

**TDM Transport over MPLS using
AAL1**

Implementation Agreement

MPLS/Frame Relay Alliance 4.0

**MPLS/Frame Relay Alliance Technical Committee
June 6, 2003**

Note: The user's attention is called to the possibility that implementation of the MPLS implementation agreement contained herein may require the use of inventions covered by patent rights held by third parties. By publication of this MPLS implementation agreement the MPLS/Frame Relay Alliance makes no representation that the implementation of the specification will not infringe on any third party rights. The MPLS/Frame Relay Alliance take no position with respect to any claim that has been or may be asserted by any third party, the validity of any patent rights related to any such claims, or the extent to which a license to use any such rights may not be available.

Editors:

Ron Insler
RAD Data Communications

For more information contact:

The MPLS/Frame Relay Alliance

Suite 307
39355 California Street
Fremont, CA 94538 USA

Phone: +1 (510) 608-3997
FAX: +1 (510) 608-5917
E-Mail: info@mplsforum.org
WWW: <http://www.mplsforum.org/>

Full Notice

Copyright © 2003 MPLS and Frame Relay Alliance.

All rights reserved.

This document and translations of it may be copied and furnished to others, and works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the MPLS Forum, except as needed for the purpose of developing MPLS implementation agreements (in which case the procedures copyrights defined by the MPLS Forum must be followed), or as required to translate it into languages other than English

This document and the information contained herein is provided on an "AS IS" basis and THE MPLS FORUM DISCLAIMS 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.

Table of Contents

1	INTRODUCTION.....	1
1.1	PURPOSE	1
1.2	SCOPE & OVERVIEW.....	1
1.3	DEFINITIONS.....	1
1.4	REFERENCE ARCHITECTURE	2
1.5	ACRONYMS.....	3
1.6	REFERENCES.....	4
2	TDM AND MPLS NETWORK INTERWORKING REQUIREMENTS	5
3	PROTOCOL STACK AND FRAME FORMAT	5
3.1	TDM-MPLS FRAME FORMAT	5
3.2	TDM-MPLS HEADER FORMAT	6
4	TDM-MPLS PAYLOAD	7
5	GENERAL PEQUIREMENTS	10
5.1	QUALITY OF SERVICE	10
5.2	JITTER AND PACKET LOSS	11
5.3	ALARMS AND INDICATION HANDLING	11
6	TDM-MPLS FRAME PROCESSING	12
6.1	GENERATING TDM-MPLS FRAMES	12
6.1.1	Setting the sequence number	12
6.1.2	Handling the R-bit	12
6.1.3	Handling the L-bit	12
6.2	RECEIVING TDM-MPLS FRAMES	13
6.2.1	Processing the sequence number	13
6.2.2	Handling the R-bit	14
6.2.3	Handling the L-bit	14
7	CLOCK RECOVERY	15
8	TDM SERVICES	15
9	TDM ATM SERVICE INTERWORKING.....	17
10	ATM PW COMPATIBILITY MODE.....	18

Revision History

Version	Change	Date
MPLS.x	Initial version	May, 2003

1 Introduction

1.1 Purpose

The purpose of this Implementation Agreement (IA) is to define network interworking between TDM circuits ($n \times 64$ kbps, E1/T1/E3/T3) over MPLS Label Switched Paths (LSPs) by using AAL1 encapsulation.

Note: SDH/SONET circuit emulation over MPLS networks is outside the scope of this document.

1.2 Scope & Overview

MPLS has the potential to consolidate service providers' networks and services such as Frame Relay, ATM, voice, TDM circuit emulation and IP services over a single infrastructure (see *Figure 1-1*).

The multiservice over MPLS reference model is shown in *Figure 1-1*. It consists of 3 distinct network elements:

- MPLS Switches (LSRs) using cell or frame transport with MPLS control protocol
- Interworking Gateways (PE) for adapting and interworking a variety of services for transport over an MPLS network
- Access / aggregation nodes that implement existing technologies.

This document suggests a method for transporting TDM ($n \times 64$ kbps, E1/T1/E3/T3) traffic across MPLS networks between two interworking nodes. In particular, the following issues are addressed:

- Method of mapping TDM traffic (including CAS) onto MPLS
- Definition of protocol data units and procedures for carrying user payloads between interworking nodes
- Exchanging alarm indications between interworking nodes
- MPLS packet loss monitoring.

The TDM services supported are defined in *Section 8, Table 2*.

1.3 Definitions

Must, Shall or Mandatory – the item is an absolute requirement of this implementation agreement.

Should – the item is desirable.

May or Optional – the item is not compulsory, and may be followed or ignored according to the needs of the implementer.

Circuit bundle – A circuit bundle is a stream of bits that have originated from the same physical interface or from different interfaces that share a common clock, and are transmitted from a specific ingress PE to a specific egress PE over a dedicated LSP. A

PE typically handles more than one circuit bundle at a time. For example, bundles may comprise several 64 kbps timeslots originating from a single E1, or an entire T3 or E3. Circuit bundles are single -direction streams, but are frequently coupled with bundles in the opposite direction to enable full-duplex communications. More than one bundle can be transmitted between two PEs, as is the case when the PW limits the bundle's packet length.

Circuit Bundle Identifier – A unique identifier of a circuit bundle associated with a certain LSP.

