregated within a bundle of transport LSPs.

The interworking labels of the ATM control channels (e.g. ATM signalling, PNNI routing control channel, and ILMI) are assigned and agreed upon by the two INEs, as specified in Section 7.3. The ATM control channels may use any of the encapsulation modes. The encapsulation mode of each ATM control channel is provisioned.

Connection admission control (CAC) is one of the control plane functions that is used to achieve ATM QoS.

7.1 Signalling

In order to set up an interworking LSP that carries an ATM switched or soft permanent connection over a transport LSP, the INE negotiates the interworking label for each direction and then binds them to the corresponding VPI/VCI values on the ATM interfaces. Signalling mechanisms to accomplish this are specified in [10].

See Section 6.6 for a discussion concerning the INE's contribution to the accumulated end-to-end maximum CTD and peak-to-peak CDV. When negotiating among multiple encapsulation modes, the QoS accumulation parameters shall account for the worst-case CTD and peak-to-peak CDV of the offered encapsulation modes.

7.2 Routing

The lowest level PNNI RCC is provisioned across a bundle of transport LSPs. The SVCC-based RCC is signaled across a bundle of transport LSPs via PNNI signalling as defined in Section 7.1. Each bundle between the two INEs is seen as a single hop link by the PNNI routing protocol. Each INE is able to advertise link state changes, as specified in PNNI [5].

7.3 LSP Maintenance

The transport LSP shall be provisioned or signaled. The signalling of transport LSPs is for further study.

An INE shall have procedures to implement at least the following:

- Transport LSP establishment, modification, and release.
- Assignment and reassignment of interworking LSPs to transport LSPs.
- Fault notification to the ATM control and management plane.
- Traffic management interworking (see Informative Appendix V for initial guidelines).

These items are for further study.

8 ATM-MPLS-ATM Management Plane Interworking Aspects

ATM OAM cells carry performance, fault, security and protection switching information for VCCs and VPCs on an end-to-end and segment basis [2] to support ATM layer management functions. The interworking function shall be capable of transferring ATM OAM information through the MPLS network by encapsulating ATM OAM cells in MPLS frames (see Section 6). In addition, the interworking function may correlate MPLS OAM information with ATM OAM information in the management plane through internal or external management interfaces. Correlation of MPLS OAM information and ATM OAM cells is for further study.

9 ATM-MPLS-ATM Traffic Management Aspects

The underlying transport LSP shall be capable of providing the required QoS for the ATM connections which are to be encapsulated. Informative Appendix V provides guidelines for the traffic engineering of transport LSPs and for the mapping of ATM service categories and conformance definition.

9.1 PNNI Resource Availability Information

The aggregate traffic characteristics of transport LSPs comprising a bundle shall be translated into PNNI resource availability information. This translation is network specific. The PNNI resource availability information per PNNI logical port is then advertised to other PNNI nodes.

9.2 ATM Service Categories

ATM supports several different service categories [7]. When a transport LSP is used to carry multiple ATM service categories, the LSP traffic parameters shall meet all ATM QoS requirements of the ATM service categories carried on that LSP.

The transport LSP may be an L-LSP or an E-LSP.

9.2.1 Label-Only-Inferred-PSC LSPs (L-LSP)

A Label-Only-Inferred- Per Hop Behavior Scheduling Class (PSC) LSP (L-LSP) is a transport LSP in which the Diffserv class of each packet is inferred from the label. The drop precedence level of each packet is inferred from the EXP field. The PSC is explicitly signaled or configured at LSP establishment, thereby enabling the LSR to infer the Diffserv class.

9.2.2 EXP-Inferred-PSC LSP (E-LSP)

An EXP-Inferred-PSC LSP (E-LSP) is a transport LSP in which both the Diffserv class of the packet and its drop precedence are inferred from the EXP field. The 3-bit EXP field of a transport LSP can represent 8 different combinations of Per Hop Behavior (PHB) and drop precedence levels. Mapping of the PSC and drop precedence to EXP fields are signaled or configured at LSP establishment thereby enabling the LSR to infer the Diffserv class and the drop precedence of the packet.

Informative Appendix I: Example of ATM-MPLS-ATM Network Interworking

In this example (Figure 8), ATM is present at the network edge as the protocol that brings multiple services into the packet core (e.g. frame relay, voice services, and circuit emulation). ATM connections are carried transparently over MPLS LSPs from one edge of the packet core to another. The LSPs are primarily used as transparent tunnels. An ATM/MPLS interworking function at each edge of the core multiplexes a number of ATM connections (VCCs, VPCs or both) into a transport LSP and originates the transport LSP. The pair of transport LSPs (providing bi-directional connectivity) between any two interworking devices at the edge of the core is either established using a network management system or is initiated through MPLS signalling.

