

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.214

(07/2007)

SERIES J: CABLE NETWORKS AND TRANSMISSION
OF TELEVISION, SOUND PROGRAMME AND OTHER
MULTIMEDIA SIGNALS

Interactive systems for digital television distribution

Cable modem TDM emulation interface

ITU-T Recommendation J.214

Document ID: 123456789

ITU-T Recommendation J.214

Cable modem TDM emulation interface

Summary

TDM emulation service (TDM-E) is a method for cable operators to deliver T1, E1 and NxDS0 emulation services that meet or exceed the quality requirement of applications that use such services. ITU-T Recommendation J.214 is part of the CableModem family of Recommendations and, in particular, defines the TDM-E architecture and components.

Source

ITU-T Recommendation J.214 was approved on 29 July 2007 by ITU-T Study Group 9 (2005-2008) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2008

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Page

1	Scope and purpose	1
1.1	Scope	1
1.2	Introduction and overview	1
2	References.....	2
2.1	Normative references.....	2
2.2	Informative references.....	4
2.3	Reference acquisition	4
3	Terms and definitions	5
3.1	Terms defined elsewhere	5
3.2	Terms defined in this Recommendation.....	5
4	Abbreviations, acronyms and conventions	6
4.1	Abbreviations and acronyms	6
4.2	Conventions.....	7
5	Technical overview.....	8
5.1	Background.....	8
5.2	Principal technical considerations	8
5.3	Devices/Interfaces	9
5.4	Architecture	13
5.5	Network environment.....	16
6	Technical specification	20
6.1	General requirements.....	20
6.2	TDM-IP network interworking.....	23
6.3	Payload formats	26
6.4	Timing aspects.....	30
6.5	Defects and monitoring	43
6.6	Security.....	44
6.7	Provisioning and service initiation	44
6.8	Management	64
Annex A	– Event, SYSLOG, and SNMP trap extensions	70
A.1	TEA events description	70
A.2	DOCSIS events extensions.....	70
Annex B	– TEA MIB definition	73
B.1	CL-PW-TC-MIB	73
B.2	CL-PW-MIB.....	77
B.3	CL-PW-TDM-MIB.....	109

Copyright International Telecommunication Union/ITU-T

	Page
Appendix I – Operator recommendations	143
I.1 Preferential treatment for TE-CMs	143
I.2 Service level agreement considerations	143
I.3 Standard DOCSIS operations	144
Appendix II – Delay aspects	146
II.1 Upstream delay	146
II.2 Downstream delay	147
Appendix III – Bandwidth usage	149
III.1 Upstream bandwidth usage.....	149
III.2 Downstream bandwidth usage.....	150
Appendix IV – eTEA configuration example	151
IV.1 Structure agnostic T1 configuration	151
IV.2 Structure agnostic E1 configuration	152
IV.3 Structure aware 5xDS0 configuration	153
IV.4 Structure aware 5xDS0 configuration with CAS	154
IV.5 Structure aware T1 with RTP configuration	155
Appendix V – Clocking options, by example	157
V.1 Background.....	157
V.2 Clocking modes	158
V.3 Clocking decision tree	161
Bibliography.....	162

ITU-T Recommendation J.214

Cable modem TDM emulation interface

1 Scope and purpose

1.1 Scope

TDM Emulation service (TDM-E) is a method for cable operators to deliver T1, E1 and NxDS0 emulation services that meet or exceed the quality requirement of applications that use such services. This Recommendation is part of the CableModem family of Recommendations and, in particular, defines the TDM-E architecture and components.

NOTE – The structure and content of this Recommendation have been organized for ease of use by those familiar with the original source material; as such, the usual style of ITU-T Recommendations has not been applied.

1.2 Introduction and overview

In legacy telecommunication networks, telephone calls are often brought into households, one at a time, over twisted pair wires. To transport many telephone calls at once (i.e., between business, wireless base stations and in the telephone network), single calls are time-multiplexed together into 'T1' signals. A single T1 signal carries 24 individual calls, and a similar 'E1' signal carries 32 calls. Since T1 and E1 (T1/E1) services have been deployed for quite some time, the performance standards, tariffs and market are well defined. A number of ITU, ETSI and ANSI standards define the various aspects of T1/E1 services. Moreover, the usage and deployment models of T1/E1 lines are well understood.

This Recommendation outlines the methods by which T1/E1 structured, unstructured, and fractional signals can be converted to IP packets, transported over a DOCSIS IP network, and converted back to T1/E1 signals with high reliability and quality. Figure 1-1 shows a simple T1 service delivered between two customer sites over DOCSIS.




Figure 1-1 – Example of T1 service delivery

2 References

2.1 Normative references

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revisions; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- [ITU-T G.703] ITU-T Recommendation G.703 (2001), *Physical/electrical characteristics of hierarchical digital interfaces*.
- [ITU-T G.704] ITU-T Recommendation G.704 (1998), *Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 kbit/s hierarchical levels*.
- [ITU-T G.706] ITU-T Recommendation G.706 (1991), *Frame alignment and cyclic redundancy check (CRC) procedures relating to basic frame structures defined in Recommendation G.704*.
- [ITU-T G.732] ITU-T Recommendation G.732 (1988), *Characteristics of primary PCM multiplex equipment operating at 2048 kbit/s*.
- [ITU-T G.810] ITU-T Recommendation G.810 (1996), *Definitions and terminology for synchronization networks*.
- [ITU-T G.823] ITU-T Recommendation G.823 (2000), *The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy*.
- [ITU-T G.824] ITU-T Recommendation G.824 (2000), *The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy*.
- [ITU-T G.1020] ITU-T Recommendation G.1020 (2006), *Performance parameter definitions for quality of speech and other voiceband applications utilizing IP networks*.
- [ITU-T G.1050] ITU-T Recommendation G.1050 (2007), *Network model for evaluating multimedia transmission performance over Internet protocol*.
- [ITU-T J.83] ITU-T Recommendation J.83 (2007), *Digital multi-program systems for television, sound and data services for cable distribution*.
- [ITU-T J.112] ITU-T Recommendation J.112 (1998), *Transmission systems for interactive cable television services*.
- [ITU-T J.122] ITU-T Recommendation J.122 (2007), *Second-generation transmission systems for interactive cable television services – IP cable modems*.
- [ITU-T J.126] ITU-T Recommendation J.126 (2007), *Embedded cable modem device specification*.
- [ITU-T J.210] ITU-T Recommendation J.210 (2006), *Downstream RF interface for cable modem termination systems*.

[ITU-T J.211]	ITU-T Recommendation J.211 (2006), <i>Timing interface for cable modem termination systems.</i>
[ITU-T J.212]	ITU-T Recommendation J.212 (2006), <i>Downstream external physical layer interface for modular cable modem termination systems.</i>
ITU-T J.222.1]	ITU-T Recommendation J.222.1 (2007), <i>Third-generation transmission systems for interactive cable television services – IP cable modems: Physical layer specification.</i>
[ITU-T J.222.2]	ITU-T Recommendation J.222.2 (2007), <i>Third-generation transmission systems for interactive cable television services – IP cable modems: MAC and upper layer protocols .</i>
[ITU-T X.690]	ITU-T Recommendation X-690 (2002) ISO/IEC 8825-1:2002, <i>Information technology – ASN.1 encoding rules: Specification of Basic Encoding rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER).</i>
[DOCSIS OSSiv2.0]	ANSI/SCTE 79-2 (2002), <i>DOCS 2.0 Operations Support System Interface.</i>
[IETF RFC 1123]	IETF RFC 1123 (1989), <i>Requirements for Internet Hosts – Application and Support.</i>
[IETF RFC 1157]	IETF RFC 1157 (1990), <i>A Simple Network Management Protocol (SNMP).</i>
[IETF RFC 1350]	IETF RFC 1350 (1992), <i>The TFTP Protocol (Revision 2).</i>
[IETF RFC 2131]	IETF RFC 2131 (1997), <i>Dynamic Host Configuration Protocol.</i>
[IETF RFC 2132]	IETF RFC 2132 (1997), <i>DHCP Options and BOOTP Vendor Extensions.</i>
[IETF RFC 2494]	IETF RFC 2494 (1999), <i>Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type.</i>
[IETF RFC 2495]	IETF RFC 2495 (1999), <i>Definitions of Managed Objects for the DS1, E1, DS2 and E2 Interface Types.</i>
[IETF RFC 2670]	IETF RFC 2670 (1999), <i>Radio Frequency (RF) Interface Management Information Base for MCNS/DOCSIS compliant RF interfaces.</i>
[IETF RFC 2863]	IETF RFC 2863 (2000), <i>The Interfaces Group MIB.</i>
[IETF RFC 3396]	IETF RFC 3396 (2002), <i>Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4).</i>
[IETF RFC 3550]	IETF RFC 3550 (2003), <i>RTP: A Transport Protocol for Real-Time Applications.</i>
[IETF RFC 3611]	IETF RFC 3611 (2003), <i>RTP Control Protocol Extended Reports (RTCP XR).</i>
[IETF RFC 3635]	IETF RFC 3635 (2003), <i>Definitions of Managed Objects for the Ethernet-like Interface Types.</i>
[IETF RFC 3636]	IETF RFC 3636 (2003), <i>Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs).</i>
[IETF RFC 3711]	IETF RFC 3711 (2004), <i>The Secure Real-time Transport Protocol (SRTP).</i>

- [IETF RFC 3895] IETF RFC 3895 (2004), *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*.
- [IEEE 1588] IEEE 1588 (2002), *Standard for a Precision Clock Synchronization Protocol for networked measurement and Control Systems*. Also *IEC 61588 (2004)*.

2.2 Informative references

This Recommendation uses the following informative references.

- [ANSI T1.403] ANSI T1.403-1999, *Network and Customer Installation Interfaces (ISDN) primary rate – Layer 1– Electrical Interfaces Specification*.
- [CMCI] *Data-Over-Cable Service Interface Specifications Cable Modem to Customer Premise Equipment Interface Specification*, SP-CMCI-I10-050408, April 8, 2005, Cable Television Laboratories, Inc.
- [IANA UDP PORTS] Internet Assigned Numbers Authority, Port Numbers, April 2006, <http://www.iana.org/assignments/port-numbers>
- [ID PW MIB] T. Nadeau, D. Zelig, *Pseudowire (PW) Management Information Base (MIB)*, draft-ietf-pwe3-pw-mib-14, July 2008.
- [ID PWTC] T. Nadeau, D. Zelig, *Definitions of Textual Conventions for Pseudowires (PW) Management*, draft-ietf-pwe3-pw-tc-mib-14, July 2008.
- [ID VCCV] T. Nadeau, C. Pignataro, *Pseudowire Virtual Circuit Connectivity Verification (VCCV), A Control Channel for Pseudowires*, draft-ietf-pwe3-vccv-15, March 2008.
- [IETF RFC 3551] IETF RFC 3551 (2003), *RTP Profile for Audio and Video Conferences with Minimal Control*.
- [IETF RFC 4553] IETF RFC 4553 (2006), *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*.
- [IETF RFC 4928] IETF RFC 4928 (2007), *Avoiding Equal Cost Multipath Treatment in MPLS Networks*.
- [MLDA] Sisalem D. and A. Wolisz, *MLDA: A TCP-friendly Congestion Control Framework for Heterogeneous Multicast Environments*, *IEEE Proc. IWQoS 2000*.
- [TIA NIM] TR41.4-03-05-029 (2003), *IP Network Impairment Simulation for VoIP Performance Testing*.
- [TR-62411] AT&T Technical Reference 62411, *AT&T ACCUNET T1.5 Service Description and Interface Specification*, December 1990.

2.3 Reference acquisition

- American National Standards Institute, Internet: <http://webstore.ansi.org>
- AT&T, <http://www.att.com>
- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone 303-661-9100; Fax 303-661-9199; Internet: <http://www.cablelabs.com>
- Institute of Electrical and Electronics Engineers (IEEE), <http://www.ieee.org>
- Internet Engineering Task Force (IETF), <http://www.ietf.org>

- International Telecommunication Union – Telecommunication Standardization Sector (ITU-T), <http://www.itu.int/itu-t/>
- International Organization for Standardization (ISO), <http://www.iso.org>
- Metro Ethernet Forum (MEF), <http://www.metroethernetforum.org>
- Society of Cable Telecommunications Engineers (SCTE), <http://www.scte.org>
- Telecommunications Industry Association (TIA), <http://www.tiaonline.org/>
- Telcordia Technologies, <http://www.telcordia.com/>

3 Terms and definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 jitter: The fluctuation in the arrival time of a regularly scheduled event such as a clock edge or a packet in a stream of packets. Jitter is defined in [ITU-T G.810] as fluctuations above 10 Hz. See wander.