Customer Edge – A Customer Edge (CE) device is the customer device connected to a Provider Edge device.

Egress – The point where the TDM service is decapsulated from a MPLS PDU (MPLS to TDM direction).

Ingress –The point where the TDM service is encapsulated into a MPLS PDU (TDM to MPLS direction).

Label Switched Path – A Label Switched Path (LSP) is the path through one or more MPLS nodes at one level of the hierarchy over which packets in a particular Forwarding Equivalence Class (FEC) are transmitted.

MPLS Node – An MPLS node is a device that is aware of MPLS control protocols, will operate one or more layer three routing protocols, and will be capable of forwarding packets based on LSP labels.

Notes – outside of Tables and Figures are informative.

PseudoWire – A Pseudo-Wire (PW) is a connection between two PEs, carried over an MPLS network.

PseudoWire Emulation Edge-to-Edge – pseudowire emulation edge-to-edge (PWE3) is a mechanism that emulates the essential attributes of a TDM service (and of some other services not covered in this document) over an MPLS network.

Provider Edge – A Provider Edge (PE) device is a Network Edge device that provides TDM service over an MPLS network. In some cases, the edge device can belong to the customer; but for consistency with other multiservice over MPLS contributions, the term PE will be used.

1.4 Reference Architecture

The TDM and MPLS network interworking Reference model is shown in *Figure 1-1*. It consists of the following elements:

- MPLS core network consisting of Label Switching Routers (LSRs)
- Provider Edge (PE) devices providing network interworking functions between TDM and MPLS
- TDM devices connected via framed or unframed TDM Interface.

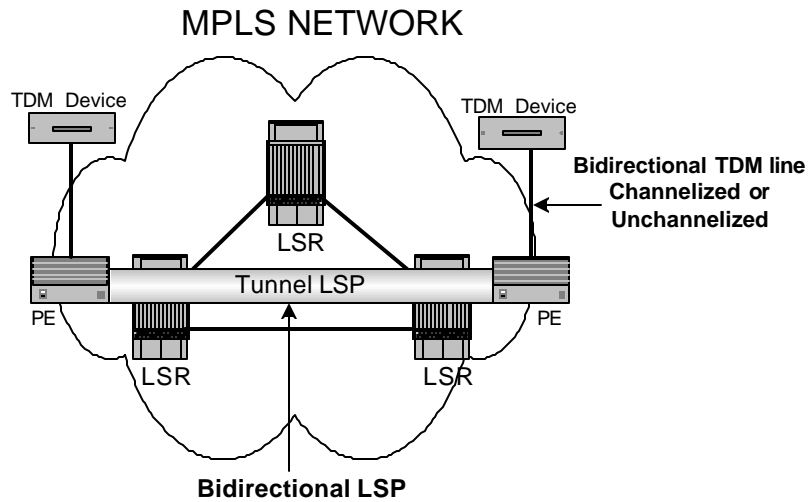


Figure 1-1 – TDM-MPLS Reference Architecture

TDM-MPLS network interworking connects two TDM networks and/or devices using an MPLS network. The TDM service is first provisioned between the TDM interface of the PE device and the TDM device. A bi-directional LSP is created either by manual provisioning or by using an MPLS signaling protocol (e.g., CR-LDP, RSVP-TE) to complete the circuit emulation connection. The use of the MPLS network by two TDM networks and/or devices is not visible to the end users. The PE provides all mapping and encapsulation functions necessary to ensure that the service provided to the TDM networks and/or devices is unchanged by the presence of an MPLS transport network.

1.5 Acronyms

AAL1	ATM Adaptation Layer 1
ATM	Asynchronous Transfer Mode
CAS	Channel Associated Signaling
CB	Circuit Bundle
CBID	Circuit Bundle Identifier
CE	Customer Edge
FEC	Forwarding Equivalency Class
IA	Implementation Agreement
IETF	Internet Engineering Task Force
IP	Internet Protocol
LDP	Label Distribution Protocol
LOS	Loss Of Signal
LSB	Least Significant Bit
LSP	Label Switched Path
LSR	Label Switching Router

MPLS	Multi Protocol Label Switching
MSB	Most Significant Bit
OOS	Out Of Service
PE	Provider Edge
PPP	Point to Point Protocol
PW	PseudoWire
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical NETwork
TBD	To Be Determined
TDM	Time Division Multiplexing
VCI	Virtual Channel Identifier (ATM networks)
VPI	Virtual Path Identifier (ATM networks)

1.6 References

- [1] ITU-T Recommendation G.704 (10/98) Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44736 kbps hierarchical levels
- [2] ITU-T Recommendation G.823 (03/00) The control of jitter and wander within digital networks which are based on the 2048 kbps hierarchy
- [3] ITU-T Recommendation G.824 (03/00) The control of jitter and wander within digital networks which are based on the 1544 kbps hierarchy
- [4] ITU-T Recommendation I.363.1 (08/96) B-ISDN ATM Adaptation Layer (AAL) specification: Type 1
- [5] ATM forum specification atm-vtoa-0078 (CES 2.0) Circuit Emulation Service Interoperability Specification Version 2.0
- [6] IETF RFC 2615, PPP over SONET/SDH, June 1999.
- [7] IETF RFC 3032, MPLS Label Stack encoding, January 2001.
- [8] ITU-T Recommendation Y.1411 ATM, MPLS Network Interworking- Cell mode User Plane Interworking
- [9] draft-ietf-pwe3-atm-encap-01.txt (02/03) Encapsulation Methods for Transport of ATM Cells/Frame Over IP and MPLS Networks, L. Martini et al., work in progress

2 TDM and MPLS Network Interworking Requirements

This section lists the requirements to be met in a TDM and MPLS network interworking configuration.

