

Network Working Group
INTERNET-DRAFT
Expiration Date:
January 2005

A. Vainshtein - Editor (Axerra Networks)
I. Sasson (Axerra Networks)
E. Metz (TNO Telecom)
T. Frost (Zarlink Semiconductor)
P. Pate (Overture Networks)

July 2004

Structure-aware TDM Circuit Emulation Service over Packet Switched
Network (CESoPSN)

[draft-ietf-pwe3-cesopsn-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, or will be disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

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 a "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

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 (NxDS0) 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. Terminology and Reference Models.....	2
2.1. Terminology.....	2
2.2. Reference Models.....	4

2.3. Requirements and Design Constraint.....	4
3. Emulated Services.....	5
4. CESoPSN Encapsulation Layer.....	5

4.1. CESoPSN Packet Format.....	5
4.2. PSN and Multiplexing Layer Headers.....	7
4.3. CESoPSN Control Word.....	7
4.4. Usage of the RTP header.....	9
5. CESoPSN Payload Layer.....	9
5.1. Common Payload Format Considerations.....	9
5.2. Basic NxDS0 Services.....	10
5.3. Extending Basic NxDS0 Services with CE Application Signaling.....	12
5.4. Trunk-Specific NxDS0 Services with CAS.....	13
6. CESoPSN Operation.....	15
6.1. Common Considerations.....	15
6.2. IWF operation.....	16
6.2.1. PSN-bound Direction.....	16
6.2.2. CE-bound Direction.....	16
6.3. CESoPSN Defects.....	18
6.4. CESoPSN PW Performance Monitoring.....	19
7. QoS Issues.....	19
8. Congestion Control (RFC 2914) Conformance.....	20
9. Security Considerations.....	20
10. Applicability Statement.....	20
11. Disclaimer of Validity.....	21
ANNEX A. A COMMON CE APPLICATION STATE SIGNALING MECHANISM.....	25
Annex B. Reference PE Architecture for Emulation of NxDS0 Services..	26

1. Introduction

This document describes a method for structure-aware encapsulation structured of NxDS0 TDM signals as pseudo-wires over packet-switching networks (PSN). In this regard, it complements similar work for for structure-agnostic emulation of TDM bit-streams [PWE3-SAToP].

Ability to emulate NxDS0 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.

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. Terminology and Reference Models

2.1. 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.

Vainshtein et al.

Expires January 2005

[Page 2]

Structured TDM Circuit Emulation Service over PSN

July 2004

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) and Remote Defect Indication (RDI) are common terms (often used as synonyms) 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. Specific rules for encoding this pattern in the TDM framing are discussed in

[\[G.775\]](#)

- o Channel-Associated Signaling (CAS) is a common term describing 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 [\[G.704\]](#).

Note: CAS can be interpreted in two different ways. A "synchronous" interpretation treats it as a set of synchronous bit-streams, while a "signaling" interpretation treats it as a method to encode signals reflecting change of state of telephony applications based upon generation and detection of certain stable bit patterns in the CAS-related bit-streams. The most commonly used patterns include "stable ones" and "stable zeroes"; (i.e., two states per bit-stream); in some cases they are augmented by a "stable alternate pattern" (providing the 3rd state of the bit-stream). The combination of these patterns allows

encoding of up to 16 different telephony application states. Most modern E1 and T1 framers support both approaches by providing:

1. For the synchronous approach – dedicated pins that allow extraction/insertion of the relevant constant-rate bit-streams into appropriate positions in the E1 or T1 trunk
2. For the signaling approach:
 - a) Dedicated memory-mapped registers which allow reading the actual stabilized CAS bits values and writing the desired combination of CAS bits values to be inserted into the resulting E1 or T1 trunk
 - b) Generation of interrupts when a de-bounced change of CAS bits has been detected.

Note: Another method of exchanging signals between telephony applications is called Common Channel Signaling (CCS). This method is not considered in this document in detail.

2.2. 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.

2.3. Requirements and Design Constraint

The CESoPSN protocol has been designed in order to meet the following design constraints:

- [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 approach simplifies 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.](#) Common data path for services with and without CE application signaling: if, in addition to TDM data, CE signaling must be transferred between a pair of CE devices for the normal operation of the emulated service, this signaling is passed in dedicated signaling packets specific for the signaling protocol while format and processing of the packets carrying TDM data remains unchanged.