3.1.2 maximum time interval error (MTIE): [ITU-T G.810].

3.1.3 time deviation (TDEV): [ITU-T G.810].

3.1.4 wander: Low speed jitter. Wander is defined in [ITU-T G.810] as fluctuations below 10 Hz. See jitter.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 alarm indication signal: Also known as the *blue alarm*. When no incoming signal is detected, a CSU/DSU transmits an unframed all-ones pattern to the network to maintain synchronization and announce its presence to the network.

3.2.2 clock unit (CU): A clock unit performs translations and distribution of TDM clocking information across (and between) physical, data and network layers.

3.2.3 clock recovery unit (CRU): A clock recovery unit exists in the IWF and is responsible for regenerating the circuits clock based on the average inter-arrival time of the packets in the adaptive clocking mode, or on the time stamp differentials received when operating in the differential clock mode. The output clock is provided to the CU.

3.2.4 channel service unit: The piece of a CSU/DSU that talks to the telco network, understands framing and line coding, and provides electrical isolation of the network from the telco network.

3.2.5 data service unit: The part of the CSU/DSU that interfaces with routers, switches and packets. It has a serial port to interface with compatible data equipment.

3.2.6 digital signal 0 (DS0): Digital signal 0 (DS0) is a basic digital signalling rate of 64 kbit/s, corresponding to the capacity of one voice-frequency-equivalent channel.

3.2.7 E1: E1 is a physical layer telephony protocol carrying data at 2048 Mbit/s. It can carry up to 32 DS0s, each of which can carry a telephone conversation.

3.2.8 eTEA: An eTEA is an eDOCSIS eSAFE; an embedded version of a TEA.

3.2.9 fractional: A fractional T1 or fractional E1 carries only a portion of the total number of DS0s that a T1 or E1 carries.

3.2.10 interworking function (IWF): An interworking function (IWF) is a logical entity. It consists of a TSP data interface on one side, an IP packet interface on the other side, and the functionalities to encode TDM data into a pseudo wire in one direction and decode TDM data from a pseudo wire in the other direction.

3.2.11 logical CPE interface (LCI): A logical CPE interface (LCI) is a logical 802.3/Ethernet MAC frame data interface.

3.2.12 loss of frame: Also called the *red alarm*. When a T1 CSU/DSU is unable to synchronize framing patterns with the remote end for 2.5 seconds, LoF is declared.

3.2.13 loss of signal: When no incoming pulses are received by a T1 CSU/DSU for a prescribed number of bit times, LoS is declared. Even if only zeros were transmitted as data, some framing bits should result in pulses on the line during that time. If LoS persists, LoF will eventually be declared because there is no incoming signal with which to synchronize.

3.2.14 out of frame: When frame synchronization is lost, an OoF event is recorded. If OoF persists, LoF is declared. OoF is cleared when frame synchronization is regained.

3.2.15 pseudo wire: A pseudo wire (PW) is a logical entity capable of the emulation of a native service (i.e., T1) over a packet switched network (like DOCSIS IP).

3.2.16 remote alarm indication: Also known as *yellow alarm*. When a CSU/DSU enters the red alarm state, an RAI is transmitted in the outgoing direction. The RAI signals to the remote end that the local end is unable to synchronize framing patterns.

3.2.17 T1: T1 is a physical layer telephony protocol carrying data at 1544 Mbit/s. It can carry up to 24 DS0s, each of which can carry a telephone conversation. It is used mainly in North America.

3.2.18 TDM interface: A physical time domain multiplex (TDM) telephony interface such as T1 or E1. Time division multiplexing (TDM) is the means by which multiple digital signals (DS0s) can be carried on a single transmission path by interleaving portions of each signal in time.

3.2.19 TDM emulation adapter (TEA): A logical entity containing various functions to provide a TDM-E circuit emulation service.

3.2.20 TDM emulation cable modem (TE-CM): A special purpose cable modem that integrates a DOCSIS eCM with an eTEA.

3.2.21 TE-CMTS: A TE-CMTS is a CMTS or M-CMTS that can support a TE-CM and can meet the minimum clocking and performance goals to support TDM-E traffic. A TE-CMTS may or may not support integrated TEAs.

3.2.22 TDM service processor (TSP): Consists of a TSP data interface on one side, a TDM interface on the other side, and the functionalities to encapsulate TSP data into TDM interface signals in one direction and extract TSP data from TDM interface signals in the other direction.

4 Abbreviations, acronyms and conventions

4.1 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ACR	Adaptive Clock Recovery
AIS	Alarm Indication Signal
A-TDMA	Advanced Time Division Multiple Access
CBR	Constant Bit Rate
CM	Cable Modem

Copyright International Telecommunication Union/ITU Telecommunication Sector

CMTS	Cable Modem Termination System
CSU	Channel Service Unit
DEPI	DOCSIS External PHY Interface
DOCSIS®	Data-Over-Cable Service Interface Specification
DS	Downstream
DSU	Data Service Unit
DTI	DOCSIS Timing Interface
HFC	Hybrid Fibre Coax
IP	Internet Protocol
LoF	Loss of Frame
LoS	Loss of Signal
MAC	Media Access Control
OoF	Out of Frame
PDCR	Prime Differential Clock Recovery
PW	Pseudo Wire
RAI	Remote Alarm Indication
RF	Radio Frequency
UDP	User Datagram Protocol
UGS	Unsolicited Grant Service

4.2 Conventions

Throughout this Recommendation, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this Recommendation.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this Recommendation.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behaviour is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behaviour described with this label.
"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, i.e., another vendor may omit the same item.

5 Technical overview

5.1 Background

Throughout the evolution of communication networks, the T1 and E1 service that is delivered to the customer has remained relatively consistent (in large part due to the strict ITU and ANSI standards); however, the method of delivering the T1 or E1 service has changed to increase efficiencies and leverage new technologies. Both SONET/SDH and ATM historically have been adapted to transport T1 and E1 services; now MPLS and IP networks are also being used. Various standards bodies and vendors have created methods for circuit emulation of T1 and E1 services across IP networks, leveraging the successful ATM circuit emulation technology developed in the 1990s. This Recommendation leverages the work done in the ITU and IETF to define circuit emulation of T1 and E1 over a DOCSIS network.

This Recommendation describes the architecture and components of a network that delivers emulated T1 or E1 or NxDS0 services over DOCSIS. It defines the T1, E1, or NxDS0 service delivered to the end customer, the requirements of a TDM-E compliant cable modem (TE-CM), the requirements of a TDM-E compliant CMTS (TE-CMTS), the requirements of the TDM emulation adapter (TEA), and pseudo wires (PW) that cross the packet network connecting two TEAs.

In developing this Recommendation, the following assumptions were made concerning the implementation and deployment of TDM-E systems:

- Interoperation with DOCSIS CMTS versions and options:
Existing CMTSs can lock their existing DOCSIS symbol clock or DOCSIS SYNC message generation to master clock inputs for both A-TDMA and S-CDMA. The chain of clock synchronization is that Stratum 1 Traceable Reference drives the DOCSIS master clock; the master clock drives the DOCSIS downstream timing; and the CM recovers timing from the DOCSIS downstream.
Some CM silicon originally developed for DOCSIS 1.1 or 2.0 residential data modems can be used in a design where the symbol clock or DOCSIS SYNC message may be recovered to generate a T1/E1 clock.
- TE-CMs are always connected to a TE-CMTS with QoS support for constant bit rate flows. The TE-CMTS complies with DOCSIS 1.1 or 2.0. A DOCSIS 1.0 CMTS will not be able to host TE-CMs. DOCSIS 1.1 may be limited in upstream bandwidth; TE-CM data flows are symmetric upstream and downstream.
- It is not in the scope of this Recommendation to cover interoperation between dissimilar services, i.e., between E1 and T1.
- The following requirements are market specific and are out of scope of this Recommendation:
 - Path redundancy and equipment redundancy (like redundant power supplies);
 - Environmental hardening, such as clock stability over temperature.
- There is not sufficient need to support the 'octet-aligned format' of [IETF RFC 4553] section 5.2, which is designed to simplify the data handling in and out of a Sonet/SDH virtual tributary/container. This Recommendation does not support that format.

5.2 Principal technical considerations

This clause describes the principle technologies and operational considerations germane to the development of a T1/E1-bearing TE-CM and its use. Major sub-topics include:

- Clocking – T1/E1 transport technology is synchronous in nature; requiring accurate and stable clocking. Since T1 transport requires that T1 signals are handed from one service provider to another, each service provider must adhere to the same clocking requirements.

Further, these synchronous clocking requirements must be very stringent to ensure reliable operation. The clause on clocking (clause 5.5) will further describe these requirements largely based on existing ITU and ATIS requirements.

- Data protocol
 - 1) Voice protocol – Phone calls at termination are completely analog, carrying linear voice signals over dedicated wires. It was not long before the phone companies realized that a wire could carry more than one phone call if the calls were time-multiplexed on the wire. Calls are now time-sampled, and the resulting digital signals are time-multiplexed over wires in T1/E1 formats (among others).
 - 2) Data protocol – It did not take long for engineers to realize that T1/E1 transport could be used for standard data (not just voice). T1 signals are now used to transport digital data in many different data formats which are beyond the scope of this Recommendation. This Recommendation accommodates the transport of either voice or data by specifying more than one data format for the IP packets that carry the T1 information over DOCSIS channels. The clauses below addressing 'structure agnostic' protocol are germane to data transport. Most of this Recommendation addresses T1 transport of voice calls, but the Recommendation accommodates data transport as well.
 - 3) Signalling data protocol – To successfully combine several voice phone calls into a T1 structure, data describing the state of each call must be included. Such call state data may include the progress of the call and the health of the transport network carrying the call. Several existing standards describe this signalling data. This Recommendation describes the format for signalling data in the IP packets that carry the T1 information over DOCSIS channels.
- Operations – The synchronous transport systems put together by telecommunications carriers for T1 data have been refined over the years to become very reliable. But the manner by which a DOCSIS transport system performs this service differs considerably from the manner in which they are performed in the PSTN by telecommunications carriers. Operators of DOCSIS systems can alter their standard operating procedures to accommodate these differences. This Recommendation makes many suggestions for operators who must configure DOCSIS systems to carry emulated T1/E1 data as specified herein.


5.3 Devices/Interfaces

5.3.1 TE-CM

A TE-CM is a standard DOCSIS cable modem that contains an embedded TDM emulation adapter (eTEA) giving it the ability to offer T1/E1 services. A TE-CM may have architecture and interfaces similar to that shown in Figure 5-1. The power supply and other functionality that is out of scope are not shown. The elements/interfaces are described as follows:

- DOCSIS eCM – This is a standard DOCSIS CM conforming to DOCSIS 1.1, 2.0, 3.0 or greater. It must be able to carry a constant bit rate service flow carrying packets at bit rates sufficient to handle T1/E1 data, management and signalling as defined herein. It also contains [ITU-T J.126] functions sufficient to support eTEA and Ethernet functionality. Per DOCSIS, it has both RF and Ethernet interfaces. The embedded DOCSIS CM chipset provides clocking to the eTEA recovered from the DOCSIS symbol clock or DOCSIS SYNC message as defined in clause 6.4.
- eTEA – The TE-CM contains an eTEA that operates as an eDOCSIS eSAFE.
- T1/E1 Interface – This interface is a standard T1/E1 connector.

- Ethernet interface – This interface is a standard Ethernet connector (RJ45). It may, in future, provide the capability to exchange IP packets with eTEA or other TDM-E entities that may exist outside the TE-CM.
- Logical CPE interface – This is an unspecified interface between the eTEA and eCM.



NOTE – [ITU-T J.126] applies to the eCM. The DOCSIS Symbol/Sync clock is Stratum 1 Traceable.

