

Network Working Group A. Vainshtein - Editor (Axerra Networks)
Internet Draft I. Sasson (Axerra Networks)
 A. Sadovski (Axerra Networks)
Expiration Date: E. Metz (TNO Telecom)
April 2004 T. Frost (Zarlink Semiconductor)
 P. Pate (Overture Networks)

October 2003

**Structured TDM Circuit Emulation Service over Packet Switched Network
(CESoPSN)**

[draft-vainshtein-cesopsn-07.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [section 10 of RFC 2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at

<http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

Abstract

This document describes a method for encapsulating structured (Nx64 kbit/s) TDM signals as pseudo-wires over packet-switching networks (PSN). In this regard, it complements similar work for structure-agnostic emulation of TDM bit-streams [PWE3-SAToP].

TABLE OF CONTENTS

[1. Introduction.....2](#)
[2. Changes from the Previous Version.....2](#)
[3. Terminology and Reference Models.....3](#)
 [3.1. Terminology.....3](#)
 [3.2. Reference Models.....4](#)
 [3.3. Generic and Specific Requirements.....4](#)
[4. Emulated Services.....5](#)
[5. CESoPSN Encapsulation Layer.....6](#)
 [5.1. CESoPSN Packet Format.....6](#)
 [5.2. PSN and Multiplexing Layer Headers.....7](#)
 [5.3. CESoPSN Control Word.....7](#)

5.4. Usage of the RTP header.....	9
6. CESoPSN Payload Layer.....	10
6.1. Common Payload Format Considerations.....	10
6.2. Basic Nx64 kbit/s Services.....	11
6.3. Trunk-Specific Nx64 kbit/s Services with CAS.....	12
6.4. Special Applications.....	13
6.4.1. An "Octet-aligned T1" Service.....	13
6.4.2. A Fractional E1/T1 Emulated Service.....	14
7. CESoPSN Operation.....	14
7.1. Common Considerations.....	14
7.2. IWF Operation.....	14
7.3. CESoPSN Defects and Performance Monitoring.....	15
8. QoS Issues.....	15
9. Congestion Control (RFC 2914) Conformance.....	15
10. Security Considerations.....	15
11. Applicability Statement.....	15
12. IANA Considerations.....	17
13. Intellectual Property Disclaimer.....	17
ANNEX A. A COMMON CE APPLICATION STATE SIGNALING MECHANISM.....	20
Annex B. Reference PE Architecture for Emulation of NX64 kbit/s Services.....	21

1. Introduction

This document describes a method for encapsulating structured (Nx64 kbit/s) TDM signals as pseudo-wires over packet-switching networks (PSN). In this regard, it complements similar work for structure-agnostic emulation of TDM bit-streams [[PWE3-SAToP](#)].

Ability to emulate Nx64 kbit/s circuits provides for saving PSN bandwidth, supports DS0-level grooming and distributed cross-connect applications. It also enhances resilience of CE devices to effects of loss of packets in the PSN and, in conjunction with standard TDM mappings, can be also used for carrying unstructured T1 bit-streams.

The CESoPSN solution presented in this document fits the PWE3 architecture described in [[PWE3-ARCH](#)], satisfies the general requirements put forward in [[PWE3-REQ](#)] and specific requirements for structured TDM emulation put forward in [[PWE3-TDM-REQ](#)].

2. Changes from the Previous Version

Note: This section will be removed from the final version of the document.

- 1. The document is limited to emulation of structured TDM services only. Unstructured (or structure-agnostic) emulation is considered in [[PWE3-SAToP](#)].**
- 2. CESoPSN operation description is derived from that described in**

[PWE3-SAToP].

- 3. The CESoPSN header formats have been aligned with these used in [PWE3-SAToP]. In particular:**

Vainshtein et al.

Expires April 2004

[Page 2]

- a) Usage of RTP header is OPTIONAL
- b) The CESoPSN control word is:
 - i) Aligned with the format defined in [[PWE3-ARCH](#)]
 - ii) Immediately precedes the RTP header (if the latter is used) when CESoPSN PWs are used across an MPLS PSN
 - iii) Indications of "normal" TDM data, invalid TDM data and remote loss of packet synchronization in the CESoPSN control word are aligned with these defined in [[PWE3-SAToP](#)].

4. The CESoPSN signaling packets:

- a) Are defined as capable for carrying arbitrary CE application state signals
- b) Can be used with or without the RTP header.

5. Common PWE3 fragmentation mechanisms are used for fragmenting the TDM trunk multiframe structures between CESoPSN packets

6. CAS signaling substructures are always appended to the last segment of the multiframe in emulation of trunk-specific Nx64 kbit/s services with CAS

7. A single mandatory default packetization latency for basic Nx64 kbit/s has been specified instead of multiple mandatory values).
The default latency is specified differently for $N < 5$ and $N \geq 5$.

3. Terminology and Reference Models

3.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

The terms defined in [[PWE3-ARCH](#)], Section 1.4 are consistently used without additional explanations.

This document uses some terms and acronyms that are commonly used in conjunction with the TDM services. In particular:

- o Loss of Signal (LOS) is a common term denoting a condition where a valid TDM signal cannot be extracted from the physical layer of the trunk. Actual criteria for detecting and clearing LOS are described in [[G.775](#)]
- o Frame Alignment Signal (FAS) is a common term denoting a special periodic pattern that is used to impose TDM structures on E1, T1, E3 and T3 circuits. Actual FAS patterns are described in [[G.704](#)] (E1, T1 and T3) and [[G.751](#)] (E3)
- o Out of Frame Synchronization (OOF) is a common term denoting the state of the receiver of a TDM signal when it failed to find valid FAS. Actual criteria for declaring and clearing OOF are described in [[G.706](#)]. Handling of this condition includes invalidation of the TDM data

- o Alarm Indication Signal (AIS) is a common term denoting a special bit pattern in the TDM bit stream that indicates presence of an upstream circuit outage. Actual criteria for declaring and clearing the AIS condition in a TDM stream are defined in [\[G.775\]](#)
- o Remote Alarm Indication (RAI) is a common term denoting a special pattern in the framing of a TDM service that is sent back by the receiver that experiences an AIS condition. This condition cannot be detected while a LOS, OOF or AIS condition is detected
- o Channel-Associated Signaling (CAS) is a common term denoting one of the methods of exchanging signals between telephony applications. CAS is based on allocation of up to 4 constant-rate synchronous bit-streams for each Voice-carrying DS0 channel in an E1 or T1 trunk. (These bit-streams are commonly denoted A, B, C and D.) The actual methods of carrying the sets of these bit streams isochronously in an E1 or T1 trunk and establishing association between a specific DS0 channel with an appropriate set of these bit streams are described in [\[G.704\]](#)
- o Common Channel Signaling (CCS) is a common term denoting an alternative method of exchanging signals between telephony applications.