[3.](#) Emulated Services

In accordance with [[PWE3-TDM-REQ](#)], structured services considered in this specification are NxDS0 services with and without CAS.

NxDS0 services are usually carried within appropriate physical trunks, and PEs providing their emulation include appropriate Native Service Processing (NSP) blocks commonly referred to as Framers.

The NSPs may also act as digital cross-connects, creating structured TDM services from multiple synchronous trunks. As a consequence, the service may contain more timeslots that could be carried over any single trunk, or the timeslots may not originate from any single trunk.

The reference PE architecture supporting these services is described in Annex B.

This document defines a single format for packets carrying TDM data regardless of the need to carry CAS or any other CE application signaling. The resulting "basic NxDS0 service" can be extended to carry CE application signaling (e.g. CAS) using separate signaling packets. Signaling packets MAY be carried in the same PW as the packets carrying TDM data or in a separate dedicated PW.

In addition, this document also defines dedicated formats for carrying NxDS0 services with CAS in signaling sub-structures in some of the packets. These formats effectively differ for NxDS0 services that originated in different trunks so that their usage results in emulating trunk-specific NxDS0 services with CAS.

PW Types for the corresponding CESoPSN PWs are defined in [[PWE3-IANA](#)].

[4.](#) CESoPSN Encapsulation Layer

4.1. 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. 1a and Fig. 1b below).

Note: Such an arrangement complies with the traditional usage of RTP for the IPv4/IPv6 PSN while allowing correct operation of CESoPSN PWs over routers that discriminate packet type on the basis of the first

four bits of the packet, as described in [PWE3-ARCH], section 5.4.3).