Figure 5-1 – TE-CM

5.3.2 TE-CMTS

A TE-CMTS is a standard DOCSIS CMTS that contains the necessary clocking functionality necessary to support a TE-CM and may also include a TDM emulation adapter (TEA) giving it the ability to terminate T1/E1 services. A TE-CMTS may have architecture and interfaces similar to that shown in Figure 5-2. The power supply and other functionality that is out of scope are not shown. The elements/interfaces are described as follows:

- RF – This is the standard DOCSIS CMTS signal as transceived on the HFC network.
- DOCSIS CMTS – This is a standard DOCSIS CMTS conforming to DOCSIS 1.1, 2.0, 3.0 or greater. It must be able to carry a constant bit rate service flow carrying packets at bit rates sufficient to handle the IP pseudo wire traffic and management as defined herein. Per DOCSIS, it has both RF and IP interfaces. The CMTS may be an M-CMTS in which case the M-CMTS Recommendation defines compliant IP, RF and synchronization interfaces.
- Clocking unit – When the TEA is not integrated within the TE-CMTS, at a minimum the clocking unit functionality must be implemented as part of the TE-CMTS. A TE-CMTS has either a [ITU-T J.211] or a BITS (building integrated timing supply) external synchronization interface as defined in clause 6.4. This is used to time and/or clock the DOCSIS symbol clock or DOCSIS SYNC message for a TE-CM to recover.
- TEA – A TE-CMTS may contain a TEA and TDM interface(s).
- Ethernet interface – The Ethernet interface will carry IP traffic to and from the managed IP network.




Figure 5-2 – TE-CMTS

L2E is an unspecified Layer 2 Ethernet interface.

5.3.3 Embedded TDM emulation adapter (eTEA)

An Embedded TDM emulation adapter (eTEA) is an eSAFE device as specified by [ITU-T J.126]. An eTEA is embedded in a TE-CM.

An eTEA has a T1 or E1 interface connector and contains the following elements/interfaces:

- L2 Ethernet – This is the Layer 2 Ethernet element where IP Pseudo wires are transceived.
- Pseudo Wires (PW1, PW2, ..., PWn) – Pseudo wires are logical IP packet flows that contain encapsulated TDM traffic. Pseudo wires are terminated and generated by IWF.
- Interworking function (IWF) – The interworking function consists of a TSP data interface on one side, an IP Packet interface on the other side, and the functionalities to encapsulate TDM data into a pseudo wire in one direction and extract TDM data from pseudo wire in the other direction. Encapsulation and transmission may involve clocking from the CU.
- TDM service flows (TS1, TS2, ..., TSn) – TDM data associated with a pseudo wire and TDM interface. The TDM service flows may be multiplexed onto one or more TDM interfaces through the TDM service processor.
- TDM service processor – The TDM service processor connects and may multiply TDM service flows onto one or more TDM interfaces.
- Ethernet interface. This Customer facing Ethernet interface has the following uses:
 - to launch the optional 1588 local master timing function on the customer Ethernet network per clause 6.4.5;
 - for DOCSIS testing of an eCM;
 - for testing of clock recovery per clause 6.4.7.1;
 - for (out of scope) use to carry TDM data per clause 5.3.1.
- TDM interface – The TDM interface is one or more T1, E1 or higher rate interface as further defined in clause 3.
- Clocking unit (CU) – This element provides high quality and reliable frequency to the IWF and TDM service processor. The CU receives clocking from the eDOCSIS module. The eDOCSIS module through DOCSIS requirements must lock to the DOCSIS symbol clock

and DOCSIS SYNC message. The eDOCSIS module must provide one of these recovered clocking signals to the CU. Since a TE-CMTS is synchronized to a Stratum 1 traceable reference, the recovered clock from DOCSIS not only has the stability needed for DOCSIS operation, it also has stability needed for T1/E1 synchronization. Clause 6.4 describes several clocking modes that may include the CU interfacing with the TDM service processor and IWF.




Figure 5-3 – eTEA

5.3.4 TDM emulation adapter (TEA)

A TDM emulation adapter (TEA) may be implemented in a TE-CMTS or other device unspecified in the network.

A TEA has a T1 or E1 TDM interface connector and contains the following elements/interfaces:

- L2 Ethernet – This is the Layer 2 Ethernet element where IP pseudo wires are transceived.
- Pseudo wires (PW1, PW2, ..., PWN) – Pseudo wires are logical IP packet flows that contain encapsulated TDM traffic. Pseudo wires are terminated and generated by IWF.
- Interworking function (IWF) – The Interworking function consists of a TSP data interface on one side, an IP packet interface on the other side, and the functionalities to encapsulate TDM data into a pseudo wire in one direction and extract TDM data from pseudo wire in the other direction. Encapsulation and transmission may involve clocking from the CU.
- TDM streams (TS1, TS2, ..., TSn) – TDM data associated with a pseudo wire and TDM interface. The TDM streams may be multiplexed onto one or more TDM interfaces through the TDM service processor.
- TDM service processor – The TDM service processor connects and may multiplex TDM streams onto one or more TDM interfaces.
- TDM interface – The TDM interface is one or more T1, E1 (or higher rate) TDM interface.
- Clocking unit (CU) – This element provides high quality and reliable frequency to the IWF and TDM service processor. The CU receives clocking from a Stratum 1 traceable reference through either a DTI or BITS interface, and a recovered clock from the clock recovery unit within the IWF. Clause 6.4 describes several clocking modes that may include the CU interfacing with the TDM service processor and IWF.




Figure 5-4 – TEA

5.4 Architecture

A TDM-E service can be deployed in several architectures. In any architecture, the TDM-E service has two end points that each contain a TEA or eTEA (TEA peers). Each device presents a TDM interface to the customer or network and a pseudo wire connecting the TEA peers logically across the IP network.

Three general architectures exist based on the existing market for T1 and E1 services. In all three architectures, one end of the TDM-E service originates in a TE-CM containing an eTEA. The three architectures differ based on where the other end of the service is terminated and thus where the other TEA/eTEA exists.

There are a variety of existing T1/E1 services including wireless backhaul, private line, PBX extension, PBX PSTN access, etc. Comprehensive T1/E1 service may include a combination of these three architectures and may include partnering with other operators for part of the implementation.

5.4.1 TE-CM to TE-CMTS w/TEA

In this architecture, the TDM-E service originates in a TE-CM (containing an eTEA) and terminates at the TE-CMTS which contains a TEA. Figure 5-5 shows this architecture with the pseudo wire flow shown as the dotted line. This is not the termination of the T1 or E1 circuit, just the two ends of the TDM-E circuit emulation. The T1 or E1 service will continue from the CMTS into the traditional TDM network.

This architecture places delay and QoS requirements on the DOCSIS network, but minimizes the impairments that the IP network may place on the TDM-E service. This architecture, however, requires TDM transport facilities at the hub/headend locations where the CMTS is located which may not be feasible or cost effective for some operators that prefer or require the TDM-E service to have a more centralized termination.




Figure 5-5 – T1 to hub architecture

5.4.2 TE-CM to centralized TEA

In this architecture, the TDM-E service originates in a TE-CM and terminates in a centralized router or other unspecified device that contains a TEA. The TEA may exist in a router, gateway, switch or other device which is out-of-scope for this Recommendation with the exception of the TEA functionality. Figure 5-6 shows this architecture with the pseudo wire flow shown as the dotted line. This is not the termination of the T1 or E1 circuit, just the two ends of the TDM-E circuit emulation. The T1 or E1 service will continue from the CMTS to its termination within traditional TDM network.

This architecture places delay, synchronization and QoS requirements on the TE-CMTS, which in this case does not contain a TEA. This architecture places delay and QoS requirements on the DOCSIS and core managed IP network. While this architecture requires a slightly more complex provisioning of the network, it centralizes the termination of the TDM-E circuit emulation in the core. This means there are fewer TDM-E devices to manage and that the TE-CMTS does not have an integrated TEA. This simplifies the TE-CMTS and reduces or eliminates modifications to existing installed CMTSs. This architecture eliminates the need for TDM transport facilities at each hub/headend since the pseudo wires will pass through the TE-CMTS (as IP packets) and be routed over the existing IP transport to the core.




Figure 5-6 – T1 to TDM core, IXC or PSTN

5.4.3 TE-CM to TE-CM through TE-CMTSs and core

In this architecture, the TDM-E service originates in a TE-CM (containing an eTEA) and in a second TE-CM (containing an eTEA). Figure 5-7 shows this architecture with the pseudo wire flow shown as the dotted line. In this architecture, the T1 service is end-to-end with transport of the circuit through the managed IP core network. If the T1 circuit requested by the customer is accessible by a single hub/headend, a single TE-CMTS may be utilized with the pseudo wire being routed locally.

This architecture places delay, synchronization and QoS requirements on the TE-CMTS, which in this case does not have a TEA. This architecture places delay and QoS requirements on the DOCSIS and core managed IP network. While this architecture requires a more complex

provisioning of the network, it can be deployed with minimal upgrades to the existing installed CMTSs and does not require any modifications to the core equipment. This architecture eliminates the need for TDM transport facilities at each hub/headend since the pseudo wires will pass through the TE-CMTSs and existing IP transport in the core between the two TE-CMs. There are, however, service limitations imposed by this architecture. This architecture only enables a provider to offer private line T1/E1 services between two customer sites. This architecture does not enable T1/E1 services to the PSTN or any other TDM network.




Figure 5-7 – T1 private line between two customer sites

5.4.4 Other architecture considerations

For many operators, a complete T1/E1 service offering will require a mixture of these three architectures. In most cases, this will be a mixture from clauses 5.4.3 and either clause 5.4.1 or 5.4.2. Figure 5-8 depicts a potential hybrid architecture.

An operator may also choose to partner with other operators creating exchange agreements where a pseudo wire may flow over both of the managed IP networks and terminate in a TE-CM in each of their networks. Moreover, a operator may choose to partner with another operator to offer T1/E1 services outside their network nationally or internationally. In most cases, this will require the operator to terminate and aggregate the TDM-E circuit emulation for handoff as a high rate TDM interface, like a DS3. This would also be the case if an operator chooses to offer wireless backhaul services. The operator would need to aggregate multiple Pseudo Wire IP flows from several TE-CMs, terminate the TDM-E circuit emulation and handoff to the wireless operator in a high speed aggregate TDM DS3 type interface. Lastly, an operator could choose the TEA defined in this Recommendation to transport TDM traffic across the managed IP core network between two routers or other devices. The details of this are out-of-scope for this Recommendation; however, it is possible with this technology.




Figure 5-8 – Comprehensive private and public T1/E1 services

5.4.5 M-CMTS & CMTS architecture considerations

The Modular-CMTS (M-CMTS) architecture offers many advantages when considering a TDM-E service. Since a T1 or E1 service has constant symmetric bandwidth of 1.5 Mbit/s to 2 Mbit/s, a traditional CMTS architecture may not scale as well as an M-CMTS. The M-CMTS architecture decouples the upstream and downstream into separate devices that can be scaled independently. Since TDM-E, residential broadband data service, VoIP and other services all utilize the DOCSIS network this independent scalability offers greater economics and flexibility as various services are deployed.

In addition to the scalability, Modular-CMTS requires the DOCSIS timing interface ([ITU-T J.211]). This is used for modular-CMTS operation, but can also be used to support TDM-E synchronization requirements as described in clause 6.4.

Existing CMTS architectures may be utilized for TDM-E as well as the M-CMTS devices that are inherently compliant. Existing CMTS may or may not have been designed with TDM-E in mind and may need some modifications. The potential for a CMTS to be made compliant to TDM-E as a TE-CMTS is vendor and platform dependent. The CMTS will need to have a synchronization interface and support QoS requirements in order to support advanced clocking modes beyond adaptive clock recovery. This may not require any modifications for some vendors, for others rather minor changes and for some it may not be possible on existing platforms. Moreover, some vendors may be able to offer upgrades to existing CMTS platforms to not only include TE-CMTS functionality, but also an integrated TEA.

5.5 Network environment

The architecture shown in clause 5.4 serves as the basis for explaining how DOCSIS can be used to carry T1/E1 using TDM-E. The architectures do not, however, show details of the environment; details that are just as important as the devices. To adequately explain these extended architectures, it is important to outline the environmental details related to clocking and QoS.

5.5.1 Clocking

5.5.1.1 Background

T1/E1 signals are designed to simultaneously connect multiple calls over a shared wire using time domain multiplexing (TDM). This imposes certain restrictions on the T1/E1 clocks used at both ends of the connection.

- **Frequency** – To make sure that each TDM circuit stays properly connected, the clocks at both ends of the circuit must be at the same frequency. This ensures that both ends of the TDM connection continue to use the same shared data time-slot, without either end slowly slipping into the next time-slot. In other words, the data rate in and the data rate out must be precisely aligned and not deviate. The responsibility for maintaining the same frequency is

shared by the two TEAs and the clock distribution system described below in clause 5.5.1.2.