3.2. Reference Models

Generic models that have been defined in Sections [4.1](#), [4.2](#) and [4.4](#) of [\[PWE3-ARCH\]](#) are fully applicable for the purposes of this document without any modifications.

The Network Synchronization reference model and deployment scenarios for emulation of TDM services have been described in [\[PWE3-TDM-REQ\]](#), Section 4.2.

Structured services considered in this document represent special cases of the structured bit stream payload type defined in Section 3.3.4 of [\[PWE3-ARCH\]](#). In each specific case the basic service structures that are preserved by a CESoPSN PW are explicitly specified (see [Section 3](#) below).

In accordance with the principle of minimum intervention ([\[PWE3-ARCH\]](#), Section 3.3.5) the TDM payload is encapsulated without any changes.

3.3. Generic and Specific Requirements

The protocol defined in this document has been designed in order to satisfy the requirements presented in [\[PWE3-REQ\]](#) and [\[PWE3-TDM-REQ\]](#).

The CESoPSN protocol has been designed with the following design parameters in mind:

- 1. Fixed amount of TDM data per packet: All the packets belonging to a given CESoPSN PW MUST carry the same amount of TDM data. This requirement allows compensation of a lost PW packet with a packet carrying exactly the same amount of "replacement" TDM data**
- 2. Fixed end-to-end delay: CESoPSN implementations SHOULD provide the same end-to-end delay between any given pair of PEs regardless of the bit-rate of the emulated service.**
- 3. Packetization latency range:**
 - a) All the implementations of CESoPSN SHOULD support packetization latencies in the range 1 to 5 milliseconds**
 - b) CESoPSN implementations that support configurable packetization latency MUST allow configuration of this parameter with the granularity which is a multiple of 125 microseconds**

4. Emulated Services

In accordance with [[PWE3-TDM-REQ](#)], structured services considered in this specification are Nx64 kbit/s with and without CAS. The reference PE architecture supporting these services is described in Annex B.

We shall use the following taxonomy of structured TDM services (adapted from one defined in [ATM-CES]):

- 1. Basic Nx64 kbit/s service:**
 - a) The structure ("frame") associated with this service that MUST be preserved in edge-to-edge emulation is an array of N octets, where the all the octets belong to the same frame of the "trunk" E1 or T1 circuit (in case of a service created by the NSP acting as a digital cross-connect from several synchronous E1 or T1 trunks, all the octets belong to the frame defined by the common frame clock pulse of these services), and i-th octet contains the data of the i-th DS0 channel (timeslot) in the bundle. The circuit generates 8000 frames per second**
 - b) This service can be optionally extended to support exchange of signals between CE applications (e.g., telephony ones) by carrying CE application signals in dedicated signaling packets**
 - c) Implementations MUST support $N \leq 31$ and MAY optionally support larger values of N**
- 2. "Trunk-specific" Nx64 kbit/s service with CAS. The structures that must be preserved by the PW are trunk multiframes sub-divided into trunk frames. Signaling information is carried appended to TDM data in the "signaling sub-structures" defines in [ATM-CES]. Since the number and bit rates of the CAS bit-streams depend on the specific framing method used with an E1 or T1 trunk, the following services are considered:**
 - a) E1-Nx64 kbit/s service with CAS, $1 \leq N \leq 30$**
 - b) T1/ESF-Nx64 kbit/s service with CAS, $1 \leq N \leq 24$**
 - c) T1/SF-Nx64 kbit/s service with CAS, $1 \leq N \leq 24$. Support of this service is OPTIONAL.**

Notes:

1. For T1 trunks using SF format (12 frames per multiframe), CESoPSN preserves the structure comprising two consecutive trunk multiframe subdivided into 24 trunk frames. This consideration is aligned with [ATM-CES].
2. It is possible to extend trunk-specific Nx64 kbit/s services to N exceeding the number of timeslots in a single E1 or T1 trunk. However, all the timeslots MUST originate in the same type of trunk.

5. CESoPSN Encapsulation Layer

5.1. CESoPSN Packet Format

The CESoPSN header MUST contain the CESoPSN Control Word (4 bytes) and MAY also contain a fixed RTP header [RFC3550]. If the RTP header is included in the CESoPSN header, it MUST immediately precede the CESoPSN control word in case of an IPv4 or IPv6 PSN, and MUST immediately follow it in the case of an MPLS PSN (see Fig. 2a and Fig. 2b below).

Note: Such an arrangement complies with the traditional usage of RTP for the IPv4/IPv6 PSN while making CESoPSN PWs ECMP-safe for the MPLS PSN (see [PWE3-ARCH], Section 5.4.4).

Note: Both UDP and L2TPv3 can provide the multiplexing mechanisms for CESoPSN PWs over an IPv4/IPv6 PSN.

