

SCTE • ISBE **S T A N D A R D S**

Data Standards Subcommittee

AMERICAN NATIONAL STANDARD

ANSI/SCTE 136-2 2019

Cable Modem TDM Emulation Interface Standard

NOTICE

The Society of Cable Telecommunications Engineers (SCTE) / International Society of Broadband Experts (ISBE) Standards and Operational Practices (hereafter called “documents”) are intended to serve the public interest by providing specifications, test methods and procedures that promote uniformity of product, interchangeability, best practices and ultimately the long-term reliability of broadband communications facilities. These documents shall not in any way preclude any member or non-member of SCTE•ISBE from manufacturing or selling products not conforming to such documents, nor shall the existence of such standards preclude their voluntary use by those other than SCTE•ISBE members.

SCTE•ISBE assumes no obligations or liability whatsoever to any party who may adopt the documents. Such adopting party assumes all risks associated with adoption of these documents, and accepts full responsibility for any damage and/or claims arising from the adoption of such documents.

Attention is called to the possibility that implementation of this document may require the use of subject matter covered by patent rights. By publication of this document, no position is taken with respect to the existence or validity of any patent rights in connection therewith. SCTE•ISBE shall not be responsible for identifying patents for which a license may be required or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

Patent holders who believe that they hold patents which are essential to the implementation of this document have been requested to provide information about those patents and any related licensing terms and conditions. Any such declarations made before or after publication of this document are available on the SCTE•ISBE web site at <http://www.scte.org>.

All Rights Reserved
© Society of Cable Telecommunications Engineers, Inc. 2019
140 Philips Road
Exton, PA 19341

Note: DOCSIS® and eDOCSIS™ are trademarks of Cable Television Laboratories, Inc., and are used in this SCTE standard with permission.

Contents

1	SCOPE AND PURPOSE	1
1.1	SCOPE	1
1.2	INTRODUCTION AND OVERVIEW	1
1.3	REQUIREMENTS AND CONVENTIONS	2
2	REFERENCES	3
2.1	NORMATIVE REFERENCES	3
2.2	INFORMATIVE REFERENCES.....	5
2.3	REFERENCE ACQUISITION	5
3	TERMS AND DEFINITIONS	6
4	ABBREVIATIONS AND ACRONYMS.....	8
5	TECHNICAL OVERVIEW	9
5.1	BACKGROUND	9
5.2	PRINCIPAL TECHNICAL CONSIDERATIONS	10
5.3	DEVICES/INTERFACES	10
5.3.1	TE-CM.....	10
5.3.2	TE-CMTS.....	11
5.3.3	Embedded TDM Emulation Adaptor (eTEA).....	12
5.3.4	TDM Emulation Adaptor (TEA).....	13
5.4	ARCHITECTURE	14
5.4.1	TE-CM to TE-CMTS w/TEA.....	14
5.4.2	TE-CM to Centralized TEA.....	15
5.4.3	TE-CM to TE-CM through TE-CMTSs and Core.....	15
5.4.4	Other Architecture Considerations.....	16
5.4.5	M-CMTS & CMTS Architecture Considerations.....	17
5.5	NETWORK ENVIRONMENT	17
5.5.1	Clocking.....	17
5.5.2	Quality of Service (QoS).....	19
6	TECHNICAL SPECIFICATION.....	21
6.1	GENERAL REQUIREMENTS.....	21
6.1.1	CU Clock Unit.....	21
6.1.2	PW Pseudo-Wire	21
6.1.3	IWF InterWorking Function.....	21
6.1.4	TSP TDM Service Processor.....	21
6.1.5	TEA TDM Emulation Adapter.....	23
6.1.6	eTEA Embedded TDM Emulation Adaptor.....	23
6.1.7	TE-CM TDM Emulation Cable Modem.....	23
6.1.8	TE-CMTS TDM Emulation CMTS.....	23
6.2	TDM-IP NETWORK INTERWORKING	24
6.2.1	IP Header Fields	25
6.2.2	UDP Header Fields.....	25
6.2.3	Control Word Header Fields.....	25
6.2.4	RTP Header Fields.....	26
6.2.5	RTCP Extended Report (XR) Option	27
6.3	PAYLOAD FORMATS.....	27
6.3.1	Structure Agnostic Transport.....	28
6.3.2	Structure Aware Transport.....	28
6.4	TIMING ASPECTS	31
6.4.1	Direct Clocking.....	31
6.4.2	Network Clock Recovery (NCR).....	32

6.4.3	<i>Prime Differential Clock Recovery (PDCR)</i>	33
6.4.4	<i>Adaptive Clock Recovery (ACR)</i>	35
6.4.5	<i>TE-CM Ethernet Timing Service</i>	35
6.4.6	<i>TE-CMTS Timing Requirements</i>	35
6.4.7	<i>Clock Recovery Performance Requirements</i>	36
6.5	DEFECTS AND MONITORING	43
6.5.1	<i>Alarm states and Alarms</i>	43
6.5.2	<i>Packet Reordering</i>	44
6.5.3	<i>Packet Loss Concealment</i>	44
6.5.4	<i>Packet Loss Statistics</i>	44
6.5.5	<i>Stray Packets</i>	44
6.5.6	<i>Packet Activity during Provisioning</i>	44
6.6	SECURITY	44
6.7	PROVISIONING AND SERVICE INITIATION	44
6.7.1	<i>Provisioning of Embedded TDM Emulation Adaptor (eTEA)</i>	44
6.7.2	<i>UGS Flow Provisioning</i>	64
6.7.3	<i>TDM Emulation Adaptor (TEA) Provisioning</i>	64
6.8	MANAGEMENT.....	64
6.8.1	<i>Management Model</i>	64
6.8.2	<i>Managed objects requirements</i>	65
6.8.3	<i>TLV to MIB Object Mapping</i>	67
ANNEX A	EVENT, SYSLOG, AND SNMP TRAP EXTENSIONS (NORMATIVE)	70
A.1	TEA EVENTS DESCRIPTION	70
A.1.1	<i>TEA event process "Init"</i>	70
A.1.2	<i>TEA event process "Oper"</i>	70
A.2	DOCSIS EVENTS EXTENSIONS.....	70
APPENDIX I	OPERATOR RECOMMENDATIONS (INFORMATIVE)	72
I.1	PREFERENTIAL TREATMENT FOR TE-CMS.....	72
I.1.1	<i>Segregation of T1 traffic</i>	72
I.1.2	<i>Preferential ranging opportunities</i>	72
I.2	SERVICE LEVEL AGREEMENT CONSIDERATIONS	72
I.3	STANDARD DOCSIS OPERATIONS.....	73
I.3.1	<i>Downloading</i>	73
I.3.2	<i>Scheduled Service</i>	73
I.3.3	<i>Provisioning</i>	73
I.3.4	<i>Balancing operations</i>	73
APPENDIX II	DELAY ASPECTS (INFORMATIVE)	74
II.1	UPSTREAM DELAY	74
II.1.1	<i>Packetization Delay</i>	74
II.1.2	<i>Processing Delay in TE-CM</i>	74
II.1.3	<i>Propagation Delay</i>	74
II.1.4	<i>Delay in CMTS</i>	75
II.1.5	<i>M-CMTS Delay</i>	75
II.2	DOWNSTREAM DELAY	75
II.2.1	<i>Packetization Delay</i>	75
II.2.2	<i>Delay in CMTS</i>	75
II.2.3	<i>Propagation Delay</i>	75
II.2.4	<i>Delay in TE-CM</i>	75
II.2.5	<i>M-CMTS delay</i>	75
APPENDIX III	BANDWIDTH USAGE (INFORMATIVE)	77
III.1	UPSTREAM BANDWIDTH USAGE	77
III.2	DOWNSTREAM BANDWIDTH USAGE	78

III.2.1	<i>Downstream RF Bandwidth Usage</i>	78
III.2.2	<i>DEPI Bandwidth Usage</i>	78
APPENDIX IV ETEA CONFIGURATION EXAMPLE (INFORMATIVE)		79
IV.1	STRUCTURE AGNOSTIC T1 CONFIGURATION	79
IV.2	STRUCTURE AGNOSTIC E1 CONFIGURATION	80
IV.3	STRUCTURE AWARE 5xDS0 CONFIGURATION	80
IV.4	STRUCTURE AWARE 5xDS0 CONFIGURATION WITH CAS	81
IV.5	STRUCTURE AWARE T1 WITH RTP CONFIGURATION	83
APPENDIX V CLOCKING OPTIONS, BY EXAMPLE (INFORMATIVE)		84
V.1	BACKGROUND	84
V.2	CLOCKING MODES	85
V.2.1	<i>Network clocking</i>	85
V.2.2	<i>Prime Differential Clock Recovery</i>	86
V.2.3	<i>Adaptive Clock Recovery</i>	86
V.3	CLOCKING DECISION TREE.....	87

List of Figures

FIGURE 1-1 - EXAMPLE OF T1 SERVICE DELIVERY	1
FIGURE 5-1 - TE-CM.....	11
FIGURE 5-2 - TE-CMTS	12
FIGURE 5-3 - eTEA	13
FIGURE 5-4 - TEA	14
FIGURE 5-5 - T1 TO HUB ARCHITECTURE	15
FIGURE 5-6 - T1 TO TDM CORE, IXC OR PSTN	15
FIGURE 5-7 - T1 PRIVATE LINE BETWEEN TWO CUSTOMER SITES	16
FIGURE 5-8 - COMPREHENSIVE PRIVATE AND PUBLIC T1/E1 SERVICES	17
FIGURE 6-1 - HEADER FORMAT (ONLY IPV4 SHOWN).....	24
FIGURE 6-2 - PAYLOAD FORMAT FOR STRUCTURE-AWARE TRANSPORT	29
FIGURE 6-3 - SIGNALING SUB-STRUCTURE	31
FIGURE 6-4 - DESCRIPTION OF PDCR "TRANSMIT" FUNCTION.....	34
FIGURE 6-5 - CLOCK RECOVERY - PERFORMANCE TEST ENVIRONMENT	36
FIGURE 6-6 - CLOCK RECOVERY INITIALIZATION TIME	41
FIGURE 6-7 - CLOCK RECOVERY SETTling TIME	42
FIGURE 6-8 - REJECTING IMPULSE NOISE	42
FIGURE 6-9 - DEVICE STATES AND STATE TRANSITIONS FOR BASIC FLOW PROVISIONING	45
FIGURE 6-10 - CONCEPTUAL LAYERING	65
FIGURE 6-11 - PW EXAMPLES.....	66
FIGURE V-1 - NETWORK CLOCKING - CMTS TO CM.....	85
FIGURE V-2 - PRIME DIFFERENTIAL CLOCK RECOVERY	86
FIGURE V-3 - ADAPTIVE CLOCK RECOVERY	87

List of Tables

TABLE 6-1 - G.823 AND G.824 CLOCK REQUIREMENTS	22
TABLE 6-2 - DOCSIS SYMBOL CLOCKS	32
TABLE 6-3 - MTIE OF CMTS MASTER CLOCK	36
TABLE 6-4 - TDEV OF CMTS MASTER CLOCK.....	36
TABLE 6-5 - NETWORK IMPAIRMENT CONSTANTS (NIM1)	39
TABLE 6-6 - NETWORK IMPAIRMENT CONSTANTS (NIM2)	40
TABLE 6-7 - NIM2 CHARACTERISTICS.....	40
TABLE 6-8 - CLOCK RECOVERY PERFORMANCE CONSTANTS.....	43
TABLE 6-9 - eTEA POWER-ON INITIALIZATION FLOW	46
TABLE 6-10 - eTEA INITIALIZATION FLOW	46
TABLE 6-11 - DHCP OPTION 43 SYNTAX	48
TABLE 6-12 - TLV DEFINITIONS	49
TABLE 6-13 - TLV TO MIB OBJECT MAPPING	67
TABLE A-1 - DOCSIS EVENTS EXTENSIONS.....	70
TABLE III-1 - UPSTREAM BANDWIDTH COMPUTED FROM DOCSIS 2.0.....	77
TABLE III-2 - AN EXAMPLE OF UPSTREAM BANDWIDTH USAGE.....	77
TABLE III-3 - AN EXAMPLE OF DOWNSTREAM BANDWIDTH USAGE	78
TABLE IV-1 - STRUCTURE AGNOSTIC T1 PW CONFIGURATION FILE EXAMPLE	79
TABLE IV-2 - STRUCTURE AGNOSTIC E1 PW CONFIGURATION FILE EXAMPLE	80
TABLE IV-3 - STRUCTURE AWARE NXDS0 PW CONFIGURATION FILE EXAMPLE	80
TABLE IV-4 - STRUCTURE AWARE NXDS0 PW CONFIGURATION FILE EXAMPLE	81
TABLE IV-5 - STRUCTURE AWARE NXDS0 PW CONFIGURATION FILE EXAMPLE	83
TABLE V-1 - BSoD CLOCKING MODES	84

1 SCOPE AND PURPOSE

This document is identical to SCTE 136-2 2013 except for informative components which may have been updated such as the title page, NOTICE text, headers and footers. No normative changes have been made to this document.

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 standard is part of the Cable Modem family of standards and in particular, defines the TDM-E architecture and components.

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 document 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.

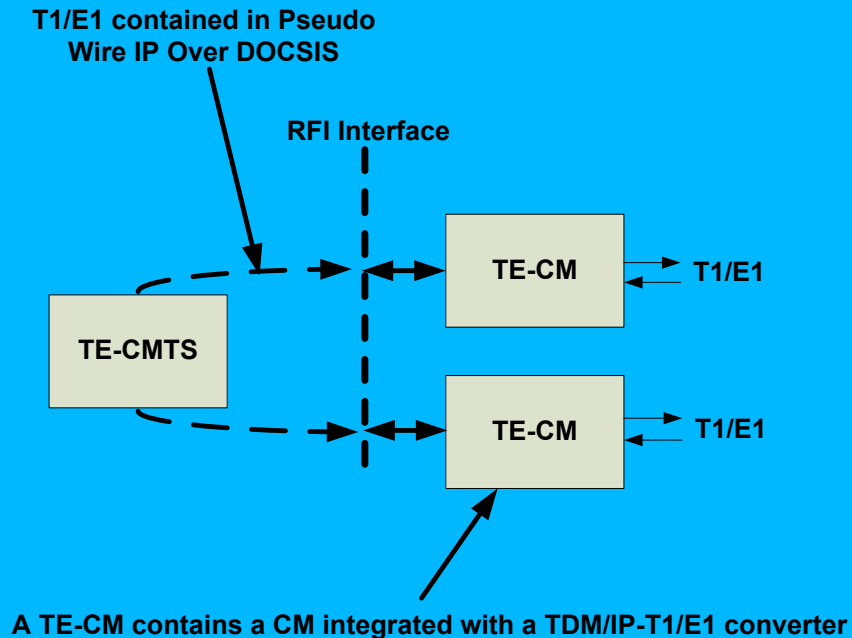


Figure 1–1 - Example of T1 Service Delivery

1.3 Requirements and Conventions

Throughout this document, 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 specification. |
| "MUST NOT" | This phrase means that the item is an absolute prohibition of this specification. |
| "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 behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior 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. |

2 REFERENCES

The following documents contain provisions, which, through reference in this text, constitute provisions of the standard. At the time of Subcommittee approval, the editions indicated were valid. All standards are subject to revision; and while parties to any agreement based on this standard are encouraged to investigate the possibility of applying the most recent editions of the documents listed below, they are reminded that newer editions of those documents may not be compatible with the referenced version.

2.1 Normative References

- [DEPI] ANSI/SCTE 137-2, DOCSIS Downstream External PHY Interface for Modular Cable Modem Termination Systems
- [DRFI] ANSI/SCTE 133, DOCSIS Downstream RF Interface for Cable Modem Termination Systems
- [CMCI] ANSI/SCTE 135-5, DOCSIS 3.0 Part 5: Cable Modem to Customer Premise Equipment Interface
- [DTI] ANSI/SCTE 137-1, DOCSIS Timing Interface for Cable Modem Termination Systems
- [MULPIv3.0] ANSI/SCTE 135-2, DOCSIS 3.0 Part 2: MAC and Upper Layer Protocols
- [OSSIv2.0] ANSI/SCTE 79-2, DOCSIS 2.0 Part 2: Operations Support System Interface
- [OSSIv3.0] ANSI/SCTE 135-4, DOCSIS 3.0 Part 4: Operations Support System Interface Specification
- [PHYv3.0] ANSI/SCTE 135-1, DOCSIS 3.0 Part 1: Physical Layer Specification
- [RFI 1.1] ANSI/SCTE 23-1, DOCSIS 1.1 Part 1: Radio Frequency Interface
- [RFI 2.0] ANSI/SCTE 79-1, DOCSIS 2.0 Part 1: Radio Frequency Interface
- [eDOCSIS] ANSI/SCTE 107, Embedded Cable Modem Devices
- [IEEE 1588] "IEC/IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", 61588-2004, November 15, 2004, IEEE.
- [ISO 8825] ISO 8825-1 (2002) Information Technology - ASN.1 encoding rules: Specification of the Basic Encoding Rules, Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER).
- [ITU-T G.703] ITU-T Recommendation G.703 (11/01) - Physical/electrical characteristics of hierarchical digital interfaces, International Telecommunication Union.
- [ITU-T G.704] ITU-T Recommendation G.704 (10/98) - Synchronous Frame Structures Used At 1544, 6312, 2048, 8488, and 44 736 Kbit/s Hierarchical Levels, International Telecommunication Union.
- [ITU-T G.706] ITU-T Recommendation G.706 (4/91) - "Frame alignment and cyclic redundancy check (CRC) procedures relating to basic frame structures defined in Recommendation G.704", International Telecommunication Union.
- [ITU-T G.732] ITU-T Recommendation G.732 (11/88) - "Characteristics of primary PCM multiplex equipment operating at 2048 kbit/s", International Telecommunication Union.
- [ITU-T G.810] ITU-T Recommendation G.810 (8/96) - "Definitions And Terminology For Synchronization Networks", International Telecommunication Union.
- [ITU-T G.823] ITU-T Recommendation G.823 (03/00) - "The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy", International Telecommunication Union.
- [ITU-T G.824] ITU-T Recommendation G.824 (03/00) - "The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy", International Telecommunication Union.
- [ITU-T G.1020] ITU-T Recommendation G.1020 (11/03) – Performance parameter definitions for quality of speech and other voiceband applications utilizing IP networks, International Telecommunication Union.
- [ITU-T G.1050] ITU-T Recommendation G.1050 (11/05) – Network model for evaluating multimedia transmission performance over internet protocol, International Telecommunication Union.