- Synchronization – To make sure that each DS0 frame is correctly multiplexed into the TDM circuit, the peer TEAs must both be synchronized to the frame of the incoming T1/E1 signal. The responsibility for maintaining frame synchronization is shared by the peer TEAs; they should be configured in a similar manner.

5.5.1.2 Clock distribution environment

There are several options available to ensure that the peer TEAs operate at the same clock frequency. From wherever the source clock exists, the frequency information about that clock must be distributed to each TEA. A few different distribution paths and methodologies exist over which a TEA or eTEA can obtain this clock information.

5.5.1.2.1 Direct clocking

Direct clocking methodology means the TEA is receiving its clocking directly from a Stratum 1 traceable reference source. This may be one of two technologies: [ITU-T J.211] or BITS.

- DTI – The DOCSIS timing interface Recommendation is specified as part of the DOCSIS M-CMTS family. [ITU-T J.211] describes a Client/Server architecture for time and frequency distribution. An M-CMTS device, CMTS or any device containing a TEA may implement a DTI client interface for connection to a DTI server. This will enable it to not only receive timing for TDM-E service, but also DOCSIS. The DTI server being used to time a device containing a TEA must implement the External Reference Mode of [ITU-T J.211] in order to comply with TDM-E. There are no additional requirements for the DTI Client.
- BITS – BITS is a generic term used to describe frequency distribution of Stratum 1 clocks in a TDM, SDH or PDH network. BITS is described in detail in ANSI T1.101-1999. In a typical deployment, a primary reference source (PRS) is containing a Stratum 1 clock or GPS (Stratum 1 traceable) connection is distributed in a Central Office, Headend or hub site directly from the PRS or through a building integrated timing supply shelf (BITS shelf). In some cases, the PRS and BITS shelf may be integrated into a single shelf called a synchronization supply unit (SSU). A network element that needs frequency synchronization will typically have a primary and secondary frequency input, typically labelled BITS input. The BITS input may have a variety of physical interfaces including T1, E1, or RS422. There are a few quality levels of BITS connections depending on the PRS and BITS topology and architecture. The same Recommendations [ITU-T G.823] and [ITU-T G.824] that govern customer T1/E1 synchronization quality also describe the PRS BITS quality used for direct clocking.

5.5.1.2.2 Network clock recovery (NCR)

Network clock recovery methodology means the TEA is receiving its clocking over the transport network (in our case DOCSIS) from a device (in our case a TE-CMTS) that has direct clocking. A device using NCR is thus traceable to a Stratum 1 clock, but has slightly lower quality clocking than a device that is directly clocked using DTI or BITS. Since DTI and BITS are both inter-office technologies, NCR is the best methodology to provide clocking to a modem. There are two methods for a TE-CM to implement NCR: DOCSIS Symbol Clock recovery and DOCSIS SYNC Message Recovery.

- DOCSIS symbol clock (SC) – The TE-CM can receive clock information encoded into the DOCSIS symbol clock. The CMTS or M-CMTS connected to the TE-CM has a clocking interface to which the DOCSIS master clock is locked. This ensures that the DOCSIS symbol clock not only meets DOCSIS requirements, but also the T1/E1 clocking requirements. The TE-CM can then provide the recovered DOCSIS symbol clock to the

eTEA. This is the most reliable and accurate method of clock distribution. One example of how this might be done is described in clause 6.4.2.2.

- DOCSIS SYNC Message Recovery – The TE-CM can receive clock information encoded in a DOCSIS SYNC Message. The CMTS or M-CMTS connected to the TE-CM has a clocking interface to which the DOCSIS master clock is locked. This ensures that the DOCSIS SYNC message not only meets DOCSIS requirements, but also T1 clocking requirements. The TE-CM can then provide a frequency reference to the eTEA disciplined by the arrival of the DOCSIS SYNC Message. One example of how this might be done is described in clause 6.4.2.3.

Direct clocking and NCR are the two most widely deployed and accurate timing methodologies deployed in communication networks today. They both are well documented and specified by the ITU and ATIS to ensure reliable and interoperable equipment and services globally. Historically, these two methodologies have been bedrock for TDM, SDH, PDH and ATM networks. However, IP/Ethernet networks are just now considering how to provide direct clocking and network clocking. DOCSIS is the first widely deployed access technology that can transport IP traffic and Stratum 1 traceable synchronization. TDM-E leverages this capability to provide a more robust IP circuit emulation service than specified by other existing Circuit Emulation standards that do not have direct clocking or NCR capabilities at this time.

5.5.1.2.3 Prime differential clock recovery (PDCR)

Prime differential clock recovery methodology leverages NCR enabling peer TEAs to precisely measure the customers T1/E1 clock and transport these measurements over the network. PDCR may be very useful for some data or cellular applications where Stratum 1 quality is needed; however, the customer has a different reference clock in their equipment.

5.5.1.2.4 Adaptive clock recovery (ACR)

Adaptive clock recovery methodology uses the customer clock to discipline a TEA (or eTEA). At the second TEA, the packet arrival rate is used to derive the original clock. In this methodology, complex algorithms and hardware are needed to filter out the jitter and wander to meet T1/E1 clocking standards. Various tradeoffs regarding cost, reliability and interoperability should be considered when using ACR. Startup time, instantaneous network delay, delay variation, packet loss, loading, and QoS all affect the clocking. The margin of safety is largely dependant on the QoS agreements that the service provider can offer.

When using ACR in the upstream direction, the quality of the recovered clock is not only affected by the IP network, it is affected by the upstream DOCSIS flow. Unsolicited grant service (UGS) is the best method to use in DOCSIS to ensure that the TDM-E service has the highest possible quality. The contention in the upstream A-TDMA or S-CDMA DOCSIS transmission will cause packet delay variation even with UGS. The effect of this will typically not have an effect on the data content or delivery; however, it may have effects on the quality of timing. To ensure high quality clocking is delivered to the customer direct clocking, NCR or PDCR should be used if available.

When using ACR across an IP network, the QoS must be provided end to end. This may impose unspecified restrictions on the combined IP network. Differential clocking (PDCR) does have some dependencies on the IP network, but is less restrictive than ACR because the actual clock is not being transmitted across the network. Rather each TEA is being timed from a Stratum 1 traceable source using NCR and the difference between the network clock and customer clock is being transmitted.

5.5.2 Quality of service (QoS)

5.5.2.1 Background

T1 services provide transport for constant bit rate (CBR) signals. Because the transport of a pseudo wire involves the sharing of resources, with the attendant requirements for reservations, IP data flows are inherently not CBR. There are two classical ways to make up for this deficiency:

- Buffering – Delay all of the IP data; buffer it into queues, and clock it out in a CBR manner. This can be done in the middle of the network IP path, but it is best done at the edge of the network, where the PW terminates in the IWF. The buffering is provided by the IWF and is configured therein. Clause 5.5.2.2.2 further describes buffering. While not strictly a 'QoS' consideration, the buffer size of the IWF must be well matched to the level of QoS so it can smooth out the IP data flow.
- QoS – Give the PW data priority over other types of data and tag it as CBR data. The IP network elements may provide QoS options allowing the operator to specify the proper priority.

5.5.2.2 QoS environment

5.5.2.2.1 QoS

The principal requirement (for a DOCSIS transport system implementing TDM-E) is to guarantee CBR transport of pseudo wire packets. In the theoretical best case, every bit presented to the network is immediately sent through the network with low and consistent delay. In the practical case, the network cannot be quite that accommodating. It may be forced by contending users to buffer (delay) traffic now and then catch up later by transporting data at a higher rate. The QoS mechanisms provided by DOCSIS can limit the amount of buffering that occurs in the network. The buffer in the TEA must be more than a match for the network buffering, or the system will break down. The better the QoS afforded to pseudo wire packets, the better the TDM-E service will function. The operators should provision the downstream and upstream of a TE-CMTS and TE-CM for the delivery of constant bit rate PWs. In the upstream, this will likely mean UGS as further defined in clause 6.7. If the pseudo wire packets leave the DOCSIS network, they must be transported with equivalent QoS. Such methods are out of scope. The weakest link in the transport network will set the expectations for remedial buffering in the TEA.

5.5.2.2.2 IWF buffering

The IWF has a configurable number of bytes to buffer each incoming pseudo wire packet stream. The IWF then performs interworking on the pseudo wire packets generating a TDM stream to the TSP.

5.5.2.2.3 Latency vs buffering errors

There is a tradeoff between the size of the buffer and the statistical chance of errors. At any level of QoS, there will be a statistical spread of packet delay. Each packet may get delayed a different amount. For the most part, a large delay would be more rare than a smaller delay. The IWF buffer must be able to accommodate almost every possible packet delay and packet delay variation. It must do so without overflow or underflow, both of which will create data errors in the T1/E1 service noticeable to the customer as BER, slips or LoS. To play it safe, the PW could always be configured with a very large buffer. This would minimize data errors. Unfortunately, many of the services carried over T1/E1 connections can only tolerate certain maximum delay times (out of scope). This max delay time puts a cap on the allowable size of the IWF buffer in that application. The cap on buffer size, in turn, places requirements for QoS on the DOCSIS network. Everything must be well matched. The goal is to simultaneously minimize the size of the buffer on each PW and stay within the statistical bounds allowable for the probability of buffer errors. Operational guidelines are out of scope for this Recommendation.

6 Technical specification

6.1 General requirements

TDM-E circuit emulation service provides emulation of T1 or E1 or nxDS0 service over a DOCSIS network. A system providing this service may consist of a TDM emulation capable CMTS (TE-CMTS), a TDM emulation cable modem (TE-CM), and possibly another router on the network side that contains a TDM emulation adapter (TEA).

6.1.1 clock unit (CU)

A CU gathers, transforms, and redistributes clocking information within a given device. There are many options as to how this may be done. The requirements for the sources of the clocking information, the transformation methods, and the redistribution methods are documented in clause 6.4.

6.1.2 pseudo wire (PW)

A pseudo wire (PW) is a flow of encapsulated TDM data within IP packets. A PW always flows between TEA peers across DOCSIS and a packet switched network.

6.1.3 Interworking function (IWF)

An IWF encodes TDM flows into IP pseudo wire flows and decodes IP pseudo wire flows into TDM flows.

An IWF **MUST** support at least one PW. An IWF **MAY** support multiple PWs. If the IWF supports multiple PWs, the configuration of each PW may be independent.

An IWF **MUST** have a configurable-sized input buffer to receive each incoming pseudo wire flow large enough to hold at least 10 ms of data at the PW line rate. For certain network conditions, larger buffers may be needed. An IWF **SHOULD** have a configurable-sized input buffer to receive incoming pseudo wire flows large enough to hold 50 ms of data at the PW line rate. The IWF should have internal control algorithms that will make sure the buffer operates near the centre of its range during normal transport (after the service is provisioned and initialized) for each PW. This will minimize buffer overflow and underflow and will help minimize the size of buffer that is needed. The buffer size should be configured to the minimum effective size to minimize latency. The algorithm necessary to control the buffer is out of scope.

An IWF **MUST** support TDM-IP interworking as defined in clause 6.2.

An IWF **MUST** support structure agnostic encapsulation as defined in clause 6.3.1.

An IWF **MUST** support structure-aware encapsulation as defined in clause 6.3.2.

An IWF **MUST** support TDM Application signalling defined in clause 6.3.2.1.

An IWF **SHOULD** support PDCR as defined in clause 6.4.3.

An IWF **MAY** support ACR as defined in clause 6.4.4.

If ACR or PDCR are supported for encapsulated PWs, the IWF **MUST** be capable of gathering clocking information from those PWs and supplying it to/from the CU.

6.1.4 TDM service processor (TSP)

A TSP **MUST** have a TDM interface for each TDM flow it can simultaneously support; the interfaces may be removable if they are only used during testing.

A TSP **MUST** accept clocking from the CU.

The TSP **MUST** provide the TDM Receive clock to the CU.

If the TSP TDM interface is a T1, it MUST comply with:

- [ITU-T G.703];
- [ITU-T G.704];
- [ITU-T G.824] as further defined in Table 6-1;
- Other National standards that are appropriate, such as [T1.403] or [TR-62411].

If the TSP TDM interface is an E1, it MUST comply with:

- [ITU-T G.703];
- [ITU-T G.704];
- [ITU-T G.706];
- [ITU-T G.732];
- [ITU-T G.823] as further defined in Table 6-1.