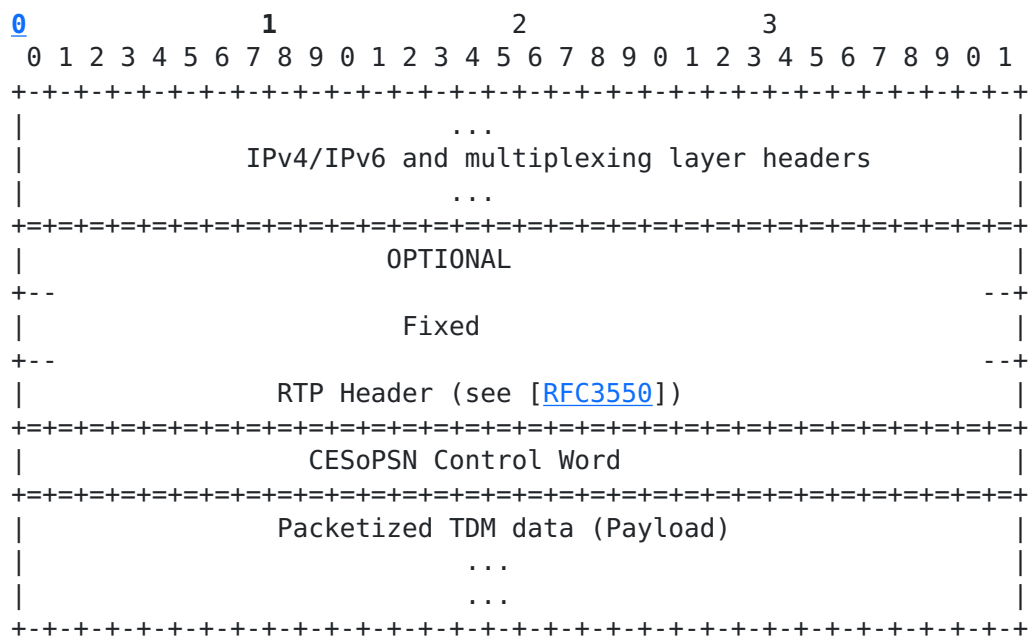


Figure 2a. CESoPSN Packet Format for an IPv4/IPv6 PSN

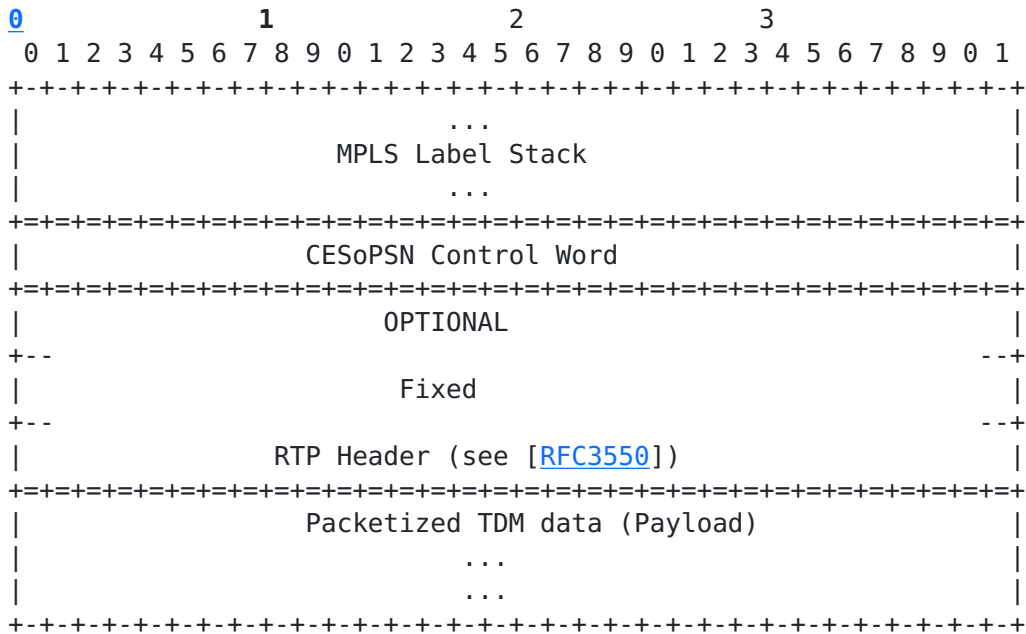


Figure 2b. CESoPSN Packet Format for an MPLS PSN

5.2. PSN and Multiplexing Layer Headers

The total size of a CESoPSN packet for a specific PW MUST NOT exceed path MTU between the pair of PEs terminating this PW. CESoPSN implementations working with IPv4 PSN MUST set the "Don't Fragment" flag in IP headers of the packets they generate.

5.3. CESoPSN Control Word

The structure of the CESoPSN Control Word is shown in Fig. 3 below.




Figure 3. Structure of the CESoPSN Control Word

Bits 0 to 3 MUST be set to 0 as described in [PWE3-ARCH], Section 5.4.4

- Bit R - If set by the PSN-bound IWF, indicates that its local CE-bound IWF is in the packet loss state. The R bit MUST be cleared by the PSN-bound IWF once its local CE-bound IWF has entered its normal operation state (see [PWE3-SAToP] for details).
- Bits L, A, X - Form a 3-bit codepoint field that is used as described in Table 1 below (all the combinations of L,R, A and X are shown).

Table 1. The CESoPSN Control Word Codepoints

L	R	A	X	Codepoint usage
0	0	0	0	CESoPSN data packet - normal operation
0	0	0	1	CESoPSN data packet - reserved for extension
0	0	1	0	CESoPSN signaling packet
0	0	1	1	Signaling packet - reserved for extension
1	0	0	0	CESoPSN data packet - invalid data. Payload MAY be omitted
1	0	0	1	CESoPSN data packet - reserved for extension Payload MAY be omitted
1	0	1	0	CESoPSN data packet, valid TDM data and trunk RAI
1	0	1	1	Reserved
0	1	0	0	CESoPSN data packet, TDM trunk OK Loss of packet synchronization
0	1	0	1	CESoPSN data packet - reserved for extension Loss of packet synchronization
0	1	1	0	Reserved
0	1	1	1	Reserved
1	1	0	0	CESoPSN data packet, invalid data. Payload MAY be omitted. Loss of packet synchronization
1	1	0	1	CESoPSN data packet - reserved for extension. Payload MAY be omitted. Loss of packet synch
1	1	1	0	CESoPSN data packet, valid data, trunk RAI Loss of packet synchronization
1	1	1	1	Reserved

Note: The structure of the CESoPSN control word and selection of the code point values are backward compatible with those used in [PWE3-SATOP].

The FRG bits are used for fragmentation and reassembly of multiframe structures between CESoPSN packets as described in [[PWE3-FRAG](#)], see [Section 6.1](#) below.

LEN (bits (10 to 15) MAY be used to carry the length of the CESoPSN packet (defined as the size of the CESoPSN header + the payload size) if it is less than 64 bytes, and MUST be set to zero otherwise.

Sequence number is used to provide the common PW sequencing function as well as detection of lost packets. Sequence numbers used in the CESoPSN data packets MUST be generated in accordance with the rules defined in [\[RFC3550\], Section 5](#) for the RTP sequence number. If signaling packets are used, their sequence numbers are generated independently from these for the CESoPSN data packets.

5.4. Usage of the RTP header

CESoPSN uses the fields of the fixed RTP header (see [\[RFC3550\], Section 5.1](#)) in the following way:

- 1. V (version) is always set to 2**
- 2. P (padding), X (header extension), CC (CSRC count) and M (marker) are always set to 0**
- 3. PT (payload type) are used as following:**
 - a) One PT value MUST be allocated from the range of dynamic values (see [\[RTP-TYPES\]](#)) for each direction of the PW. The same PT value MAY be reused for both directions of the PW and also reused between different PWs
 - b) The PE at the PW ingress MUST set the PT field in the RTP header to the allocated value
 - c) The PE at the PW egress MAY use the received value to detect malformed packets
- 4. Sequence number in the RTP header MUST be equal to the sequence number in the CESoPSN CW**
- 5. Timestamps are used for carrying timing information over the network:**
 - a) Their values are generated in accordance with the rules established in [\[RFC3550\]](#)
 - b) Frequency of the clock used for generating timestamps MUST be an integer multiple of 8 kHz. All implementations of CESoPSN MUST support the 8 kHz clock. Other frequencies that are integer multiples of 8 kHz MAY be used if both sides agree to that
 - c) Possible modes of timestamp generation are discussed below
- 6. The SSRC (synchronization source) value in the RTP header MAY be used for detection of misconnections.**

The RTP header in CESoPSN can be used in conjunction with at least the following modes of timestamp generation:

- 1. Absolute mode: the ingress PE sets timestamps using the clock recovered from the incoming TDM circuit. As a consequence, the timestamps are closely correlated with the sequence numbers. All CESoPSN implementations that support RTP MUST support this mode**
- 2. Differential mode: PE devices connected by the PW have access to the same high-quality synchronization source, and this synchronization source is used for timestamp generation. As a**

consequence, the second derivative of the timestamp series

Vainshtein et al.