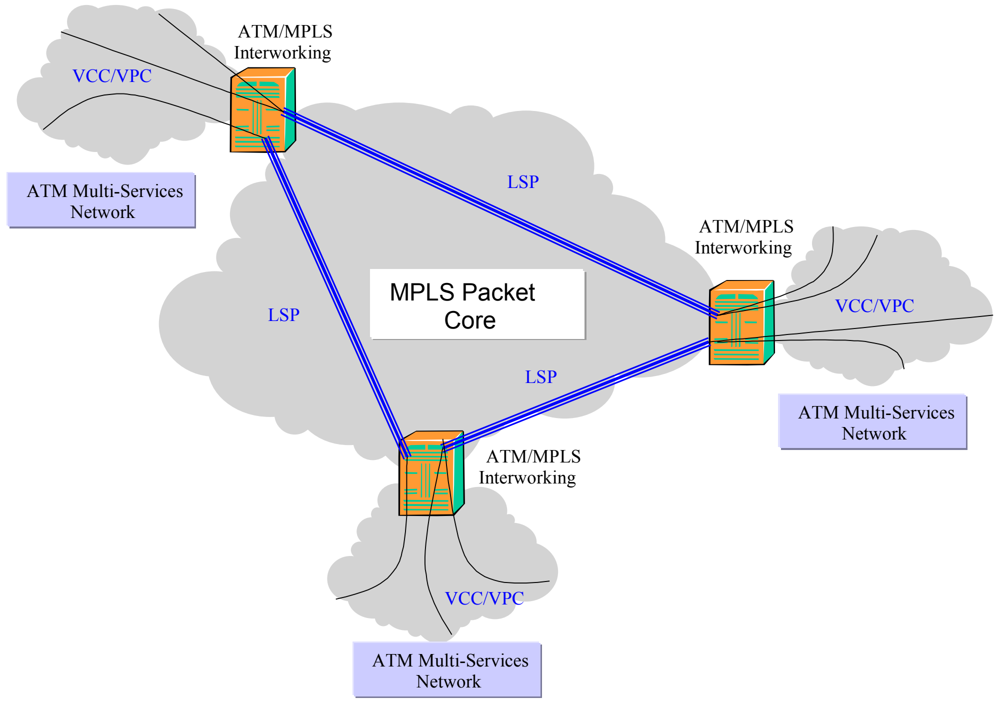

Figure 8: Example of ATM-MPLS-ATM Network Interworking

Transparency in this context means that ATM based services should be carried over the packet core unaffected. In this example, the interworking function is characterized by:

i. ATM and MPLS user data planes are interworked.

ii. ATM control plane information (signalling channels, routing control channels) and layer management plane information (OAM cells) are tunneled through the LSP from one ATM/MPLS edge to the other edge.

In this example, the recovery mechanisms of the ATM network are operational in the event a non-recoverable failure occurs in the MPLS core network. The following is an example of such a recovery operation. If a LSP fails in the MPLS network, the MPLS network is responsible for redirecting the traffic over a backup LSP. However, in some scenarios, the recovery may not be fast enough or may not be possible at all. In such cases, the MPLS network will have to generate a path failure indication back to the LSP originating points (i.e. the nodes where the interworking function resides). If the LSP between the 2 edges was originally signaled, it will be released back to its originating point via MPLS signalling procedures.

The LSP originating node, where the IWF resides, may act on the received failure indication to:

i. Generate the proper OAM cell type and content over the ATM PVCs towards the ATM network.

ii. Release ATM SVC calls back to their sources. ATM SPVC calls may be rerouted normally by their source nodes through an alternate path, which may or may not traverse the MPLS core. From the perspective of the PNNI protocol, this operation is completely transparent and does not require any additional routing or signalling procedures.

Informative Appendix II: Per-Bundle Interworking Labels

This Appendix provides an example of a functional model for an INE that would allow interworking labels to be unique over a bundle without being unique across the entire INE.

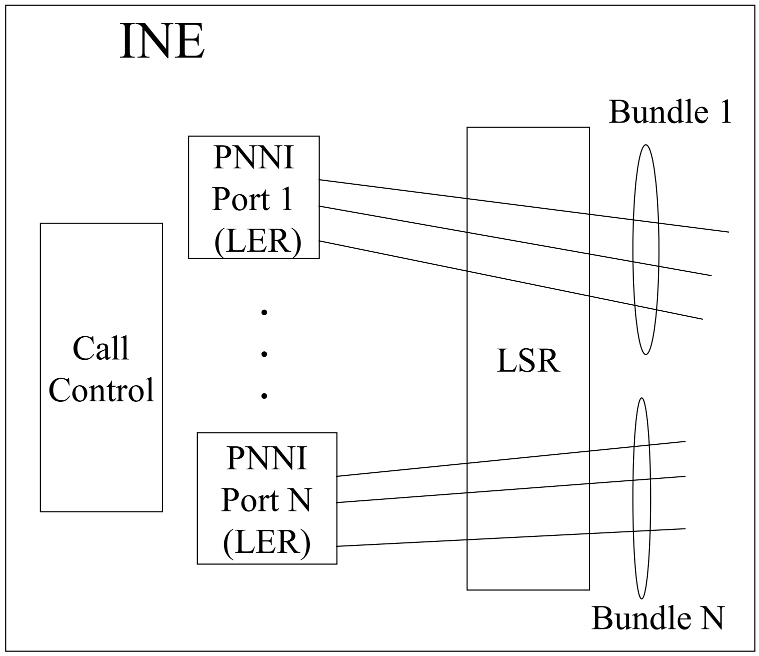

Figure 9 Example Funmctional Model For an INE

Figure 9 shows a possible functional model for an INE. There are 3 types of functional entities shown: LSR, PNNI port, and Call Control. A bundle of transport LSPs is modelled as a logical ATM port to ATM signalling and routing. For each logical PNNI port, there is a PNNI port functional entity. The PNNI port functional entity acts as an LER, which terminates/originates all of the transport and interworking LSPs for this PNNI port. The LSR switches the transport LSPs between the INE's external MPLS interfaces and the PNNI port functional entities. The call control entity provides the PNNI node call and routing control.