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

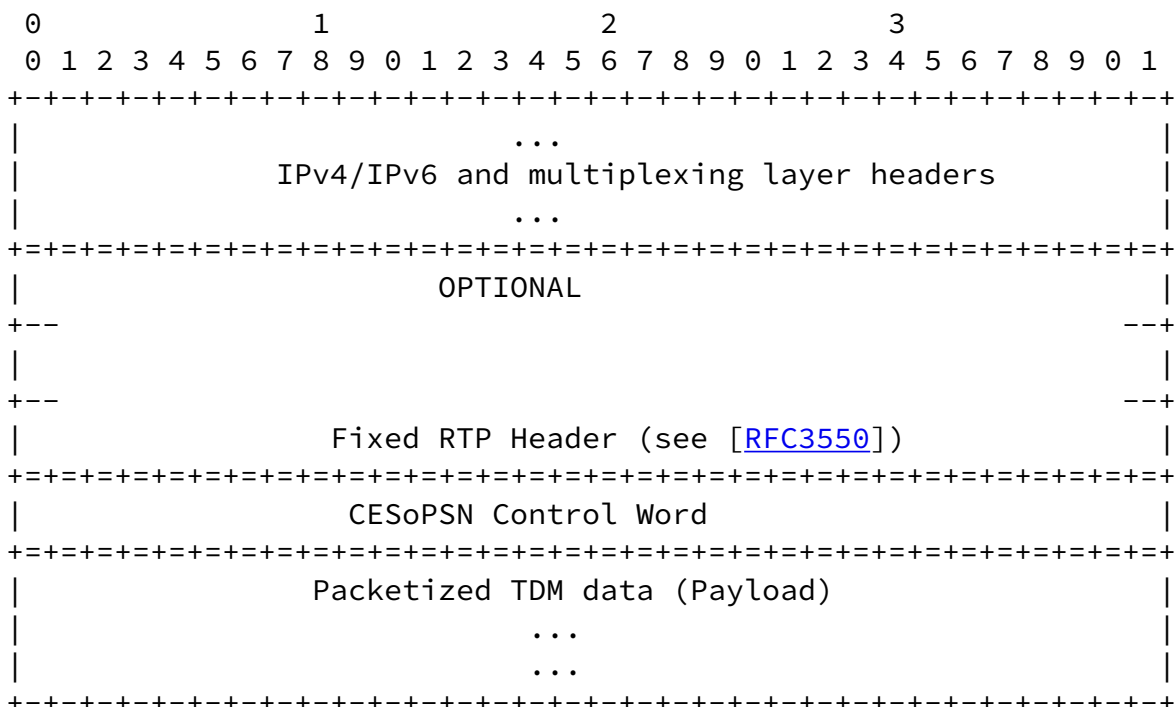


Figure 1a. CEsPSN Packet Format for an IPv4/IPv6 PSN

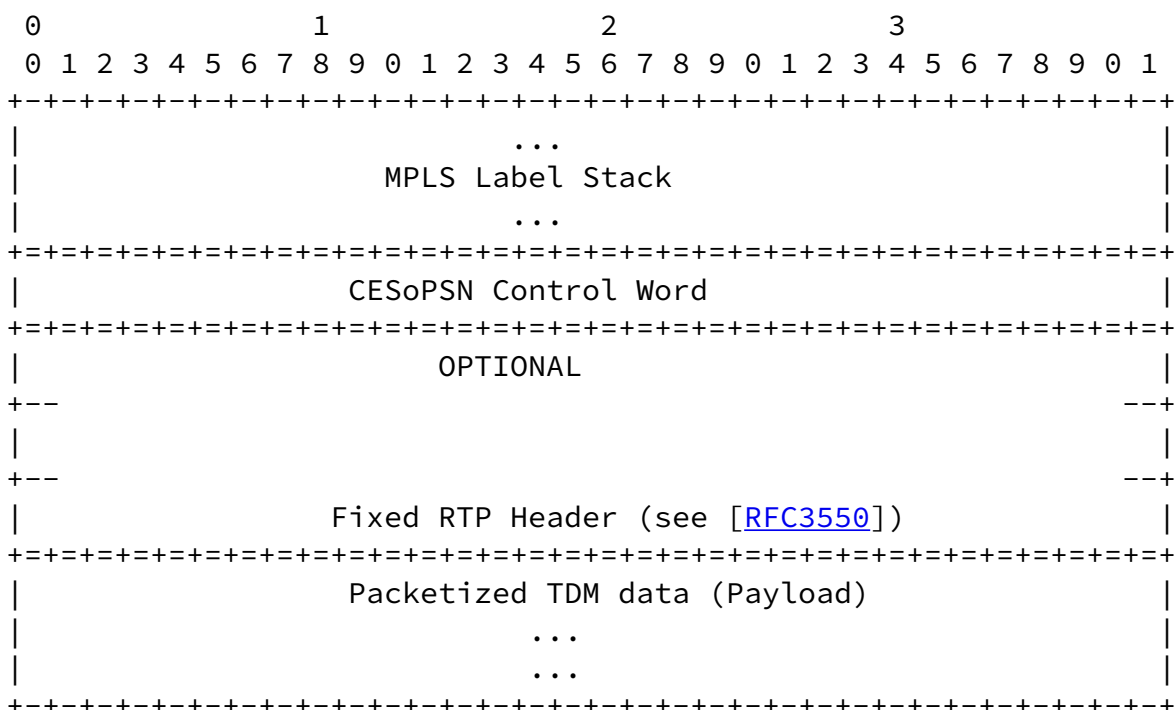


Figure 1b. CESoPSN Packet Format for an MPLS PSN

- L - if set, indicates some abnormal condition of the attachment circuit.
- M - a 2-bit modifier field. In case of L cleared allows discrimination of signaling packets, carrying RDI of the attachment circuit across the PSN etc. In case of L set

only the '00' value is currently defined, other values are reserved for future extensions. L and M bits can be treated as a 3-bit code point space that is described in detail in Table 1 below

- R - if set by the PSN-bound IWF, indicates that its local CE-bound IWF is in the packet loss state, i.e. has lost a preconfigured number of consecutive packets. The R bit MUST be cleared by the PSN-bound IWF once its local CE-bound IWF has exited the packet

loss state, i.e. has received a preconfigured number of consecutive packets.

=====		
L	M	Code Point Interpretation
=====		
0	00	CESoPSN data packet - normal situation. All CESoPSN implementations MUST recognize this code point. Payload MUST be played out "as received".
0	01	Reserved for future extensions.
0	10	CESoPSN data packet, RDI condition of the AC. All CESoPSN implementations MUST support this codepoint: I payload MUST be played out "as received", and, if so configured, the receiving CESoPSN IWF instance SHOULD be able to command the NSP to force the RDI condition on the outgoing TDM trunk.
0	11	Reserved for CESoPSN signaling packets.
1	00	TDM data is invalid, payload MAY be omitted. All implementations MUST recognize this code point and insert appropriate amount of the configured "idle code" in the outgoing attachment circuit. In addition, if so configured, the receiving CESoPSN IWF instance SHOULD be able to force the AIS condition on the outgoing TDM trunk.
1	01	Reserved for future extensions
1	10	Reserved for future extensions
1	11	Reserved for future extensions
=====		

Table 1. Interpretation of bits L and M in the CESoPSN CW.