Expires April 2004

[Page 9]

represents the difference between the common timing source and the clock of the incoming TDM circuit. Support of this mode is OPTIONAL. If supported, it SHOULD be used in conjunction with the timestamping clock frequencies above the TDM line clock rates. Specification of the recommended timestamping clock frequencies to be used in conjunction with the differential timestamping mode is left for further study.

6. CESoPSN Payload Layer

6.1. Common Payload Format Considerations

All the services considered in this document are treated as sequences of "basic structures" (see [Section 3](#) above). The payload of a CESoPSN packet always consists of a fixed number of octets filled, octet by octet, with the data contained in the corresponding consequent basic structures preserving octet alignment between these structures and the packet payload boundaries in accordance with the following rules:

- 1. The order of the payload octets corresponds to their order on the TDM AC.**
- 2. Consecutive bits coming from the TDM AC fill each payload octet starting from its most significant bit to the least significant one.**
- 3. All CESoPSN data packets MUST carry the same amount of valid TDM data in both directions of the PW. In other words, the time needed to fill a CESoPSN packet with valid TDM data MUST be the same for the duration of the PW. If there no valid TDM data to fill the packet, the payload may be omitted altogether as described in [Section 5.3](#) above.**
- 4. CESoPSN MUST use alignment of the basic structures with the packet payload boundaries in order to carry the structures across the PSN.** This means that:
 - a) In case of "short" basic structures ("frames"), payload of each CESoPSN frame comprises an integer number of frames and the first octet of the first frame is the first octet of the payload
 - b) In case of "long" basic structures (multiframe sub-divided into frames):
 - i) Alignment of frame sub-structures with the packet payload boundaries is preserved as described above
 - ii) The multiframe structure may be fragmented between multiple CESoPSN packets using the common PWE3 fragmentation mechanism (see [[PWE3-FRAG](#)])
 - c) CESoPSN PWs MAY carry CE signaling information appended to packets carrying valid TDM data. In this case:
 - i) The amount of the former does not affect the amount of the latter
 - ii) The CE signaling information MUST be appended to the packet carrying the last fragment of the multiframe.

Note: The terms "short" and "long" basic structures do not refer to the size of these structures in bytes but, rather, to the times required to fill them with the TDM data. The frame structure always requires 125

microseconds to fill it (regardless of the number of timeslots), while

Vainshtein et al.

Expires April 2004

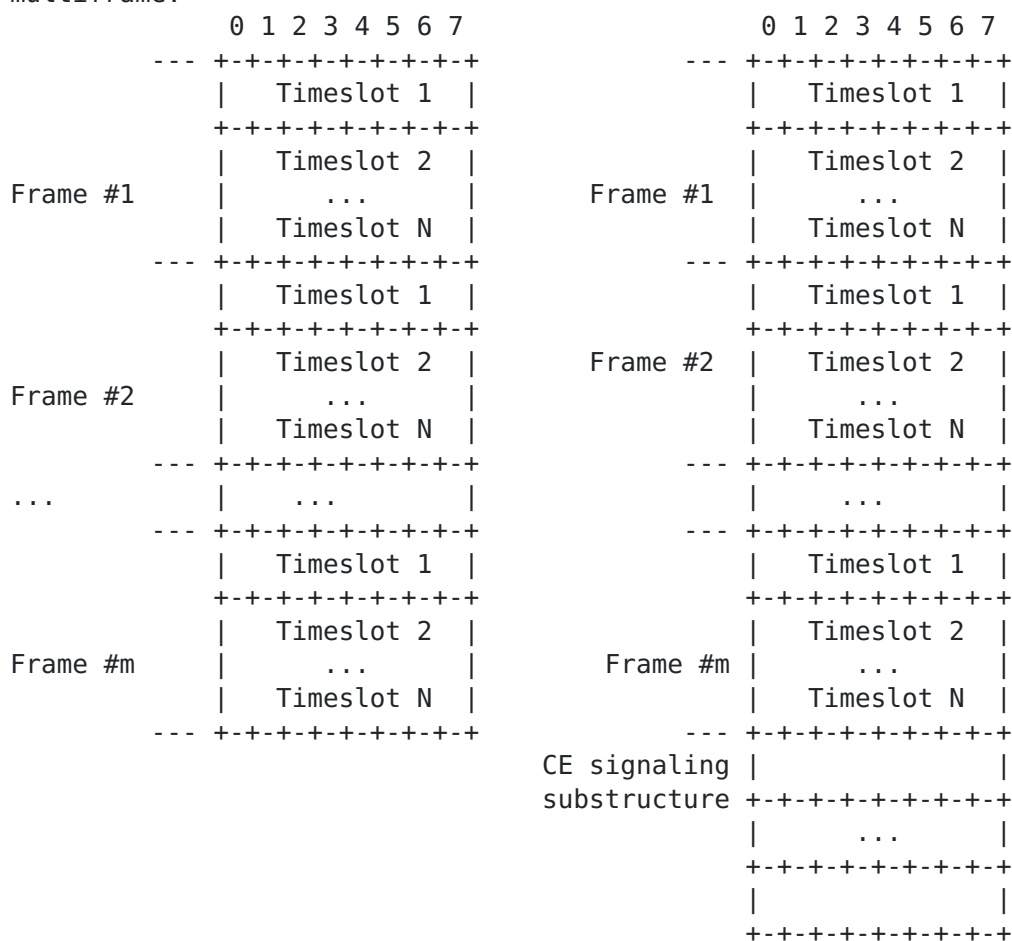
[Page 10]

the multiframe structure requires 2 or 3 milliseconds depending on the type of the trunk.

The resulting payload formats are shown in Fig 4 (a) and (b) below.

This mode of operation complies with the recommendation in [\[PWE3-ARCH\]](#) to use similar encapsulations for structured bit stream and cell generic payload types.

Note: CESoPSN implementations preserving multiframe structures cannot carry more TDM data than is contained in a single trunk multiframe.