- [ITU-T J.83] ITU-T Recommendation J.83 (12/07) - Digital multi-program systems for television, sound and data services for cable distribution, International Telecommunication Union.
- [MEF8] Metro Ethernet Forum, Technical Specification MEF8, Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004.
- [RFC 1123] IETF RFC 1123/STD0003, Braden, R., Requirements for Internet Hosts – Application and Support, October 1989, Internet Engineering Task Force.
- [RFC 1157] IETF RFC1157/STD0015, J.D. Case, M. Fedor, M.L. Schoffstall, J. Davin, Simple Network Management Protocol (SNMP), May 1990. Internet Engineering Task Force.
- [RFC 1350] IETF RFC 1350, K. Sollins, The TFTP Protocol (Revision 2), July 1992, Internet Engineering Task Force.
- [RFC 2131] IETF RFC 2131, R. Droms, Dynamic Host Configuration Protocol, March 1997, Internet Engineering Task Force.
- [RFC 2132] IETF RFC 2132, S. Alexander, DHCP Options and BOOTP Vendor Extensions, March 1997, Internet Engineering Task Force.
- [RFC 2863] IETF RFC 2863, K. McCloghrie, F. Kastenholz, The Interfaces Group MIB, June 2000 Internet Engineering Task Force.
- [RFC 2494] IETF RFC 2494, D. Fowler, Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type, January 1999, Internet Engineering Task Force.
- [RFC 3396] IETF RFC 3396, T. Lemon, Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4), November 2002, Internet Engineering Task Force.
- [RFC 3550] IETF RFC 3550, H. Schulzrinne, RTP: A Transport Protocol for Real-Time Applications, July 2003, Internet Engineering Task Force.
- [RFC 3611] IETF RFC 3611, T. Friedman, R. Caceres, A. Clark RTP Control Protocol Extended Reports (RTCP XR), November, 2003, Internet Engineering Task Force.
- [RFC 3635] IETF RFC 3635, J. Flick, Definitions of Managed Objects for the Ethernet-like Interface Types, September 2003, Internet Engineering Task Force.
- [RFC 3636] IETF RFC 3636, J. Flick, Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs), September 2003, Internet Engineering Task Force.
- [RFC 3711] IETF RFC 3711, M. Baugher, and D. McGrew, The Secure Real-time Transport Protocol (SRTP), March 2004, Internet Engineering Task Force.
- [RFC 3895] IETF RFC 3895, O. Nicklass, Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types, September 2004, Internet Engineering Task Force.
- [RFC 5542] IETF RFC 5542, T. Nadeau, D. Zelig, O. Nicklass, Definitions for Textual Conventions for Pseudo-wire (PW) Management, May 2009, Internet Engineering Task Force.
- [RFC 5601] IETF RFC 5601, T. Nadeau, et al., Pseudowire (PW) Management Information Base (MIB), July 2009, Internet Engineering Task Force.
- [RFC 5604] IETF RFC 5604, O. Nicklass, Managed Objects for Time Division Multiplexing (TDM) over Packet Switched Networks (PSNs), July 2009, Internet Engineering Task Force.
- [T1-403] ANSI T1.403-2005, "Network and Customer Installation Interfaces - DS1 - Electrical Interface", American National Standards Institute.
- [TIA NIM] TR41.4-03—05-029 (9/05) – IP Network Impairment Simulation for VoIP Performance Testing, May 7, 2003, Telecommunications Industry Association.
- [TR-62411] AT&T Technical Reference 62411, "AT&T ACCUNET T1.5 Service Description and Interface Specification", December 1990, AT&T.

2.2 Informative References

This standard uses the following informative references.

- [DSG] ANSI/SCTE 106, DOCSIS Set-Top Gateway (DSG) Specification
- [IANA UDP PORTS] Internet Assigned Numbers Authority, Port Numbers, January 4 2010, Refer to <http://www.iana.org/assignments/port-numbers>.
- [MLDA] Sisalem D. and A. Wolisz, MLDA: A TCP-friendly Congestion Control Framework for Heterogeneous Multicast Environments, IEEE Proc. IWQoS 2000.
- [RFC 4553] IETF RFC 4553, Vainshtein, A. and Stein, Y., Structure-Agnostic Time Division Multiplexing over Packet (SAToP), June 2006, Internet Engineering Task Force.
- [RFC 4928] IETF RFC 4928. George Swallow, G., Bryant, S., and Andersson, L., Avoiding Equal Cost Multipath Treatment in MPLS Networks, June 2007, Internet Engineering Task Force.
- [RFC 5085] IETF RFC 5085, T. Nadeau, C. Pignataro, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires, December 2007, Internet Engineering Task Force.
- [RFC 5086] IETF RFC 5086. Vainshtein, A. et al, Structure-aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN), December 2007, Internet Engineering Task Force.
- [RFC 5087] IETF RFC 5087, Stein, Y., et al., Time Division Multiplexing over IP (TDMoIP), December 2007, Internet Engineering Task Force.

2.3 Reference Acquisition

- American National Standards Institute, Internet: <http://webstore.ansi.org>.
- AT&T, <http://www.att.com>.
- Institute of Electrical and Electronic Engineers (IEEE), <http://www.ieee.org>.
- Internet Engineering Task Force (IETF), Internet: <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 Associations (TIA), <http://www.tiaonline.org/>.
- Telcordia Technologies, <http://www.telcordia.com/>.

3 TERMS AND DEFINITIONS

This standard uses the following terms:

Alarm Indication Signal	Also known as the <i>blue alarm</i> . 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.
CU	A Clock Unit (CU) performs translations and distribution of TDM clocking information across (and between) physical, data, and network layers.
CRU	A Clock Recovery Unit (CRU) exists in the IWF and is responsible for regenerating the circuit's 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.
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.
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.
DS0	Digital signal 0 (DS0) is a basic digital signaling rate of 64 kbit/s, corresponding to the capacity of one voice-frequency-equivalent channel.
E1	E1 is a physical layer telephony protocol carrying data at 2.048 Mbps. It can carry up to 32 DS0s, each of which can carry a telephone conversation.
eTEA	An eTEA is an eDOCSIS eSAFE; an embedded version of a TEA.
Fractional	A Fractional T1 or Fractional E1 carries only a portion of the total number of DS0s that a T1 or E1 carries.
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 Pseudo Wire in the other direction.
J1	J1 is the Japanese version of a T1; it has some alarm and framing differences as compared with T1.
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.
LCI	A Logical CPE Interface (LCI) is a logical 802.3/Ethernet MAC frame data interface.
Loss of Frame	Also called the <i>red alarm</i> . When a T1 CSU/DSU is unable to synchronize framing patterns with the remote end for 2.5 seconds, LOF is declared
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.
MTIE	Maximum Time Interval Error as defined in [ITU-T G.810].
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.
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).

Remote alarm indication	Also known as <i>yellow alarm</i> . 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.
T1	T1 is a physical layer telephony protocol carrying data at 1.544Mbps. It can carry up to 24 DSOs, each of which can carry a telephone conversation. It is used mainly in North America.
TDEV	Time Deviation as defined in [ITU-T G.810].
TDM Interface	A physical Time Domain Multiplex (TDM) telephony interface such as T1, E1, or J1. Time Division Multiplexing (TDM) is the means by which multiple digital signals (DSOs) can be carried on a single transmission path by interleaving portions of each signal in time.
TEA	A TDM Emulation Adapter (TEA) is a logical entity containing various functions to provide a BSoD-TE circuit emulation service.
TE-CM	The TDM Emulation Cable Modem (TE-CM) is a special purpose cable modem that integrates a DOCSIS eCM with an eTEA.
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 BSoD-TE traffic. A TE-CMTS may or may not support integrated TEAs.
TSP	A 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.
Wander	Low speed jitter. Wander is defined in [ITU-T G.810] as fluctuations below 10 Hz. See Jitter.

4 ABBREVIATIONS AND ACRONYMS

This standard 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
CMTS	Cable Modem Termination System
CSU	Channel Service Unit
DEPI	DOCSIS External PHY Interface
DOCSIS®	Data-Over-Cable Service Interface Specifications
DS	Downstream
DSU	Data Service Unit
DTI	DOCSIS Timing Interface
HFC	Hybrid Fiber Coax
IP	Internet Protocol
LOF	Loss of frame
LOS	Loss of signal
MAC	Media Access Control
MPLS	Multiprotocol Label Switching
OOF	Out of frame
PDCR	Prime Differential Clock Recovery
PW	Pseudo Wire
RAI	Remote alarm indication
RF	Radio Frequency
S-CDMA	Synchronous Code Division Multiple Access
UDP	User Datagram Protocol
UGS	Unsolicited Grant Service

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 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 90s. This standard leverages the work done in the ITU and IETF to define circuit emulation of T1 and E1 over a DOCSIS network.

This standard 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 BSoD-TE compliant cable modem (TE-CM), and the requirements of a BSoD-TE compliant CMTS (TE-CMTS), the requirements of the TDM Emulation Adaptor (TEA), and Pseudo Wires (PW) that cross the packet network connecting two TEAs.

In developing this standard, the following assumptions were made concerning the implementation and deployment of BSOD-TE 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 [RFI 2.0] and S-CDMA [PHY v3.0]. 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.
- From time to time this document may refer to the voice communications capabilities of a TE-CM device / TDM-IP protocol in terms of "IP Telephony." The legal/regulatory classification of IP-based voice communications provided over cable networks and otherwise, and the legal/regulatory obligations, if any, borne by providers of such voice communications, are not yet fully defined by appropriate legal and regulatory authorities. Nothing in this standard is addressed to, or intended to affect, those issues. In particular, while this document uses standard terms such as "call," "call signaling," "telephony," etc., it should be recalled that, while a TE-CM device / TDM-IP protocol performs activities analogous to these PSTN functions, the manner by which it does so differs considerably from the manner in which they are performed in the PSTN by telecommunications carriers, and that these differences may be significant for legal/regulatory purposes. Moreover, while reference is made here to "IP Telephony," it should be recognized that this term embraces a number of different technologies and network architectures, each with different potential associated legal/regulatory obligations. No particular legal/regulatory consequences are assumed or implied by the use of this term.
- 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.
- J1 can be partially supported, as outline in Section 6.3.1.
- It's not in the scope of this standard to cover interoperation between dissimilar services, i.e., between E1 and T1.
- The following requirements are market specific and are out of scope in this standard:
 - 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 [RFC 4553] section 5.2, which is designed to simplify the data handling in and out of a Sonet/SDH virtual tributary/container. This standard does not support that format.

5.2 Principal technical considerations

This section 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 section on clocking (Section 5.5.1) 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 wasn't 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 spec. This standard 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 sections below addressing 'structure agnostic' protocol are germane to data transport. Most of this standard addresses T1 transport of voice calls, but the standard accommodates data transport as well.
 3. Signaling 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 signaling data. This standard describes the format for signaling 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 standard 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 Adaptor (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 below. 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 signaling as defined herein. It also contains [eDOCSIS] 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 Section 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 BSoD-TE entities that may exist outside the TE-CM.
- Logical CPE Interface - This is an unspecified interface between the eTEA and eCM.

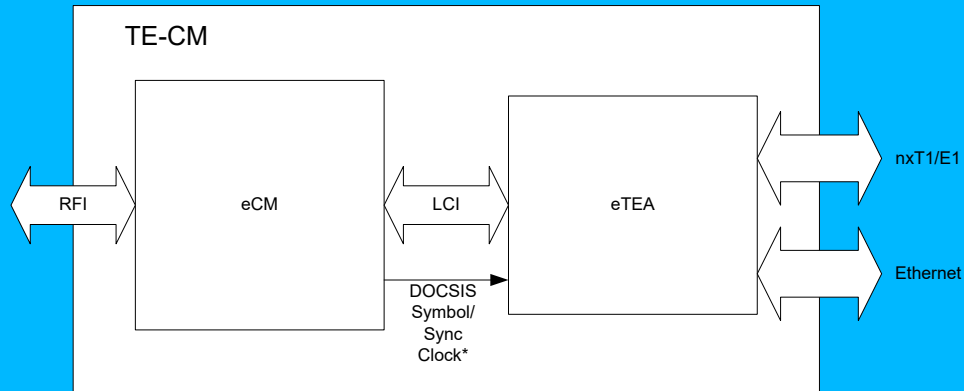


Figure 5-1 - TE-CM

NOTE: [eDOCSIS] specifications apply to the eCM. The DOCSIS Symbol/Sync clock is Stratum 1 Traceable.

5.3.2 TE-CMCTS

A TE-CMCTS is a standard DOCSIS CMTS that contains the necessary clocking functionality necessary to support a TE-CM and may also include a TDM Emulation Adaptor (TEA) giving it the ability to terminate T1/E1 services. A TE-CMCTS may have architecture and interfaces similar to that shown in Figure 5-2 below. 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 specification defines compliant IP, RF, and synchronization interfaces.
- Clocking Unit - When the TEA is not integrated within the TE-CMCTS, at a minimum the Clocking Unit functionality must be implemented as part of the TE-CMCTS. A TE-CMCTS has either a [DTI] or a BITS external synchronization interface as defined in Section 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-CMCTS may contain a TEA and TDM interface(s).
- Ethernet Interface - The Ethernet interface will carry IP traffic to and from the managed IP network.

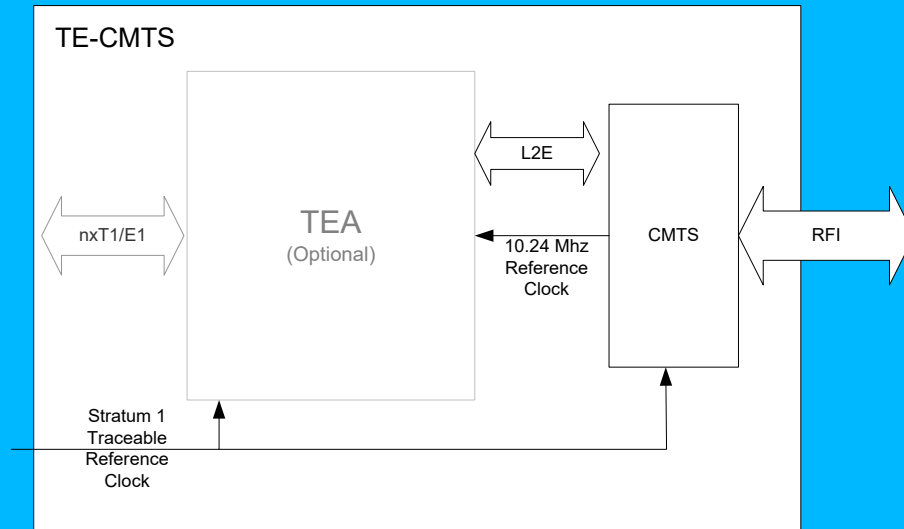


Figure 5-2 - TE-CMTS

L2E is an unspecified Layer 2 Ethernet Interface.

5.3.3 Embedded TDM Emulation Adaptor (eTEA)

An Embedded TDM Emulation Adaptor (eTEA) is an eSAFE device as specified by [eDOCSIS]. An eTEA is embedded in a TE-CM.

An eTEA has a T1, E1, or J1 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 multiple 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 Section 6.4.5;
 - For DOCSIS testing of an eCM;
 - For testing of clock recovery per Section 6.4.7.1;
 - For (out of scope) use to carry TDM data per Section 5.3.1.

- TDM Interface - The TDM interface is one or more T1, E1, J1 or higher rate interface as further defined in Section 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. Section 6.4 describes several clocking modes that may include the CU interfacing with the TDM Service Processor and IWF.

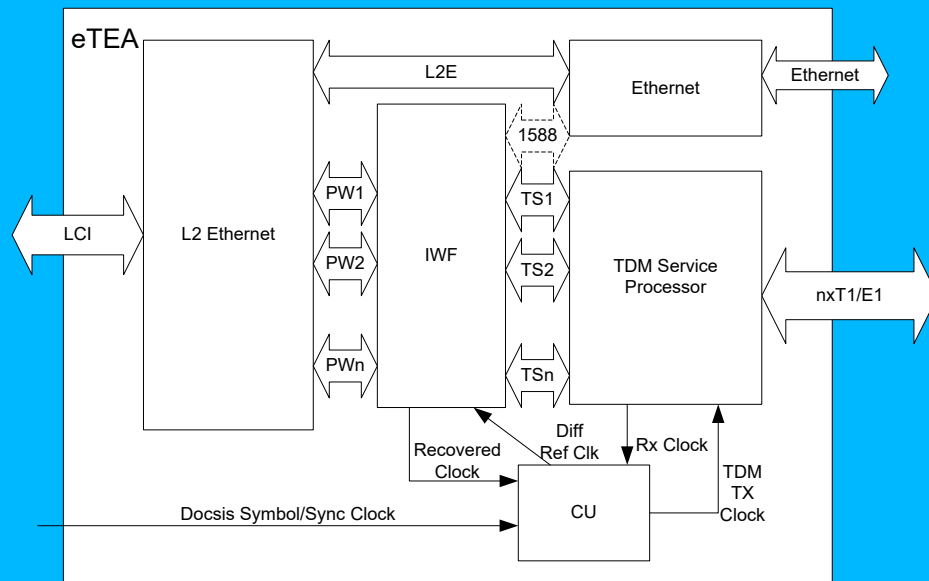


Figure 5-3 - eTEA

5.3.4 TDM Emulation Adaptor (TEA)

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

A TEA has a T1, E1, or J1 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, J1 (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. Section 6.4 describes several clocking modes that may include the CU interfacing with the TDM Service Processor and IWF.

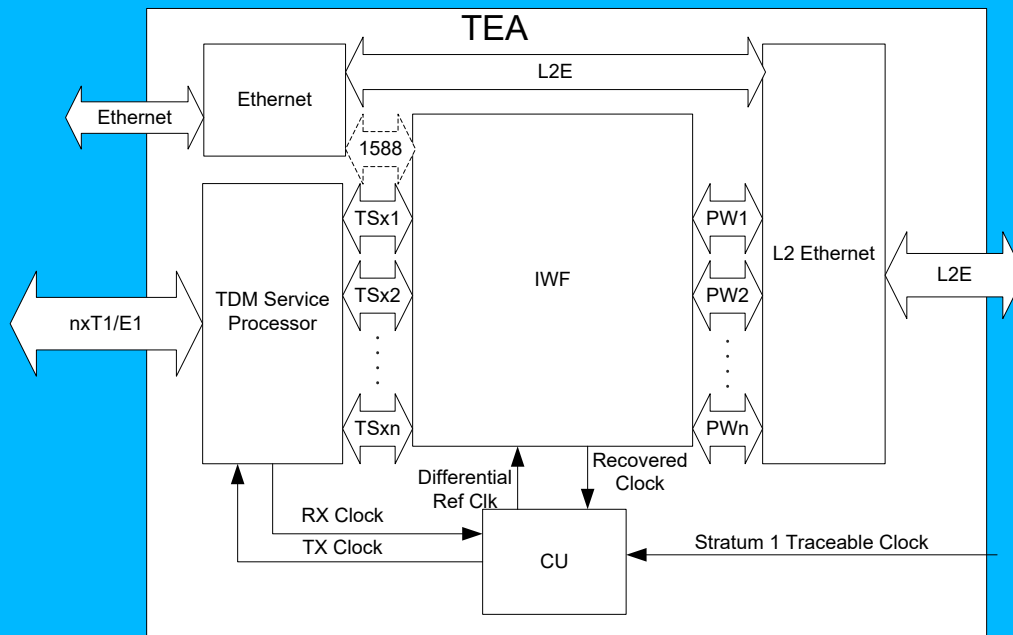


Figure 5-4 - TEA

5.4 Architecture

A BSoD-TE service can be deployed in several architectures. In any architecture the BSoD-TE 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 BSoD-TE 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 BSoD-TE 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 BSoD-TE 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 BSoD-TE 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 BSoD-TE 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 BSoD-TE 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 standard 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 BSoD-TE 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 BSoD-TE circuit emulation in the core. This means there are fewer BSoD-TE 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.