Each PNNI port has its own label space for interworking labels and can assign these independently of all the other PNNI ports. This allows for interworking labels that are unique over a bundle without being unique across the entire INE. In the context of RFC 3031, each PNNI port and the LSR is a platform that implements per-platform labels.

An advantage of this approach is that it allows an INE to scale to more than 2^{20} LSPs. Additionally, each PNNI port can be isolated from the other PNNI ports. For example, a PNNI port would detect and take appropriate action if another INE (in the same network or another network) used an interworking LSP that has not been assigned to it. However, if only per-platform labels are used for an INE, the downstream INE might not be able to determine which upstream INE had sent a particular MPLS frame. Therefore the downstream INE could not determine if the upstream INE was assigned the interworking LSP.

Informative Appendix III: Relationship Between ATM and LSP Connection Identifiers

Transport of ATM traffic over the MPLS network is performed using a two-level LSP stack. The outer label is the Transport LSP and provides the equivalent functionality of a transport link. The inner label is the Interworking LSP and provides support for VP/VC switching functions.

The transport LSP exists only within the MPLS network. It is used by the INE to provide a pseudo transmission path between interconnected INE devices. This provides the capability to perform ATM signalling and routing between INE devices.

Virtual Path Connections

When the INE is operating at the Virtual Path Connection (VPC) level, translation is performed between the ATM VPI and the IWF LSP label. This mode of operation is equivalent to performing virtual path switching in an ATM switch. The relationship between the ATM side of the INE and the MPLS side of the INE is shown in Figure 10.

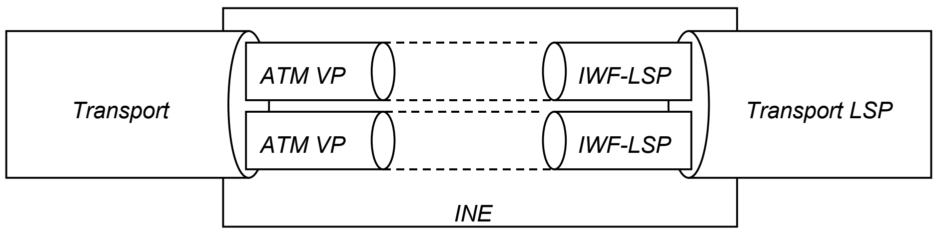

Figure 10: Relationship between VPC on ATM Side and LSP on MPLS Side of INE

An example of switching at the virtual path level is show in Figure 11.

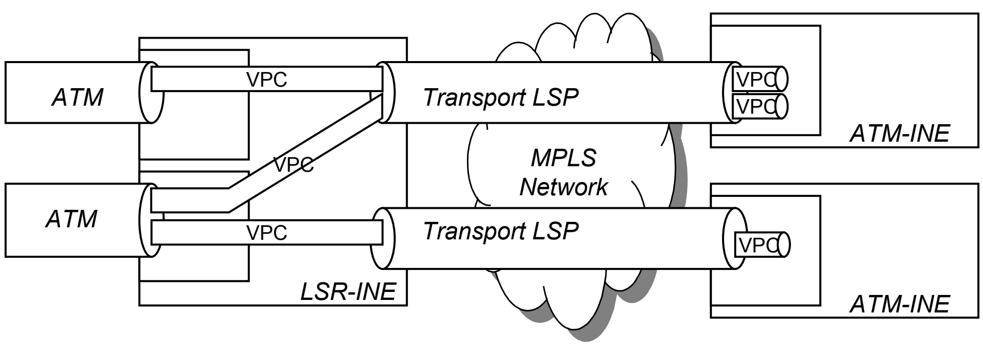

Figure 11: Switching at Virtual Path Level

Virtual Channel Connections

When the INE is operating at the Virtual Channel Connection (VCC) level, translation is performed between the ATM VPI/VCI and the IWF LSP label. This mode of operation is equivalent to performing virtual channels switching in an ATM switch. The relationship between the ATM side of the INE and the MPLS side of the INE is shown in Figure 12.

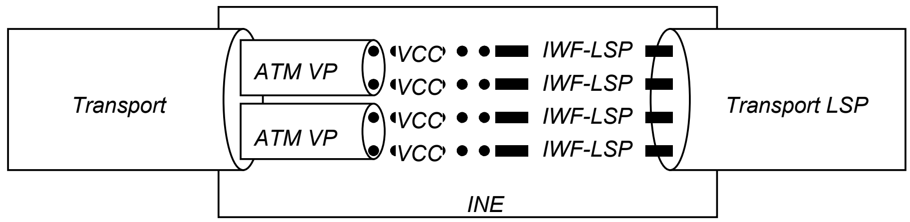

Figure 12: Relationship between VCC on ATM Side and LSP on MPLS Side of INE

An example of switching at the virtual channel level is show in Figure 13.