Note: Implementations that do not support the reserved code points MUST silently discard the corresponding packets upon reception.

The FRG bits in the CESoPSN control word MUST be cleared for all services excluding trunk-specific NxDS0 with CAS. In case of these services they MAY be used to denote fragmentation of the multiframe structures between CESoPSN packets as described in [[PWE3-FRAG](#)], see also 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. It MUST be generated in accordance

Vainshtein et al.

Expires January 2005

[Page 8]

Structured TDM Circuit Emulation Service over PSN

July 2004

with the rules defined in [[RFC3550](#)], [Section 5](#) for the RTP sequence number.

4.4. Usage of the RTP header

When a fixed RTP header (see [[RFC3550](#)], [Section 5.1](#)) is used with CESoPSN, its fields are used 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) is 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 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 represents the difference between the common timing source and the clock of the incoming TDM circuit. Support of this mode is OPTIONAL.

[5.](#) CESoPSN Payload Layer

5.1. 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 the CESoPSN packets MUST carry the same amount of valid TDM data in both directions of the PW. In other words, the time that is required to fill a CESoPSN packet with the TDM data must be constant. The PE devices terminating a CESoPSN PW MUST agree on the number of TDM payload octets in the PW packets for both directions

of the PW at the time of the PW setup.

Notes:

1. CESoPSN packets MAY omit invalid TDM data in order to save the PSN BW. If the CESoPSN packet payload is omitted, the L bit in the CESoPSN control word MUST be set
2. CESoPSN PWs MAY carry CE signaling information either in separate packets or appended to packets carrying valid TDM data. If signaling information and valid TDM data are carried in the same CESoPSN packet, the amount of the former does not affect the amount of the latter.


5.2. Basic NxDS0 Services

As mentioned above, the basic 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.

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:

1. The amount of TDM data in a CESoPSN packet MUST be either an integer multiple of the basic structure size
2. The first structure in the packet MUST start immediately at the beginning of the packet payload.

The resulting payload format is shown in Fig. 3 below.






Figure 3. The CESoPSN Packet Payload Format for the Basic NxDS0 Service

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

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

$$L = 8 \times N \times D$$

where:

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

CESoPSN implementations supporting NxDS0 services MUST support the following set of configurable packetization latency values:

- o For N = 1: 8 milliseconds (with the corresponding packet payload size of 64 bytes)
- o For 2 ≤ N ≤ 4: 4 millisecond (with the corresponding packet payload size of 32×N bytes)
- o For N ≥ 5: 1 millisecond (with the corresponding packet payload size of 8×N octets).

Support of 5 ms packetization latency for N = 1 is RECOMMENDED.

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

5.3. Extending Basic NxDS0 Services with CE Application Signaling

Implementations that have chosen to extend the basic NxDS0 service to support CE application state signaling carry encoded CE application state signals in separate signaling packets.

Format of the CESoPSN signaling packets over IPv4/IPv6 and MPLS PSN for the case when the CE maintains a separate application state per DS0 channel (e.g., CAS for the telephony applications) is shown in Fig. 4a and 4b below respectively.

Signaling packets SHOULD be carried a separate dedicated PW. However, implementations MAY carry them in the same PW as the TDM data packets for the basic NxDS0 service. The methods of "pairing" the PWs carrying TDM data and signaling packets for the same extended NxDS0 service are out of scope of this document.

Regardless of the way signaling packets are carried across the PSN, the following rules apply:

1. The CESoPSN signaling packets MUST:
 - a) Use their own sequence numbers in the control word
 - b) Set the flags in the control word like following:
 - i) L = 0
 - ii) M = '11'
 - iii) R = 0
2. If RTP header is used in the data packets, it MUST be also used in the signaling packets with the following restrictions:
 - a) An additional RTP payload type (from the range of dynamically allocated types) MUST be allocated for the signaling packets.
 - b) In addition, the signaling packets MUST use their own SSRC value.

The protocol used to assure reliable delivery of signaling packets is discussed in Annex A.