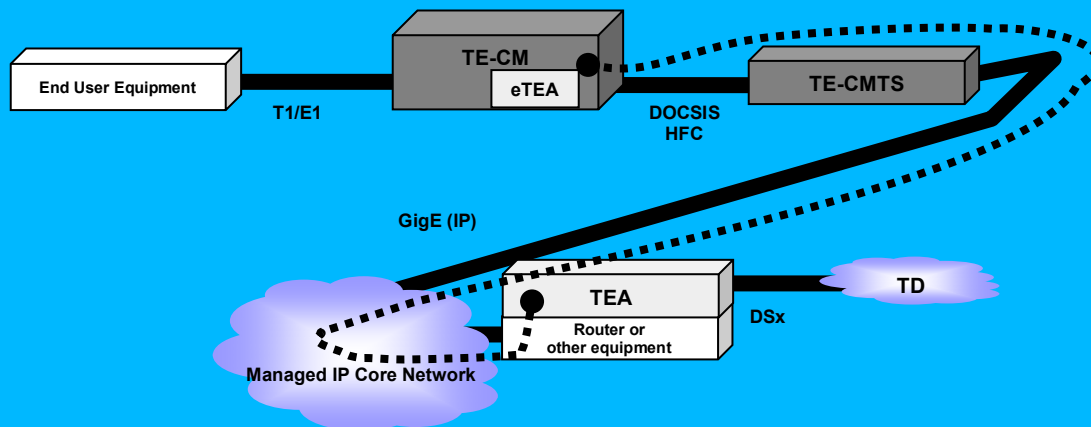


Figure 5-6 - T1 to TDM core, IXC or PSTN

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

In this architecture the BSoD-TE 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.

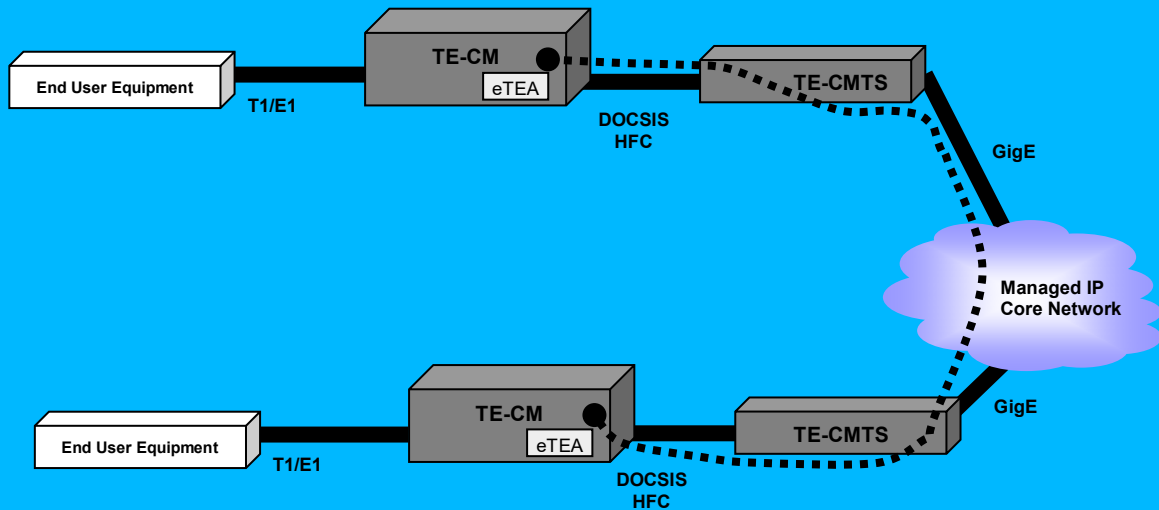


Figure 5-7 - T1 private line between two customer sites

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 Sections 5.4.3 and either Section 5.4.1 or 5.4.2. Figure 5-8 depicts a potential hybrid architecture.

An MSO may also choose to partner with other MSOs creating exchange agreements where a Pseudo Wire may flow over both of their managed IP networks and terminate in a TE-CM in each of their networks. Moreover, a MSO 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 MSO to terminate and aggregate the BSoD-TE circuit emulation for handoff as a high rate TDM interface, like a DS3. This would also be the case if an MSO chooses to offer wireless backhaul services. The MSO would need to aggregate multiple Pseudo Wire IP flows from several TE-CMs, terminate the BSoD-TE 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 standard 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 standard; 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 BSoD-TE service. Since a T1 or E1 service has constant symmetric bandwidth of 1.5M.bits to 2M.bits, 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 BSoD-TE, 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 ([DTI]). This is used for Modular-CMTS operation, but can also be used to support BSoD-TE synchronization requirements as described in Section 6.4.

Existing CMTS architectures may be utilized for BSoD-TE as well as the M-CMTS devices that are inherently compliant. Existing CMTS may or may not have been designed with BSoD-TE in mind and may need some modifications. The potential for a CMTS to be made compliant to BSoD-TE 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 Section 5.4 serves as the basis for explaining how DOCSIS can be used to carry T1/E1 using BSoD-TE. 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's 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 Section 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: [DTI] or BITS.

- DTI - The DOCSIS Timing Interface standard is specified as part of the DOCSIS M-CMTS family. [DTI] 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 BSoD-TE service, but also DOCSIS. The DTI Server being used to time a device containing a TEA must implement the External Reference Mode of [DTI] in order to comply with BSoD-TE. 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+ T11 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 labeled 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 specifications [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 Section 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 Section 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. BSoD-TE 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 dependent 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. UGS is the best method to use in DOCSIS to ensure that the BSoD-TE 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. Section 5.5.2.2.1 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 BSoD-TE) 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 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 BSoD-TE service will function. The MSOs 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 Section 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 rarer 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 document.

6 TECHNICAL SPECIFICATION

6.1 General Requirements

BSoD-TE 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 CU Clock Unit

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 Section 6.4.

6.1.2 PW Pseudo-Wire

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 IWF InterWorking Function

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 msec 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 msec of data at the PW line rate. The IWF should have internal control algorithms that will make sure the buffer operates near the center 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 Section 6.2.

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

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

An IWF **MUST** support TDM Application signaling defined in Section 6.3.2.1.

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

An IWF **MAY** support ACR as defined in Section 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 TSP TDM Service Processor

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 or a J1, it **MUST** comply with:

- [TR-62411];

- [ITU-T G.703];
- [ITU-T G.704];
- [T1-403];
- [ITU-T G.824] as further defined in Table 6–1.

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 Recommendations [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 standard places normative requirements on TDM Interfaces to conform to these ITU Recommendations and Table 6–1 below is required to fully specify which clauses of the ITU Recommendations apply. Table 6–1 is used to specify different requirements for different TDM clocking methods which are defined in the sections cited in the table.

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

- [ITU-T G.824] table 1 for 1544kbit/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 1544kbit/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 2048kbit/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	Section	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] Recommendations SHOULD comply as a 'Synchronization Interface'.

6.1.5 TEA TDM Emulation Adapter

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 (Section 6.4.3), PDCR (Section 6.4.3), or SC (Section 6.1.6).

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 Section 6.2.2.,
- the TEA MUST meet this specification 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 eTEA Embedded TDM Emulation Adaptor

An eTEA MUST meet the normative requirements of a TEA.

6.1.7 TE-CM TDM Emulation Cable Modem

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 [RFI 1.1], DOCSIS 2.0 ([RFI 2.0]) or greater.

A TE-CM MUST contain an eCM as defined in [eDOCSIS].

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

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

A TE-CM MUST be compliant to [CMCI], including an interface with 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 Section 6.4.

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 ([RFI 2.0]) symbol clock or DOCSIS SYNC message and provide it to the TSP and IWF as described in Section 6.4.

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

6.1.8 TE-CMTS TDM Emulation CMTS

A TDM Emulation Cable Modem Termination System (TE-CMTS) is a CMTS (or M-CMTS) with additional functionalities to enable TDM emulation services.

A TE-CMTS MUST support either DOCSIS 1.1, DOCSIS 2.0 ([RFI 2.0]), or greater.

A TE-CMTS MUST contain a CU supporting Direct Clocking mode.

A TE-CMTS must be able to distribute clocking information downstream to the TE-CMs.

- A TE-CMTS SHOULD be able to lock the DOCSIS ([RFI 2.0]) SYNC message to the CU as described in Section 6.4 and keep the DOCSIS SYNC Message compliant with Section 6.4.6.
- A TE-CMTS SHOULD be able to lock the DOCSIS ([RFI 2.0]) symbol clock to the CU as described in Section 6.4 and keep the DOCSIS Symbol Clock compliant with Section 6.4.6.
- A TE-CMTS MUST be able to lock the symbol clock or the SYNC message to the CU as described in this section. See Section 6.4.2.1 for further comments on this requirement.

A TE-CMTS MAY contain a TEA.

The TE-CMTS should provide low latency service to PW service flows in both the upstream and downstream directions. A BSoD-TE system should provide a round-trip latency of less than 14 msec, including the latency of the all equipment involved, including TE-CMTSs, TE-CMs, and TEAs.

6.2 TDM-IP Network Interworking

This section specifies the format of packet headers for PW packets. The format of the payload data for PW packets is specified in Section 6.3.

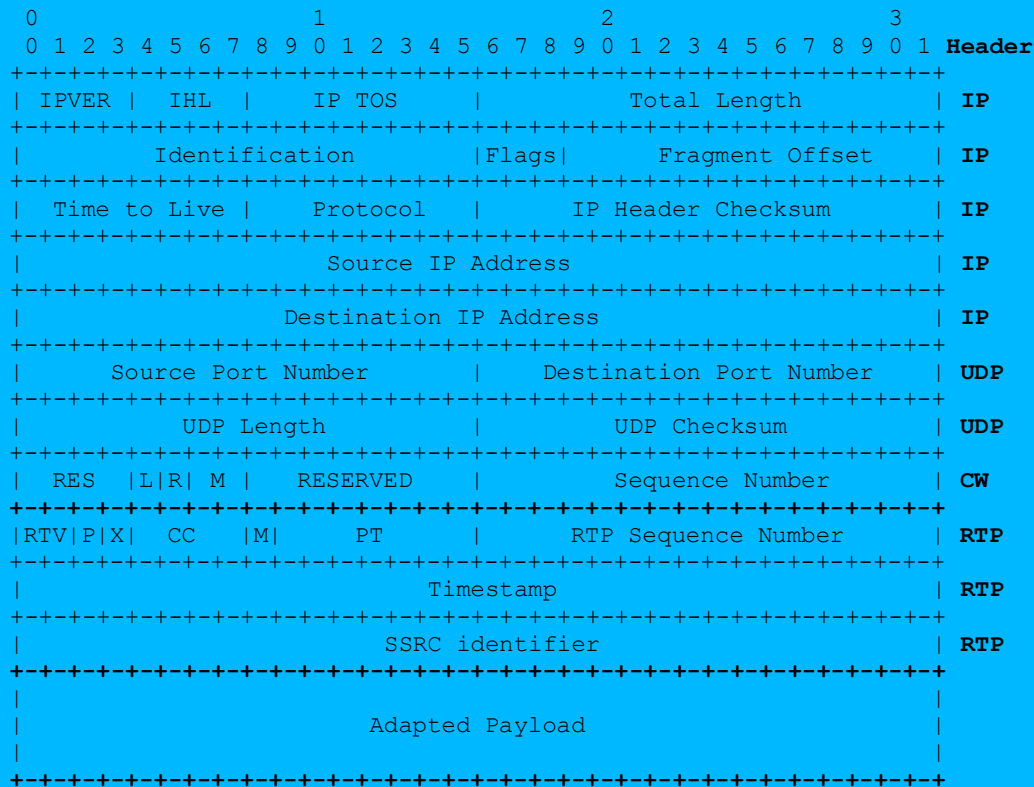


Figure 6–1 - Header Format (Only IPv4 shown)

The first five rows are the IP header (in the case of IPv4), the sixth and seventh rows are the UDP header. Row 8 is the TDMoIP control word. Rows 9 through 11 are the optional RTP header.

An IWF MUST implement the header fields (and payload) shown in Figure 6–1 in the order shown, with the possible exception of the RTP header which is optional for each PW.

6.2.1 IP Header Fields

The IWF MUST support IPv4. The IWF MAY support IPv6.

The IWF MUST be configurable for the IPv type for each PW. The IWF MUST use the same IPv type for transmission and reception of PW flows. An IWF SHOULD use the same IPv type for both PW data flows and management.

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 mis-connections as defined in section 6.6.1 of [MEF8].

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 specifications may differ from this specification 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 standard was written to avoid a problem where optional [RFC 3711] encryption would obliterate the Control Word if implemented. This standard moves the Control Word outside the encryption range of [RFC 3711].

The Control Word for MPLS PWE was moved ahead of the RTP header (in [RFC 4553]) to avoid problems with [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 specification because DOCSIS systems don't use MPLS internally. The reader is referred to section 4.3 of [RFC 4553] and [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 [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 RDaction. This is applicable only to structure aware emulation.

11 - Indicates that the packet contains only the signaling 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. Signaling 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 [RFC 3550]. See Section 6.2.4 below.

6.2.4 RTP Header Fields

An IWF MAY support [RFC 3550]. If the IWF supports PDCR, then it must support RTP per Section 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 [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 Section 6.7.3 and "eTEA PW RTP Payload Type" in Section 6.7.1.10) 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 Section 6.7.3 and "eTEA PW RTP Peer Payload Type" in Section 6.7.1.10). See section 3 of [RFC3551] for more information on assignment of PT values in RTP. The following (decimal) PT value may be used in this specification:

- 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 [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 [RFC 3550] if PDCR is supported.

SSRC (Synchronization Source) Identifier (16 bits). This is the SSRC 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

[RFC 3611] specifies 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 signaling data is in-band, it is possible to make one-way and round trip latency and latency variation measurements.

- An IWF MAY support [RFC 3611].
- An IWF should not support, or should not be configured to use SDP as it is applied in [RFC 3611]. Peer PWs must be configured similarly for inter-operation using [RFC 3611].
- An IWF that supports [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 [RFC 3611] section 4.5, to report latency measurements. Further informative text may be found in [MLDA] and [RFC 3550] Figure 2. The IWF transmitting Block Type 4 MAY substitute any time stamp 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 section 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 signaling packets (per Section 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 paragraphs. 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.
- J1 will only be accommodated using T1 structure agnostic mode. This means fractional J1 will not be supported.

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 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., $K \times 125 \text{ usec}$). The resulting payload format is shown in Figure 6–2 that follows.




Figure 6–2 - Payload Format for Structure-aware transport

6.3.2.1 Signaling Transport

TDM applications exchange signaling 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 signaling. Signaling 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 Signaling (CCS) may be achieved by carrying the signaling channel together with the emulated service (i.e., channel 23 for DS1, or channel 16 for E1). However, Channel Associated Signaling (CAS) (i.e., DS1 Robbed Bit Signaling 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 x 64 kbit/s basic service.

This section describes a generic method for extending the Nx64kbit/s basic service by carrying CAS signaling for each pseudo-wire in separate signaling packets that is independent of the TDM circuit type. It may be used in situations where the individual 64kbit/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 Signaling Frames

The generic format of the PW signaling 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 signaling 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 and TOS bits. This makes it possible to give special treatment to signaling frames.

- M bits (signaling frame have a value of 'b11).
- TOS field - 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 signaling frames.
- If an IWF uses the [RFC 3550] header:
 - The IWF MUST use a payload type value provided from configuration (see Section 6.7.3 and "eTEA PW RTP CAS Signaling Payload Type" in Section 6.7.1.10.1) to set the PT field for associated signaling frames. The IWF MUST reject a configuration that assigns the same payload type value for data frames and signaling frames. The IWF MAY use a configured PT value in incoming packets to detect malformed packets (see Section 6.7.3 and "eTEA PW RTP CAS Signaling Peer Payload Type" in Section 6.7.1.10.1).
 - The IWF MUST assign identical timestamp values for the data and associated signaling frames at any given time.

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

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




Figure 6–3 - Signaling sub-structure

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

- Setup of the PW.
- A change in the signaling state of any of the 64 Kbps 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 signaling change is detected, multiple PW signaling packets are sent. The IWF MUST send a PW signaling packet within *sp_int* ms. The IWF MUST send PW signaling 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 signaling frame every *sp_idle* sec in the absence of any signaling change event described above. The default value of *sp_idle* is 5 seconds. The TSP MUST preserve the ABCD signaling bits at their previous state whenever packets are lost.

6.4 Timing Aspects

The timing aspects of a BSoD-TE 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] specifications.

As required in Sections 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 section 7.2.1 of [DTI] 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 section 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 specification only requires a TE-CMTS to support one method. The [PHY v3.0] standard 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 table is from [DRFI].

Table 6–2 - DOCSIS Symbol Clocks

Downstream mode	Nominal Specified Symbol Rate, fb (MHz)	M/N	Master Clock Rate, fm' (MHz)	Downstream Symbol Rate, fb' (MHz)	Offset from Nominal
64-QAM ([ITU-T J.83] Annex B)	5.056941	401/812	10.239990	5.056945	0.95 ppm
256-QAM ([ITU-T J.83] Annex B)	5.360537	78/149	10.240000	5.360536	0.02 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 SCMDA recovery of 10.24 MHz, in which case the S-CDMA value is to be used. See [RFI 2.0], Table 8-18, 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 [RFI 2.0]. 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 BSoD-TE. 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 Section 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 customer's 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.

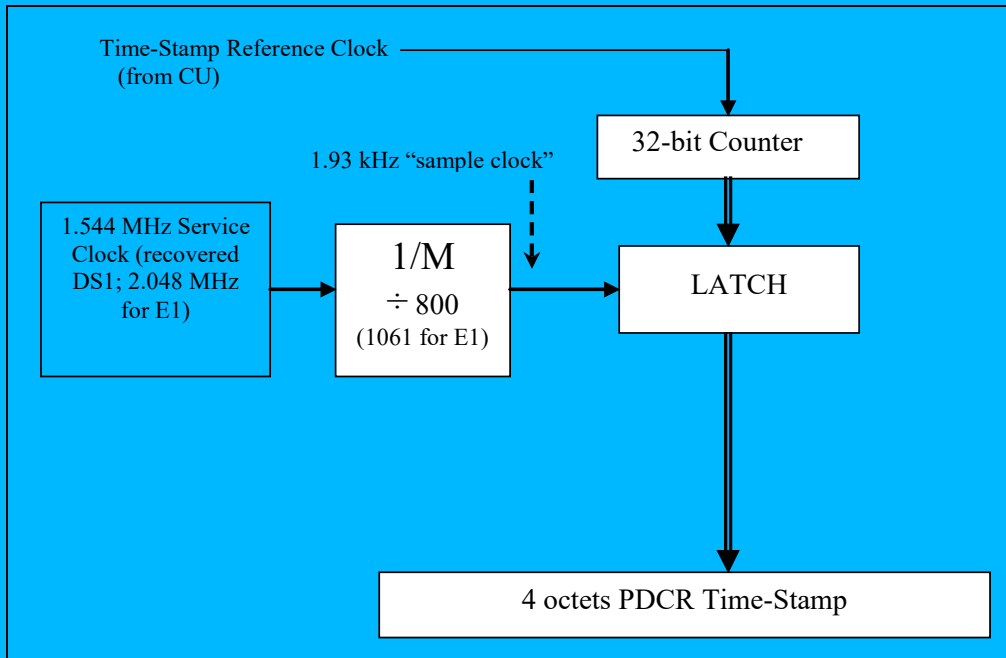


Figure 6–4 - Description of PDCR "transmit" function

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., 1.93kHz for T1 and 2048/1061kHz for E1) and $TSRC$ is the time stamp reference clock (i.e., 10240kHz).

For example, assume the DS1 CES packets are sent at a rate of 1kHz, $N=256$, $SC=1.93\text{kHz}$, $TSRC=10240\text{kHz}$ and the offset of the service clock is 10ppm. Each sample the Timestamp advance 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 100ppm, 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:

19.44MHz 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 24.704MHz clock (16*1.544MHz) and 32.768MHz clock (16*2.048MHz) is left for future study.