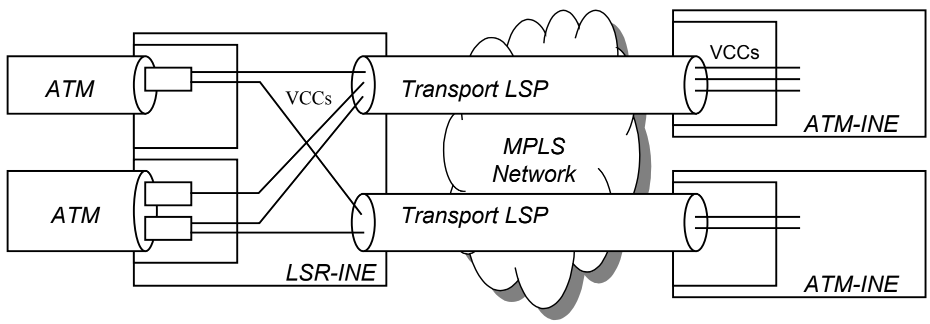

Figure 13: Switching at Virtual Channel Level

Relationship of ATM Connections and LSPs

An ATM connection is a bi-directional pair of virtual circuits that use the same identifiers for both directions of a communication channel. ATM switching can be viewed by examining only one of these directions. The ATM VP or VC arrives at the ingress side of an ATM switch. The input port/VPI/VCI is translated to the output port/VPI/VCI and switched based on the resulting values. Bi-directional communications is achieved by ensuring that switching of the output port/VPI/VCI maps back to the input port/VPI/VCI.

An LSP does not maintain this same relationship. The LSP label used in bi-directional communications is not required to have the same value. The relationship between an ATM switch, the INE and an LSP switch is shown in Figure 14.

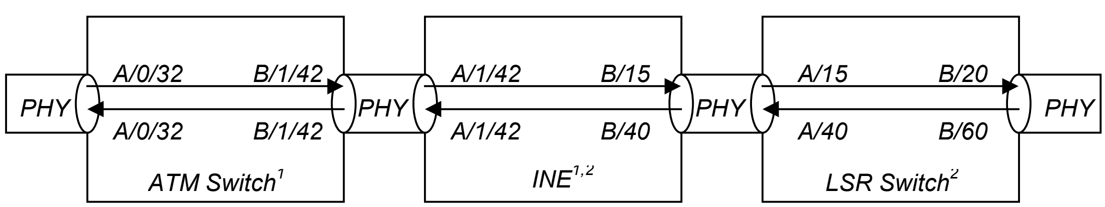

Notes:
1. x/y/z refers to port/vpi/vci
2. p/q refers to port/label

Figure 14: Relationship between an ATM switch, the INE and an LSP switch

Informative Appendix IV: INE Functional Decomposition

The Interworking Network Element includes three functional elements:

1. ATM functions interacting with an ATM network
2. MPLS functions interacting with an MPLS network
3. Interworking functions interacting between ATM functions and MPLS functions.

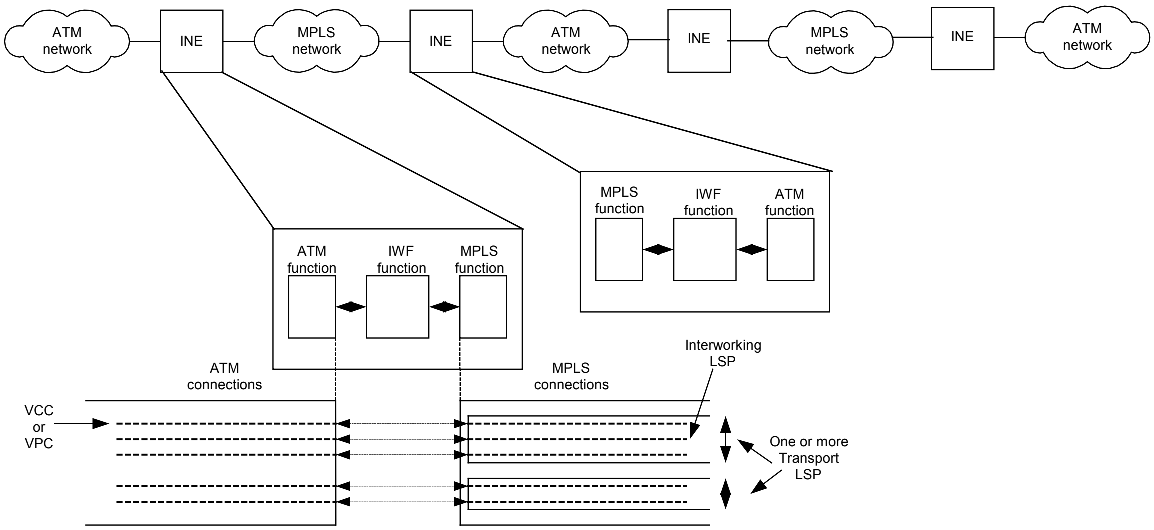

Figure 15 INE High Level Functional Model

Interworking functions between ATM and MPLS are managed at the following functional levels:

1. Data plane
2. Fault Management
3. Connection Management which includes binding connections, signalling interworking and traffic management
4. Control Plane Routing function