(a) The packet without the the signaling structure

(b) The packet with the signaling structure.

Figure 4. The CESoPSN Packet Payload Formats

6.2. Basic Nx64 kbit/s Services

The structure preserved across the PSN for this service consists of N octets filled with the data of the corresponding DS0 channels belonging to the same frame of the originating trunk(s), and the service

generates 8000 such structures per second.

Vainshtein et al.

Expires April 2004

[Page 11]

Packetization latency, number of timeslots and payload size are linked by the following obvious relationship:

$$L = 8 \cdot N \cdot D$$

where:

- o D is packetization latency, milliseconds
- o L is packet payload size, octets
- o N is number of DS0 channels.

CESoPSN implementations supporting basic Nx64 kbit/s services MUST support the following packetization latency values:

- o For $N \geq 5$: 1 millisecond (with the corresponding packet payload size of $8 \cdot N$ octets). Support of packetization latencies of 2, 3 and 5 milliseconds is RECOMMENDED
- o For $1 \leq N \leq 4$: 5 milliseconds (with the corresponding packet payload size of $40 \cdot N$ octets).

Usage of any other packetization latency (packet payload size) that is compatible with the restrictions given in [Section 6.2.1](#) above is OPTIONAL.

Note: Dependency between the number of timeslots and the default latency is introduced to avoid mandating very short CESoPSN packets (and involved low bandwidth efficiency).

6.3. Trunk-Specific Nx64 kbit/s Services with CAS

The structure preserved by CESoPSN for this group of services is the trunk multiframe sub-divided into the trunk frames, and signaling information is carried appended to the TDM data using the signaling substructures defined in [ATM-CES]. These substructures comprise N consecutive nibbles, so that the i-th nibble carries CAS bits for the i-th DS0 channel, and are padded with a dummy nibble for odd values of N (see Fig. 5 below).




Figure 5: The CAS Signaling Substructure (for an odd N).

All CESoPSN implementations supporting trunk-specific Nx64 kbit/s with CAS MUST support the default mode where a single CESoPSN packet carries exactly one trunk multiframe aligned with the packet payload. In this

case:

Vainshtein et al.

Expires April 2004

[Page 12]

- 1. Packetization latency is:**
 - a) 2 milliseconds for E1 Nx64 kbit/s
 - b) 3 milliseconds for T1 Nx64 kbit/s
- 2. The packet payload size is:**
 - a) $16*n + \text{floor}((n+1)/2)$ for E1-Nx64 kbit/s
 - b) $24*n + \text{floor}((n+1)/2)$ for T1/ESF-Nx64 kbit/s and T1/SF-Nx64 kbit/s
- 3. The packet payload format coincides with the special AAL1 structure with signaling defined in [ATM-CES] for transport of trunk-specific Nx64 kbit/s with CAS.**

In order to provide lower packetization latency, CESoPSN implementations for trunk-specific Nx64 kbit/s with CAS SHOULD support fragmentation of multiframe structures between multiple CESoPSN packets as described in [Section 5.1](#) above. In this case, the signaling substructure MUST be appended to the last fragment of the multiframe structure.

Notes:

- 1. In case of T1-Nx64 kbit/s with CAS, the signaling bits are carried in the TDM data as well as in the signaling substructure.** However, the receiver MUST use the CAS bits as carried in the signaling substructures
- 2. It is possible to emulate trunk-specific Nx64 kbit/s services with CAS by just carrying the trunk multiframe structures over the PSN (and, in case of an E1 trunk, Nx64 kbit/s, including timeslot 16 in the end service).** Such an approach would be fully consistent with the Principle of Minimum Intervention. Its applicability is left for further study
- 3. In case of trunk-specific Nx64 kbit/s with CAS originating in a T1-SF trunk, each nibble of the signaling substructure contains A and B bits from two consecutive trunk multiframes as described in [ATM-CES].**

6.4. Special Applications

6.4.1. An "Octet-aligned T1" Service

Support of Nx64 kbit/s services provides an additional option for transferring unstructured T1 in an "octet-aligned" way:

- o First, it is mapped into a structured E1 using 25 timeslots (e.g., timeslots 1 to 15 and 17 to 26) in accordance with the rules described in [\[G.802\]](#), Annex B) by an appropriate NSP
- o The bundle of 25 timeslots is then carried over the PSN as an appropriate CESoPSN PW.

Support of the "octet-aligned" T1 service allows convenient integration with a wide range of NSPs commonly used with T1 (e.g., integrated E1/T1 framers, SONET/SDH mappers etc.) that inherently use [\[G.802\]](#)-based (or

similar) T1-to-E1 mapping.

Vainshtein et al.

Expires April 2004

[Page 13]

6.4.2. A Fractional E1/T1 Emulated Service

In some cases a pair of CE devices uses only a subset of timeslots of an E1 or T1 trunk but expects some aspects of native trunk behavior to be preserved by the emulated service, e.g.:

- o If the TDM data is invalid (e.g., because the transmitting CE TDM interface fails), the receiving CE TDM interface experiences an AIS condition
- o If the receiving CE TDM interface experiences an AIS condition, the CE TDM interface at the other side of the service experiences an RAI condition.

CESoPSN facilitates deployment of these applications while saving the PSN bandwidth as following:

- 1. Only the actually used timeslots (and, if necessary, their associated CE application signaling) are carried by the CESoPSN PW across the PSN**
- 2. The CE-bound IWFs at both ends of the PW are locally configured:**
 - a) To force AIS transmission on the local NSP if CESoPSN packets without valid TDM data are received
 - b) To force RAI transmission on the local NSP if CESoPSN packets with valid TDM data and RAI indication are received.

See also Annex B for details of the reference PE architecture.

7. CESoPSN Operation

7.1. Common Considerations

Edge-to-edge emulation of a TDM service using CESoPSN is only possible when the two PW attachment circuits are of the same type (basic Nx64 kbit/s or one of the trunk-specific Nx64 kbit/s with CAS) and bit rate. The service type and bit rate are exchanged at PW setup as described in [[PWE3-CONTROL](#)].

7.2. IWF Operation

Operation of CESoPSN IWF in both PSN-bound and CE-bound directions is similar to that described in [[PWE3-SATOP](#)] with the following differences:

- 1. In the situations where the SAToP CE-bound IWF is REQUIRED to send "all-ones" pattern to the CE, the CESoPSN CE-bound IWF MUST send a locally configurable constant bit pattern ("idle code") on all timeslots**
- 2. If some form of CE signaling is used (CAS or CCS), the CE-bound IWF must also send a locally configured "Channel Idle" signal on all timeslots**
- 3. CESoPSN CE-bound IWF MAY use application-specific techniques for generation of "replacement packets", e.g., as described in**

[[PACKETLOSS](#)].