A TEA compliant with PDCR MUST be compliant with Section 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 Section 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 Table 6–3, 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 sample ¹

¹ Credit [DTI]

Table 6–3 - MTIE of CMTS Master Clock

MTIE Master Clock	
Observation Time, \mathbf{T} (seconds)	MTIE (nanoseconds)
$0.1 \leq \mathbf{T} < 280$	$30 + 10 * \mathbf{T} + .05 * \mathbf{T}^2 + 8.75e-5 * \mathbf{T}^3$
$280 \leq \mathbf{T}$	$800 + 0.01 * \mathbf{T}$


The TE-CMTS symbol (master) clock output MUST be compliant with Table 6–4, where Time Deviation (TDEV) is defined in [ITU-T G.810].

Table 6–4 - TDEV of CMTS Master Clock

TDEV Master Clock	
Observation Time, \mathbf{T} (seconds)	TDEV (nanoseconds)
$0.1 \leq \mathbf{T} < 200$	6
$200 \leq \mathbf{T} < 1000$	$0.03 * \mathbf{T}$
$1000 \leq \mathbf{T} < 10000$	30

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 section 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 section also defines a model of the intervening network; defines the characteristics of that model; and states the performance requirements in the context of the model characteristics. Figure 6–5 shows the model system within which the clock recovery system must meet performance requirements.

**Figure 6–5 - 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 Sections 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 Section 6.1.4.