TDM Transport – Must support transfer of TDM traffic: framed (including fractional/channelized) and unframed E1/T1, $n \times 64$ kbps, and E3/T3 over MPLS

Frame Ordering – Must preserve frame ordering end-to-end

Signaling – Must provide for transparent transfer of CAS

Timing – May provide a mechanism for reconstruction of the TDM clock per ANSI/T1.101-1999 or appropriate ITU-T G.81x/G.82x recommendations

Alarm Mapping – Must support the transport of standard alarms to and from the physical interface across the MPLS network

Interworking – Should be amenable to connection to existing $n \times 64$ kbps, E1/T1, E3/T3, AAL1 and TDMoIP equipment without requiring extensive processing capabilities. May support compatibility with ATM PW N:1 mode as per [9], [8].

Overhead vs. Latency Tradeoff – The option to configure, by user or by signaling protocol, the number of AAL1 cells per frame must be given in order to balance overhead and latency.

3 Protocol Stack and Frame Format

3.1 TDM-MPLS Frame Format

Figure 3-1 shows the frame format and encapsulation of TDM-MPLS frames expressed in number of bytes per component. The TDM-MPLS frame format consists of a TDM-MPLS header followed by the payload. An MPLS LSP label stack [7] precedes the TDM-MPLS frame. This frame is encapsulated in a link layer frame (e.g., Ethernet, PPP in HDLC framing [6]).

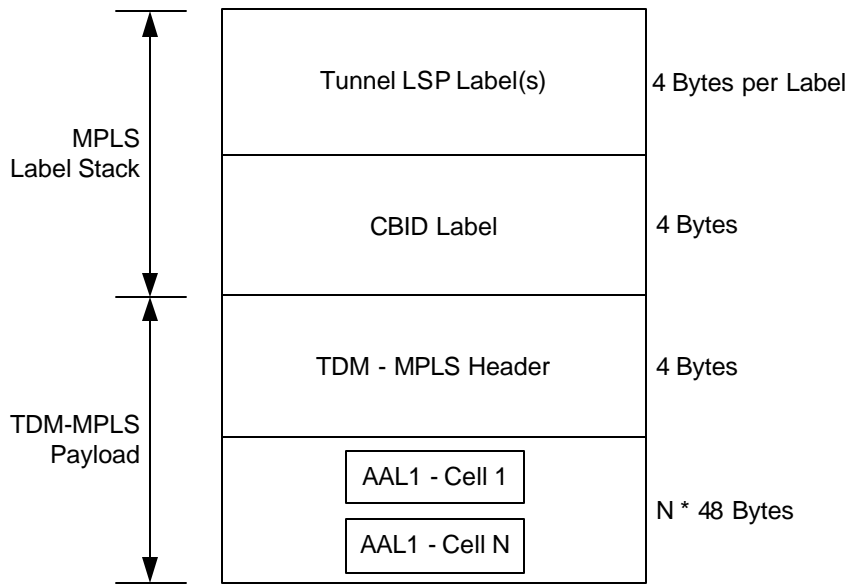


Figure 3-1 – Format and General Encapsulation of TDM-MPLS Frame

Note: Overhead vs. latency tradeoff – A largervalue of n leads to increased latency and smaller overhead; and vice versa, a smaller value of n results in decreased latency and bigger overhead. The number of cells per frame must be user-configurable or signaling protocol configurable.

3.2 TDM-MPLS Header Format

The TDM-MPLS header is shown below. *Table3-1* defines the role of each field.

BIT	0-3	4	5	6-9	10-15	16-31
	Reserved	L	R	Reserved	Length	Sequence Number

Figure 3-2 – TDM-MPLS Specific Header

Field	Meaning
Bits 0-3: Reserved	Reserved bits. Set to zero on transmission and ignored on reception.
Bit 4: L-bit	The L-bit being set indicates that the source has detected or has been informed of a TDM physical layer or ATM VC/VP fault impacting the data to be transmitted. This bit can be used to indicate physical layer LOS that should trigger AIS generation at the far end. When interworking with ATM the L-bit initiates VC/VP-AIS in the ATM environment. L-bit support is mandatory.
Bit 5: R-bit	The R-bit being set indicates that the source is not receiving packets at its TDM-MPLS receive port. R-bit support is mandatory.

Bits 6 – 9: Reserved	its TDM-MPLS receive port. R-bit support is mandatory. Reserved bit. Set to zero on transmission and ignored on reception.
Length: (Bits 10 – 15)	The length field is used to indicate the length of the TDM-MPLS packet (TDM-MPLS header and payload) in case padding is employed to meet minimum transmission unit requirements of layer 2 of the MPLS network. It must be used if the TDM-MPLS packet length (including layer 2 overhead) is less than 64 bytes, and must be set to zero if this length exceeds 64 bytes.
Sequence Number: (Bits 16 – 31)	The 16 bits sequence number may be used to guarantee ordered packet delivery. Processing of the sequence number field is optional. The sequence number space is a 16 bit, unsigned, circular space. The sequence number must be set to zero when this field is not used.