Vainshtein et al.

Expires April 2004

[Page 14]

7.3. CESoPSN Defects and Performance Monitoring

CESoPSN uses the same defects, error events and error block definitions from [PWE3-SAToP].

8. QoS Issues

If the PSN providing connectivity between PE devices is Diffserv-enabled and provides a PDB [[RFC3086](#)] that guarantees low-jitter and low-loss, the CESoPSN PW SHOULD use this PDB in compliance with the admission and allocation rules the PSN has put in place for that PDB (e.g., marking packets as directed by the PSN).

9. Congestion Control ([RFC 2914](#)) Conformance

CESoPSN PWs represent a special case of PWs carrying constant bit rate (CBR) services across the PSN. These services, by definition, cannot behave in a TCP-friendly manner prescribed by [[RFC2914](#)] under congestion while retaining any value for the user.

CESoPSN will use the generic PWE3 approach for handling congestion in PWs carrying CBR services when such an approach has been specified.

10. Security Considerations

This document does not affect the underlying security issues of specific PSN.

In addition, it defines misconnection detection capabilities of CESoPSN. These capabilities increase resilience of CESoPSN to misconfiguration and some types of DoS attacks.

11. Applicability Statement

CESoPSN is an encapsulation layer intended for carrying circuits Nx64 kbit/s services with or without CAS over PSN.

CESoPSN allows, within reasonable limits, to emulate end-to-end delay properties of TDM networks. In particular, in most cases the edge-to-edge delay introduced by CESoPSN PWs does not depend upon the type and bit-rate of the emulated service.

CESoPSN fully complies with the principle of minimal intervention minimizing overhead and computational power required for encapsulation.

CESoPSN can be used in conjunction with various clock recovery techniques and does not presume availability of a global synchronous clock at the ends of a PW. However, if the global synchronous clock is available at both ends of a CESoPSN PW, using RTP and differential mode of timestamp generation improves the quality of the recovered clock.

CESoPSN allows carrying CE application state signaling that requires synchronization with data in-band in separate signaling packets. A special combination of flags in the CESoPSN control word is used to distinguish between data and signaling packets, while the Timestamp field in the RTP headers is used for synchronization. This makes CESoPSN extendable to support different types of CE signaling without affecting the data path in the PE devices.

CESoPSN also allows emulation of Nx64 kbit/s services with CAS carrying the signaling information appended to (some of) the packets carrying TDM data.

CESoPSN allows the PSN bandwidth conservation by carrying only AIS and/or Idle Code indications instead of data.

CESoPSN allows deployment of BW-saving Fractional point-to-point E1/T1 applications.

Being a constant bit rate (CBR) service, CESoPSN cannot provide TCP-friendly behavior under network congestion. If the service encounters congestion, it should be temporarily shut down.

CESoPSN allows collection of TDM-like faults and performance monitoring parameters hence emulating 'classic' carrier services of TDM circuits (e.g., SONET/SDH). Similarity with these services is increased by the CESoPSN ability to carry 'far end error' indications.

CESoPSN provides for a carrier-independent ability to detect misconnections and malformed packets. This feature increases resilience of the emulated service to misconfiguration and DoS attacks.

CESoPSN completely hides effects packet loss on the receiving CE TDM interface by locally regenerating all the information (framing, checksums) used for detection of error events and defects at the TDM level.

CESoPSN provides for detection of lost packets and allows using CE application-specific techniques for generation of replacement packets.

CESoPSN carries indications of outages of incoming attachment circuit across the PSN thus providing for effective fault isolation.

Faithfulness of a CESoPSN PW may be increased if the carrying PSN is Diffserv-enabled and implements a PDB that guarantees low loss and low jitter.

CESoPSN does not provide any mechanisms for protection against PSN outages. As a consequence, resilience of the emulated service to such outages is defined by the PSN behavior. On the other hand:

- o The jitter buffer and packets' reordering mechanisms associated with CESoPSN increase resilience of the emulated service to fast PSN rerouting events

- o Remote indication of lost packets is carried backward across the PSN from the receiver (that has detected loss of packets) to transmitter. Such an indication MAY be used as a trigger for activation of proprietary service-specific protection mechanisms.

12. IANA Considerations

This specification requires assignment of new PW Types for CESoPSN PWs listed in [Section 4.1](#).

13. Intellectual Property Disclaimer

This document is being submitted for use in IETF standards discussions. Axerra Networks, Inc. has filed one or more patent applications relating to the CESoPSN technology outlined in this document. Axerra Networks, Inc. will grant free unlimited licenses for use of this technology to the users who will register and sign up at the Axerra web site.

ACKNOWLEDGEMENTS

We express deep gratitude to Stephen Casner who reviewed this document in detail, corrected some serious errors and provided many valuable inputs.

The present version of the text of the QoS section has been suggested by Kathleen Nichols.

We thank Maximilian Riegel, Sim Narasimha, Tom Johnson, Ron Cohen and Yaron Raz for valuable feedbacks.

We thank Alik Shimelmits for many fruitful discussions.

MANDATORY REFERENCES

[RFC3550] H. Schulzrinne et al, RTP: A Transport Protocol for Real-Time Applications, [RFC 1889](#), IETF, 1996

[RFC2119] S. Bradner, Key Words in RFCs to Indicate Requirement Levels, [RFC 2119](#), IETF, 1997

[RFC2833] H. Schulzrinne, S. Petrack, RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals. [RFC 2833](#), IETF, 2000

[RFC2914] S. Floyd, Congestion Control Principles, [RFC 2914](#), IETF, 2000

[RFC3086] K. Nichols, B. Carpenter, Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification, RFC 3086, IETF, 2001

[RFC3095] C. Bormann (Ed.), Robust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed, [RFC 3095](#), IETF, 2001

[RTP-TYPES] RTP PARAMETERS, <http://www.iana.org/assignments/rtp-parameters>

[G.704] ITU-T Recommendation G.704 (10/98) - Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 Kbit/s hierarchical levels

[G.706] ITU-T Recommendation G.706 (04/91) - Frame Alignment and Cyclic Redundancy Check (CRC) Procedures Relating to Basic Frame Structured Defined in Recommendation G.704

[G.751] ITU-T Recommendation G.751 (xx/93) - Digital Multiplex Equipments Operating at the Third Order Bit Rate of 34368 kbit/s and the Fourth Order Bit Rate of 139264 kbit/s and Using Positive Justification

[G.775] ITU-T Recommendation G.775 (10/98) - Loss of Signal (LOS), Alarm Indication Signal (AIS) and Remote Defect Indication (RDI) Defect Detection and Clearance Criteria for PDH Signals

[G.802]

[ATM-CES] The ATM Forum Technical Committee. Circuit Emulation Service Interoperability Specification version 2.0 af-vtoa-0078.000, January 1997.