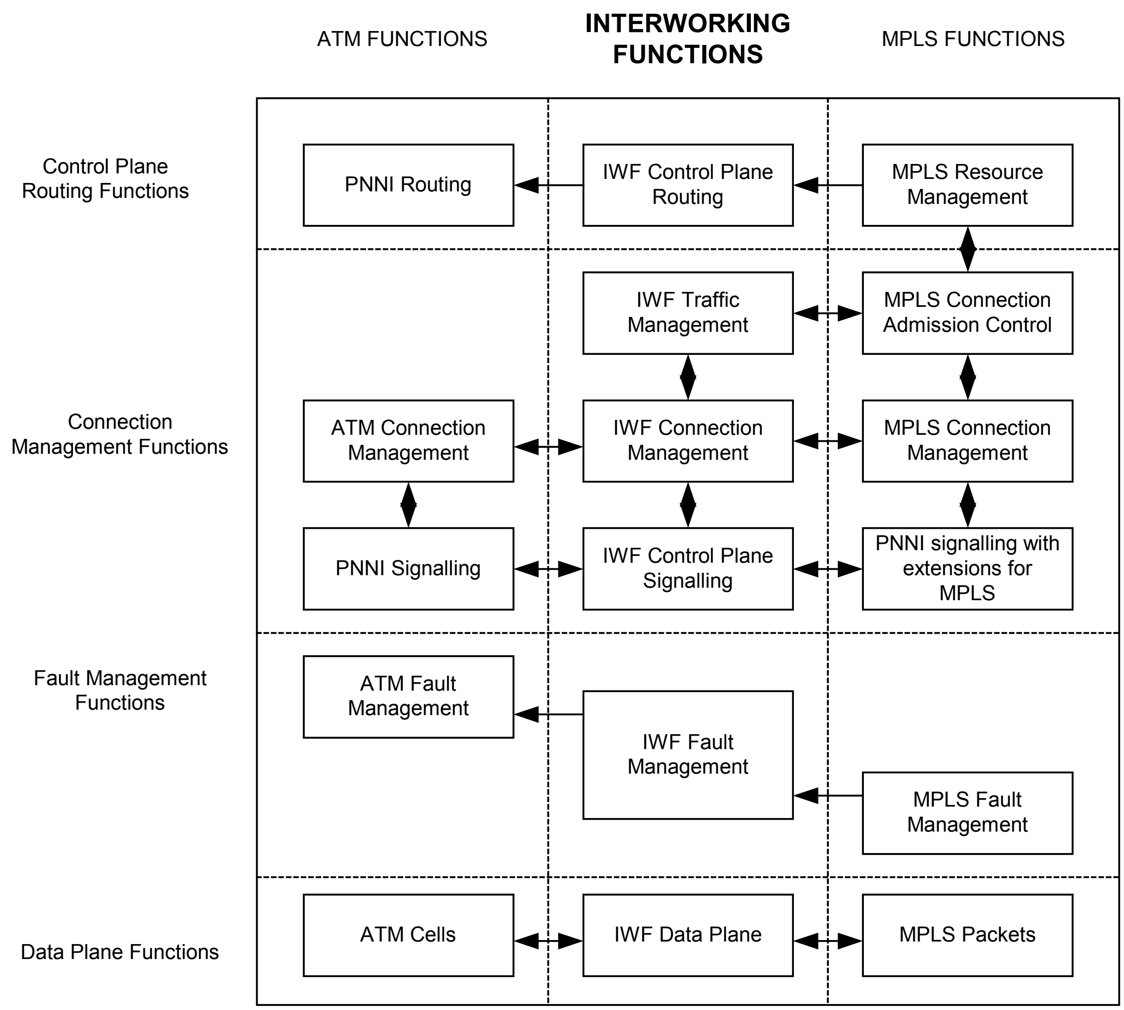

Figure 16 Interworking Functions

Data Plane

The IWF Data Plane function is responsible for:

1. ATM cells encapsulation into an interworking LSP
2. ATM cells extraction from an interworking LSP

This function can support multiple encapsulation modes.

Fault Management

ATM and MPLS networks have a client/server relationship as far as fault management is concerned. The MPLS fault management entity provides MPLS fault detection. The IWF fault management entity notifies MPLS layer faults to the ATM fault management entity.

Connection Management

Connections have to be managed on both ATM and MPLS sides. The IWF connection management function binds an ATM connection with an MPLS connection. The MPLS connection consists of a combination of an inter-working LSP and a transport LSP.

The IWF traffic management function translates ATM traffic contract, service category and QoS parameters into similar MPLS elements to allow connection admission control to be performed on the MPLS side.

The MPLS signalling functions setup, teardown and restart inter-working LSPs. It selects interworking labels from the INE-wide or interface-wide pool of available MPLS labels. The IWF signalling entity provides conversion between PNNI signalling on the ATM side and the extended PNNI signalling protocol used on the MPLS side.

Negotiation of configurable parameters of the ATM-MPLS interworking function for a given Interworking LSP occurs between peer INEs. This negotiation does not span multiple INEs that are not peers.