6.4.7.2 Network Impairment Model (NIM1)

Two different Network Impairment Models are defined in this specification. This section defines NIM1; NIM2 is defined in Section 6.4.7.3. Either NIM may be used to demonstrate compliance as outlined in Section 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 [ITU-T G.1020] Appendix I and in [TIA NIM]. NIM1 will be configured according to Section 6.4.7.2.1 as configured according to the NIM1 constants (Section 6.4.7.2.2). It is an assumption that any TE-CMTS involved in the transport of IP packets in Figure 6-5 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 modeled 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 %4.1f\tlastd %4.1f \t reorder ", d, lastd);
        }
        else {
            printf("d %4.1f\tlastd %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 */

```

Notes:

1. The FL is used to filter and smooth the delay impulses as they become delays. The term $(ds*(FL - 1) + impulse) / FL$ 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.
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 6–5 serves to define and limit the possible variations of the Gilbert-Elliott model.

The following constants pertain to the Performance requirements in Section 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 6–5. Each NIM1 constant in the pseudo code of Section 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 6–5 - 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 6–6. This set of constants should generate a [ITU-T G.1050] NIM with the characteristics outlined in Table 6–7. The normative requirements are defined by Table 6–6 whereas Table 6–7 is informative only.

Table 6–6 - 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 6–7 - NIM2 Characteristics

NIM2 Characteristic	Value
Latency	16 msec
Jitter	3 msec
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 section 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 Sections 6.4.3 and 6.4.4. Figure 6–6 through Figure 6–8 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 modeled 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 section in the face of network impairments described in Section 6.4.7 and generated by either NIM1 specified in Section 6.4.7.2 or NIM2 specified in Section 6.4.7.3.
- Table 6–8 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 T_{init} seconds during the initialization of the TEA.




Figure 6–6 - Clock recovery initialization time

- Step response settling time - A CRU within the IWF **MUST** be capable of the following clock recovery sequence: settle for T_{init} 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 Section 6.1.4 within T_{settle} seconds after the step input.




Figure 6-7 - Clock recovery settling time

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

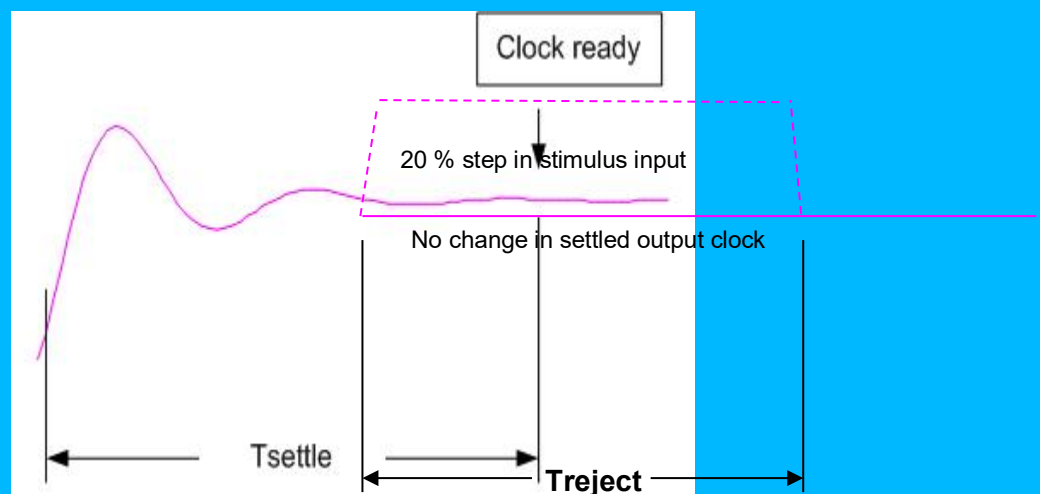


Figure 6-8 - Rejecting impulse noise

- 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 Section 6.4.7.4.1.

Table 6–8 - Clock Recovery Performance Constants

Item	Description	ACR Value	PDCR Value	Units
Tinit	The time between initialization (reset) and clock recovery.	700	700	sec
Tsettle	The settling time of the output frequency.	900	900	sec
Treject	The length of stimulus noise impulse that must be rejected.	1	1	sec
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	Sec
Tloss	The duration of an excessive packet loss condition.	1	1	Sec

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 signaling 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, [RFC 3711] seems to be a very efficient and tight security mechanism for this application.

- The IWF MAY implement [RFC 3711] security for end-to-end security between peer IWFs in a service.
- If an IWF implements [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 [RFC 3711] and is out of scope herein.

6.7 Provisioning and Service Initiation

6.7.1 Provisioning of Embedded TDM Emulation Adaptor (eTEA)

This section describes the Embedded TDM Emulation Adaptor (eTEA), its attributes and its configuration.

6.7.1.1 eTEA Device Provisioning

The eTEA is defined as a eSAFE device as specified by [eDOCSIS]. 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 Section 6.7.1.9.4.

Figure 6–9 illustrates the eTEA provisioning process.

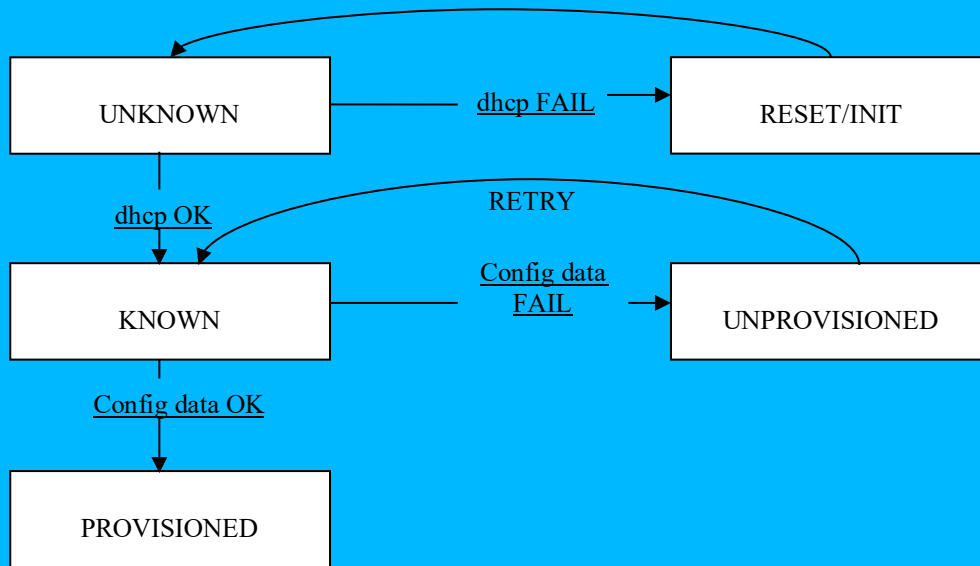


Figure 6–9 - Device States and State Transitions for Basic Flow Provisioning

6.7.1.2 eTEA Endpoint Provisioning

After an eTEA is successfully provisioned, it has to establish transport connectivity with its peer TEA or eTEA to transition to an operational state and ready to carry traffic. Such end-to-end provisioning and monitoring protocol is at the discretion of the operators, and [RFC 5085] may be used for such purpose.

6.7.1.3 Security

An eTEA is provisioned through a DOCSIS-like provisioning flow and should use standard DOCSIS security methods.

6.7.1.4 Backoff, Retries and Timeouts

The backoff mechanisms help the network to throttle device registration during a typical or multiple client registration condition when requests are not serviced within the protocol specified timeout values. This specification does not delve into the details of these requirements but outlines the following as a general guideline:

- The recommendation for the throttling of registration may be based on DOCSIS 1.1 CM registration.
- The eTEA must follow DHCP specifications for the timeout and retry mechanisms.
- The eTEA must use an adaptive timeout for TFTP as specified in the DOCSIS 1.1 specification.

6.7.1.5 Power-on Initialization Flow

The eTEA first obtains its IP address along with other information including its config file name and address, and then download the eTEA config file from the TFTP server. This process is illustrated in Table 6–9 and described in more detail in Table 6–10.

Table 6–9 - eTEA Power-on Initialization Flow

Flow	eTEA	CMTS	DOCSIS DHCP	DOCSIS TFTP	SYSLOG SERVER
eTEA-1			→		
eTEA-2	←				
eTEA-3			→		
eTEA-4	←				
eTEA-5				→	
eTEA-6	←				
eTEA-7					→

Table 6–10 - eTEA Initialization Flow

Step	eTEA Initialization Description	Normal Flow Sequencing	Must proceed to here if this step fails
eTEA-1	DHCP Broadcast DISCOVER The eTEA MUST send a broadcast DHCP DISCOVER message. This message MUST include option code 60 (vendor specific option) in the format "bsod-te1.0:xxxxx". The eTEA MUST include the DHCP option code 43 in the DHCP DISCOVER message as defined in Section 6.7.1.8.3. The eTEA MUST request in DHCP option 55 the following: option 1 (subnet mask). The eTEA MAY request in DHCP option 55 the following: option 7 (syslog server).	eTEA-1 MUST NOT occur before CM registration completion.	eTEA-1
eTEA-2	DHCP OFFER The eTEA MUST only accept a valid DHCP OFFER message. A valid DHCP OFFER MUST include option 1, the configuration file location in the "siaddr" field and the configuration file name in the "file" field.	eTEA-2 MUST occur after eTEA-1 completion.	eTEA-1
eTEA-3	DHCP broadcast REQUEST The eTEA MUST send a DHCP REQUEST after selecting a valid DHCP OFFER.	eTEA-3 MUST occur after eTEA-2 completion.	eTEA-1
eTEA-4	DHCP ACK The DHCP server sends a DHCP ACK message to the eTEA. The DHCP ACK message MUST include all options and sub-options which had been sent in eTEA-2 (DHCP OFFER). If the option and sub-option values of this DHCP ACK differ with the preceding DHCP OFFER (eTEA-2), the option and sub-option values of this DHCP ACK MUST be treated as authoritative (per [RFC 2131]). If the DHCP ACK is not valid as per the criteria established in eTEA-2, the eTEA MUST fail this step.	eTEA-4 MUST occur after eTEA-3 completion.	eTEA-1
eTEA-5	TFTP Configuration File Request The eTEA MUST perform a TFTP protocol exchange to download its configuration file. The 'siaddr' and 'file' fields of the DHCP ACK are used to locate the configuration file. Specific details of the TFTP protocol can be found in [RFC 1350].	eTEA-5 MUST occur after eTEA-4 completion.	eTEA-1
eTEA-6	TFTP Configuration File Response The TFTP server MUST send the requested configuration file to the eTEA. Specific details of the TFTP protocol can be found in [RFC 1350]. The number of TFTP retries is governed per [RFI 2.0].	eTEA-6 MUST occur after eTEA-5 completion.	If the configuration file download failed per TFTP protocol, go to eTEA-1. If the configuration file itself is in error, log the error and end the process.
eTEA-7	SYSLOG NOTIFICATION If the Syslog server IP address has been received, the eTEA MUST send a provisioning complete notification. This should include the PASS/FAIL result of the provisioning operation.	eTEA-7 MUST occur after the completion of eTEA-6.	

The eTEA configuration file should contain all the necessary parameters to provision every entity including the PW, TSP Ethernet and CU, within the eTEA and the eTEA itself. These entities inside eTEA may be enabled or disabled individually via SNMP.

An eCM in an eDOCSIS device implementing an eTEA supports the eSAFE-MIB as indicated in [eDOCSIS].

6.7.1.5.1 eSAFE MIB esafeProvisioningStatusTable

Support for the esafeProvisioningStatusTable are identified in this section.

An eCM in an eDOCSIS device implementing an eTEA MUST report its current provisioning state as indicated in the eSAFE MIB object esafeProvisioningStatusProgress.

An eCM in an eDOCSIS device implementing an eTEA MUST report any provisioning errors as indicated in the eSAFE MIB object esafeProvisioningStatusFailureFound.

An eCM in an eDOCSIS device implementing an eTEA MUST support the eSAFE MIB object esafeProvisioningStatusFailureFlow with the flow steps defined in Table 6–10 'Step' column.

An eCM in an eDOCSIS device implementing an eTEA MUST report the Event ID value from the 'Event ID' column of the DOCSIS Events Extensions table Table A–1 as indicated in the eSAFE MIB object esafeProvisioningStatusFailureEventID.

An eCM in an eDOCSIS device implementing an eTEA MUST report the value from the 'Event Message' column of the DOCSIS Events Extensions table Table A–1 as indicated in the eSAFE MIB object esafeProvisioningStatusFailureErrorText.

6.7.1.5.2 eSAFE MIB esafeDevStatusTable

Support for the esafeDevStatusTable are identified in this section.

An eCM in an eDOCSIS device implementing an eTEA MUST report a value of 'significant' for the eSAFE MIB object esafeDevServiceIntImpact.

An eCM in an eDOCSIS device implementing an eTEA MUST support the esafeDevServiceIntImpactInfo object as specified in the eSAFE MIB. An eCM in an eDOCSIS device implementing an eTEA MAY report a vendor-specified string for the eSAFE MIB object esafeDevServiceIntImpactInfo.

6.7.1.6 Behavior during a disconnected state

Changes in the state of the connection can be reported by each eTEA.

- The eTEA MUST define 'starting' and 'stopping' for Signaling and Data Trap activity over a configurable interval with a default of 30 seconds.
- The eTEA MUST limit the combined TDM Signaling and TDM Data Traps to a configurable number in any one hour with a default setting of 24. The eTEA should not control the rate until the limit is reached; the total sent may be reset each hour.
- TDM Signaling Traps -
 - When a TEA senses that its peer has stopped sending TDM Signaling packets, the TEA SHOULD send a TDM Signal Link Down Trap.
 - When a TEA senses that its peer has started sending TDM Signaling packets, the TEA SHOULD send a TDM Signal Link Up Trap.
- TDM Data Traps -
 - When a TEA senses that its peer has stopped sending TDM Data packets, the TEA SHOULD send a TDM Data Link Down Trap.
 - When a TEA senses that its peer has started sending TDM Data packets, the TEA SHOULD send a TDM Data Link Up Trap.

6.7.1.7 Temporary Signal Loss

If the eCM or DOCSIS reset for any reason, the eTEA MUST reset and reinitialize, this will result in the TDM service disruption till the configuration parameters and the operational status of the eTEA is restored.

6.7.1.8 DHCP Options

6.7.1.8.1 Standard DHCP Options

DHCP options 1,2,3,4 and 7 are defined in [RFC 2132].

- Option code 1 (Subnet Mask).
- Option code 2 (Time Offset).
- Option code 3 (Router Option/Default Gateway).
- Option code 4 (Time Server Option).
- Option code 7 (Log Server Option).

6.7.1.8.2 DHCP Option 60: Vendor Client Identifier

Option code 60 contains a string identifying Capabilities of the eTEA. The eTEA MUST send the following ASCII Coded String in DHCP Option code 60: "bsod-te1.0:zzzzz". Where zzzzzz MUST be an ASCII representation of the hexadecimal encoding of the eTEA TLV Encoded Capabilities, as defined in Section 6.7.1.11.

6.7.1.8.3 DHCP Option 43

The eTEA MUST send the DHCP Option 43 in the DHCP DISCOVER and DHCP REQUEST.

DHCP Option 43 contains the number of sub-options defined to provide the eTEA device specific information to the back-office systems. The DHCP option 43 sub-options 1 through 10 and 31 are specified by BSoD-TE. The BSoD-TE DHCP option 43 sub-options MUST be present in the format of "Encapsulated vendor-specific extensions" ([RFC 2131]).

The following table contains the sub-options of the DHCP Option-43, which the eTEA MUST use. The eTEA MUST send all required sub-options listed in the table below unless explicitly stated otherwise. If the total number of octets in all DHCP option 43 sub-options exceeds 255 octets, the eTEA MUST follow [RFC 3396] to split the option into multiple smaller options.

Table 6–11 - DHCP Option 43 Syntax

eTEA DHCP Option 43	Required / Not used	Value	Description
Sub-option 1	Not Used		The request sub-option vector is a list of sub-options (within option 43) to be returned to client by the server upon reply to the request. None defined. The DHCP option 43 sub-option 1 MUST NOT be used by the eTEA.
Sub-option 2	Required	"ETEA"	The sub-option 2 contains the device type of the component making the DHCP request. The eTEA MUST send the DHCP option 43 sub-option 2. For BSoD-TE, this is: "ETEA"= embedded TEA.
Sub-option 3	Not Used		The sub-option 3 contains a colon separated list of all components in the eDOCSIS device. It is used by the eDOCSIS eCM device. The DHCP option 43 sub-option 3 MUST NOT be sent by the eTEA.
Sub-option 4	Required	<device serial number>	The sub-option 4 contains the device serial number represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 4. The DHCP option 43 sub-option 4 value MUST be identical to the value of the docsDevSerialNumber MIB object [eDOCSIS].
Sub-option 5	Required	<Hardware version>	The sub-option 5 contains the hardware version number represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 5. The DHCP option 43 sub-option 5 value MUST be identical to the <Hardware version> field in MIB II object sysDescr.

eTEA DHCP Option 43	Required / Not used	Value	Description
Sub-option 6	Required	<Software version>	The sub-option 6 contains the software version number represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 6. The DHCP option 43 sub-option 6 value MUST be identical to the <Software version> field in MIB II object sysDescr.
Sub-option 7	Required	<Boot ROM version>	The sub-option 7 contains the Boot ROM Version represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 7. The DHCP option 43 sub-option 7 value MUST be identical to the <Boot ROM version> field in MIB II object sysDescr.
Sub-option 8	Required	<OUI>	The sub-option 8 contains the Organizational Unique Identifier (OUI) represented as a hexadecimal-encoded 3-byte octet string. The eTEA MUST send the DHCP option 43 suboption 8. If omitted, the DHCP Server SHOULD use the eTEA MAC address as the eTEA OUI.
Sub-option 9	Required	<Model number>	The sub-option 9 contains the eTEA Device Model Number represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 9. The DHCP option 43 sub-option 9 value MUST be identical to <Model Number> field in the MIB-II object sysDescr.
Sub-option 10	Required	<Vendor name>	The sub-option 10 contains the Vendor Name represented as an ASCII string. The eTEA MUST send the DHCP option 43 suboption 10. The DHCP option 43 sub-option 10 value MUST be identical to <Vendor Name> field in the MIB-II object sysDescr.
Sub-option 31	Required	<eTEA MAC Address>	The sub-option 31 contains the eTEA MAC Address encoded as a 6 byte octet string. The eTEA MUST send the DHCP option 43 suboption 31. The DHCP option 43 sub-option 31 value MUST be identical to the content of the MacAddress MIB object.

6.7.1.9 eTEA Configuration File

The eTEA configuration file follows the CM configuration format specified in [RFI 2.0].

6.7.1.9.1 eTEA Binary Configuration File Format

The eTEA-specific configuration data MUST be contained in a file which is downloaded to the eTEA via TFTP. This is a binary file in the same format defined for DHCP vendor extension data [RFC 2132].

The Configuration File MUST consist of a number of configuration settings, each of the form "Type Length Value (TLV)". Definitions of these terms are provided in Table 6–12.

Table 6–12 - TLV Definitions

Type	A single-octet identifier, which defines the parameter.
Length	A single-octet field specifying the length of the value field (not including Type and Length fields).
Value	A set of octets Length long containing the specific value for the parameter.

The configuration settings MUST follow each other directly in the file, which is a stream of octets (no record markers).

6.7.1.9.2 eTEA Configuration File Settings

An eTEA MUST be capable of processing all standard configuration settings. An eTEA MUST ignore any configuration setting present in the configuration file which it cannot interpret.

An eTEA configuration file MUST include:

- the eTEA Message Integrity Check (MIC) (Type 53).
- the End of Data Marker (Type 255).

An eTEA configuration file MAY include the other TLVs defined in Section 6.7.1.10.

Authentication of the provisioning information is provided by a message integrity check (MIC), eTEA MIC. It is a digest which ensures that the data sent from the provisioning server were not modified en route. This is NOT an authenticated digest (it does not include any shared secret).

6.7.1.9.3 Configuration File Creation

The sequence of operations required to create the configuration file is as follows:

1. Create the type/length/value entries for all the parameters required by the eTEA.
2. Calculate the eTEA message integrity check (MIC) configuration setting as defined in Section 6.7.1.9.4 and add to the file following the last parameter using code and length values defined for this field.
3. Add the end of data marker.

6.7.1.9.4 eTEA MIC Calculation

The eTEA message integrity check configuration setting MUST be calculated by performing an MD5 digest over the bytes of the configuration setting fields. It is calculated over the bytes of these settings as they appear in the TFTP image, without regard to TLV ordering or contents. There is one exception to this disregard of the contents of the TFTP image:

The bytes of the eTEA MIC TLV itself are omitted from the calculation. This includes the type, length, and value fields.

On receipt of a configuration file, the eTEA MUST recompute the digest and compare it to the eTEA MIC configuration setting in the file. If the digests do not match then the configuration file MUST be discarded.

6.7.1.10 eTEA Provisionable Attributes

The eTEA provisionable attributes are defined as TLVs in the following subsection.

6.7.1.10.1 Configuration File Format Requirements

The configuration settings MUST follow each other directly in the file, which is a stream of octets. The eTEA MUST be capable of properly receiving and processing the configuration file.

An eTEA configuration file contains the parameters to configure in the TDM interface using TLV 11 on the standard DS1 and DS0 MIBs, configuration of the cable modem via TLV8, and configuration of the PWs using TLV 9 and 10.

An eTEA configuration file MUST contain the eTEA Message Integrity Check (MIC) (Type 53) and the End of Data Marker (Type 255).

eTEA Symbol Clock

This object configures the parameters for dividing down the reference clock and contains the M, N values.

Type	Length	Value
8	4	m,n

This TLV contains the values for the M/N PLL.

eTEA IWF configuration encoding

This object configures the Interworking function of the eTEA. It consists of a number of encapsulated type/length/value fields. The encapsulated fields define the individual parameters for the IWF being configured. Within the IWF, each supported PW is configured separately as shown below. Parameters that are typically configured identically for many PWs are grouped together in TLV type 10.

Type	Length	Value
9	n	

eTEA PW Index setting

This object configures the index number of the PW. The settings for each individual PW are preceded by explicitly setting the index number of the PW.

Type	Length	Value
9.1	2	<PW Index>

eTEA PW Emulation Type

This object configures the emulation type of the PW. The default value is t1Satop(17).

Type	Length	Value
9.2	1	<e1Satop(17), t1Satop(18), basicCesPsn(21), tdmCasCesPsn(23)>

eTEA PW Peer Address

This object configures the IPv4 address used by the PW as the destination IP address of the peer.

Type	Length	Value
9.3	4	ip1.ip2.ip3.ip4

eTEA PW Peer Address IPv6

This object configures the peer destination IP address when IPv6 is used. This object is required only when IPv6 is used.

Type	Length	Value
9.4	16	ip1....ip16

eTEA PW Destination Port

This object configures the destination UDP port to be used by the currently indexed PW. The UDP port number is to be chosen from the range of dynamically allocated UDP ports numbers (49152 through 65535). For a given IWF, a unique port number should be assigned for each PW.

Type	Length	Value
9.5	2	<Valid port number>

eTEA PW Peer Destination Port

This object configures the UDP port used by the peer PW for the currently indexed PW. The UDP port number is to be chosen from the range of dynamically allocated UDP ports numbers (49152 through 65535). This value may also be used as the source UDP port for packets sent to the Peer.

Type	Length	Value
9.6	2	<Valid port number>

eTEA PW Name

This object stores the canonical name assigned to the PW.

Type	Length	Value
9.7	1-31	<string>

eTEA PW Description

This object initializes textual string containing information about the PW. This is an optional parameter. If there is no description this object contains a zero length string.

Type	Length	Value
9.8	1-31	<string>

eTEA PW Admin Status

This object is used to alter the administrative status of this PW. Testing is used to route the packet side of the PW to the externally accessible Ethernet port, as opposed to the RFI port. This is useful for testing PW operations such as the clock recovery without an intervening DOCSIS network. In some implementations this may be a global parameter. Application loopback is used to route packets on this PW to itself so that the PW can be looped back on itself. This is useful for testing the TDM interface through the internal data path of the TEA.

Type	Length	Value
9.9	2	up(1),down(2),testing(3), app_loopback(4)

eTEA PW Status Change Notification Enable

This object configures whether notifications are generated on changes in the status of this PW from up to down or down to up. The default value is Disabled.

Type	Length	Value
9.10	1	Disable (0), Enable (1)

eTEA PW TDM Type

This object configures the type of TDM link being emulated. When a line type with CAS is selected, this enables the CAS application signaling as described in Section 6.3.2.2.

Type	Length	Value
9.11	1	other(1), ds1(2),e1(3), ds3(4),e3(5),octetAlignedT1(6) ,nXds0(7),nXds0WithCASe1(8), nXds0WithCASds1Esf(9),nXds0WithCASds1Sf(10) Note: J1 interfaces will use a value of (1) other.

eTEA PW TDM Configuration Table Index

This object configures the index to the TDM configuration table. The TDM configuration table contains parameters that are typically common across many PWs of an IWF in an eTEA. This is described by TLV 10 below.

Type	Length	Value
9.12	2	index of PW TDM configuration table

eTEA PW CESoPSN Configuration Table Index

This object configures the index to the CESoPSN configuration table. This is described by TLV 13 below.

Type	Length	Value
9.13	2	index of CESoPSN specific configuration table

eTEA PW RTP SSRC

In the case where RTP headers are being used on this PW, this object configures the value of the SSRC field to use in the RTP header.

Type	Length	Value
9.14	4	<SSRC value>

eTEA PW RTP Peer SSRC

In the case where RTP headers are being used on this PW, this object configures the value of the SSRC field to expect in the RTP header. When this value is non-zero it may be used to check incoming data frames for validity. Default: 0.

Type	Length	Value
9.15	4	<SSRC value>

eTEA PW TDM Circuit Map

This object configures the TDM circuit transported by this PW. This TLV includes two suboptions to declare the port number and the timeslots. For Structure Agnostic operation it configures the port transported by this PW. For Structure Aware Emulation the 64K timeslots transported by this PW are configured. It is possible for a PW to support timeslots from multiple ports.

Type	Length	Value
9.16	N	

eTEA PW TDM Port Number

This object references the physical E1/T1 port number. The default value is 1, such that in implementations supporting a single port this parameter is not required.

Type	Length	Value
9.16.1	1	Port number (1-255)

eTEA PW TDM Timeslot Map

This object selects the timeslots transported by this PW.

Type	Length	Value
9.16.2	4	Selected timeslots. 32-bit bit map for E1, 24 bit for T1

Example: The value 0x00000605 configures time slots 1, 4, 10, 11; a value of 0x00ffffff configures all 24 timeslots of a T1 interface.

eTEA PW TDM Configuration Table

This object configures the common parameters of a Pseudo Wire configured for TDM Emulation. Table entries may be indexed by multiple Pseudo Wires.

Type	Length	Value
10	n	

eTEA PW TDM Configuration Table Index

This object configures the current entry of the PW TDM configuration table.

Type	Length	Value
10.1	2	<TDM config table index>

eTEA PW TDM Payload Size

This object configures the current payload size used for this PW. In Structure Agnostic transport the default value for this parameter is 192 bytes for T1 and 256 bytes for E1. For structure aware transport, this value is the number of timeslots times the number of frames in the payload.

Type	Length	Value
10.2	2	Payload size

eTEA PW RTP header Used

This optional object configures the use of RTP header.

Type	Length	Value
10.3	1	true(1), false(2)

eTEA jitter buffer setting

This object configures the jitter configuration and is an estimate of maximum jitter in a given network. The jitter buffer is used to counter some of the network congestion and managed network changes which may cause delay variations in the packet arrival rate.

Type	Length	Value
10.5	2	<Value in msec>

eTEA PW Payload Suppression

This optional object configures the PW to not send out payload bytes when the L bit is set. The L bit indicates that there is a problem at the TDM interface. Default value is disabled (2).

Type	Length	Value
10.6	1	enable(1), disable(2)

eTEA PW LOPS Exit Criteria

The required number of packets with sequential sequence numbers that are available to be played out of the jitter buffer in order to exit the LOPS state. If set to 0, the LOPS state is exited when the excessive packet loss rate falls below the excessive packet threshold value for a complete window time (AvePktLossTimeWindow).

Type	Length	Value
10.7	2	# consecutive frames

eTEA PW LOPS Entrance Criteria

This object configures the criteria for entering the Loss of Packet Synchronization state. Indicates the number of consecutive packets that are late or missing on playout to be considered the criteria for entering the LOPS state. When the parameter is set to 0, the LOPS state is entered when the excessive packet loss threshold has been reached. Default value is 0.

Type	Length	Value
10.8	2	# consecutive frames

eTEA PW Packet Replace Policy

This parameter determines the information played out on the TDM interface instead of a lost packet assuming the pseudo wire is not in loss of packet synchronization state. Default value is 0.

Type	Length	Value
10.10	1	filler (0), AIS (1), other (2)

eTEA PW Packet Loss Window

This object configures the size of the Window to be used for calculating the packet loss rate.

Type	Length	Value
10.11	2	Ms

eTEA PW Excessive Loss Threshold

This object configures the % threshold to be considered excessive. A second in which this is counted is considered an error second.

Type	Length	Value
10.12	1	%

eTEA PW Severe Loss Threshold

This object configures the % threshold to be considered severe. A second in which this is counted is considered a severe error second.

Type	Length	Value
10.15	1	%

eTEA PW RTP Timestamp Mode

This parameter determines the timestamp generation mode. Default value is 1.

Type	Length	Value
10.16	1	absolute (1), differential (2), prime (3), other(4)

eTEA PW Default Fill Pattern

This object configures the fill pattern to play out on the TDM interface when a packet is not present to be played out and the replace policy is to use the configured filler byte.

Type	Length	Value
10.17	1	Fill (default 0x00)

eTEA PW L Flag Payload Policy

This object determines whether the payload received in packets.

With L bit set should be ignored or used. Default value is ignore(0).

Type	Length	Value
10.18	1	ignore(0), use(1)

eTEA PW TOS

This object configures the IP TOS value used for packets sent on this PW. The default value is 184 (EF class).

Type	Length	Value
10.19	1	TOS value

eTEA PW RTP Payload Type

This object configures the payload type to use in the RTP header of the data frames.

Type	Length	Value
10.20	1	PT value

eTEA PW RTP Peer Payload Type

This object configures the payload type that the peer uses in the RTP header of the data frames. When this value is non-zero it may be used to check incoming data frames for validity. Default: 0.

Type	Length	Value
10.21	1	PT value

eTEA PW RTP Timestamp Reference

This object configures the timestamp reference to use. This is the number clocks in a 125 usec period. For example the following frequencies are represented as follows: 8k(1), 32k(4), 256k(64), 1.544M(193), 2.048(256), 10.24(1280), 19.44 (2430).

Type	Length	Value
10.22	2	TS rate

eTEA PW RTP Peer Timestamp Reference

This object configures the timestamp reference use by peer. This is the number clocks in a 125 usec period. For example the following frequencies are represented as follows: 8k(1), 32k(4), 256k(64), 1.544M(193), 2.048(256), 10.24(1280), 19.44 (2430).

Type	Length	Value
10.23	2	TS rate

eTEA PW SRTP Enable

This optional object configures the use of SRTP.

Type	Length	Value
10.24	1	true(1), false(2)

eTEA PW TDM CESoPSN Configuration Table

This object configures the common parameters of a Pseudo Wire configured for CESoPSN. Table entries may be indexed by multiple Pseudo Wires.

Type	Length	Value
13	N	

eTEA PW CESoPSN Configuration Table Index

This object configures the current entry of the PW CESoPSN configuration table.

Type	Length	Value
13.1	2	<CESoPSN config table index>

eTEA PW CESoPSN Default Idle Pattern

This object configures the default Idle Pattern played out on the TDM interface on various defects, according to policy, including receipt of L bit, packet loss and when the pseudo wire is administratively disabled.

Type	Length	Value
13.2	1	idle (default 0xFF)

eTEA PW CESoPSN L Flag Policy

This object configures the action to perform on the TDM interface when packets are received with the L flag set. Default value is idle(0).

Type	Length	Value
13.3	1	idle(0), TrunkAIS(1), idle_code(2)

eTEA PW CESoPSN R Flag Policy

This object configures the action to perform, if any, on the TDM interface when packets are received with the R flag set. For Structure Agnostic applications this parameter should be set to 0. Default value is (0).

Type	Length	Value
13.4	1	nothing(0), RAI(1), idle_code(2)

eTEA PW CESoPSN Remote Defect Policy

This object configures the action to perform on the TDM interface when packets are received with the L flag as 0 and the M bits as 01 indicating a Remote Defect. Default value is 0.

Type	Length	Value
13.5	1	Nothing(0), RAI(1), Channel Idle(2)

eTEA PW CESoPSN LOPS Policy

This parameter determines the information played out on the TDM interface while in the LOPS state. This determines if the idle pattern or AIS pattern is played out. For structured aware TDM emulation 'channel idle' indication on the trunk interface can also be specified.

default: idle(1)

Type	Length	Value
13.6	1	idle (1), TrunkAIS(2), Channel Idle(3)

eTEA PW Application Signaling TOS

This object configures the value to use in the TOS field when configured to send the CAS signaling in a separate packet stream. The default value is 0.

Type	Length	Value
13.7	1	TOS value

eTEA PW RTP CAS Signaling Payload Type

This object configures the payload type to use in the RTP header of the signaling frames.

Type	Length	Value
13.8	1	PT value

eTEA PW RTP CAS Signaling Peer Payload Type

This object configures the payload type to use in the RTP header of the signaling frames. When this value is non-zero it may be used to check incoming data frames for validity. Default: 0.

Type	Length	Value
13.9	1	PT value

eTEA PW Application Signaling Idle

This object configures the CAS signaling Idle pattern. This is the pattern sent when a failure is detected on the TDM interface, including multi-frame failures. This may need to be different per PW. Default 0x0F.

Type	Length	Value
13.10	1	Idle value

eTEA PW Application Signaling Interval

This object configures the application signaling interval to use when sending the CAS signaling in a separate packet stream. The default value is 5 ms.

Type	Length	Value
13.11	1	Ms

eTEA PW Application Maximum Signaling Interval

This object configures the application signaling time in which to refresh the peer when no application signaling has changed in this interval. This is only valid when configured to send the CAS signaling in a separate packet stream. The default value is 5 seconds.

Type	Length	Value
13.12	1	S

eTEA dsx1 configuration table

This object configures the parameters of each dsx1 TDM interface. The parameters correspond to parameters in RFC3895. Status of the interface may be read out using SNMP gets.

Type	Length	Value
14	n	

eTEA dsx1 Port ID

This object configures the PORT ID for the physical dsx1 TDM port that is being configured.

Type	Length	Value
14.1	1	<Port ID>

eTEA dsx1 Line Type

This object configures the Line Type for the physical TDM port that is being configured.

Type	Length	Value
14.2	1	other(1), dsx1ESF(2), dsx1D4(3), dsx1E1(4), dsx1E1CRC(5), dsx1E1MF(6), dsx1E1CRCMF(7), dsx1Unframed(8), dsx1E1Unframed(9)

eTEA dsx1 Line Coding

This object configures the Line Coding for the physical TDM port that is being configured.

Type	Length	Value
14.3	1	dsx1JBZS (1), dsx1B8ZS (2), dsx1HDB3 (3), dsx1ZBTSI (4), dsx1AMI (5), other(6), dsx1B6ZS(7)

eTEA dsx1 Circuit ID

This object configures the transmission vendor's circuit identifier to facilitating troubleshooting. This identifier should include the physical port ID.

Type	Length	Value
14.4	1-255	string

eTEA dsx1 Loopback Configuration

This object configures the Loopback configuration for the physical TDM port that is being configured. This is for diagnostic and commissioning purposes.

Type	Length	Value
14.5	1	dsx1NoLoop(1), dsx1PayloadLoop(2), dsx1LineLoop(3), dsx1OtherLoop(4), dsx1InwardLoop(5), dsx1DualLoop(6)

eTEA dsx1 Signal Mode

This object configures the signaling mode used for DS0s of this physical port. The default value is none(1).

Type	Length	Value
14.6	1	none (1), robbedBit (2), (CAS Signaling for T1) bitOriented (3), (CAS Signaling for E1) messageOriented (4), (CCS Signaling) other (5)

eTEA dsx1 Transmit Clock Source

This object configures the clock source for this physical port. Specification of the clock source includes the mode that this port works in along with a reference to the source of the clock, if necessary.

Type	Length	Value
14.7	n	

eTEA dsx1 Transmit Clock Source Mode

This object configures the clock mode for this physical port. This suboption must always be present. Other suboptions are optional as described below. All PWs attached to this port work in this timing mode.

Type	Length	Value
14.7.1	1	loopTiming(1), localTiming(2), throughTiming(3), adaptive (4)

When working in looptiming, the clock source is taken from the received TDM clock of this port. When working in local timing, the clock source is taken from an unspecified internal clock source. For throughTiming, the clock source is taken from the cable modem interface unless otherwise specified with TLVs 14.7.2 and 14.7.3. For adaptive timing, the clock source is taken from the PW attached to this port. If multiple PWs are associated with this port - the clock comes from the PW attached to the lowest numbered DS0 of this port, unless otherwise specified with TLV 14.7.4.

eTEA dsx1 Transmit Clock Source Interface Type

This object configures the interface type. The default is 0.

Type	Length	Value
14.7.2	1	0 Cable Modem port, 1 Ethernet port, 2 e1/t1 port

eTEA dsx1 Transmit Clock Source Port Number

This object configures the index of the clock source. The default is 1.

Type	Length	Value
14.7.3	1	Port Number

eTEA dsx1 Transmit Clock Source PW Index

This object configures the PW index of the adaptive timing source.

Type	Length	Value
14.7.4	2	PW Index

eTEA dsx1 Fdl

This object configures the set of capabilities desired the facilities data link on this port. This parameter is only relevant for T1 operation. For structure agnostic emulation, all values can be supported. For structure aware emulation, only a value of dsx1FdlNone is supported.

Type	Length	Value
14.8	1	other(1), dsx1AnsiT1403(2), dsx1Att54016(4), dsx1FdlNone(8)

eTEA dsx1 Line Length

The length of the ds1 line in meters. This objects provides information for line build out circuitry. This object is only useful if the interface has configurable line build out circuitry.

Type	Length	Value
14.9	2	Line length (Meters)

eTEA dsx1 Line Status Trap Enable

Indicates whether a change in dsx1LineStatus should generate an snmp trap.

Type	Length	Value
14.10	1	enabled(1), disabled(2)

eTEA dsx1 Channelization

Indicates whether this DS1 is channelized into separate DS0s.

Type	Length	Value
14.11	1	disabled(1), enabledDS0(2)

eTEA dsx1 Line Mode

Indicates whether this DS1 is working as a long or short haul. This is only applicable for T1 and may be useful for configuring the dsx1 interface.

Type	Length	Value
14.12	1	csu(1), (Long Haul) dsu(2) (short Haul)

eTEA dsx1 Line Build Out

Indicates the required line build out for this interface port. This is only applicable for T1 and may be useful for configuring the dsx1 interface.

Type	Length	Value
14.13	1	notApplicable (1), neg75dB (2), neg15dB (3), neg225dB (4), zerodB (5)

eTEA SNMP MIB Object

This object allows arbitrary SNMP MIB objects to be Set via the TFTP-Registration process.

Type	Length	Value
11	n	variable binding

The value is an SNMP VarBind as defined in [RFC 1157]. The VarBind is encoded in ASN.1 Basic Encoding Rules, just as it would be if part of an SNMP Set Request.

The cable modem MUST treat this object as if it were a part of an SNMP Set Request with the following caveats:

- It MUST treat the request as fully authorized (it cannot refuse the request for lack of privilege).
- SNMP Write-Control provisions do not apply.
- No SNMP response needs to be generated.

This object MAY be repeated with different VarBinds to "Set" a number of MIB object. All such SETS MUST be treated as simultaneous.

Each VarBind MUST be limited to 255 bytes.

eTEA SNMP Write-Access Control

This object makes it possible to disable SNMP "Set" access to individual MIB objects. Each instance of this object controls access to all of the writeable MIB objects whose Object ID (OID) prefix matches. This object may be repeated to disable access to any number of MIB objects.

Type	Length	Value
12	n	OID prefix plus control flag

Where n is the size of the ASN.1 Basic Encoding Rules [ISO 8825] encoding of the OID prefix plus one byte for the control flag.

The control flag may take values:

- 0 - allow write access.
- 1 - disallow write access.

Any OID prefix may be used. The Null OID 0.0 may be used to control access to all MIB objects. When multiple instances of this object are present and overlap, the longest (most specific) prefix has precedence.

Example:

- anyTable disallow write-access.
- anyTable.1.3 allow write-access.

eTEA Vendor Specific Extensions

This object allows vendor specific extensions to the eTEA. This MUST include the Vendor ID. The Vendor ID MUST be the first TLV. The eTEA MUST ignore any DOCSIS extensions that it cannot interpret.

Type	Length	Value
43	n	

eTEA Message Integrity Check (eTEA MIC)

This parameter contains a MD5 Digest over the bytes of the config file settings, as outlined in Section 6.7.1.9.4. This TLV appears once in the configuration file immediately before the End of Data marker.

Type	Length	Value
6	16	A 128-bit (16 octet) MD5 Digest

eTEA SNMP Access Control Configuration

The eTEA MUST support the following TLVs from Annex C of [MULPIv3.0] for SNMPv1v2c access configuration in SNMP v3 coexistence mode of the eTEA management entity:

- TLV-53 SNMPv1v2c Coexistence Configuration.
- TLV-54 SNMPv3 Access View Configuration.

The eTEA MUST expand the TLVs 53 and 54 as described in [OSSIV3.0] with the purpose of configuring the SNMPv1v2c access rules. The eTEA uses the key '@eTEA' instead of @CM while populating those entries.

eTEA End of Data

This is a special marker for end of data. It has no Length or value fields.

Type	Length	Value
255		

6.7.1.11 eTEA Device Capabilities

No eTEA device capability is defined but future enhancement is possible. The hexadecimal fields in DHCP option 60 should all be set at 0.

6.7.2 UGS Flow Provisioning

For each PW, a UGS flow must be provisioned. The classifier for each UGS flow must include the parameters that uniquely associate the UGS flow to the PW. Each UGS flow should be activated when the PW module in eTEA is enabled, and be deactivated when the PW module in eTEA is disabled.

6.7.3 TDM Emulation Adaptor (TEA) Provisioning

A TEA may be located either in a CMTS or a router or a switch inside the network. It should have similar parameter sets as an eTEA, but its provisioning is out of the scope of this standard. The TEA MUST support the MIB modules as required in the “Managed objects requirements” Section 6.8.2. The TEA is configured via SNMP sets or optionally via the configuration file.

6.8 Management

Pseudo Wires provided by the IWF of the TEA are managed by setting up the configuration, monitoring the statistics provided and responding to the alarms.

Setup and teardown of Pseudo Wires is based on proper configuration of the peer TEAs that terminate the Pseudo Wire. The OSS is responsible for insuring that parameters of the peer TEAs are compatible.

6.8.1 Management Model

Several MIB modules are used for providing a standard way of monitoring the operation of the device containing a TEA. The TDM management model consists of several MIB modules, according to the conceptual layering and relationship shown in Figure 6–10 below. The TEA is managed using the MIB modules described in [RFC 2494],

[RFC 3895], [RFC 2863], [RFC 5601], [RFC 5604], and the textual conventions defined in [RFC 5542] and [RFC 5601].




Figure 6-10 - Conceptual Layering

A TDM connection will be Pseudo Wire connection. It will not be treated as an interface and therefore is not represented in the ifTable.

6.8.2 Managed objects requirements

The following sections detail implementation requirements for the RFCs listed.

6.8.2.1 Requirements for DS0 MIB and DS0 Bundle MIB [RFC 2494]

A TEA supporting Structure Aware Transport MUST implement [RFC 2494] with read-only access (i.e., SNMP GET).

If the TEA implements [RFC 2494], the TEA MAY implement the following objects with read-only access:

```

dsx0RobbedBitSignalling
dsx0CircuitId
dsx0IdleCode
dsx0SeizedCode
dsx0TransmitCodesEnable

```

6.8.2.2 Requirements for DS1 MIB [RFC 3895]

A TEA MUST implement [RFC 3895] with read-only access (i.e., SNMP GET).

A TEA MAY implement the DS2/E2 interfaces of [RFC 3895] with read-only access.

A TEA MAY implement the following objects from [RFC 3895] with read-only access:

```

dsx1SendCode
dsx1LoopbackConfig
dsx1SignalMode
dsx1Fdl
dsx1LineLength
dsx1LineCoding
ds1ChanMappingGroup
ds1FarEndGroup
ds1DS2Group

```

If Structure Aware Transport is supported, the TEA MUST support the following parameters from [RFC 3895] with read-only access:

dsx1Channelization

The TEA MUST support the dsx1TransmitClockSource object from [RFC 3895] using the parameters loop and through. Support for internal is not required.

6.8.2.3 Requirements for Ethernet Interface MIB [RFC 3635] and MAU MIB [RFC 3636]

If the TEA provides an external Ethernet interface, the TEA MUST implement [RFC 3635] and [RFC 3636]. As an eDOCSIS device, the packetized Ethernet interface is a logical interface and its configuration and management is beyond the scope of this document.

6.8.2.4 Requirements for Interfaces Group MIB [RFC 2863]

A TEA MUST implement [RFC 2863].

A TEA MUST support the ifAdminStatus object as read-write (RW) to provide administrative control over both TDM and Packet interfaces.

The ifType object has been assigned the following enumerated values for each instance of a TEA interface:

ethernetCsmacd(6)
ds1(18)
ds0 (81),
ds0Bundle (82)

6.8.2.4.1 Example use of the ifStackTable

This section describes by example how to use ifStackTable to represent the relationship of ds0 and ds0Bundles with ds1 interfaces and PWs. Implementers of the stack table for ds0 and ds0Bundle interfaces should use the appropriate RFC for the service being stacked on ds0s and ds0Bundles. Examples given below are for illustration purposes only. Implementation of the ifStackTable and ifStack should be implemented for Gets.

Example: Two PWs, one directly on a ds1, and one PW being carried on 4 ds0s of a ds1.




Figure 6–11 - PW Examples

The assignment of the index values could, for example, be:

ifIndex	Description	
1	ds1 #1	(type 18)
2	ds0Bundle	(type 82)
3	ds0 #1	(type 81)
4	ds0 #2	(type 81)
5	ds0 #3	(type 81)


```

6          ds0 #4          (type 81)
7          ds1 #2          (type 18)

```

The ifStackTable is then used to show the relationships between the various interfaces.