Structured TDM Circuit Emulation Service over PSN

July 2004

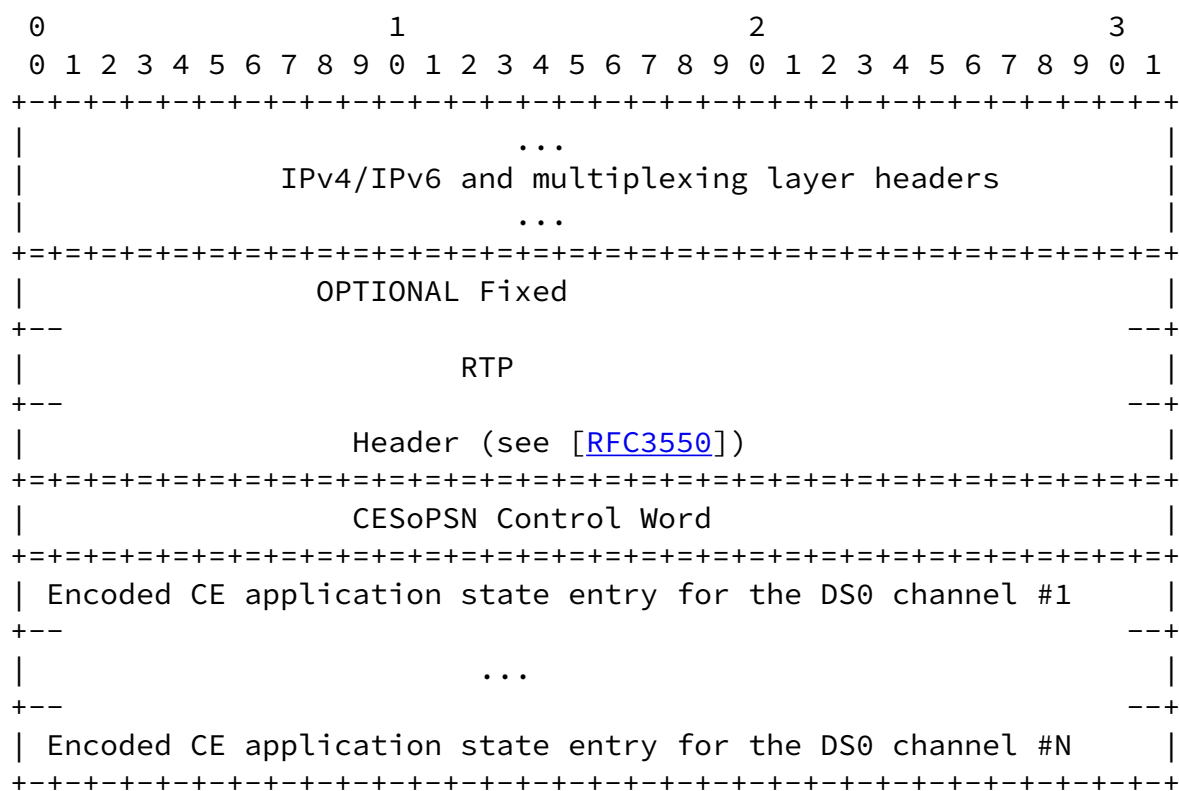
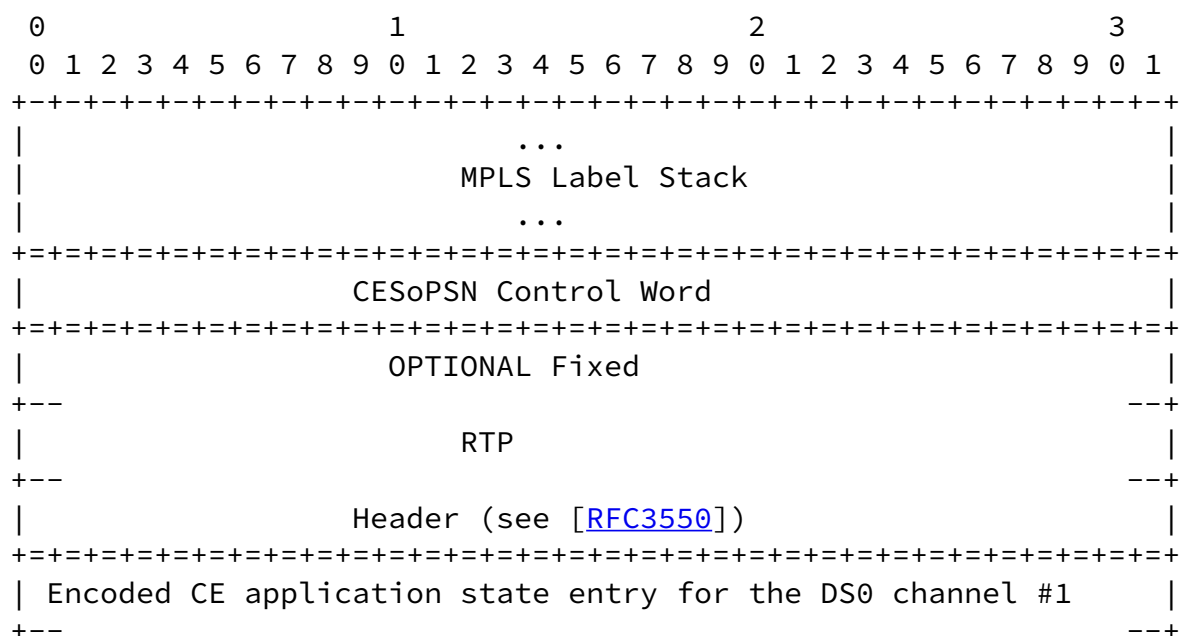