Control Plane - Routing

The IWF control plane routing function advertises into the ATM network link state changes resulting from the modification of transport LSPs characteristics. The type of routing information exchanged between the INE and the ATM network will depend on the type of interface between the INE and the ATM network (i.e., PNNI inside link, PNNI outside link, AINI, UNI or IISP).

Informative Appendix V: Traffic Engineering and Traffic Management Guidelines

This appendix provides guidelines for the provision of QoS for the individual interworking LSPs over the MPLS network. The objective is to provide the ability to support the traffic contracts and the QoS commitments made to the ATM connections.

ATM connections associated with different ATM service categories can either be multiplexed over a single Transport LSP or transported over several Transport LSPs. The mapping of ATM connections to Transport LSPs must take into consideration the Traffic Engineering properties of Transport LSPs. Two LSP types are defined in this appendix:

- Class Multiplexed LSP. This transports several ATM service categories.
- Class Based LSP. This transports a single service category.

Class Multiplexed LSP

ATM connections with different ATM service categories are multiplexed over one Transport LSP. The following are two examples:

- All ATM service categories multiplexed over one Transport LSP
- Real-time traffic (CBR and rt-VBR) connections are multiplexed over one Transport LSP while non-real-time traffic (nrt-VBR, ABR and UBR) connections are multiplexed over a second Transport LSP.

Class multiplexed LSPs can be implemented using either the E-LSP or L-LSP methods.

Class Multiplexed L-LSP Method

In this method, the LSP must meet the most stringent QoS requirements of the ATM connections transported by the LSP. LSRs in the core MPLS network only look at the transport label. They cannot decide which scheduling treatment to apply on a per-packet or per-ATM connection basis. All packets in one L-LSP receive the same scheduling treatment.

For instance, if the LSP carries two 20Mbit/s connections of type CBR and UBR, the LSP needs to be created with 40Mbit/s and it should provide CBR type of QoS for both UBR and CBR traffic.

The following figure illustrates the implementation of the Class Multiplexed L-LSP method.

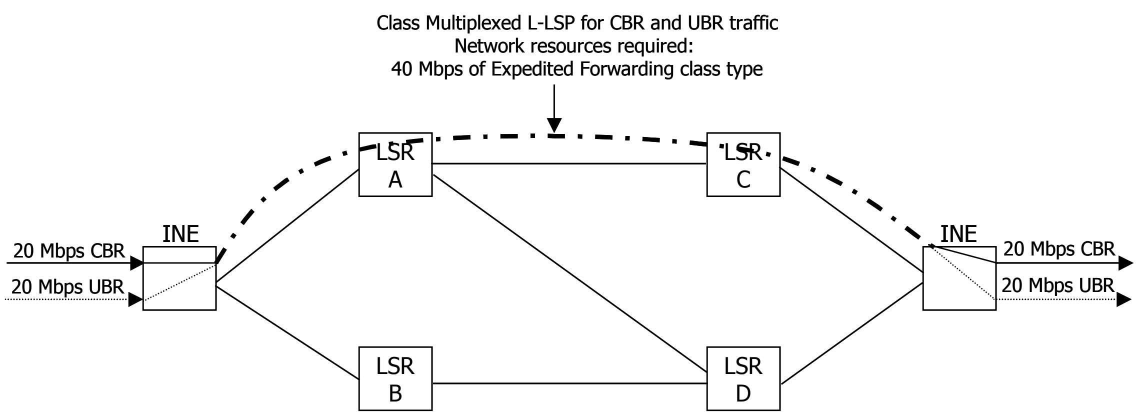

Figure 17: Class Multiplexed L-LSP

Class Multiplexed E-LSP Method

The MPLS network identifies the scheduling and dropping treatments to be applied to a packet based on the value of the EXP field inside the transport label. Each LSR can then apply a different scheduling treatment for each packet transported over an E-LSP.

The Class Multiplexed E-LSP method requires a mapping between ATM service categories and the four possible values of the EXP-bits field. An example mapping is given in Table 1.

The Class Multiplexed E-LSP method requires the MPLS network routing function to find a path for the E-LSP that meets the requirements of all of the ATM service categories transported in that LSP. This method increases the complexity of the path-searching algorithm by increasing the number of constraints to be satisfied. It can also lead to an under-utilization of network resources because one single path must meet the constraints of all ATM service categories transported over that E-LSP.

The following figure illustrates an implementation of the Class Multiplexed E-LSP method for the previously described example.