INFORMATIONAL REFERENCES

[PWE3-REQ] XiPeng Xiao et al, Requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3), Work in Progress, June 2003, [draft-ietf-pwe3-requirements-06.txt](#)

[PWE3-TDM-REQ] Maximilian Riegel et al, Requirements for Edge-to-Edge Emulation of TDM Circuits over Packet Switching Networks (PSN), Work in Progress, June 2003, [draft-ietf-pwe3-tdm-requirements-01.txt](#)

[PWE3-ARCH] S. Bryant, P. Pate et al, Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3), Work in progress, October 2003, [draft-ietf-pwe3-framework-06.txt](#)

[PWE3-CONTROL] L. Martini et al, Pseudowire Setup and Maintenance using LDP, Work in progress, October 2003, [draft-ietf-pwe3-control-protocol-04.txt](#)

[PWE3-IANA] L. Martini, M. Townsley, IANA Allocations for pseudo Wire Edge to Edge Emulation (PWE3), Work in progress, October 2003, [draft-ietf-pwe3-iana-allocation-02.txt](#)

[PWE3-FRAG] A. Malis, M. Townsley, PWE3 Fragmentation and Reassembly,
Work in progress, June 2003, [draft-ietf-pwe3-fragmentation-02.txt](#)

[PWE3-SATOP] A. Vainshtein, Y. Stein, Structure-agnostic TDM over
Packet, Work in Progress, September 2003, [draft-ietf-pwe3-satop-00.txt](#)

[PACKETLOSS] Y. Stein, I Druker, The Effect of Packet Loss on Voice
Quality for TDM over Pseudowires, Work in Progress, October 2003,
[draft-stein-pwe3-tdm-packetloss-01.txt](#)

Authors' Addresses

Alexander ("Sasha") Vainshtein
Axerra Networks
[24 Raoul Wallenberg St.](#),
Tel Aviv 69719, Israel
email: sasha@axerra.com

Israel Sasson
Axerra Networks
[24 Raoul Wallenberg St.](#),
Tel Aviv 69719, Israel
email: israel@axerra.com

Akiva Sadovski
Axerra Networks
[24 Raoul Wallenberg St.](#)
Tel Aviv 69719, Israel
email: akiva@axerra.com

Eduard Metz
TNO Telecom,
email: eduard.metz@hetnet.nl

Tim Frost
Zarlink Semiconductor
Tamerton Road, Roborough, Plymouth, PL6 7BQ, UK
email: tim.frost@zarlink.com

Prayson Pate
Overture Networks
[507 Airport Boulevard](#)
Building 111 Morrisville, North Carolina, 27560
Email: prayson.pate@overturenetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind,

provided that the above copyright notice and this paragraph are

Vainshtein et al.

Expires April 2004

[Page 19]

included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English. The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

ANNEX A. A COMMON CE APPLICATION STATE SIGNALING MECHANISM

A PW that requires conveyance of CE application state signals MAY carry encoded CE application signals in CESoPSN signaling packets using:

- 1. A dedicated "signaling CESoPSN packet" codepoint in the CESoPSN control word (see [Section 4.3.](#) above)**
- 2. A separate numbering sequence to distinguish between loss of a CESoPSN data packet and loss of a signaling packet**
- 3. If RTP is used, then:**
 - a) An additional PT value MUST be allocated for the CESoPSN signaling packets from the range of unused values (see [\[RTP-TYPES\]](#)). This value MUST be different from one allocated for the data packets for the same PW**
 - b) An additional SSRC value MUST be allocated for the CESoPSN signaling packets. This value MUST be different from one used for the data packets in order to allow a separate numbering sequence for the signaling packets**
 - c) The sequence number in the CESoPSN control word MUST be generated in accordance with the rules used for generation of RTP sequence numbers, and the sequence numbers in the CESoPSN control word and RTP header of the signaling packets MUST be the same**
 - d) M (marker) will be set to 1 to for "urgent" signaling packets. The CE application state carried in these packets will be conveyed to the CE at the egress of the PW immediately, without any re-synchronization with the data. Signals carried in**

"normal" signaling packets will be conveyed to the

Vainshtein et al.

Expires April 2004

[Page 20]

- CE at the PW egress after re-synchronization with the TDM data
- e) PEs terminating the PW SHOULD send RTCP sender reports (see [\[RFC3550\]](#), [Section 6.3.1](#)) for synchronization sources used for timestamping both data and signaling CESoPSN packets
 - f) Timestamps SHOULD be used for re-synchronization between TDM data and CE application state signals at the PW egress.

Signaling packets are generated by the ingress PE in accordance with the following logic (adapted from [\[RFC2833\]](#)):

- 1. The CESoPSN signaling packet with the same information is sent 3 times at an interval of 5 ms under one of the following conditions:**
 - a) The CESoPSN PW has been set up. These packets MUST be marked as "urgent"
 - b) Failure of the TDM trunk has been detected. The packets must be filled with the (locally configured) signals indicating idle state of the application. These packets MUST be marked as "urgent"
 - c) A change in the CE application state has been detected. If another change of the CE application state has been detected during the 15 ms period, this process continues
 - d) The CE-bound CESoPSN IWF has entered its normal operation state (see [PWE3-SAToP] for details). These packets MUST reflect the current CE application state and SHOULD be marked as "urgent"
 - e) Remote Loss of Packets indication has been cleared (after previously being set) These packets MUST reflect the current CE application state and SHOULD be marked as "urgent"
- 2. Otherwise, the CESoPSN signaling packets reflecting the current CE application state SHOULD be sent every 5 seconds.**

These rules allow fast probabilistic recovery after loss of a single signaling packet as well as deterministic (but, possibly, slow) recovery following PW setup and PSN outages.

Encoding of CE application state signals for various common applications will be considered in separate documents. For CAS this encoding has been defined in [\[RFC2833\]](#).

ANNEX B. REFERENCE PE ARCHITECTURE FOR EMULATION OF NX64 KBIT/S SERVICES

Structured TDM services do not exist as physical circuits. They are always carried within appropriate physical TDM trunks, and the PE providing their emulation always includes a Native Processing Block

(NSP) commonly referred to as Framer. As a consequence, the

Vainshtein et al.

Expires April 2004

[Page 21]

architecture of a PE device providing edge-to-edge emulation for these services includes the Framer and Forwarder blocks.

In case of Nx64 kbit/s services (the only type of structured services considered in this document), the trunk is either an E1 or a T1 trunk, and bundles of Nx64 kbit/s are cut out of it using one of the framing methods described in [\[G.704\]](#).