[ITU-T G.823] and [ITU-T G.824] differ in the requirements placed on a 'Traffic Interface' and 'Synchronization Interface' as defined therein. This Recommendation places normative requirements on TDM interfaces to conform to these ITU-T Recommendations, and Table 6-1 below is required to fully specify which clauses of the ITU-T Recommendations apply. Table 6-1 is used to specify different requirements for different TDM clocking methods which are defined in the clauses cited in the table.

A TSP TDM T1 'Traffic' interface MUST comply with:

- [ITU-T G.824] Table 1 for 1544 kbit/s digital rate interfaces;
- [ITU-T G.824] Table 2.

A TSP TDM T1 'Synchronization' interface MUST comply with:

- [ITU-T G.824] Table 1 for 1544 kbit/s digital rate interfaces;
- [ITU-T G.824] Table 5, Figure 3;
- [ITU-T G.824] Table 6, Figure 4.

A TSP TDM E1 'Traffic' interface MUST comply with:

- [ITU-T G.823] Table 1 for 2048 kbit/s digital rate interfaces;
- [ITU-T G.823] Table 2, Figure 1.

A TSP TDM E1 'Synchronization' interface MUST comply with:

- [ITU-T G.823] Table 5 for PDH synchronization;
- [ITU-T G.823] Table 12, Figure 10;
- [ITU-T G.823] Table 13, Figure 11.

Table 6-1 – G.823 and G.824 clock requirements

Clocking method	Clause	Traffic/Synchronization
Direct	6.4.1	Synchronization
NCR	6.4.2	Synchronization
PDCR	6.4.3	Synchronization
ACR	6.4.4	Traffic

A TSP interface that complies with [ITU-T G.823] or [ITU-T G.824] SHOULD comply as a 'Synchronization Interface'.

6.1.5 TDM emulation adapter (TEA)

A TEA MUST contain an IWF.

A TEA MUST contain a TSP.

A TEA MUST contain a CU supporting a synchronized clock input to support NCR or direct clocking modes.

A TEA SHOULD contain an IWF supporting PDCR mode.

A TEA MAY contain an IWF supporting ACR.

A TEA SHOULD indicate when the CU, IWF or TSP are not locked to a valid clock, including periods when it is acquiring a clock signal using ACR (clause 6.4.4), PDCR (clause 6.4.3), or SC (clause 6.4.2.2).

A TEA must be associated with a static IP address.

A TEA MUST have an LCI interface capable of accepting IP packets.

If a TEA contains an IWF that contains multiple PWs:

- the TEA MUST be capable of supporting UDP addressing per clause 6.2.2.
- the TEA MUST meet this Recommendation on all TDM interfaces and PWs simultaneously.
- the TEA MUST be configurable for each PW without disturbing configuration on any other PW (exclusive of intentionally shared parameters).

A TEA MUST make available every active TDM interface of the embedded TSP.

A TEA MUST support the provisioning of the CU, the TSP, the IWF, and the PWs it contains.

6.1.6 embedded TDM emulation adapter (eTEA)

An eTEA MUST meet the normative requirements of a TEA.

6.1.7 TDM emulation cable modem (TE-CM)

A TDM emulation cable modem (TE-CM) is a cable modem with additional functionalities to provide TDM emulation services.

A TE-CM MUST support either DOCSIS 1.1 [ITU-T J.112], DOCSIS 2.0 [ITU-T J.122] or greater.

A TE-CM MUST contain an eCM as defined in [ITU-T J.126].

A TE-CM MUST contain an eTEA as defined in clause 6.1.5.

A TE-CM MUST make available every TDM interface made available by its embedded eTEA.

A TE-CM MUST support a CPE interface that includes 10BASE-T capability. A TE-CM SHOULD have an interface with 100BASE-T capability. The 10BASE-T or 100BASE-T interface MAY be compliant to clause 6.4. One example of an acceptable industry interface is CMCI.

A TE-CM MUST contain a CU that accepts a synchronized clock input in order to support NCR mode. This allows the TE-CM to lock to the DOCSIS ([ITU-T J.122]) symbol clock or DOCSIS SYNC message and provide it to the TSP and IWF as described in clause 6.4.

A TE-CM MAY contain an Ethernet timing service based on [IEEE 1588] as outlined in clause 6.4.5.

6.1.8 TDM emulation CMTS (TE-CMTS)

A TDM emulation cable modem termination system (TE-CMTS) is a CMTS (or M-CMTS) with additional functionalities to enable TDM emulation services.

The IWF MUST discard datagrams with zero Time to Live.

The IWF must not fragment transmitted data so it MUST set the IPv4 IP Header Flags (3 bits) to 010.

6.2.2 UDP header fields

Source and destination port numbers (16 bits each). Each TDM-bound PW of a given IWF must be given a unique UDP destination port number chosen from the range of dynamically allocated and private UDP ports numbers (49152 through 65535).

An IWF MUST use the destination UDP port field to differentiate between multiple PW streams. The source UDP port MAY be used to detect misconnections. An example can be found in section 6.6.1 of [b-MEF 8].

The IWF MUST NOT be configurable for 'well known' UDP ports (0 through 1023) per [IANA UDP PORTS]. UDP length (16 bits) is the length in bytes of the UDP header and data.

UDP checksum (16 bits) is the checksum of the UDP/IP header and data. If an IWF does not compute the UDP checksum, it MUST set the UDP checksum to 0.

6.2.3 Control word header fields

Emerging IETF and ITU-T Recommendations may differ from this Recommendation in the use of the control word header. In particular, the control word field and the optional RTP field may be in different order within the packet. This Recommendation was written to avoid a problem where optional [IETF RFC 3711] encryption would obliterate the control word if implemented. This Recommendation moves the control word outside the encryption range of [IETF RFC 3711].

The control word for MPLS PWE was moved ahead of the RTP header (in [IETF RFC 4553]) to avoid problems with [IETF RFC 4928] when MPLS is used. ECMP is a shortcut taken by MPLS routers that can cause continuous path shifting and a considerable increase in jitter. This is not a concern for this Recommendation because DOCSIS systems do not use MPLS internally. The reader is referred to section 4.3 of [IETF RFC 4553] and [IETF RFC 4928] for further reading. During transport of a PW over MPLS, users are advised to strip off the IP and UDP headers so the RESERVED field of the control word will prevent MPLS ECMP rerouting. This may create problems however for port labeled PWs. Alternatively, the headers can be left on if ECMP is turned off in the MPLS routing path(s). It is recommended that any transport of this PW data over MPLS be carefully done and follow the guidelines in [IETF RFC 4928] to avoid problems.

The control word header has the following fields:

RES (4 bits) is RESERVED. The IWF MUST set the four control word RESERVED bits to 0.

L (1 bit) is the Local TDM failure indicator bit.

- For structure-agnostic emulation:
The IWF MUST set the L bit when LoS is detected on the TDM interface associated with the PW.
The IWF MAY set the L bit when the AIS defect is detected on the TDM interface associated with the PW.
- For structure-aware emulation, the IWF MUST set the L bit when LoS, OOF or AIS defects are detected on a TDM interface associated with the PW.
- When the TDM fault is cleared, the IWF MUST clear the L bit.

R (1 bit) is the Remote Receive Failure bit:

- The IWF MUST set the R bit when it detects that the loss of packet state (LoPS) has been entered for the PW.

- The IWF MUST clear the R bit when it detects an exit from the loss of packet state for the PW.

M (2 bits)

When the L bit is 0, the M bits (bit 6, bit 7) have the following meaning:

00 – Indicates regular TDM payload.

01 – Reserved.

10 – Reports the receipt of a defect such as RAI at the TDM input to the IWF. When this value is received by the peer IWF, RAI may be generated on the TDM interface according to the current configuration of RD action. This is applicable only to structure aware emulation.

11 – Indicates that the packet contains only the signalling sub-structure.

When L bit is 1, the M bits have the following meaning:

00 – Indicates there may not be valid TDM payload data in the packet.

01 – Reserved.

10 – Reserved.

11 – Reserved

RESERVED (8 bits). The IWF MUST set the 8 control word RESERVED bits to 0.

Sequence number (16 bits). The IWF SHOULD use this field to detect lost and reordered packets of each PW. The context of the sequence number is different for different values of the M bit. Signalling packets have their own sequence space. If RTP is not used, the IWF MUST process the sequence number (for each sequence space as selected by the M bits) as described in section 5.1 of [IETF RFC 3550]. See clause 6.2.4 below.

6.2.4 RTP header fields

An IWF MAY support [IETF RFC 3550]. If the IWF supports PDCR, then it must support RTP per clause 6.4.3.

The RTP header field is optional and may be eliminated, but peer IWFs must be configured the same way. For reference, consult [IETF RFC 3550].

RTV (2 bits) is the RTP Version. The IWF MUST set the 2 RTV bits to 2.

P (1 bit) is the padding. The IWF MUST set the P bit to 0.

X (1 bit) is the header extension. The IWF MUST set the X bit to 0.

CC (4 bits) is the CSRC count. The IWF MUST set the 4 CC bits to 0.

M (1 bit) is the marker. The IWF MUST set the M bit to 0.

Accordingly, RTP header extensions, padding, and contributing synchronization sources are not used.

PT (7 bits) is the Payload Type. The IWF transmitting a PW MUST utilize a PT value provided from configuration (see clause 6.7.3 and "eTEA PW RTP Payload Type" in clause 6.7.1.10.1) to set the PT field for each packet transmitted for the PW. The IWF MAY use a configured PT value in incoming packets to detect malformed packets (see clause 6.7.3 and "eTEA PW RTP Peer Payload Type" in clause 6.7.1.10.1). See section 3 of [IETF RFC 3551] for more information on assignment of PT values in RTP. The following (decimal) PT value may be used in this Recommendation:

0 – For structure agnostic mode or if the PT is not used to detect malformed packets.

RTP sequence number (16 bits). The IWF MUST set and process the RTP sequence number as described in section 5.1 of [IETF RFC 3550]. The IWF MUST set the RTP sequence number identical to the sequence number in the control word.

RTP timestamps (32 bits). These are used for carrying optional timing information over the network. The IWF MUST generate timestamps for each PW in accordance with [IETF RFC 3550] if PDCR is supported.

SSRC (synchronization source) identifier (16 bits). The SSRC (synchronization source) identifier field in the RTP. This field should be used to indicate each unique clock domain and may be used for detection of misconnections.

6.2.5 RTCP extended report (XR) option

[IETF RFC 3611] specifies the method by which RTCP headers can be used as an in-band conduit for data that is pertinent to management, operation and performance. In this particular application, the principal benefit is that operators may be able to retrieve management information from remote TEAs by simply querying the local peer TEA. In addition, since the XR signalling data is in-band, it is possible to make one-way and round-trip latency and latency variation measurements.

- An IWF MAY support [IETF RFC 3611].
- An IWF should not support, or should not be configured to use SDP as it is applied in [IETF RFC 3611]. Peer PWs must be configured similarly for inter-operation using [IETF RFC 3611].
- An IWF that supports [IETF RFC 3611]:
 - MUST support XR Block Types 4 and 5 and be configurable to disable the sending of each Block Type. An IWF can use the Receiver Reference Time Report Block (Block Type 4) and the DLRR Report Block (Block Type 5) to make and report round-trip latency measurements.
 - MUST use the method outlined in [IETF RFC 3611] section 4.5, to report latency measurements. Further informative text may be found in [MLDA] and [IETF RFC 3550] Figure 2. The IWF transmitting Block Type 4 MAY substitute any timestamp in place of the NTP timestamp specified for Block Type 4 as long as it can adequately compute the round-trip delay from the echoed information received in the received Block Type 5 report.
 - MUST issue Block Type 4 and 5 reports at least every 10 minutes if they are enabled.
 - MUST NOT support any XR Block Type that it cannot disable or report as supported.
 - SHOULD be configurable to limit the total size of the RTCP header including all XR extension blocks. This may help prevent DoS attacks.
 - MUST be configurable as to the rate of transmission of each Block Type. This will allow operators to accurately predict data bandwidth.

It may be desirable to send XR messages on a regular basis. XR messages should be sent with UGS-AD scheduling and care should be taken to make sure the XR packets travel the same physical path as the payload packets.

6.3 Payload formats

This clause specifies the format of payload data for PW packets.

The IWF MUST preserve the order of payload bytes that correspond to their order in the TDM circuit.

The IWF MUST follow network bit order when creating packets from the TDM data and vice versa for each PW.