Table 3-1 – TDM-MPLS Frame Header Fields

4 TDM-MPLS Payload

The AAL1/2 family of protocols is a natural choice for trunking applications. Although originally developed to adapt various types of application data to the rigid format of ATM, the mechanisms are general solutions to the problem of transporting constant or variable bandwidth data streams over a byte-oriented packet network. For the prevalent case for which the timeslot allocation is static and no activity detection is performed, the payload can be most efficiently encoded using constant bit rate AAL1 adaptation. The AAL1 format is described in ITU-T standard I.363.1 [4] and its use for TDM circuit emulation is explained in ATM Forum specification atm-vtoa-0078 (CES 2.0)[5].

The TDM-MPLS payload consists of between one and thirty-one 48-octet sub-frames. The number of sub-frames which can be inferred by the receiving side from the total length of the payload is preconfigured and typically chosen according to latency and bandwidth constraints. Using a single sub-frame reduces latency to a minimum, but incurs the highest overhead. Using eight sub-frames, for example, reduces the overhead percentage, while increasing the latency by a factor of eight.

TDM-MPLS header	48 Octet Sub-frame
-----------------	--------------------

Figure 4-1 – Single TDM-MPLS-AAL1 Sub-Frame per TDM-MPLS Frame

TDM-MPLS header	48 Octet Sub-frame (1)	48 Octet Sub-frame (n)
-----------------	------------------------	------------------------

Figure 4-2 – Multiple TDM-MPLS-AAL1 Sub-Frames per TDM-MPLS Frame

Bit 0	Bits 1-3	Bits 4-6	Bit 7	
C	SN	CRC	P	47 octets of payload

Figure 4-3 – Sub-Frame Format

where:

- C** – (1 bit) convergence sub-layer indication; its use here is limited to indication of the existence of a pointer (see below)
C=0 means no pointer; C=1 means a pointer is present.
- SN** – (3 bits) sequence number, increments from sub-frame to sub-frame
- CRC** – (3 bits) is a 3-bit error cyclic redundancy code on C and SN
- P** – (1 bit) even byte parity

As can be readily inferred, this octet can only take on eight different values, and incrementing the sequence number forms an eight sub-frame sequence number cycle, the importance of which will become clear shortly.

The MPLS frame is protected by layer 2 frame protection (e.g CRC), hence for TDM network interworking the receiver **should not** check the content of the CRC and P bits in the AAL1 cell header. On the other hand, for ATM service interworking, the CRC and P bits of the AAL1 sub-frame **must** be checked, since the sub-frame had originated in the ATM network. Hence checking the CRC and P bits in the AAL1 sub-frame **shall** be configurable by the user or by a signaling protocol.

The structure of the remaining 47 octets in the TDM-MPLS AAL1 sub-frame depends on the sub-frame type, of which there are three, corresponding to the three types of AAL1 circuit emulation service defined in ATM Forum specification af-vtoa-0078 [5]. These are known as unstructured circuit emulation, structured circuit emulation and structured circuit emulation with CAS.

The simplest sub-frame is the unstructured one, which is used for transparent transfer of whole trunks (T1, E1, T3, E3), and serial Nx64Kbps interfaces. The 47 octets after the sequence number octet contain 376 bits from the TDM bit stream. No frame synchronization is supplied or implied; framing is the sole responsibility of the end-user equipment. Hence the unstructured mode can be used for leased lines which carry data rather than $n \times 64$ kbps timeslots, even for trunks with non-standard frame synchronization. For the T1 case, the raw frame consists of 193 bits, and hence one 183/193 T1 frame fits into each TDM-MPLS AAL1 sub-frame. The E1 frame consists of 256 bits, and so one 15/32 E1 frame fits into each sub-frame. In either case, if a single sub-frame is used per TDM-MPLS frame, the overhead percentage is 12 header octets when using a stack of 2 labels + 1 octet of AAL1 subframe header for every 47 TDM payload octets, i.e. 21.6%. However, for eight TDM-MPLS AAL1 sub-frames per TDM-MPLS frame, the payload is $8 \times 47 = 376$ octets, resulting in an overhead of about 3.3%. These overhead figures do not include layer 2 overhead.

When the TDM trunk is segmented into timeslots according to ITU-T G.704 [1], and it is desired to transport $n \times 64$ kbps circuit where n is only a fraction of the full E1 or T1, it is more efficient to use one of the structured AAL1 circuit emulation services. Structured AAL1 views the data not merely as a bit stream, but as a bundle of timeslots. Furthermore, when CAS is used it can be formatted such that it can be readily detected and manipulated.

In the structured circuit emulation mode without CAS, n octets from the n timeslots to be transported are first arranged in order of timeslot number. Thus, if timeslots 2, 3, 5, 7 and 11 are to be transported, the corresponding five octets are placed in the sub-frame immediately after the sequence number octet.