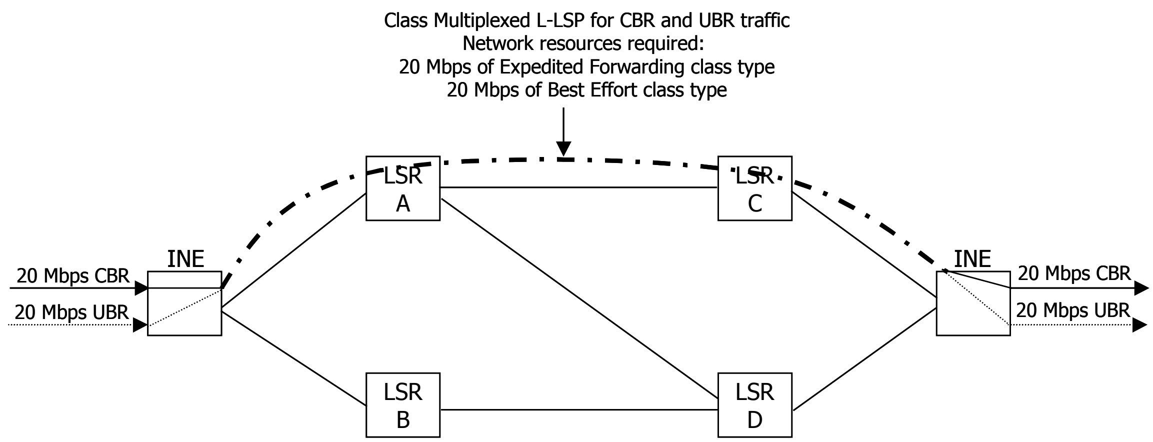

Figure 18: Class Multiplexed E-LSP

Class-Based LSP

In the Class-based LSP scenario a Transport LSP carries ATM connections that all belong to the same ATM service category. The MPLS network routing function needs only to find a path inside the MPLS network that meets the requirements for the one ATM service category carried over that LSP. This method provides a more granular approach than the Class Multiplexed LSP approach and may allow a better optimization of network resources.

This scenario can be implemented using either L-LSP or E-LSP methods. The E-LSP version of this method only supports one class of service per LSP.

The following figure describes a typical implementation of the Class Based LSP method.

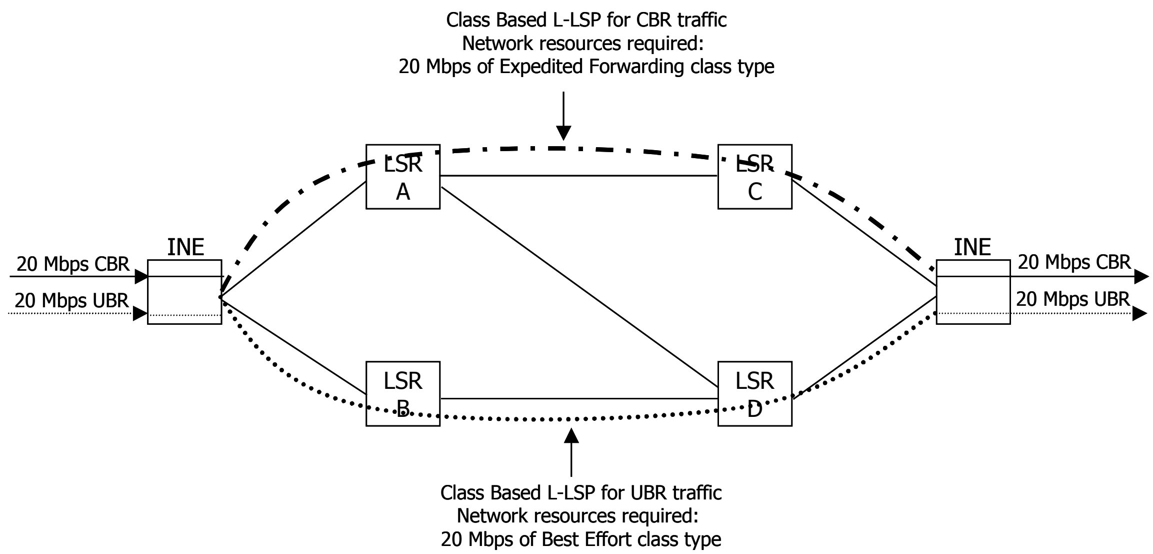

Figure 19: Class Based L-LSP

Diff-Serv Support of ATM Service Categories

When an ATM service is configured over an MPLS network, each ATM service category should be mapped to a compatible class of service in the MPLS network. A compatible class of service maintains the integrity of the end-to-end service. For example, the CBR service category should be mapped to a class of service with stringent loss and delay objectives. If the MPLS network uses the IP Diff-Serv framework to provide QoS classes, a class of service based on the Expedited Forwarding (EF) Per-Hop Behavior (PHB) is a good candidate.

Furthermore, ATM service categories have support for multiple conformance definitions. The conformance definition specifies the conformance of cells of a connection at an interface, e.g., public UNI, in relation to the conformance algorithm and corresponding parameters specified in the connection traffic descriptor [V.1]. For example, the conformance definition specifies if the requested QoS parameters, e.g., CLR, apply to the aggregate CLP0+1 conforming cell flow or to the CLP0 conforming flow only. Thus, the conformance definition should be respected in the selected PSN class of service. For example, a connection CLP1 cell flow in a VBR.3 conformance definition is treated as excess traffic in the ATM network and thus has no QoS guarantees associated with it. This flow should be provided a treatment no better than the treatment of the CLP0 cell flow in the MPLS network. This does not mean however that the MPLS network should mirror the various conformance definitions of the ATM service categories.