All packets belonging to a given PW in both directions must carry the same number of payload bytes. Accordingly, the IWF MUST process received PW packet payloads based on its configuration for the number of payload bytes used in transmitting PW packets.

The IWF MUST support a configurable packet payload size for each PW. The configuration of the packet size must remain unchanged for the lifetime of the PW. The total size of a PW packet for a specific PW flow must not exceed path MTU between the pair of IWFs.

An IWF MUST NOT support payloads smaller than 20 bytes.

Except for the possible presence of signalling packets (per clause 6.3.2), PW flows do not contain specific bits describing the data organization as being 'Structure Agnostic' or 'Structure Aware' as outlined in the following subclauses. This information is contained in the configuration of the pair of IWFs.

6.3.1 Structure-agnostic transport

Structure-agnostic transport completely disregards any TDM structure and simply transports all data received from the TDM network in fixed length packets. No byte or frame alignment is implied. However, network bit order must be followed when creating packets from the TDM data and vice versa.

An IWF MUST be capable of supporting at least the following default payload sizes:

- E1 – 256 bytes.
- T1 – 192 bytes.

6.3.2 Structure-aware transport

Structure-aware transport removes the structure overhead associated with the TDM data at the ingress IWF and regenerates it at the egress IWF. Only the payload extracted from the TDM frames is carried via the PW across the packet switched network. Structure-aware transport allows channelized/fractional T1/E1 operation.

In order to preserve the integrity of the TDM structure, for each PW the IWF MUST encode all non-signal packets to carry a fixed amount of data corresponding to the integral multiple (K) of TDM frames for which it is provisioned. Thus, the packetization delay for structure aware transport is $K \times \text{frame interval}$ (i.e., $K \times 125 \mu\text{s}$). The resulting payload format is shown in Figure 6-2 below.




Figure 6-2 – Payload format for structure-aware transport

6.3.2.1 Signalling transport

TDM applications exchange signalling information in addition to TDM data. The typical example is telephony applications that exchange their state (i.e., off-hook/on-hook) in addition to TDM data carrying PCM-encoded voice.

When delivering the TDM services over a structure-agnostic pseudo wire, it is not necessary to intercept or process CE signalling. Signalling is embedded in the packetized TDM data stream, and hence it is carried end-to-end across the pseudo wire.

With structure-aware emulation, transport of common channel signalling (CCS) may be achieved by carrying the signalling channel together with the emulated service (i.e., channel 23 for DS1, or channel 16 for E1). However, channel associated signalling (CAS) (i.e., DS1 Robbed Bit Signalling or E1 CAS) requires knowledge of the relationship of the timeslot to the trunk multi-frame structure. This is indicated by the framing bits, which may not be preserved by $N \times 64$ kbit/s basic service.

This clause describes a generic method for extending the $N \times 64$ kbit/s basic service by carrying CAS signalling for each pseudo-wire in separate signalling packets that is independent of the TDM circuit type. It may be used in situations where the individual 64 kbit/s channels are selected from multiple TDM circuits, or picked off a TDM bus rather than from a specific TDM circuit. It also saves bandwidth, since only changes in the CE application state are carried.

6.3.2.2 PW signalling frames

The generic format of the PW signalling frames corresponds to the format shown in Figure 6-3. In DOCSIS systems, these frames are likely to be sent in a UGS-AD flow. The following additional requirements apply:

- An IWF MUST ensure that the following items are consistent between the PW data frames and their associated signalling frames:
 - Usage of the optional RTP header (that is, either both use it or both do not use it).
 - Destination IP Address and port.
 - Source IP Address.

The two types of frames differ only in the value of the M bits (signalling frame have a value of 'b11) and ToS bits (these are configurable separately for each frame type).

- An IWF MUST use a separate sequence number space for the PW data frames and their associated signalling frames.
- If an IWF uses the [IETF RFC 3550] header:
 - The IWF MUST use a payload type value provided from configuration (see clause 6.7.3 and "eTEA PW RTP CAS Signalling Payload Type" in clause 6.7.1.10.1) to set the PT field for associated signalling frames. The IWF MUST reject a configuration that assigns the same payload type value for data frames and signalling frames. The IWF MAY use a configured PT value in incoming packets to detect malformed packets (see clause 6.7.3 and "eTEA PW RTP CAS Signalling Peer Payload Type" in clause 6.7.1.10.1).
 - The IWF MUST assign identical timestamp values for the data and associated signalling frames at any given time.

NOTE – This enables synchronization of the signalling and data information using the standard RT-based mixing procedures described in [IETF RFC 3550].

- The IWF MUST assign the *ith* nibble of the payload the current "ABCD" value of the CAS signal corresponding to the *ith* timeslot of the PW as shown in Figure 6-3.
- The IWF MUST assign to the payload of each signalling frame an integral number of octets with at most padding of a single 4-bit nibble.




Figure 6-3 – Signalling sub-structure

The IWF MUST send PW packets with the signalling sub-structure as payload on a signalling change event including:

- Setup of the PW.
- A change in the signalling state of any of the 64 kbit/s channels of the PW.
- The loss of frames defect has been cleared.
- The remote loss of frames (R-bit) indication has been cleared.

In order to protect against packet loss, when a signalling change is detected, multiple PW signalling packets are sent. The IWF MUST send a PW signalling packet within *sp_int* ms. The IWF MUST send PW signalling packets out every *sp_int* ms until three identical packets are sent. The default value of *sp_int* is 5 ms.

The IWF SHOULD send a PW signalling frame every *sp_idle* second in the absence of any signalling change event described above. The default value of *sp_idle* is 5 seconds. The TSP MUST preserve the ABCD signalling bits at their previous state whenever packets are lost.

6.4 Timing aspects

The timing aspects of a TDM-E service are bounded by the existing DOCSIS, ITU and ATIS requirements for DOCSIS transport, and T1 or E1 service. The primary specifications that govern timing and synchronization for a T1 or E1 are the [ITU-T G.823] and [ITU-T G.824] Recommendations.

As required in clauses 6.1.4 and 6.1.5, the T1/E1 interface will be tested for compliance to [ITU-T G.823] or [ITU-T G.824].

6.4.1 Direct clocking

A device containing a CU compliant with direct clocking MUST have at least one DTI client interface to a DTI server compliant with [ITU-T J.211] or at least one BITS input for connection to a PRS compliant with [ITU-T G.823] or a PRS compliant with [ITU-T G.824].

A CU embedded in a TEA with direct clocking MUST lock the IWF and TSP to the CU.

6.4.2 Network clock recovery (NCR)

A CU compliant with NCR MUST lock the IWF and TSP to the DOCSIS SYNC message.

A CU compliant with NCR MUST lock the IWF and TSP to the DOCSIS Symbol Clock.

6.4.2.1 Support for NCR

This clause documents the strategy for maintaining long-term interoperability between TE-CMs and TE-CMTSs that use NCR.

6.4.2.1.1 DOCSIS 1.1 and 2.0

Some DOCSIS 1.1 and 2.0 CMTSs may only be able to support the SYNC method. For this reason, all TE-CMs are required to support the SYNC method.

6.4.2.1.2 DOCSIS 3.0

Going forward, the Symbol Clock method is increasingly preferred over the SYNC method. For this reason, all TE-CMs are required to support the Symbol Clock method. TE-CMTSs should support both methods if possible. This Recommendation only requires a TE-CMTS to support one method. [ITU-T J.222.1] requires a CMTS to support both methods.

6.4.2.2 Symbol clock (SC) recovery

One method of locking to the symbol clock (SC) is described as follows. A DTI or BITS source is used to derive the CMTS master clock of 10.24 MHz. The symbol clock is derived from (and locked to) the master clock via an M/N. The TE-CM recovers the 10.24 MHz master clock from the symbol clock via an N/M PLL. The 1.544 MHz T1 clock is recovered from the 10.24 MHz via a 193/1280 PLL. The 2.048 MHz E1 clock is recovered from the 10.24 MHz via a 1/5 PLL. An Ethernet clock (like 25 MHz) could be recovered from 10.24 MHz via a 625/256 PLL. Every clock is in sync.

N/M values – While the values of M and N are configurable at the CMTS, most CMTSs have symbol clocks operating at specific rates. The following information is extracted from [ITU-T J.210].

Table 6-2 – DOCSIS symbol clocks

Downstream mode	Nominal specified symbol rate, f_b (MHz)	M/N	Master clock rate, f_m' (MHz)	Downstream symbol rate, f_b' (MHz)	Offset from nominal
64QAM ([ITU-T J.83] Annex B)	5.056941	401/812	10.239990	5.056945	0.95 ppm
256QAM ([ITU-T J.83] Annex B)	5.360537	78/149	10.240000	5.360536	0.02 ppm
64QAM or 256 QAM ([ITU-T J.83] Annex A)	6.952	869/1280	10.240000	6.952	0 ppm

- SC TLV – The TE-CM MUST be capable of receiving values for M and N through an 'SC' TLV specific to the TE-CM configuration. If the TE-CM does not receive the SC TLV, it SHOULD use default M/N values from Table 6-2 as selected by the configured QAM modulation.

- S-CDMA override – The TE-CM MUST use the SC-TLV value (or its default value) unless the TE-CM receives an M/N value in the UCD for S-CDMA recovery of 10.24 MHz, in which case the S-CDMA value is to be used. See Table 8-18 of [ITU-T J.122] Channel TLV parameters, Types 12 and 13. The TE-CM MUST report an event if the S-CDMA M/N value overrides the SC TLV.
- Supported values – The TE-CM MUST support the M/N values listed in Table 6-2. The TE-CM SHOULD support all values of M/N as defined in [ITU-T J.122]. The TE-CM MUST report an event if it receives an M/N value it cannot support and revert to the default value from Table 6-2 as selected by the configured QAM modulation.

6.4.2.3 SYNC clock recovery

The existing DOCSIS 1.1, 2.0 and higher describe the DOCSIS SYNC message and the DOCSIS performance requirements. The TE-CMTS using direct clocking uses a frequency reference that meets the [ITU-T G.823] and [ITU-T G.824] requirements to generate the DOCSIS SYNC message and the contained timestamps. The DOCSIS SYNC message not only meets the performance bounds for DOCSIS, it also has long term stability needed for TDM-E. A TE-CM can recover this direct reference clock from the DOCSIS SYNC message timestamps using a PLL. A typical implementation would have a PLL that is disciplined by the DOCSIS SYNC message timestamps to generate an 8 kHz reference for the eTEA that meets the [ITU-T G.823] and [ITU-T G.824] requirements detailed in clause 6.1.4.

6.4.3 Prime differential clock recovery (PDCR)

The IWF SHOULD support prime differential clock recovery.

The principle underlying PDCR is to "encode" the timing attributes of the ingress DS1 service clock using a PRS-traceable reference clock. The peer IWF "decodes" this information using its PRS-traceable reference clock as reference to recreate the service clock for the egress DS1. (In PDCR each TEA has direct clocking or NCR available, but the customer prefers to transport their timing across the network. To do this, the IWF uses the clock from the CU in direct clocking or NCR mode to measure the customers Tx clock in the TSP. The offset is encoded and sent in an RTP PW. The PW is then received at the second TEA when the offset is added to the clocking from the CU using direct clocking or NCR to time the TSP. Since direct clocking and NCR are traceable to a Stratum 1 reference clock they are equivalent at both TEAs, so the offset measured at one TEA is precisely recreated at the second TEA.)

The PDCR encoding method is described with reference to Figure 6-4 below.

A 32-bit counter forms the base of the PDCR time-stamp (PDCR-TS). The counter operates continually based on the PRS-traceable clock ("Time Stamp reference") with the normal rollover between "all-1s" and "all-0s".

- In PDCR implementations, the IWF MUST use a counter of 32 bits for the PDCR-TS.
- In PDCR implementations, the IWF MUST support RTP.
- In PDCR implementations for T1, the IWF MUST divide the incoming 1.544 MHz "service clock" by $M = 800$ to generate a 1.93 kHz "sample clock" (PDCR time-stamp-clock).
- In PDCR implementations for E1, the IWF MUST divide the incoming 2.048 MHz clock by $M = 1061$ to generate a 1.93 kHz "sample clock" (PDCR time-stamp-clock). The generated value will not be precisely 1.93 kHz, but it is close enough.
- The PDCR time-stamp is therefore 4 octets.
- Because two clocks are involved, meta stable states may occur at the output of the latch (if a latch is used). Multiple methods should be used to avoid meta-stable states.