This placement is repeated until all 47 octets in the sub-frame are taken:

Octet	1	2	3	4	5	6	7	8	9	10	...	41	42	43	44	45	46	47
Timeslot	2	3	5	7	11	2	3	5	7	11	...	2	3	5	7	11	2	3

The next sub-frame commences where the present sub-frame left off:

Octet	1	2	3	4	5	6	7	8	9	10	...	41	42	43	44	45	46	47
Timeslot	5	7	11	2	3	5	7	11	2	3	...	5	7	11	2	3	5	7

and so forth. The set of timeslots 2, 3, 5, 7, 11 is called a structure, and the point where one structure ends and the next commences is a structure boundary.

The problem with this arrangement is the lack of explicit indication of the octet identities. As can be seen in the above example, each TDM-MPLS AAL1 sub-frame starts with a different timeslot, so a single lost frame will result in misidentifying timeslots from that point onwards, without possibility of recovery. The solution to this deficiency is the periodic introduction of a pointer to the boundary of the next structure. This pointer need not be used too frequently, as the timeslot identification is uniquely inferable unless frames are lost.

The particular method used in AAL1 is to insert a pointer once every sequence number cycle of length eight sub-frames. The pointer is seven bits and protected by an even parity MSB, and so occupies a single octet. Since seven bits are sufficient to represent offsets larger than 47, the placement of the pointer octet can be limited to sub-frames with an even sequence number. Unlike the usual TDM-MPLS AAL1 sub-frames with 47 octets available for payload, sub-frames which contain a pointer, called P-format sub frames, have the following format:

Bit 0	Bits 1-3	Bits 4-6	Bit 7	Bit 0	Bits 1- 7	
C	SN	CRC	P	E	Pointer	46 octets of payload

where

C – (1 bit) convergence sub-layer indication, C=1 for P-format sub frames

SN – (3 bits) even sequence number

CRC – (3 bits) is a 3 bit error cyclic redundancy code on C and SN

P – (1 bit) even byte parity LSB for sequence number octet

E – (1 bit) even byte parity MSB for pointer octet

Pointer – (7 bits) byte pointer to next structure boundary

The MPLS frame is protected by layer 2 frame protection (e.g., CRC); hence in case of TDM and MPLS network interworking, the MPLS receiver **should not** check the content of CRC, P and E bits in the AAL1 cell to detect errors in the AAL1 sub-frame. On the other hand, in

an application of ATM to MPLS service interworking, the CRC, P and E bits of the AAL1 sub-frame **must** be checked since the sub-frame had originated in the ATM network. Hence checking the CRC, P and E bits in the AAL1 sub-frame **shall** be configurable by user or signaling protocol.

P-format sub-frames have 46 octets of payload and the next sub-frame has 47 octets. Viewed as a single entity, the pointer needs to indicate one of 93 octets. If P=0, it is understood that the structure commences with the following octet (i.e. the first octet in the payload belongs to the lowest numbered timeslot).

P=93 means that the last octet of the second sub-frame is the final octet of the structure, and the following sub-frame commences with a new structure. The special value P=127 indicates that there is no structure boundary to be indicated (needed when extremely large structures are being transported).

The P-format sub-frame is always placed at the first possible position in the sequence number cycle that a structure boundary occurs, and can only occur once per cycle.

The only difference between the structured circuit emulation format and structured circuit emulation with CAS is the definition of the structure. Whereas in structured circuit emulation the structure is composed of the n timeslots, in structured circuit emulation with CAS the structure encompasses the super-frame consisting of multiple repetitions of the n timeslots and then the CAS bits. The CAS bits are tightly packed into octets and the final octet is padded with zeros if required.

For example, for E1 trunks the CAS bits are updated once per super frame of 16 frames. Hence the structure for $N \times 64$ derived from an E1 with CAS consists of 16 repetitions of n octets, followed by n sets of the four ABCD bits, and finally four zero bits if n is odd. For example, the structure for timeslots 2, 3 and 5 will be as follows:

```
2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5
2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 [ABCD2 ABCD3] [ABCD5 0000]
```

Similarly for T1 ESF (Extended Super frame) trunks, the super frame is 24 frames, and the structure consists of 24 repetitions of n octets, followed by the ABCD bits as before. For T1 SF (Super frame, D4), the signaling bits will in general appear twice, in their regular (bit-robbled) positions and at the end of the structure.

5 General Pequirements

5.1 Quality of Service

TDM-MPLS does not provide a mechanism to ensure timely delivery or provide other quality of service guarantees. Hence it is required that the MPLS layer and/or Layer 2 do it.

TDM-MPLS assumes a relatively benevolent environment. On the MPLS side, this means a network with prioritization and sufficient bandwidth, and low probability of bit error, packet order interchange or lost packets.

Layer 2 priority can be bestowed upon a TDM-MPLS stream by using the VLAN priority field. The MPLS layer can provide priority by using EXP bits or by using MPLS traffic engineering protocol extensions. Switches and routers, which the TDM-MPLS stream must

traverse, should be configured to respect these priorities. Sufficient bandwidth can be guaranteed by under-subscription and/or traffic engineering.

5.2 Jitter and Packet Loss

In order to compensate for packet delay variation, which exists in any MPLS network, a jitter buffer **shall** be provided. The length of this buffer **should** be configurable and **may** be dynamic (i.e., grow and shrink in length according to the statistics of the delay variation).