```

ifStackTable  Entries
HigherLayer  LowerLayer
0            1
1            0
0            2
2            3
2            4
2            5
2            6
3            7
4            7
5            7
6            7
7            0

```

The association between PWs and DS0s is configured by the PW circuit map TLV. This allows an implementation to build the ifStackTable and show the relationship based on ifIndexes.

The ifStackTable can be used to define more complicated relationships such as where DS0s from multiple DS1s are grouped together in a DS0 bundle. Application based restrictions are beyond the scope of this document.

6.8.2.5 Requirements for PW-STD-MIB [RFC 5601]

A TEA MUST implement the PW-STD-MIB as defined in [RFC 5601], utilizing the textual conventions defined in PW-TC-STD-MIB [RFC 5542] and IANA-PWE3-MIB [RFC 5601].

6.8.2.6 Requirements for PW-TDM-MIB [RFC 5604]

A TEA MUST implement the PW-TDM-MIB as defined in [RFC 5604], utilizing the textual conventions defined in PW-TC-STD-MIB [RFC 5542].

6.8.3 TLV to MIB Object Mapping

The following table shows the mapping between the TLV options and the MIB Object:

Table 6–13 - TLV to MIB Object Mapping

TLV Number	TLV Name	MIB Object
PW Table		
9.1	PW Index	pwIndex
9.2	PW Emulation Type	pwType
		pwPeerAddrType
9.3	PW Peer Addr	teaPwPeerAddr
9.4	PW Peer Addr IPv6	
9.5	PW Dest Port	pwOutboundLabel
9.6	PW Peer Dest Port	pwInboundLabel
9.7	PW Name	pwName
9.8	PW Description	pwDescr
9.9	PW Admin Status	pwAdminStatus
9.10	PW Change Notification Enable	pwUpDownNotifEnable
9.11	PW TDM Type	Deprecated
9.12	PW TDM config table Index	pwGenTDMCfgIndex
9.13	PW TDM CESoPSN table Index	pwRelTDMCfgIndex

TLV Number	TLV Name	MIB Object
9.14	PW TDM RTP SSRC	Deprecated
9.15	PW TDM RTP SSRC Peer	Deprecated
9.16	Circuit Map	mapped to DS0 MIB, ifTable and ifStackTable
PW TDM configuration Table		
10.1	PW TDM Config Table Index	pwTDMCfIndex
10.2	Payload Size	pwTDMCfPayloadSize
10.3	RTP header used	pwTDMCfRtpHdrUsed
10.5	Jitter Buffer Depth	pwTDMCfJtrBfrDepth
10.6	Payload Suppression	pwTDMCfPayloadSuppression
10.7	LOPS Exit Criteria	pwTDMCfConsecPktsInSynch
10.8	LOPS Entrance Criteria	pwTDMCfConsecMissPktsOutSynch
10.10	Packet Replace Policy	pwTDMCfPktReplacePolicy
10.11	Packet Loss Window	pwTDMCfAvePktLossTimeWindow
10.12	Excessive Loss Threshold	pwTDMCfExcessivePktLossThreshold
10.15	Severe Loss Threshold	pwTDMCfMissingPktsToSes
10.16	RTP Timestamp Mode	pwTDMCfTimestampMode
10.17	Default Fill Pattern	pwTDMCfPktFiller
10.18	L bit Payload Policy	Deprecated
10.19	PW IP TOS	Deprecated
10.20	RTP Payload Type	Deprecated
10.21	RTP Peer Payload Type	Deprecated
10.22	RTP TS Reference	Deprecated
10.23	RTP Peer TS Reference	Deprecated
10.24	SRTP enable	Deprecated
CESoPSN Configuration Table		
13.1	PW CESoPSN Config Table Index	Deprecated
13.2	Default Idle Pattern	Deprecated
13.3	LFlag Policy	Deprecated
13.4	RFlag Policy	Deprecated
13.5	Remote Defect Policy	Deprecated
13.6	LOPS Policy	Deprecated
13.7	App Sig TOS	Deprecated
13.8	RTP CAS PT	Deprecated
13.9	RTP CAS Peer PT	Deprecated
13.10	App Sig Idle	Deprecated
13.11	App Sig Interval	Deprecated
13.12	App Sig Max Interval	Deprecated
dsx1 Configuration Table		
14.1	dsx1 port ID	physical port number of the ds1 interface
14.2	dsx1 Line type	dsx1LineType
14.3	dsx1 Line Coding	dsx1LineCoding
14.4	dsx1 Circuit ID	dsx1CircuitIdentifier
14.5	dsx1 Loopback Config	dsx1LoopbackConfig
14.6	dsx1 Signal Mode	dsx1SignalMode
14.7	dsx1 Transmit Clock Source	dsx1TransmitClockSource

TLV Number	TLV Name	MIB Object
14.8	dsx1 Fdl	dsx1Fdl
14.9	dsx1 Line Length	dsx1LineLength
14.10	dsx1 Line Status Trap Enable	dsx1LineStatusChangeTrapEnable
14.11	dsx1 Channelization	dsx1Channelization
14.12	dsx1 Line Mode	dsx1LineMode
14.13	dsx1 Line Build Out	dsx1LineBuildOut

Annex A Event, SYSLOG, and SNMP Trap Extensions (Normative)

The TEA MUST support the DOCSIS Event extensions defined in this section.

This section follows the format of Annex D (Format and Content for Event, SYSLOG, and SNMP Trap) of [OSSv2.0]. The TEA MUST conform to the requirements of [OSSv2.0] section 7.4, Fault Management, pertaining to these events, unless otherwise explicitly indicated in this section.

A.1 TEA Events Description

"Event" is used in this part to reference Annex D [OSSv2.0].

A.1.1 TEA event process "Init"

The event process "Init" refers to the initial provisioning process. The Event extensions associated with the "Init" process are divided into two sub-processes, PROVISIONING and CLOCK. The extensions for PROVISIONING use Error Code Set P01, while the extensions for CLOCK use Error Code Set P02.

A.1.2 TEA event process "Oper"

The Event extensions herein designated as "Oper" cover events generated during operation. The only defined sub-process is LINK. The Error Code Sets used for these events are P03.

A.2 DOCSIS Events Extensions

Table A-1 - DOCSIS Events Extensions

Process	Sub-Process	TEA Priority	Event Message	Message Notes And Details	Error Code Set	Event ID	SNMP Notification
INITIALIZATION							
Init	PROVISIONING	Notice	DHCP Broadcast DISCOVER (eTEA-1)				
Init	PROVISIONING	Notice	DHCP OFFER (eTEA-2)				
Init	PROVISIONING	Notice	DHCP broadcast REQUEST (eTEA-3)				
Init	PROVISIONING	Notice	DHCP ACK (eTEA-4)	For SYSLOG only, append: Status = Pass (or Fail).			
Init	PROVISIONING	Notice	TFTP config file request (eTEA-5)				
Init	PROVISIONING	Error	TFTP config file download error (eTEA-6)		P01.0	8000100	
Init	PROVISIONING	Error	TFTP config file error (eTEA-6)		P01.1	8000101	
Init	PROVISIONING	Notice	Provisioning complete (eTEA-7)	For SYSLOG only, append: Status = Pass (or Fail).	P01.3	8000103	linkUp

Process	Sub-Process	TEA Priority	Event Message	Message Notes And Details	Error Code Set	Event ID	SNMP Notification
Init	CLOCK	Notice	Override of clock M/N value by S-CDMA	For SYSLOG only, append: M/N = <P1>/<P2>. P1= M, P2 = N	P02.0	8000200	
Init	CLOCK	Warning	Cannot support clock M/N value	For SYSLOG only, append: M/N = <P1>/<P2>. P1= M, P2 = N	P02.1	8000201	
OPERATION							
Oper	LINK	Error	TDM Data link Status Change.		P03.0	8000300	TDMdataLinkDown
Oper	LINK	Notice	TDM Data link Up.		P03.1	8000301	TDMdataLinkUp
Oper	LINK	Error	TDM Signal link Status Change.		P03.2	8000302	TDMsignalLinkDown
Oper	LINK	Notice	TDM Signal link Up.		P03.3	8000303	TDMsignalLinkUp
Oper	IWF	Error	PW has gone down		P03.4	8000304	PWdown
Oper	IWF	Notice	PW now operational		P03.5	8000305	PWup
Oper	IWF	Notice	PW was deleted		P03.6	8000306	PWdeleted
Oper	IWF	Notice	PW Excessive Loss		P03.7	8000307	PWexcessiveLoss
Oper	IWF	Error	PW Severe Loss		P03.8	8000308	PWsevereLoss

Appendix I Operator Recommendations (Informative)

This section pulls together information operators can use to construct and manage systems carrying T1/E1 bearing traffic. This specification outlines performance requirements for error rates, availability, etc., that may be difficult to achieve without attention to many operating details. This section is not meant to contain an exhaustive list of the operating details germane to successful operation. Rather this section is meant to highlight relevant methods and operating practices that could be overlooked.

I.1 Preferential treatment for TE-CMs

The DOCSIS RFIv2.0 [RFI 2.0] and DSG [DSG] standards together describe two TLVs that can be used to improve the performance of T1-bearing systems. The paragraphs in the subsections just below are taken from the DSG standard and slightly modified. See the two referenced standards for technical details.

I.1.1 Segregation of T1 traffic

Traditional Telco T1 signals must function with very low error rates. To further pave the way for the transport of T1 signals over DOCSIS, it is desirable to be able to segregate upstream T1 signals on a specifically "chosen upstream" (excluding all other CMs). The chosen US may, for instance, have less noise on it. DOCSIS RFIv2.0 [RFI 2.0] defines TLV 19 (for UCD messages) which regulates the usage of an US channel. MSOs may use this new capability in the following way.

MSOs can arrange it so that non-T1 bearing CMs (which do monitor TLV19) can be held off the chosen US. CMs that do not monitor TLV 19 and choose to register on the chosen upstream may be issued an "Abort Ranging" message. Non-DSG CMs will move to try another US. DSG CMs will drop into One-way mode until Tdsg3 times out after five minutes (default) before trying again.

I.1.2 Preferential ranging opportunities

Further, traditional Telco T1 signals must function with very high availability. It is desirable to allow T1-bearing CMs to range before other CMs in the event of power outages.

[RFI 2.0] defines a mechanism whereby classes of devices can be preferentially delayed in initial ranging. Cable Devices will have a default Ranging Class ID, based on the device type, which could be used to cause them to inhibit initial ranging after an outage. A bit-field TLV (18) is added to the UCD message to indicate which device classes are to inhibit ranging. MSOs can help T1-bearing CMs to range first by delaying other CMs.

I.2 Service Level Agreement considerations

Providers of T1 services often maintain Service Level Agreements (SLAs) with subscribers. These agreements are fairly standard but will be determined by individual MSOs as part of their offering. This section outlines the typical requirements as compiled from 3-4 existing SLAs from non-MSO T1 service suppliers. The general details mentioned herein are only meant to serve as a guideline for the control of a T1-bearing system, with an eye towards minimizing impact to revenue. Said another way, the following list of performance parameters may be directly important to revenue and merit particular attention. Some parameters are operational, and some are more technical. In particular, operators should strive to maximize service availability and minimize maximum latency (in each calendar month).

- Operational considerations:
 - Delivery time for the installation of service - in days After Receipt of Order.
 - Time to restore services in the event of outages - in hours.
- Technical considerations:
 - Maximum latency roundtrip (in any one month) - in milliseconds.
 - Availability (in any one month) - in percent.

The measurement of this parameter is out of scope.

It is a common practice to schedule routine service intervals in advances. The SLA often provides for such coordination and excuses the supplier for such outages. Further, the time used in such scheduled service intervals is often subtracted so the statistical parameters (used to judge SLA performance) are not affected.

I.3 Standard DOCSIS Operations

Some standard DOCSIS operations might adversely affect T1 availability and error rates. While CMTS systems have traditionally supported service flows that may be interrupted, voice and T1 services cannot be interrupted casually. Any change to the modem or services provided to subscribers must be done without impeding the ability of the two TE-CMs to forward packets over their link. Further attention may need to be paid to the manner in which the following operations are carried out.

I.3.1 Downloading

Many different downloading mechanisms have been specified for CMs. Unless a TE-CM can maintain continuous, error-free T1 service throughout a downloading cycle, then a mechanism should be found to perform the operation at a time that minimizes disruption to the service.

I.3.2 Scheduled Service

As mentioned at the end of the last section, scheduled maintenance on the under-lying DOCSIS system should be coordinated with subscribers to T1-bearing services.

I.3.3 Provisioning

To maximize the availability of the service, initialization delays should be minimized so the service recovers quickly from power outages.

I.3.3.1 IP addresses

Delays should be minimized as much as possible while obtaining IP addresses. Section 6.7 covers this in more detail.

I.3.3.2 Delays

Try to minimize other delays that may occur during initialization during:

- Ranging.
- Establishing Time of Day.
- Transfer of Operational Parameters.
- Registration.
- Baseline Privacy Initialization.

I.3.4 Balancing operations

Operations designed to balance the network should be avoided if they would disrupt service. Such operations might best be performed during scheduled maintenance intervals per the SLA (see Section I.2). If such operations are performed, some attention should be paid to using the quickest method available.

I.3.4.1 Load balancing

Load balancing operations affecting TE-CMs should be avoided. The provisioning of T1 services may require the static assignment of downstream and upstream to a TE-CM. This may be done, for instance, to decrease the bit-error rates of the TE-CM data communications (see Section I.1.1).

I.3.4.2 Node splitting and maintenance

In a similar manner, plant maintenance, like node splitting, should not be carried out without considering the effect on T1-bearing services.

Appendix II Delay Aspects (Informative)

As commercial services are frequently used by delay sensitive applications such as voice, it is critical that its delay through the network be kept within the acceptable bounds of such applications. While such delay constraint comes often as an end-to-end requirement, it may be translated into delay constraints for each network segments along the traffic path, to facilitate the network planning and management.

In this appendix, delays experienced by TDM emulation traffic through a BSoD-TE system are described. For delay sensitive applications such as cell site backhaul, it is desirable to keep the one-way delay through the BSoD-TE system, for example, the one illustrated in Figure 5–5, below 7 ms. Therefore, it is important to understand and minimize the delay within each network element and at their interfaces. Excessive delay may result in degradation of quality of service (e.g., incurring voice echo) and limit the range of applications BSoD-TE can support.

In cell site backhaul applications, calls may be dropped (as cell phones move between cells) if the delay is too large. At this time, it's believed that there are no definitive specifications governing the maximum allowable delay. Cell site operators may each have their own requirements. Anecdotally, the tightest spec mentioned is that the one way delay must be kept below 7 milliseconds; others cite 10 milliseconds. Additionally, it may be important to keep the delays as close to equal as possible for all cell sites.

In an M-CMTS or DOCSIS 3.0 network, the DOCSIS timestamp may be common across the network if DTI Servers are deployed with GPS. If this is true, then delay may be measured between any two DOCSIS devices using DEPI Latency Measurement as defined in the [DEPI] standard. This may be useful to understand what segments of the network are contributing delays.

II.1 Upstream Delay

II.1.1 Packetization Delay

Each TDM circuit is sampled periodically and the bits form the payload of a packet. The packetization delay is the time between the arrival of the first bit and that of the last bit in the same packet from the TDM circuit. Thus this delay is the same as the sample interval.

II.1.2 Processing Delay in TE-CM

The eTEA inside the TE-CM packages the payload bits into a packet for transmission to the CMTS. The delay inside TE-CM includes the UGS scheduling delay (see [RFI 2.0]), the interleaver delay and the transmission delay.

II.1.2.1 UGS Scheduling Delay

The UGS grants for TDM emulation traffic are periodical. In the best case, if the packet is ready right before the UGS grant time there is no additional delay. In the worst case, if the packet just misses the UGS grant time there is an additional delay of one UGS grant interval. The TE-CM should minimize this delay.

II.1.2.2 TDMA Byte Interleaver Delay

For TDMA, interleaver may be used in the upstream direction to reduce the packet error rate. The Interleaver delay depends on the interleaver depth and packet size. This delay is not likely to be significant as it only involves packet processing internal to the TE-CM.

II.1.2.3 Transmission Delay

The transmission delay onto the wire is:

$$\text{raw packet size} / \text{upstream bandwidth},$$

where the raw packet size includes all the protocol and physical layer overhead.

II.1.3 Propagation Delay

The propagation delay is up to 0.8 ms between the most distant CM and CMTS (see [RFI 2.0]) but typically much less.

II.1.4 Delay in CMTS

All the TDM emulation traffic should be of high priority. Still a TDM emulation flow may have to contend with other high priority traffic such as high priority control traffic and other TDM emulation flows. Hence a TDM emulation packet may have to be queued behind these types of packets and possibly also a low priority traffic already in transmission. The packet delay will also depend on the queuing architecture and the scheduling algorithm.

The TEA, if residing in a CMTS, may also have a play-back buffer with a depth of more than one packet to absorb the jitter. In such cases, additional buffering delay occurs. See Sections 5.5.2.2.2 and 5.5.2.2.3.

II.1.5 M-CMTS Delay

If the upstream RF receiver is physically separated from the core M-CMTS, there will be additional delay due to packet serialization, queuing, transmission, etc.

II.2 Downstream Delay

II.2.1 Packetization Delay

Similar to the upstream direction, each TDM circuit is sampled periodically and the bits form the payload of a packet. The packetization delay is the time between the arrival of the first bit and that of the last bit in the same packet from the TDM circuit. Thus this delay is the same as the sample interval.

II.2.2 Delay in CMTS

The delay in CMTS includes the queuing delay, the downstream interleaver delay and the transmission delay.

II.2.2.1 Queuing Delay

Although similar in essence to the upstream queuing delay as discussed in Appendix II.1, the queuing delay in the downstream direction may differ from that of the upstream due to different queuing architecture, scheduling algorithm and contending traffic.

II.2.2.2 Interleaver Delay

The downstream interleaver delay is specified in [RFI 2.0] Table 6-15.

II.2.2.3 Transmission Delay

The transmission latency onto the wire is raw packet size / downstream bandwidth where the raw packet size includes all the protocol overhead. The downstream bandwidth is specified in [ITU-T J.83] Table B.3.

II.2.3 Propagation Delay

Same as the upstream direction, the propagation delay from RF output to the most distant CM is up to 0.8 ms but typically much less.

II.2.4 Delay in TE-CM

The IWF function, if residing in CMTS, may also have a play-back buffer with a depth of more than one packet to absorb the jitter. In such cases, additional buffering delay occurs.

II.2.5 M-CMTS delay

In the M-CMTS architecture, the core M-CMTS is physically separated from the EQAM and the interface is defined in [DEPI]. Hence, additional delays exist compared to that in the traditional CMTS architecture. More discussions can be found in [DEPI] Appendix I.

II.2.5.1 Queuing Delay in Core M-CMTS

This is similar to the queuing delay in traditional CMTS as described in Appendix II.2.2.

II.2.5.2 PSP Packet Serialization Delay

PSP mode is recommended for BSoD-TE because its traffic may be prioritized over other traffic.