In addition to detecting the FAS and imposing associated structure on the TDM trunk, E1 and T1 framers commonly support some additional functionality including:

- 1. Detection of special states of the incoming TDM trunk**
(e.g., AIS, 00F or RAI)
- 2. Forcing special states (e.g., AIS and RAI) on the outgoing AC upon an explicit request**
- 3. Extraction and insertion of CE application signals that may accompany specific DS0 channel(s).**

The resulting PE architecture for Nx64 kbit/s services is shown in Fig. [B.1](#) below. In this diagram:

- 1. Only the case of services originating in a single TDM trunk is represented**
- 2. In the PSN-bound direction:**
 - a) The Framer:
 - i) Detects frame alignment signal (FAS) and splits the incoming ACs into separate DS0 channels
 - ii) Detects special AC states
 - iii) If necessary, extracts CE application signals accompanying each of the separate DS0 services
 - b) The Forwarder:
 - i) Creates one or more Nx64 kbit/s bundles
 - ii) Sends the data received in each such bundle to the PSN-bound direction of a respective CESoPSN IWF instance
 - iii) If necessary, sends the current CE application state data of the DS0 services in the bundle to the PSN-bound direction of the respective CESoPSN IWF instance
 - iv) If necessary sends the AC state indications to the PSN-bound directions of all the CESoPSN instances associated with the given AC
 - c) Each PSN-bound PW IWF instance encapsulates the received data, application state signal and the AC state into PW PDUs and sends the resulting packets to the PSN
- 3. In the CE-bound direction:**
 - a) Each CE-bound instance of the CESoPSN IWF receives the PW PDUs from the PSN, extracts the

TDM data, AC state and CE application state signals and sends them

- b) The Forwarder sends the TDM data, application state signals and, if necessary, a single command representing the desired AC state, to the Framer
- c) The Framer accepts all the data of one or more NX64 kbit/s bundles possibly accompanied by the associated CE application state and commands referring to the desired AC state, and generates a single AC accordingly with correct FAS.

Notes:

1. This model is asymmetric:

- a) **AC state indication can be forwarded from the framer to multiple** instances of the CESoPSN IWF
- b) No more than one CESoPSN IWF instance should forward AC state-affecting commands to the framer.

2. This model can be easily extended to the case of multiple TDM trunks. In this case:

- a) All the trunks must be synchronous
- b) The Forwarder must operate as a digital cross-connect with regard to both the TDM data and CE application signals
- c) Sending the trunk state indications to the CESoPSN IWF and forced transmission of trunk states according the CESoPSN IWF commands is not required.

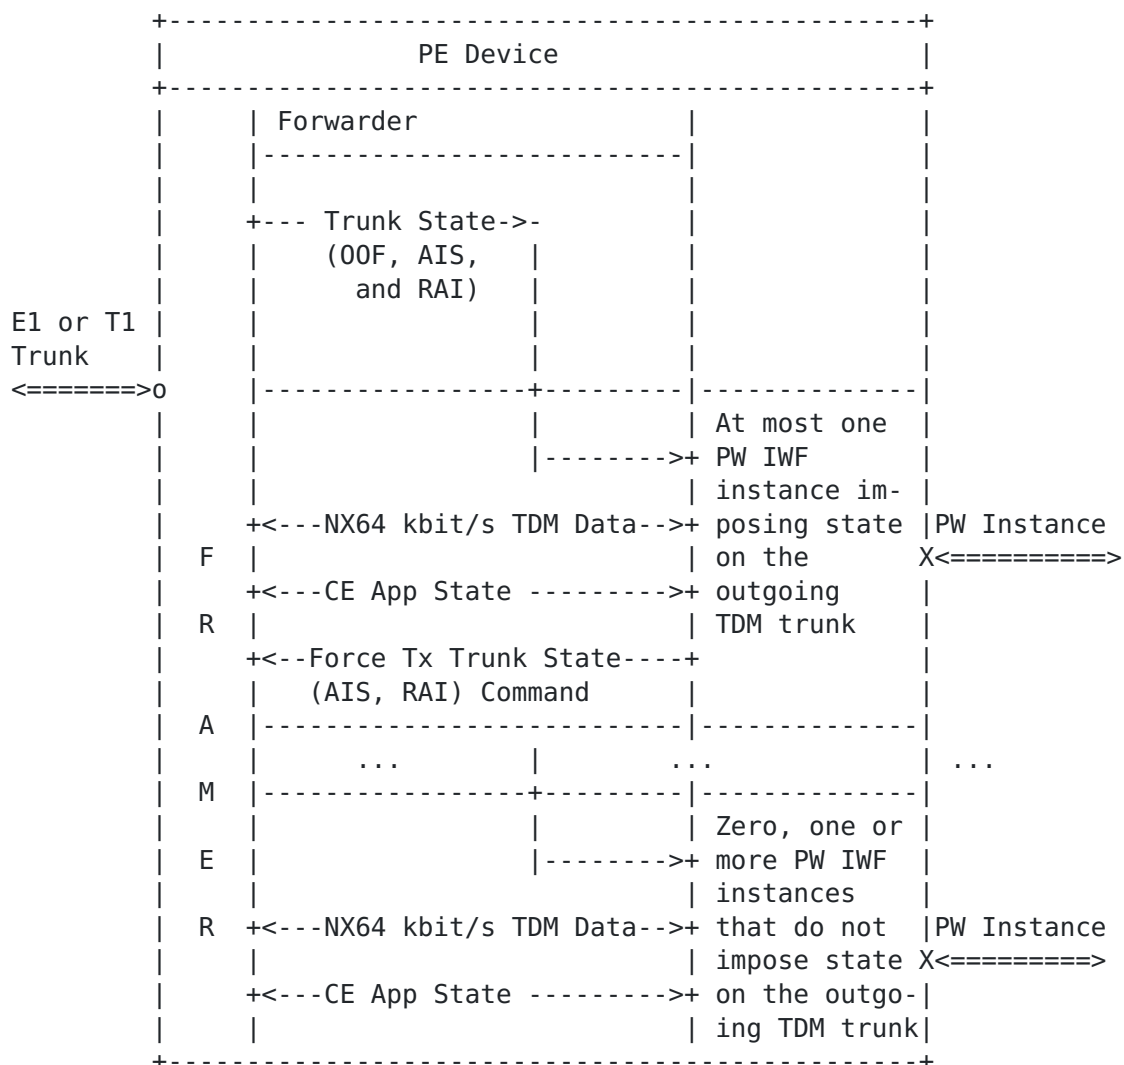


Figure B.1. Reference PE Architecture for Nx64 kbit/s Services