In order to handle (infrequent) packet loss and misordering, a packet order integrity mechanism **shall** be provided. This mechanism **shall** track the sequence numbers of packets in the jitter buffer and **must** take appropriate action when faults are detected. When missing packets are detected the mechanism **shall** output interpolation packets in order to retain TDM timing. Packets with incorrect sequence numbers or other detectable header errors **may** be discarded. Packets arriving in incorrect order **may** be reordered. Whenever possible, interpolation packets **should** ensure that proper synchronization bits are sent to the TDM network.

Packet loss causes degradation of the perceived audio quality of voice services. Hence, for TDM timeslots known to be transporting voice services, the contents of the interpolation packets can be important. In the simplest implementation, predefined constant values **may** be substituted for the lost TDM data, resulting in preservation of TDM timing but rapid deterioration of voice quality. A better alternative replays TDM data, so that each timeslot is fed with data previously received; if the TDM system utilizes CAS, then such replay **must** be performed in such fashion as to ensure the integrity of the CAS. In optimal implementations, each timeslot known to be carrying voice services is analyzed and packet loss concealment techniques utilized.

5.3 Alarms and Indication Handling

The TDM-MPLS PE **must** be able to transfer traffic-affecting alarms and indications, and map them among TDM, ATM and MPLS networks. Support of alarms and indications is critical in order to have a working I.A.

The egress PE **should** be able to generate OOS code towards the TDM interface, based on an MPLS failure (such as MPLS tunnel failure or LOS on the MPLS port) or L-bit in the TDM-MPLS header. When the Circuit Bundle contains CAS, OOS code **should** also be generated. In addition an update of the R-bit in the TDM-MPLS header, or the sending of an MPLS OAM message, is needed.

The ingress PE **should** be able to stop transmission of TDM traffic towards the MPLS network upon reception of a corresponding lower layer defect (such as LOS or AIS from the TDM interface) and update the L-bit in the TDM-MPLS header or send an MPLS OAM message.

Two types of indication, L (Local defect) and R (Remote defect) **must** be supported:

L-Bit

When interworking with TDM networks, the L-bit indicates that the local trunk (e.g. E1/T1) indicates a fault condition. Although it may be the case that several trunks are multiplexed together, these must be related trunks that share a common timing source, and so may indeed simultaneously fail. When interworking with ATM, the L-bit indicates VC/VP-AIS/LOC in the ATM environment.

R-Bit

The R-bit indicates that the source is not receiving packets at its TDM-MPLS receiving port, indicating failure of that direction of the bi-directional connection.

This bit is the equivalent of VC/VP-RDI in the ATM environment.

6 TDM-MPLS Frame Processing

6.1 Generating TDM-MPLS Frames

The generation process of a TDM-MPLS frame consists of generating AAL1 cells and aggregating several cells in one packet. Thereafter the following TDM-MPLS frame header fields, including the L- and R-bits, are prepended.

The length field will be generated according to the rules specified in Table 1. The sequence number field is set according to the current sequence number of the circuit bundle.

The AAL1 cells are generated as per ITU-T I363.1 [4] and atm-vtoa-0078 (CES 2.0) [5], the number of cells per packet being set according to the configuration or signaling protocol.

Once the TDM-MPLS frame has been generated, lower layers (such as ETHERNET) will perform additional processing.

6.1.1 Setting the sequence number

The sequence number is defined separately for each CB and incremented by one for each TDM-MPLS frame sent for that CB. It **may** be used by the receiver to detect packet loss and to restore packet sequence. In addition, since the basic clock rate for each CB is constant, the sequence number **may** be used as an approximate timestamp for synchronization purposes.

The following procedure **must** be used by the ingress PE if sequencing is desired for a given TDM-MPLS service:

- The initial PDU transmitted on the LSP **must** use sequence number 1.
- The PE **must** increment the sequence number by one for each subsequent PDU.
- When the transmit sequence number reaches the maximum 16 bit value (65535), the sequence number **must** wrap to 1.
- If the ingress PE does not support sequence number processing, then the sequence number field in the control word **must** be set to 0.

6.1.2 Handling the R-bit

The following procedure **must** be used by the egress PE in a given TDM-MPLS service:

- The R-bit **must** be set after a preconfigured number of consecutive packets are not received on a specific CB, and **must** be cleared once packets are once again received.

6.1.3 Handling the L-bit

The following procedure **must** be used by the egress PE in a given TDM-MPLS service:

If the Ingress PE detects or has been informed of a TDM physical layer or ATM VC/VP fault impacting the data to be transmitted, such as physical layer LOS/LOF/AIS in the TDM

or F1/F2/F3 LOS/LOP/LOF/AIS or F4/F5 AIS/LOC in the ATM environment, the L-bit **must** be set to 1, otherwise the L-bit **must** be set to 0.

When the L-bit is set, the contents of the packet may not be meaningful, and the payload size **may** be reduced in order to conserve bandwidth. Once set, the L-bit **must** be cleared if the TDM fault is rectified.