If the BSoD-TE packets are sent to DEPI with PSP mode, there is possible packet serialization delay depending on the packet size. PSP should not serialize BSoD-TE traffic to avoid incurring this delay.

II.2.5.3 Transmission Delay onto DEPI

The transmission latency onto the wire is raw packet size / DEPI bandwidth where the raw packet size includes all the protocol overhead as described in [DEPI].

II.2.5.4 CIN Delay

Delay added by CIN may be noticeable and is under the control of the network operator.

II.2.5.5 Latency in EQAM

Additional queuing delay is possible due to traffic scheduling as illustrated in [DEPI], Figure 6-1.

Appendix III Bandwidth Usage (Informative)

In choosing the appropriate packet payload size, one usually has to make tradeoffs between delay and bandwidth utilization. The packet size related delays have been discussed in Appendix II. The bandwidth utilization is discussed here for both upstream and downstream directions.

As described in Section 6.2, the fixed protocol overhead for a TDMoIP packet includes: TDMoIP Control Word (4 bytes), optional RTP header (12 bytes), UDP header (8 bytes), IPv4 header (20 bytes) or IPv6 header (40 bytes). So the fixed per packet overhead added by IWF is 44 bytes for IPv4 and 64 bytes for IPv6.

III.1 Upstream Bandwidth Usage

The raw upstream bandwidth over the DOCSIS interface is a function of the channel width (as given in [RFI 2.0] Table 6.5) and the modulation mode as shown in Table III-1 below.

Table III-1 - Upstream Bandwidth computed from DOCSIS 2.0

Channel Width(MHz)	QPSK(Mbps)	8-QAM (Mbps)	16-QAM(Mbps)	32-QAM (Mbps)	64-QAM (Mbps)
0.2	0.32	0.48	0.64	0.8	0.96
0.4	0.64	0.96	1.28	1.6	1.92
0.8	1.28	1.92	2.56	3.2	3.84
1.6	2.56	3.84	5.12	6.4	7.68
3.2	5.12	7.68	10.24	12.8	15.36
6.4	10.24	15.36	20.48	25.6	30.72

In addition to the fixed per-packet overhead added by the eTEA, DOCSIS in the upstream direction adds an Ethernet header (18 byte) and a MAC header (6 bytes), which brings the overhead per packet overhead to P_hdr=68 bytes for IPv4 and P_hdr=88 bytes for IPv6. This overhead may be significantly reduced by deploying PHS but the saving depends on the implementation.

For the structure-agnostic transport, if the packetization interval is T_p ms, the TDM payload for a T1 circuit is $P_{tdm}=1.544\text{Mbps} \cdot T_p \text{ ms}/8=193 \cdot T_p$ bytes, for a E1 circuit is $P_{tdm}=2.048\text{Mbps} \cdot T_p \text{ ms}/8=256 \cdot T_p$ bytes.

The physical layer overhead depends on the configured FEC K byte size, FEC T byte size, preamble length $P_{preamble}$ and guard time byte P_{guard} when using TDMA. For the physical layer payload size of $T_{hdr}+T_{tdm}$, the number of whole FEC codeword is $CW=\text{floor}((T_{hdr}+T_{tdm})/K)$, the number of bytes in the shortened codeword is $(T_{hdr}+T_{tdm}) - CW \cdot K$. The number of shortened codeword is $CW_{short} = \text{ceiling}((T_{hdr}+T_{tdm})/K) - CW$ and the total number of bytes is $CW \cdot (K+2 \cdot T) + ((T_{hdr}+T_{tdm}) - CW \cdot K) + 2 \cdot T \cdot CW_{short} + P_{preamble} + P_{guard}$.

As an example, with $T=8$, $K=20$, $P_{preamble}=20$ and $P_{guard}=5$, the raw byte per packet and the raw bandwidth for a UDP/IPv4/RTP encapsulated flow is shown in the following table, assuming using shortened codeword. Along with table III-7, the result of such calculation can be used by the operator for bandwidth provisioning.

Table III-2 - An Example of Upstream Bandwidth Usage

Packetization Interval (ms)	T1		E1	
	Raw Byte per Packet	Bandwidth Usage (Mbps)	Raw Byte per Packet	Bandwidth Usage (Mbps)
1	510	4.08	621	4.97
2	847	3.39	1069	4.28
4	1537	3.07	1997	3.99

Based on the upstream cable plant conditions, a network operator may choose the appropriate modulation mode, the PHY layer parameters, and then decide on the packetization interval based on the calculations illustrated above along with the delay budget.

III.2 Downstream Bandwidth Usage

III.2.1 Downstream RF Bandwidth Usage

The raw downstream RF bandwidth is 26.97035 Mbps for 64-QAM and 38.81070 Mbps for 256-QAM, as shown in Table B.3 in [ITU-T J.83].

The downstream PDU and MAC format is the same as the upstream so the overhead is still P_hdr=68 bytes for IPv4 and P_hdr=88 bytes for IPv6. The DOCSIS payload is then encapsulated in fixed size MPEG frames (183 byte).

Table III-3 - An Example of Downstream Bandwidth Usage

Packetization Interval (ms)	T1		E1	
	MPEG Packet Size (Byte)	Bandwidth Usage (Mbps)	MPEG Packet Size (Byte)	Bandwidth Usage (Mbps)
1	376	3.008	376	3.008
2	564	2.256	752	3.008
4	940	1.880	1128	2.256

III.2.2 DEPI Bandwidth Usage

For M-CMTS, the DEPI bandwidth usage can be calculated based on the encapsulation described in [DEPI] and how the implementation carries the payload using the recommended PSP mode.

Appendix IV eTEA Configuration Example (Informative)

IV.1 Structure Agnostic T1 Configuration

The following is an example of an eTEA configuration file used to configure a structure agnostic T1 PW.

Table IV-1 - Structure Agnostic T1 PW Configuration File Example

Type	Length	Value (subtype)	Length	Value	Description
9	43				Pseudo wire configuration
		1	2	1	PW index 1
		2	1	18	Structure agonistic TDM over IP T1 PW
		3	4	1.2.3.4	Peer's IPv4 address
		5	2	50000	Destination UDP port number
		6	2	50001	Peer Destination UDP port number
		7	4	T1-1	PW name
		9	2	2	PW administratively down
		10	1	0	Notifications not generated when the PW status changes
		12	2	1	This PW uses TDM configuration table 1
		16	3		
		16.1	1	2	This PW is associated with port 2
10	39				Common parameters for TDM configurations
		1	2	1	TDM configuration table 1
		2	2	192	Payload size is 192 bytes
		3	1	0	Optional RTP header is not used
		5	2	5	The jitter buffer is configured to be 5 msec
		6	1	2	Disable payload suppression
		7	2	10	Set the threshold to be 10 for exiting the loss of packet synchronization state
		8	2	4	Set the threshold to be 4 for entering the loss of packet synchronization state
		11	2	5	Window for calculating packet loss rate is set to 5 msec
		12	1	5	Excessive loss occurs if the loss exceeds 5%
		15	1	1	Severe loss occurs if the loss exceeds 1%
		17	1	0xFF	Fill pattern is defined as 0xFF
6	16	Digest value			MD5 Digest
255					End of data marker

IV.2 Structure Agnostic E1 Configuration

The following is an example of an eTEA configuration file used to configure a structure agnostic E1 PW, with IPv6 addresses for the PW endpoints.

Table IV-2 - Structure Agnostic E1 PW Configuration File Example

Type	Length	Value (subtype)	Length	Value	Description
9	55				Pseudo wire configuration
		1	2	1	PW index 1
		2	1	17	Structure agonistic TDM over IP E1 PW
		4	16	2001::1.2.3.4	Peer's IPv6 address
		5	2	50000	Destination UDP port number
		6	2	50001	Peer Destination UDP port number
		7	4	E1-1	PW name
		9	2	2	PW administratively down
		10	1	0	Notifications not generated when the PW status changes
		12	2	1	This PW uses TDM configuration table 1
		16	3		
		16.1	1	2	This PW is associated with port 2
10	35				Common parameters for TDM configurations
		1	2	1	TDM configuration table 1
		2	2	256	Payload size is 256 bytes
		3	1	0	Optional RTP header is not used
		5	2	5	The jitter buffer is configured to be 5 msec
		6	1	2	Disable payload suppression
		7	2	10	Set the threshold to be 10 for exiting the loss of packet synchronization state
		8	2	5	Window for calculating packet loss rate is set to 5 msec
		12	1	5	Excessive loss occurs if the loss exceeds 5%
		15	1	1	Severe loss occurs if the loss exceeds 1%
		17	1	0xFF	Fill pattern is defined as 0xFF
6	16	Digest value			MD5 Digest
255					End of data marker

IV.3 Structure Aware 5xDS0 Configuration

The following is an example of an eTEA configuration file used to configure a structure aware 5xDS0.

Table IV-3 - Structure Aware nxDS0 PW Configuration File Example

Type	Length	Value (subtype)	Length	Value	Description
9	53				Pseudo wire configuration
		1	2	1	PW index 1
		2	1	21	Structure aware PW (CESoPSN)
		3	4	1.2.3.4	Peer's IPv4 address
		5	2	50000	Destination UDP port number
		6	2	50001	Peer Destination UDP port number

Type	Length	Value (subtype)	Length	Value	Description
		7	4	5DS0	PW name
		9	2	2	PW administratively down
		10	1	0	Notifications not generated when the PW status changes
		12	2	1	This PW uses TDM configuration table 1
		13	2	1	This PW uses CESoPSN configuration table 1
		16	9		
		16.1	1	2	This PW is associated with port 2
		16.2	4	0x1F	The first five timeslots are used
10	39				Common parameters for TDM configurations
		1	2	1	TDM configuration table 1
		2	2	40	Payload size is 40 bytes
		3	1	0	Optional RTP header is not used
		5	2	5	The jitter buffer is configured to be 5 msec
		6	1	2	Disable payload suppression
		7	2	10	Set the threshold to be 10 for exiting the loss of packet synchronization state
		8	2	4	Set the threshold to be 4 for entering the loss of packet synchronization state
		11	2	5	Window for calculating packet loss rate is set to 5 msec
		12	1	5	Excessive loss occurs if the loss exceeds 5%
		15	1	1	Severe loss occurs if the loss exceeds 1%
		17	1	0xFF	Fill pattern is defined as 0xFF
13	16				CESoPSN parameters for TDM configurations
		1	2	1	TDM CESoPSN configuration table 1
		2	1	0xFF	idle pattern is defined as 0xFF
		3	1	1	Alarm indication signal on entire trunk is sent if packets are received with L flag set
		4	1	1	When packets are received with the R flag set, remote alarm indication is set
		5	1	0	Do nothing if received packet has L=0 and M=01
6	16	Digest value			MD5 Digest
255					End of data marker

IV.4 Structure Aware 5xDS0 Configuration with CAS

The following is an example of an eTEA configuration file used to configure a structure aware 5xDS0.

Table IV-4 - Structure Aware nxDS0 PW Configuration File Example

Type	Length	Value (subtype)	Length	Value	Description
9	53				Pseudo wire configuration
		1	2	1	PW index 1
		2	1	23	Structure aware PW (CESoPSN) with CAS
		3	4	1.2.3.4	Peer's IPv4 address
		5	2	50000	Destination UDP port number
		6	2	50001	Peer Destination UDP port number

Type	Length	Value (subtype)	Length	Value	Description
		7	4	5DS0	PW name
		9	2	2	PW administratively down
		10	1	0	Notifications not generated when the PW status changes
		12	2	1	This PW uses TDM configuration table 1
		13	2	1	This PW uses CESoPSN configuration table 1
		16	9		
		16.1	1	2	This PW is associated with port 2
		16.2	4	0x1F	The first five timeslots are used
10	39				Common parameters for TDM configurations
		1	2	1	TDM configuration table 1
		2	2	40	Payload size is 40 bytes
		3	1	0	Optional RTP header is not used
		5	2	5	The jitter buffer is configured to be 5 msec
		6	1	2	Disable payload suppression
		7	2	10	Set the threshold to be 10 for exiting the loss of packet synchronization state
		8	2	4	Set the threshold to be 4 for entering the loss of packet synchronization state
		11	2	5	Window for calculating packet loss rate is set to 5 msec
		12	1	5	Excessive loss occurs if the loss exceeds 5%
		15	1	1	Severe loss occurs if the loss exceeds 1%
		17	1	0xFF	Fill pattern is defined as 0xFF
13	25				CESoPSN parameters for TDM configurations
		1	2	1	TDM CESoPSN configuration table 1
		2	1	0xFF	Fill pattern is defined as 0xFF
		3	1	1	Idle pattern is sent if packets are received with L flag set
		4	1	1	When packets are received with the R flag set, remote alarm indication is set
		5	1	0	Do nothing if received packet has L=0 and M=01
		7	1	5	TOS value for the CAS signaling traffic is 5
		10	1	4	Application signaling idle interval is 4 sec
		11	1	4	Application signaling interval is 4 msec
6	16	Digest value			MD5 Digest
255					End of data marker

IV.5 Structure Aware T1 with RTP Configuration

The following is an example of an eTEA configuration file used to configure a structure aware T1 with optional RTP.

Table IV-5 - Structure Aware nxDS0 PW Configuration File Example

Type	Length	Value (subtype)	Length	Value	Description
9	48				Pseudo wire configuration
		1	2	1	PW index 1
		2	1	21	Structure aware PW (CESoPSN)
		3	4	1.2.3.4	Peer's IPv4 address
		5	2	50000	Destination UDP port number
		6	2	50001	Peer Destination UDP port number
		7	4	5DS0	PW name
		9	2	2	PW administratively down
		10	1	0	Notifications not generated when the PW status changes
		12	2	1	This PW uses TDM configuration table 1
		13	2	1	This PW uses CESoPSN configuration table 1
		18	4	0xABCD	SSRC value
10	35				Common parameters for TDM configurations
		1	2	1	TDM configuration table 1
		2	2	40	Payload size is 40 bytes
		3	1	1	Optional RTP header is used
		5	2	5	The jitter buffer is configured to be 5 msec
		6	1	2	Disable payload suppression
		7	2	10	Set the threshold to be 10 for exiting the loss of packet synchronization state
		8	2	5	Window for calculating packet loss rate is set to 5 msec
		12	1	5	Excessive loss occurs if the loss exceeds 5%
		15	1	1	Severe loss occurs if the loss exceeds 1%
		17	1	0xFF	Fill pattern is defined as 0xFF
13	16				CESoPSN parameters for TDM configurations
		1	2	1	TDM CESoPSN configuration table 1
		2	1	0xFF	idle pattern is defined as 0xFF
		3	1	1	Alarm indication signal on entire trunk is sent if packets are received with L flag set
		4	1	1	When packets are received with the R flag set, remote alarm indication is set
		5	1	0	Do nothing if received packet has L=0 and M=01
6	16	Digest value			MD5 Digest
255					End of data marker

Appendix V Clocking Options, by example (Informative)

This appendix discusses the various clocking modes supported in the spec and details several examples.

V.1 Background

A central element in emulating TDM services is the clocking methodology. This specification allows many different options for clocking the system. This appendix builds on information contained in Sections 5.4 and 5.5 and uses the same terminology. In terms of clocking the TDM links, there are two options: either these links can be tied directly to a Stratum-1 Traceable Reference Source (PRS) or the links provide their own independent timing. When links are tied to a PRS, the TEA at the CMTS uses what is called Direct Clocking to take the clock from a BITS interface or a DTI interface. The DTI interface can be used in DOCSIS 3.0 and the M-CMTS to provide intra-box synchronization of all of the modes, and this timing reference often includes a PRS reference. On the Cable Modem side of the network, the DOCSIS 10.24 MHz clock can be recovered in one of two ways: from SYNC messages or from the Symbol Clock. When TDM links have independent timing, again there are two options. If a reference clock with a common source is available at both ends of the connection, this reference can be used as a reference to timing the link clocks. If no common reference is available, then adaptive clock recovery can be used.

The table below summarizes legitimate clock modes between the ends of the TDM emulated link; between the ends of the Pseudowire. Each column represents a possible configuration between either two TE-CMs or between a TEA associated with a CMTS and a CM.

Table V-1 - BSoD Clocking Modes

Device	Clocking Mode												
	NCR		PDCR					ACR					
TE-CMTS	DC		DC					Any	Any	DC	Any	Any	Any
TEA	DC		Adaptive	Loop	Adaptive			DC	Loop				Adaptive
TE-CM	Through	Through	Adaptive	Adaptive	Loop	Adaptive	Loop	Adaptive	Adaptive	Through	Adaptive	Loop	Loop
TE-CM(2)		Through				Adaptive	Adaptive			Adaptive	Adaptive	Adaptive	
	tested	tested		tested			tested		tested				

Notes:

1. When two TE-CMs are used in the configuration, the clock mode of each TDM port is shown.
2. The Through mode implies that the clock is coming from the CU and is the DOCSIS 10.24 MHz clock derived from either the SYNC messages and/or the Symbol clock.
3. DC - Direct Clocking: used in the TE-CMTS and assumes that the clock comes from BITS or from DTI.
4. When PDCR is enabled (together with RTP) this indicates that packets are generated with PDCR timestamps.

V.2 Clocking Modes

V.2.1 Network clocking

With Network clocking, the network clock is used directly as the reference for clocking all of the interfaces. A PLL is required for converting between the frequencies of the references. Figure V-1 below, depicts the configuration of the first column of NCR where a PRS clock source is used for clocking all of the TDM interfaces directly. It is recommended that the PRS clock source be a Stratum 1 traceable clock.

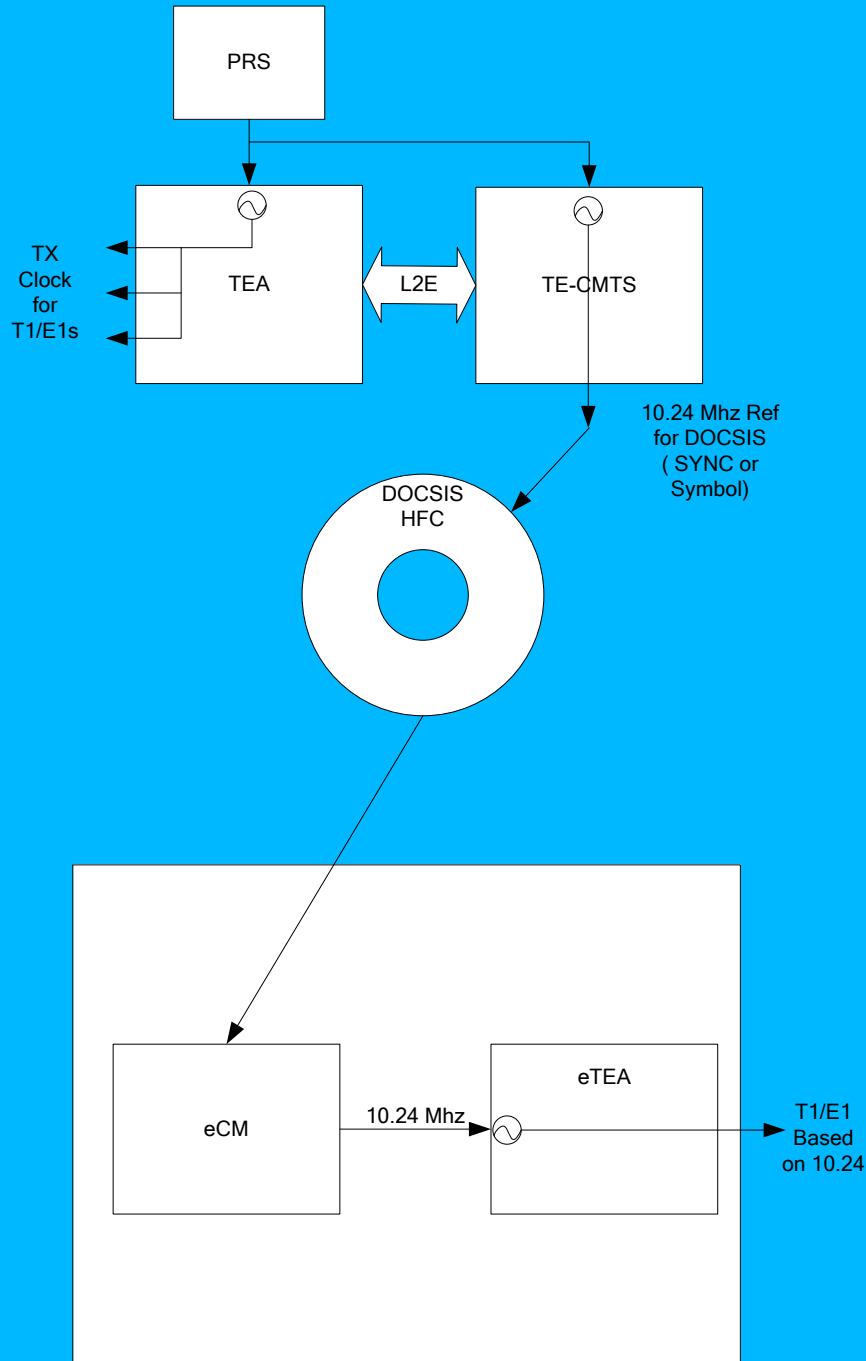


Figure V-1 - Network Clocking - CMTS to CM

The second column of NCR has a similar picture, except that in place of the T1 being terminated in the TEA; it is terminated in the TE-CM. The clock reference is still taken from the PRS in the core network.

V.2.2 Prime Differential Clock Recovery

With Prime Differential Clock Recovery, a PRS is used as the reference clock and conveyed to the TE-CMs over the DOCSIS network either using the SYNC messages or using the Symbol Clock. In this way both ends of the network are using a common reference clock. Figure V-2 below, illustrates the second column of Figure V-1, under PDCR. In this case, the T1 device attached to the TEA provides the timing reference for the T1 service. Information to discern the difference between the two clocks is passed in the RTP, such that the clock can be recovered in the Clock Recovery Unit in the TE-CM.

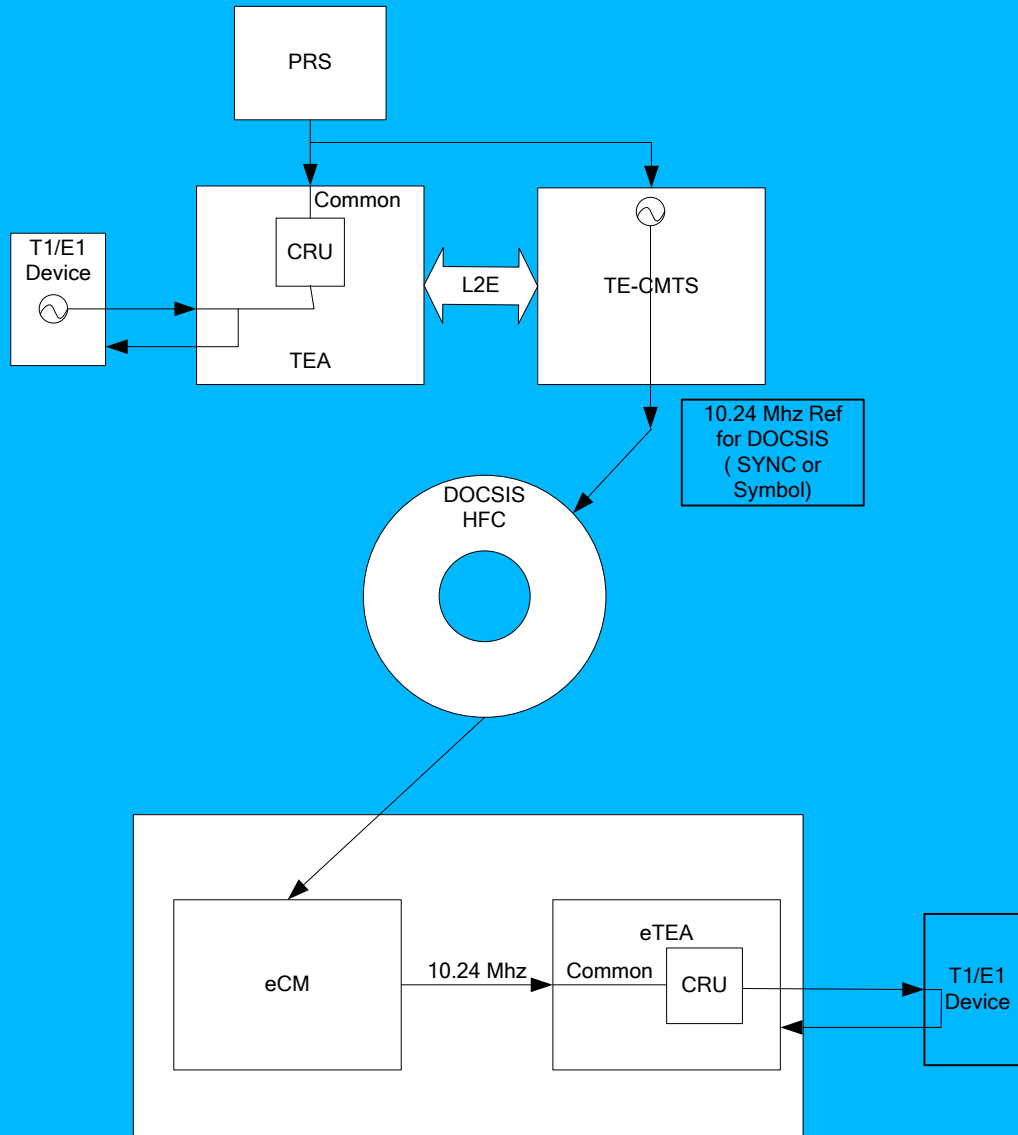


Figure V-2 - Prime Differential Clock Recovery

With PDCR, since the operation can be symmetric, other clocking options are supported, as shown in Table V-1. For example, separate clock sources can be supported in each direction, and connection between two TE-CMs can easily be accommodated.

V.2.3 Adaptive Clock Recovery

Adaptive clock recovery operates with no common clock between the two ends of the link.

Figure V-3 below, depicts the configuration of the second column of Figure V-1 of Adaptive Clock Recovery, where the TEA operates in a loop mode and takes the timing source from the attached T1/E1 device. The CRU within the IWF of the eTEA implements the Adaptive Clock Recovery. The clock is recovered based only on the inter-arrival time of the packets and the recovered clock is sent to the attached E1/T1 device.

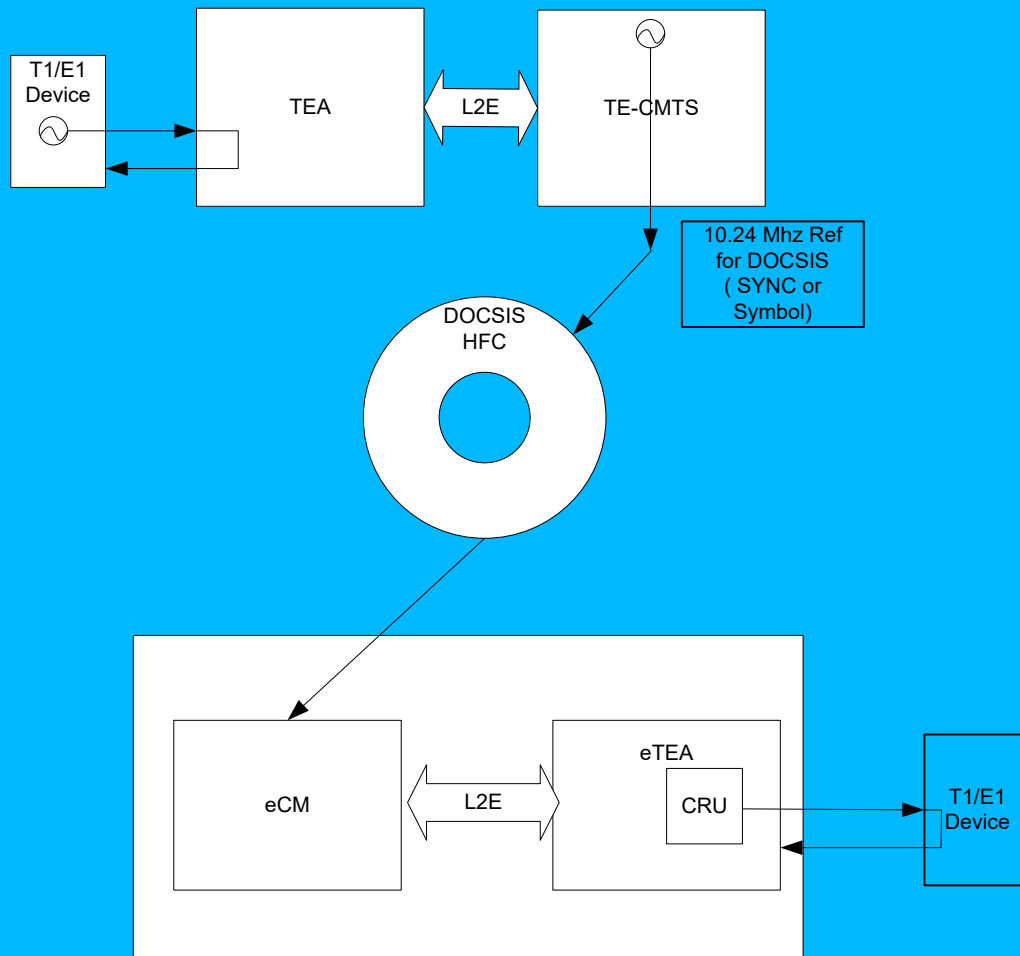


Figure V-3 - Adaptive Clock Recovery

Other configuration options of Adaptive Clock Recovery are possible, however due to the adaptive nature of ACR; care must be taken in setting performance expectations for different configurations. In particular, care must be taken to match the PW packet sizes and rates with the scheduling of the DOCSIS Service flows.

Performance of ACR can be affected by Bit Error Rates and scheduling variations in upstream and downstream data transport. For this reason, planning a service should be done with consideration of the characteristics of both downstream and upstream connections and its relation to the ACR algorithm used.

V.3 Clocking Decision Tree

When selecting between the various options, the following simplified decision tree is useful for selecting the best option. These are arranged in a prioritized order:

1. Can all nodes in system be synchronized to a single reference source clock (e.g., PRS)?
Yes - use Network Clocking.
2. (else) Do the TE-CMTS and TEA support and have access to a common reference clock (e.g., PRS)?

Yes - Use Prime Differential Clock Recovery.

3. (else) Does the TE-CM support Adaptive Clock Recovery?

Yes - use Adaptive Clock Recovery. (Preferred mode of operation is to recover clock from TE-CMTS in TE-CM.)