- One method is already specified herein. The occurrence of meta-stable states is minimized because the 1.93 kHz "sample clock" is relatively prime to the 10.24 MHz CM master clock.
- Additional methods SHOULD be devised to avoid meta-stable states.

A PDCR time-stamp is thus generated at the 1.93 kHz rate; and, if, for example, the payload contains 192/193 information octets (i.e., DS1 data), there will be 1 or 2 PDCR time-stamps generated for each DS1-payload unit.

In PDCR implementations, the latest PDCR-time-stamp MUST be sent in the payload.




Figure 6-4 – Description of PDCR "transmit" function

The PDCR implementation at the receiving IWF uses the packet sequence numbers carried in the CES control word or RTP header to identify missing or out-of-sequence packets.

The PDCR implementation at the receiving IWF uses the incoming PDCR time-stamps and its locally available PRS traceable clock to recover the original service clock.

The PDCR implementation at the receiving IWF can determine the number of PDCR timestamp (service clock) samples using the following formula:

$$\text{Number of samples} = \text{round} \left(\frac{(T_{i+N} - T_i) * SC}{TSRC} \right)$$

Where T_i and T_{i+N} are the value of the PDCR time stamps received in packets with sequence number i and consecutive packet with sequence number $i+N$ respectively, SC is the nominal rate of the sample clock (i.e., 1.93 kHz for T1 and 2048/1061 kHz for E1) and $TSRC$ is the time stamp reference clock (i.e., 10'240 kHz).

For example, assume the DS1 CES packets are sent at a rate of 1 kHz, $N = 256$, $SC = 1.93$ kHz, $TSRC = 10'240$ kHz and the offset of the service clock is 10 ppm. For each sample, the timestamp advances by $10.24\text{M}/(1.93\text{k} + 10 \text{ ppm})$ ticks or equivalently by $(10.24\text{M}-10 \text{ ppm})/1.93\text{k}$ ticks. The number of samples of the PDCR timestamp between the two packets is either 494 or 495 samples. Therefore, the difference in timestamps would be either $((10.24\text{M}-10 \text{ ppm})/1.93\text{k}) * 494$ or

$((10.24\text{M}-10\text{ ppm})/1.93\text{k})*495$. Therefore, as long as $10\text{ ppm}*495$ is much smaller than one, the formula above applies. The DS1 clock offset from the PRS traceable clock is bounded by Telecom standards to be much less than 100 ppm, and therefore the formula above applies as long as N is small enough ($N < 5000$).

When a DS1 is terminated within the DOCSIS network at both ends, both IWFs have access to the 10.24 MHz common time-stamp reference clock. However in situations where the DS1 is transported over the DOCSIS and terminated outside of the DOCSIS network, the external IWF would have access to a PRS traceable clock usually using a dedicated interface to BITS.

In order to support such topologies, the PDCR receiver implementation SHOULD support recovery of DS1 and E1 clocks using PDCR timestamps generated by different master PRS traceable clocks, different sample clocks generated by different divisor M. In particular:

19.44 MHz PRS traceable master clock (DS1 and E1 emulation) with sampling clock as defined by $M = 800$ for DS1 and $M = 1061$ for E1 signals.

Selection of the divisor M for other useful frequencies such as 24.704 MHz clock ($16*1.544\text{ MHz}$) and 32.768 MHz clock ($16*2.048\text{ MHz}$) is left for future study.

A TEA compliant with PDCR MUST be compliant with clause 6.4.7.

6.4.4 Adaptive clock recovery (ACR)

A clock recovery unit (CRU) in the IWF compliant with ACR MUST regenerate the clock provided to the TSP based on the packet arrival rate.

The RTP header in ACR can optionally be used in conjunction with the Absolute RTP mode of timestamp generation. Devices operating in the ACR mode work opposite devices operating in the internal or looped timing modes. In the loopback clocking mode, the IWF sets timestamps using the clock received from the incoming TDM circuit. As a consequence, the timestamps are closely correlated with the sequence numbers. All ACR implementations that support RTP must work opposite a TEA operating in this mode.

A CRU compliant with ACR MUST support the use of "absolute mode" timestamps, where the clock used to generate the timestamp is that recovered from the incoming TDM circuit.

A CRU compliant with ACR MUST support the use of "absolute mode" timestamps generated using an 8 kHz clock.

A CRU compliant with ACR MUST support the use of "absolute mode" timestamps generated using clocks with frequencies that are an integer multiples of 8 kHz.

A CRU compliant with ACR MUST be compliant with clause 6.4.7.

6.4.5 TE-CM Ethernet timing service

A TE-CM MAY support [IEEE 1588] on the Ethernet interface.

- A TE-CM that supports [IEEE 1588] SHOULD use the DOCSIS symbol clock as a Time Source.
- If the TE-CM supports [IEEE 1588], it SHOULD lock the [IEEE 1588] timestamp to the DOCSIS timestamp.

6.4.6 TE-CMTS timing requirements

The TE-CMTS symbol (master) clock output MUST be compliant with Figure 6-5, where MTIE (defined in [ITU-T G.810] at observation time (S) is:

$$MTIE(S) = \max_{j=1}^{N-n+1} \left[\max_{i=j}^{n+j-1} (x_i) - \min_{i=j}^{n+j-1} (x_i) \right]$$

where:

τ = sample period

N = number of samples in the sequence

$n = \lceil S/\tau \rceil + 1$

S = observation time

x_i = time delay sample¹

MTIE master clock	
Observation time, τ (s)	MTIE (ns)
$0.1 \leq \tau < 280$	$30 + 10 * \tau + 0.05 * \tau^2 + 8.75e-5 * \tau^3$
$280 \leq \tau$	$800 + 0.01 * \tau$

Figure 6-5 – MTIE of CMTS master clock

The TE-CMTS symbol (master) clock output MUST be compliant with Figure 6-6, where time deviation (TDEV) is defined in [ITU-T G.810].

TDEV master clock	
Observation time, τ (s)	TDEV (ns)
$0.1 \leq \tau < 200$	6
$200 \leq \tau < 1000$	$0.03 * \tau$
$1000 \leq \tau < 10000$	30

Figure 6-6 – TDEV of CMTS master clock

6.4.7 Clock recovery performance requirements

Clock recovery methods (ACR and PDCR) are designed to take the TDM clock from at one peer TEA and reproduce that clock at the other peer TEA. This clause outlines the performance requirements of the clock recovery methods. Further, since the performance of these methods may be affected by the characteristics of the intervening IP network, this clause also defines a model of the intervening network; defines the characteristics of that model; and states the performance

¹ Credit [ITU-T J.211].

requirements in the context of the model characteristics. Figure 6-7 shows the model system within which the clock recovery system must meet performance requirements.




Figure 6-7 – Clock recovery – Performance test environment

6.4.7.1 Definitions

- **Stimulus** – The stimulus to the performance test environment consists of IP packets containing clock recovery data. These may be generated directly by a PC, or by a 'golden' TE-CM driven by a top quality T1 clock. The IP packets entering the TEA DUT may come over an RF port (associated with a TE-CM), or over the Ethernet port in a special test mode.
- **NIM** – The stimulus packet stream will be delayed (and otherwise altered for test purposes) by the network impairment simulator operating according to one of the NIM models (see clauses 6.4.7.2 and 6.4.7.3). The NIM simulator will be concatenated with an unloaded DOCSIS system to mimic DOCSIS effects like the upstream grant process where the T1 clock may beat against the DOCSIS clock.
- **DUT** – The TEA device under test must reproduce the T1/E1 clock within the quality requirements of a T1/E1 clock and within the performance requirements of clause 6.1.4.

6.4.7.2 Network impairment model (NIM1)

Two different network impairment models are defined in this Recommendation. This clause defines NIM1; NIM2 is defined in clause 6.4.7.3. Either NIM may be used to demonstrate compliance as outlined in clause 6.4.7.4.1. The NIM chosen is then referred to herein only as "NIM".

This network impairment model (NIM1) will be the Gilbert-Elliott model as outlined in Appendix I of [ITU-T G.1020] and in [TIA NIM]. NIM1 will be configured according to clause 6.4.7.2.1 as configured according to the NIM1 constants (clause 6.4.7.2.2). It is an assumption that any TE-CMTS involved in the transport of IP packets in Figure 6-7 does not add significant jitter, packet loss, or packet reordering to the IP stream when compared to the NIM effects. It is not expected that the NIM will sufficiently replicate the characteristics (not yet identified) of the real network. It is only expected that the NIM will sufficiently stress a clock recovery circuit for testing purposes.

6.4.7.2.1 NIM1 state machine

The Gilbert-Elliott NIM1 can be modelled per [ITU-T G.1020]. The Gilbert-Elliott model behaves according to a state machine with two states:

- State 0 is a low loss state with loss probability $PL(0)$ and a state transition probability $PT0$.
- State 1 is a lossy state with loss probability $PL(1)$ and a state transition probability $PT1$. It should be noted that $PL(1) \gg PL(0)$.

Packet loss, state transition, and packet delay are driven by a Markov process. There is an option allowing (or disallowing) packet loss. The following is a pseudocode implementation of the state machine. It contains initialization, a loop that computes the delay for each packet, and a final reporting section.

```
/* Version NIM4.C */
#include <c:\Progra~1\turboc\total\stdio.h>
```

```

#include <c:\Progra~1\turboc\total\time.h>

/* input constants (to be filled in from the NIM constants table) */
float PL[2] = {.01,.01};      /* the prob of packet loss in @ state */
float PT[2] = {.04,.5};      /* the prob of state transition in @ state */
float PI = .2;                /* the prob of delay impulse */
float HI = 60.;              /* the delay driver, impulse height */
float FL = 20.;              /* the delay impulse smoothing filter */
int AllowPacketLoss = 1;      /* is packet loss allowed ? */
int AllowPacketReorder = 0;   /* is packet reordering allowed ? */

/* program constants */
int numPackets = 1000;
float lost = 32000.;          /* an encoding delay signifying 'lost' */
int states = 2;               /* a two state Gilbert-Elliott model */
float packetTime = 1.;       /* in ms (estimated), for packet ordering */

/* variables */
int state;                    /* 0 for low loss, 1 for high loss */
float d;                      /* the delay the packet will undergo */
float ds;                     /* the accumulated & filtered delay state */
float impulse;                /* the raw height of a delay impulse */
float lastd;                  /* tracks prev delay, reordering needed? */

int numDiscarded = 0;         /* number of discarded packets */
int numNotDiscarded = 0;      /* number not discarded */
float cumDelay = 0.;          /* Constants to compute variance of delay */
float cumDelaySquare = 0.;
float fnum, fdenom, Variance;
int PDiscarded, PNotDiscarded;

FILE *stream;
int i,s,packet;

main() {

    stream = fopen("myfile.txt","w");
    /* The following code is executed before packets are processed */
    printf(" PL[0] %4.3f\n",PL[0]);
    printf(" PL[1] %4.3f\n",PL[1]);
    printf(" PT[0] %4.3f\n",PT[0]);
    printf(" PT[1] %4.3f\n",PT[1]);
    printf(" PI %4.3f\n",PI);
    printf(" HI %4.3f\n",HI);
    printf(" FL %4.3f\n",FL);

    /* initialization */
    srand( (unsigned int) time( NULL )); /* seed the random generator */
    ds = d = lastd = 0; /* initialize the delay state */
    /* ----- */
    /* after initialization, for every packet... */
    for (packet = 0; packet < numPackets; packet++) {
        if (rand()/32768. < PI) { impulse = HI; } /* determine delay */
        else { impulse = 0.; }
        ds = (ds*(FL - 1) + impulse) / FL; /* filter the raw delay */
        d = ds; /* delay */
        if ((AllowPacketLoss == 1) && (rand()/32768. < PL[state])) {d += lost;}
        if (rand()/32768. < PT[state]) {state = ((state + 1) & 1);}
        if ((AllowPacketReorder == 0) && (d < (lastd - packetTime))) {
            d = lastd - packetTime;
            printf("d \t%4.1f\tlastd \t%4.1f \t reorder ",d,lastd);
        }
        else {
            printf("d \t%4.1f\tlastd \t%4.1f \tno reorder ",d,lastd);
        }
        if (d < lost) {lastd = d;} /* not discarded */
        else {
            if (lastd >= packetTime) {lastd -= packetTime;} /* discarded */
        }
        if ((AllowPacketLoss == 1) && (d >= lost)) {
            /* {discardPacket();} */
        }
    }
}

```

```

        printf("Discarded \n");
        numDiscarded++;
    }
    else {
        /* ScheduleToForwardPacketAfterDelay(d); */
        printf("ForwardAfterDelay %8.1f\n",d);
        numNotDiscarded++;
        cumDelay += d;
        cumDelaySquare += d*d;
    }
    fprintf(stream,"%4.1f\n",d);
} /* -----*/
printf("\nDone\n");
fclose(stream);
PDiscarded = (100 * (double)numDiscarded)/numPackets;
PNotDiscarded = (100 * (double)numNotDiscarded)/numPackets;
printf("Number Discarded\t= %6i %3i Percent \n"
        ,numDiscarded, PDiscarded);
printf("Number Not Discarded\t= %6i %3i Percent \n"
        ,numNotDiscarded, PNotDiscarded);
printf("Average Delay\t=\t %4.1f\n", (cumDelay/numNotDiscarded));

/* StdDev is sqrt( (n* sum(x^2) - sum(x)^2) / ((n(n-1)) ) */
fnum = ((float)numNotDiscarded * cumDelaySquare) - (cumDelay * cumDelay);
fdenom = (float)numNotDiscarded * ((float)numNotDiscarded - 1);
Variance = (fnum/fdenom);
printf("Delay Variance\t=\t%5.1f\n", Variance);
return;
} /* end */

```

NOTE 1 – The FL is used to filter and smooth the delay impulses as they become delays. The term $(ds*(FL - 1) + \text{impulse}) / FL$ shows that the entries in FL should be picked to be larger than 1. Variable ds provides delay memory between packets and decays because of the -1 . If FL is set high, it will decay (in the absence of impulses) slowly. If FL is set to 2, it will decay rapidly. FL should be determined first. Then HI can be set to determine the amplitude of the impulses and delays. Note that impulses are divided by FL. This is the reason to set FL first, then determine the size of HI needed to drive the impulses.