6.2 Receiving TDM-MPLS Frames

When a PE receives a TDM-MPLS frame, it checks the layer 2 checksum; if the checksum is incorrect, the frame is discarded. Otherwise it processes the TDM-MPLS frame header fields and AAL1 cell contents, processing as per ITU-T I.363.1 [4] and atm-vtoa-0078 (CES 2.0) [5] in order to restore the TDM traffic for transmission through one of its TDM interfaces. The PE performs the following actions (not necessarily in the order shown):

- The PE processes the length and sequence number fields.
- It processes the AAL1 header fields per cell and generates the TDM traffic towards the TDM interface.
- **The CRC, P and E bits in the subframe are checked if the receiver was configured to check them.**

6.2.1 Processing the sequence number

If the egress PE supports receive sequence number processing, then the following procedure **must** be used:

When TDM-MPLS service is initially created, the “expected sequence number” associated with it **must** be initialized to 1. When a PDU is received on the LSP associated with the TDM-MPLS service, the sequence number must be processed as follows:

If the sequence number on the packet is 0, then the PDU passes the sequence number check

Otherwise, if the PDU sequence number \geq the expected sequence number and the PDU sequence number – the expected sequence number < 32768 , then the PDU is in order.

Otherwise, if the PDU sequence number $<$ the expected sequence number and the expected sequence number – the PDU sequence number ≥ 32768 , then the PDU is in order.

Otherwise the PDU is not in order.

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

$$\text{Expected sequence number} = \text{PDU sequence number} + 1 \text{ mod } 2^{**16}$$

if (expected sequence number = 0), then expected sequence number = 1.

LSP PDUs that are received out of order **may** be dropped or reordered at the discretion of the egress PE.

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

6.2.2 Handling the R-bit

The following procedure **Must** be used by the egress PE in a given TDM-MPLS service:

- If the R-bit is set to 1 on a specific TDM-MPLS PW, then the egress PE **must** generate an alarm/notification to its network management system.

6.2.3 Handling the L-bit

The following procedure **must** be used by the egress PE in a given TDM-MPLS service:

- If the L-bit is set to 1 on a specific TDM-MPLS PW, then the egress PE **must** transmit “out of service” code in the relevant timeslots or send AIS for the TDM environment and F4/F5 AIS in the ATM environment.

7 Clock Recovery

TDM networks are inherently synchronous; somewhere in the network there will always be at least one extremely accurate primary reference clock, with long-term accuracy of one part in $10E^{-11}$. This node, whose accuracy is called "stratum 1", provides reference timing to secondary nodes with lower "stratum 2" accuracy, and these in turn provide reference clock to "stratum 3" nodes. This hierarchy of time synchronization is essential for the proper functioning of the network as a whole; for details see ANSI/T1.101-1999 or ITU-T G.81x series recommendations. The use of time standards less accurate than stratum 3 is NOT RECOMMENDED, as it may result in service impairments.

Packets in MPLS networks reach their destination with delay that has a random component, known as jitter. When emulating TDM on an MPLS network, it is possible to overcome this randomness by using a "jitter buffer" on all incoming data, assuming the proper time reference is available. The problem is that the original TDM time reference information is not disseminated through the MPLS network.

In broadest terms, there are two methods of overcoming this difficulty; in one method, the timing information is provided by some means independent of the MPLS network, while in the other the timing must be transferred over the MPLS network.

For example, in the "whole network" scenario, the entire TDM infrastructure (or at least major portions of it) is replaced by TDM-MPLS. Here timing information **must** be delivered over the MPLS network, and the reconstructed TDM stream **should** conform to ITU-T recommendations [2] for E1 and [3] for T1 trunks. When the provision of accurate local time references is not practical, then clock recovery **should** be performed based on the rate of arrival of incoming packets, using an appropriate 'averaging' process that negates the effect of zero mean random jitter. Conventionally, a phase locked loop (PLL) is used for this purpose.

The regenerated clock **should** conform to ITU-T recommendations [2] for E1 and [3] for T1 trunks.

However, TDM-MPLS is frequently used in a "toll-bypass" scenario, where an MPLS link connects two existing TDM networks. In such a case, both TDM-MPLS devices **shall** receive accurate timing from the TDM networks to which they connect and **must** use this local timing.

8 TDM Services

TDM-MPLS is designed to transport TDM services over an MPLS network. The TDM service maps several TDM circuits to an MPLS LSP. As such, it enables 0/1/3 cross-connect over the MPLS network. Table 2 describes supported TDM services.

Specific implementations may support all or any subset of these services.

Service	Standard	Rate
$N \times 64K$	Serial interface	$N \times 64$ kbps
Unframed T1	G.703	1.544 Mbps
Unframed E1	G.703	2.048 Mbps
Unchannelized/ Unframed E3	G.703	34.368 Mbps
Unchannelized/ Unframed T3	G.703	44.736 Mbps
Channelized/ Fractional E1	G.703, G.704	$N \times 64$ kbps $1 \leq N \leq 31$
Channelized/ Fractional T1	G.703, G.704	$N \times 64$ kbps, $1 \leq N \leq 24$
Channelized/ Fractional T1/E1 with CAS	G.703, G.704	
AAL1	ITU-T I.363.1	All of the above

Table 8-2 – TDM-MPLS Services

Both data and clock information must be transferred edge to edge.

When present, CAS is transparently transferred edge to edge.

Trunk-associated CCS is transferred transparently edge to edge.

Standard TDM alarms are generated when required and transferred edge to edge

Note: Service interworking between different TDM service types is outside the scope of this I.A.