In the remainder of this section, it is assumed that the MPLS network implements the IP Diff-Serv framework to provide QoS. All ATM traffic management functions specified in [V.1] are applicable for the ATM part of the interworking LSP in the upstream and downstream INE. In the ATM-to-MPLS direction, each ATM connection may be policed and/or shaped to conform to its traffic descriptor in the ATM endpoint of the interworking LSP in the INE. Whenever allowed by the conformance definition, a non-conforming CLP0 cells may be tagged. Connection admission control (CAC) should be applied to make sure sufficient resources are available to carry the interworking LSP over the transport LSP. The mapping of ATM service category and conformance definition to the Diff-Serv PHB determines the outgoing PHB. This is the PHB to be applied to the cells or packets of the interworking LSP in the upstream INE, downstream INE, and by the MPLS network. The transport label of the outgoing MPLS packet should be marked to identify the selected PHB. This consists of marking the EXP field in the transport shim header.

Table 1 provides an example of mapping ATM service category and conformance definition to a Diff-Serv PHB. The Expedited Forwarding (EF) and Assured Forwarding (AF) PHBs are defined in [V.3] and [V.2] respectively.

ATM Service Category	Conformance Definition	CLP Setting	Diff-Serv PHB[1]	Diff-Serv Code Point
CBR	CBR.1	0/1	EF	101110
Rt-VBR	VBR.1	0/1	EF	101110
	VBR.2/VBR.3	0	AF41	100010
	VBR.2/VBR.3	1	AF42	100100
Nrt-VBR	VBR.1	0/1	AF31	011010
	VBR.2/VBR.3	0	AF31	011010
	VBR.2/VBR.3	1	AF32	011100
ABR	ABR	0	AF21	010010
UBR+MDCR	UBR.1/UBR.2	0/1	AF21	010010
GFR	GFR.1/GFR.2	0	AF21	010010
	GFR.1/GFR.2	1	AF22	010100
UBR	UBR.1/UBR.2	0/1	DF	000000

1. RFC2597 does not mandate relative levels of priority between AF classes. In typical implementations, AF4x is treated as the highest and AF1x the lowest.

Table 1: Example ATM Service Category to Diff-Serv Class Mapping

Note that an alternative mapping may not distinguish between the conformance definitions in a given ATM service category. In this case, the CLP0 and CLP1 flows of an ATM connection are provided with the same QoS in the MPLS network. As an example, all conformance definitions of the nrt-VBR service category may be mapped to the AF31 PHB.

For an MPLS network, the selected PHB for the interworking LSP is conveyed in different ways depending if the transport LSP is an L-LSP or an E-LSP. In the case of an L-LSP, the PHB Scheduling Class is signaled at the transport LSP establishment and is therefore inferred from the transport label value. The Drop Precedence of the individual MPLS packets is conveyed in the EXP field of the transport LSP shim header. In the case of an E-LSP, the PHB is conveyed in the EXP field of the transport LSP shim header. The actual values to be marked in the EXP field to reflect the example in Table 1 are outside the scope of this document.

In the presence of congestion, the INE may mark the EFCI bit and may perform selective cell discard based on CLP setting, if allowed by the conformance definition, and in accordance with TM4.1 [V.1]. The method used to transfer the CLP and EFCI information of the individual cells into the ATM-MPLS interworking specific header of the interworking LSP packet is described in detail in Section 6.

In the MPLS-to-ATM direction, the ATM service category and conformance definition is obtained from the context accessed via the interworking label. The information needed to reconstruct the ATM header, including the setting of the CLP and EFCI fields, for the individual cells is contained in the ATM specific information field of the MPLS packet. The method used to determine the CLP and EFCI information of the individual cells from the ATM specific information field of the MPLS packet is described in detail in Section 6.3.

References

[V.1] ATM Forum: af-tm-0121.000 (March 1999): Traffic Management Specification Version 4.1

[V.2] IETF: RFC 2597 (June 1999): Assured forwarding PHB Group

[V.3] IETF: RFC: 3246 (March 2002): An Expedited Forwarding PHB

Informative Appendix VI: Multiple Bundles

This specification allows for multiple bundles to exist between a pair of INEs. Each bundle is a logical PNNI port. This has some potential advantages versus all of the transport LSPs being in a single bundle and being treated as a single logical PNNI port. The following are some cases where this may be the case:

1. The ATM PNNI model is based on a facility associated model – i.e. the signalling channel is in the facility (e.g. VP or physical port). For ATM/MPLS this same model would have the logical PNNI port be a pair of transport LSPs.
2. An INE may be connected to 2 LSRs and to support diverse routing capabilities (e.g. some ATM customers use a pair of ATM Soft PVCs to provide redundancy and require that these 2 connections to be diversely routed to prevent a single point of failure) it may be desirable to have two logical PNNI ports: one with all of its transport LSPs routed to one of the LSRs and the other logical PNNI port with all of its transport LSPs routed to the other LSR.
3. It may be useful to have a logical PNNI port dedicated to specific applications (e.g. to support voice over ATM with a VPN). One advantage is that it can be used to protect that application, not only by bandwidth reservation but also by protection against call overload.
4. The QoS parameters of a bundle, e.g. delay, must be the worst of all the transport LSPs in a bundle. So if there are alternate routes between two INEs with significantly different delay characteristics, it may be advantageous to have different bundles so that transport LSPs taking the “short” path are in one bundle which then would have “better” QoS than would be if transport LSPs of the “long” path were also included in the bundle.
5. In some implementations call control is distributed among multiple processors and having multiple logical PNNI ports may make it easier to share the call load among the processors.

END OF DOCUMENT