NOTE 2 – The function ScheduleToForwardPacketAfterDelay(d) engenders the need for multi-threading in the software. No pseudo code is included for this function (or for discardPacket()).

6.4.7.2.2 NIM1 constants

The NIM1 constants configure NIM1 according to the Gilbert-Elliott model as outlined in [ITU-T G.1020]. The specific enumeration of the constants in Table 6-3 serves to define and limit the possible variations of the Gilbert-Elliott model.

The following constants pertain to the performance requirements in clause 6.4.7.4.1. To further simulate network characteristics that vary during the day (ToD), two separate columns of NIM1 constant values are specified below in Table 6-3. Each NIM1 constant in the pseudo code of clause 6.4.7.2.1 may be faded from one ToD value to the other (and back) on a time scale simulating real ToD variations of the network.

Table 6-3 – Network impairment constants (NIM1)

Item	Description	Value ToD1	Value ToD2	Value ToD3 (Excessive packet loss)	Format	Model variable name
Probability of packet loss	The probability of packet loss in each state. (The lossy state is second)	.01/.01	.01/.01	.01/.99	float	PL[state]
Probability of state transition	The probability of state transition in each state	.04/.5	.04/.5	.5/.001	float	PT[state]
Probability of delay	The probability of impulse	.2	.2	.2	float	PI
Amount of delay	The delay impulse height	60	60	60	float	HI
Delay smoothing filter	The smoothing factor for delays	20	6	6	float	FL
Packet Loss	Determines whether the model is allowed to lose packets	1 (true)	1 (true)	1 (true)	int	AllowPacketLoss

6.4.7.3 Network impairment model (NIM2)

NIM2 is defined according to [ITU-T G.1050] and configured according to the constants in Table 6-4. This set of constants should generate a [ITU-T G.1050] NIM with the characteristics outlined in Table 6-5. The normative requirements are defined by Table 6-4 whereas Table 6-5 is informative only.

Table 6-4 – Network impairment constants (NIM2)

G.1050 constant	Value
Service Profile	B
Test Case	127
LAN A Occupancy*	0
Access A Occupancy*	0
MTU A*	512 bytes
Route flap interval*	None
Route flap duration*	0
Delay (regional)*	0
Delay (intercontinental)*	0
Link fail interval*	None
Link fail duration*	0
Packet loss*	0
Reordered packets*	0
Access B Occupancy*	0
MTU B*	512 bytes
LAN B Occupancy*	0
* This minimizes network effects that might be generated by the [ITU-T G.1050] model, leaving only generated delays and packet losses.	

Table 6-5 – NIM2 characteristics

NIM2 characteristics	Value
Latency	16 ms
Jitter	3 ms
Loss	0.1%
Average Sequential Loss	3.3
Max Sequential Loss	7
Sequential Loss Rate	0.2
OoS	0.003%
Discards	0%

6.4.7.4 Performance requirements

This clause specifies that clock recovery should function in the presence of network impairments like jitter and packet loss. It should be noted that PDCR should not be susceptible to network impairments like delay or jitter because they are subtracted out.

6.4.7.4.1 Specific requirements

The following requirements only apply to the clock recovery unit within the IWF when operating in PDCR or ACR mode implemented according to clauses 6.4.3 and 6.4.4. Figures 6-8 through 6-10 are informative only, serving only to illustrate measurements. The tests which follow, and the informative figures, are based on the assumption that the clock recovery mechanism can be modelled as a linear, second order control loop. In that context, there are no requirements on ringing

frequency, overshoot, or damping; the control loop may be over or under damped and the measurement techniques still apply.

- An IWF MUST meet the requirements in this clause in the face of network impairments described in clause 6.4.7 and generated by either NIM1 specified in clause 6.4.7.2 or NIM2 specified in clause 6.4.7.3.
- Table 6-6 is part of the following requirements; it contains the constants to which they refer. A TEA MUST have a configurable test mode so that the PW packets may be looped back directly to the TEA (as a replacement for packets received at the TE-CM RF port).
- Initialization – A CRU MUST acquire an existing stream of pseudo wire IP packets and lock the TDM output clock within Tinit seconds during the initialization of the TEA.




Figure 6-8 – Clock recovery initialization time

- Step response settling time – A CRU within the IWF MUST be capable of the following clock recovery sequence: settle for Tinit seconds to a stable locked frequency; accept a 20% step up from nominal of the stimulus frequency; and re-settle its TDM clock output to acceptable TDM clock requirements in clause 6.1.4 within Tsettle seconds after the step input.




Figure 6-9 – Clock recovery settling time

- Rejection of impulse noise – A CRU MUST be capable of rejecting a Treject second long, 20% (from nominal) amplitude deviation of stimulus frequency (after it has settled) without disturbing its clock output.




Figure 6-10 – Rejecting impulse noise

- Capture range – A CRU MUST successfully acquire, recover, and track TDM input stimulus frequencies within T_{capture} Hz of the nominal TDM output clock frequency, where T_{capture} may have different values for T1 and E1.
- Rejection of temporary change in average delay – A CRU MUST be capable of rejecting a T_{vary} second long, 20% (from nominal) change of average delay without disturbing its clock output. This temporary change of average delay may be created by the NIM model.
- Rejection of permanent change in average delay – A CRU MUST be capable of rejecting a permanent 20% (from nominal) change of average delay without disturbing its clock output. This change of average delay may be created by the NIM model.
- Rejection of excessive packet loss condition – A CRU MUST be capable of keeping the clock output stable throughout a single excessive packet loss rate condition of duration T_{loss} seconds. This excessive packet loss rate condition may be created by the NIM model using the ToD3 constants for T_{loss} seconds.

6.4.7.4.2 Performance requirement constants

The following constants pertain to the performance requirements in clause 6.4.7.4.1.

Table 6-6 – Clock recovery performance constants

Item	Description	ACR value	PDCR value	Units
Tinit	The time between initialization (reset) and clock recovery.	700	700	Seconds
Tsettle	The settling time of the output frequency.	900	900	Seconds
Treject	The length of stimulus noise impulse that must be rejected.	1	1	Seconds
Tcapture	The maximum deviation of input frequency over which recovery functions properly.	± 50 (T1) ± 60 (E1)	± 50 (T1) ± 60 (E1)	Hz
Tvary	The length of a temporary change in average delay.	1	1	Seconds
Tloss	The duration of an excessive packet loss condition.	1	1	Seconds

6.5 Defects and monitoring

6.5.1 Alarm states and alarms

- *LoPS (loss of packet state)*

An IWF MUST cause a PW to enter the LoPS state when more than the configurable percentage of packets are lost over a configurable packet loss time window or when a configurable number of consecutive packets are missing.

If a PW is in the LoPS state, the IWF MUST set the R bit in transmitted packets.

If a PW is in the LoPS state, the IWF MUST cause the PW to exit the state when a configurable number of consecutive packets are played out of the jitter buffer with no errors or when a complete jitter buffer has been played out with no missing packets.

- *Excessive packet loss alarm (EPL)*

If the IWF stays in the LoPS state for more than 2.5 seconds, it MUST set the EPL Alarm.

The IWF MUST clear the EPL alarm after 10 seconds of no errors.

- *Errored seconds*

The IWF MUST detect and count both errored seconds, severely errored seconds, and unavailable seconds for each PW.

The IWF MUST detect an error defect whenever a filler data is played out of the jitter buffer. For structure agnostic operation, a filler packet is considered as one error defect.

The IWF MUST count an errored second as any second in which one or more error defect is detected.

The IWF MUST count a severely errored second as a second in which the number of error defects is equal to or greater than the configurable severely errored second threshold number of error defects.

The IWF MUST count an unavailable second when ten consecutive seconds qualify as severely errored seconds, and it continues to be unavailable until ten consecutive seconds occur that do not qualify as severely errored seconds.

- *Fill values*

For missing data or packets, or when in the LoPS state, the IWF MUST fill transmitted TDM signals with the configured fill value.

In structure-agnostic mode, the IWF MUST use the configured fill value. The IWF MAY use the AIS code as the configured fill value in structure agnostic mode.

In structure-aware mode, the IWF MUST use the configured fill value for the applicable timeslots and maintain transport framing. The IWF MAY use the AIS code as the configured fill value in structure-aware mode.

6.5.2 Packet reordering

The IWF MUST detect mis-ordered PW packets. The IWF SHOULD reorder mis-ordered PW packets. If the IWF cannot reorder a mis-ordered packet, it MUST discard the packet. All mis-ordered packets that are not reordered MUST be counted.

6.5.3 Packet loss concealment

When occasional packet loss occurs on a PW, the IWF MUST maintain framing and fill the TDM stream with a configurable fill octet; default all ones.

In extended superframe (ESF) mode on T1, the IWF MUST set the ESF data link signalling bits to 1's in a PW packet loss condition.

6.5.4 Packet loss statistics

The IWF MUST maintain a MIB for each PW with 24 hours/96 intervals of 15-minute packet loss statistics. In each 15-minute interval, the IWF MUST update MIB objects maintaining a count of lost packets and a count of seconds where the SPLA alarm condition has been declared.

6.5.5 Stray packets

An IWF MUST discard stray (non-PW) packets. An IWF MUST NOT alter its packet loss statistics based on discarding stray packets.

6.5.6 Packet activity during provisioning

An IWF MUST output an 'all ones' pattern as its TDM flow to the TSP while it is being set up (before it is ready to re-transmit received data) and while it is being torn down (after it no longer can retransmit). These time periods may be defined by the fill state of the IWF buffer for a given PW.

6.6 Security

PW traffic between IWF peers may cross networks where security is an issue. While latency is an issue with any further processing, [IETF RFC 3711] seems to be a very efficient and tight security mechanism for this application.

- The IWF MAY implement [IETF RFC 3711] security for end-to-end security between peer IWFs in a service.
- If an IWF implements [IETF RFC 3711], the IWF MUST be able to optionally disable it.
- The IWF may use other security protocols like IPSec.
- The IWF MUST NOT use security protocols that cannot be disabled.
- Key exchange is not fully specified in [IETF RFC 3711] and is out of scope herein.

6.7 Provisioning and service initiation

6.7.1 Provisioning of embedded TDM emulation adapter (eTEA)

This clause describes the embedded TDM emulation adapter (eTEA), its attributes and its configuration.

6.7.1.1 eTEA device provisioning

The eTEA is defined as a eSAFE device as specified by [ITU-T J.126]. Device provisioning is the process by which the eTEA device is configured to support TDM Emulation service. The device provisioning involves the eTEA obtaining its IP configuration required for basic network connectivity, announcing itself to the network, and downloading its configuration data from its provisioning server. A basic content integrity verification check MUST be conducted on the configuration file by the eTEA as defined in clause 6.7.1.9.4.