9 TDM ATM Service Interworking

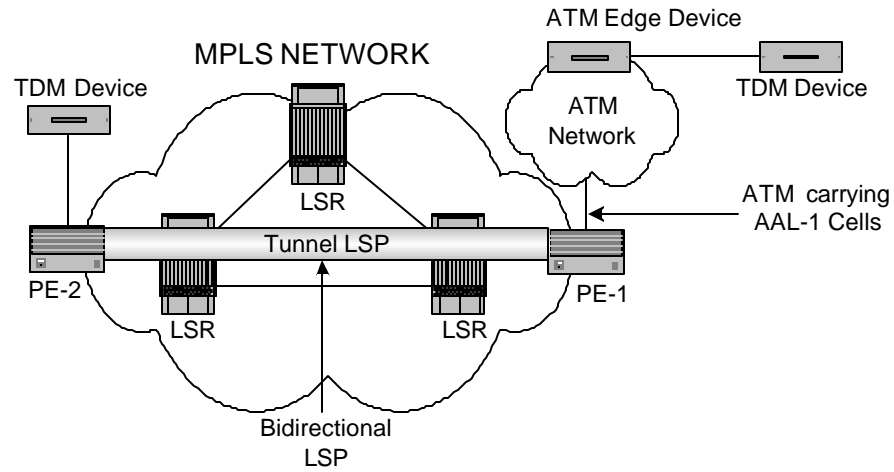


Figure 9-1 – ATM-AAL1 Service Interworking

Figure 9-1 shows a case where the TDM data is transmitted across an ATM and MPLS network. It is carried over the ATM network using AAL1 circuit emulation. PE-1 extracts the AAL1 cells from the ATM stream, and maps them (according to VCI/VPI) to CBID, appends a TDM over MPLS header, and builds a TDM-MPLS packet. At the remote side, PE-2 removes the headers, reassembles a bit stream from the AAL1 cells and restores the TDM traffic. This model supports bidirectional service interworking.

10 ATM PW Compatibility Mode

Since TDM traffic can be carried over ATM circuit emulation services using AAL1, the protocols described in [8], [9] **may** be used to indirectly transport TDM over pseudo-wires. In such a case the TDM is first converted into an AAL1 ATM flow according to [4] and [5], and thereafter this ATM flow is encapsulated as described in [8], [9].

The TDM-MPLS control word is compatible with the control word of [8], [9] for the mandatory N:1 mode.

The N:1 mode concatenates ATM cells, including their cell headers, with the exception of the HEC. Hence, a valid and locally unique VPI/VCI must be allocated to the TDM bundle before this mode can be utilized.

The format of the control word and payload are as follows:

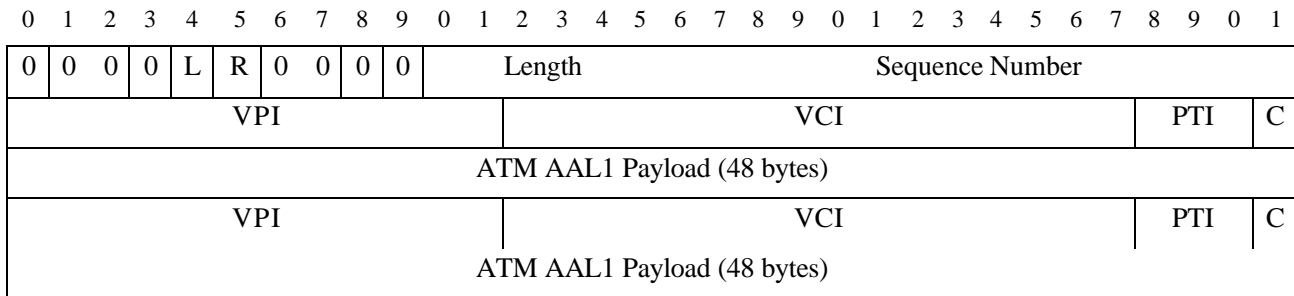


Figure 10-1 – Multiple Cell ATM Encapsulation (N:1 mode)

L- and R-bit s are described in section 3.2 above

VPI (12 bits), VCI (16 bits) are the ATM labels taken from the ATM cell header.

PTI (3 bits) is the ATM payload type identifier copied from the ATM header and indicates congestion as well as differentiating between cells containing user data and those used for maintenance.

C (1 bit) is the cell loss priority field copied from the ATM header, with C=1 indicating lower priority.

When using N:1 mode with n greater than one, MPLS OAM signaling **MUST** be employed to signal local and remote defects.

While ATM PW compatibility mode enables utilization of network devices designed according to [8], [9] and facilitates service interworking with existing ATM circuit emulation systems, it has higher overhead (an additional 4 bytes per 48 byte cell) and its use impedes exploitation of some features of the intrinsic AAL1 mode. For example, due to the separation of the TDM processing from the edge devices, access to timing-related information may be lost, resulting in jitter and wander attenuation inferior to that obtainable via the intrinsic AAL1 mode. Packet interpolation (see section 5.2) may suffer as compared with other TDM-MPLS modes.

TDM-MPLS mode can be configured by manual provisioning or by using some signaling protocol (outside the scope of this I.A.)

End of Document