Figure 4a. CESoPSN Signaling Packet Format over an IPv4/IPv6 PSN






Figure 4b. CESoPSN Signaling Packet Format over an MPLS PSN

5.4. Trunk-Specific NxDS0 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

Structured TDM Circuit Emulation Service over PSN

July 2004

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.

CESoPSN implementations supporting trunk-specific NxDS0 services with CAS MUST NOT carry more TDM data per packet than is contained in a single trunk multiframe.

All CESoPSN implementations supporting trunk-specific NxDS0 with CAS MUST support the default mode where a single CESoPSN packet carries exactly the amount of TDM data contained in exactly one trunk multiframe and appended with the signaling sub-structure. The TDM data is aligned with the packet payload. In this case:

1. Packetization latency is:
 - a) 2 milliseconds for E1 NxDS0
 - b) 3 milliseconds for T1 NxDS0
2. The packet payload size is:
 - a) $16 \times n + \text{floor}((N+1)/2)$ for E1-NxDS0
 - b) $24 \times n + \text{floor}((N+1)/2)$ for T1/ESF-NxDS0 and T1/SF-NxDS0
3. The packet payload format coincides with the multiframe structure described in [\[ATM-CES\]](#) ([Section 2.3.1.2](#)).

In order to provide lower packetization latency, CESoPSN implementations for trunk-specific NxDS0 with CAS SHOULD support fragmentation of multiframe structures between multiple CESoPSN packets. In this case:

1. The FRG bits MUST be used to indicate first, intermediate and last fragment of a multiframe as described in [\[PWE3-FRAG\]](#)
2. The amount of the TDM data per CESoPSN packet must be constant.

3. Each multiframe fragment MUST comprise an integer multiple of the trunk frames
4. The signaling substructure MUST be appended to the last fragment of each multiframe.

Nibbles 1,2	A B C D A B C D
	+--+--+--+--+--+--+
Nibbles 3,4	A B C D A B C D
	+--+--+--+--+--+--+
Nibble n	A B C D (pad)
(odd) & pad	+--+--+--+--+--+--+

(a) The packet with the signaling structure (the last fragment of the multiframe)

(b) The packet without the signaling structure (not the last fragment of the multiframe)

Figure 5. The CESoPSN Packet Payload Format for Trunk-Specific NxDS0 with CAS

Notes:

1. In case of T1-NxDS0 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. In case of trunk-specific NxDS0 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. CESoPSN Operation

6.1. Common Considerations

Vainshtein et al.

Expires January 2005

[Page 15]

Structured TDM Circuit Emulation Service over PSN

July 2004

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 NxDS0 or one of the trunk-specific NxDS0 with CAS) and bit rate. The service type and bit rate are exchanged at PW setup as described in [PWE3-CONTROL].

6.2. IWF operation

6.2.1. PSN-bound Direction

Once the PW is set up, the PSN-bound CESoPSN IWF operates as follows:

TDM data is packetized using the configured number of payload bytes per packet.

Sequence numbers, flags, and timestamps (if the RTP header is used) are inserted in the CESoPSN headers and, for trunk-specific NxDS0 with CAS, signaling substructures are appended to the packets carrying the last fragment of a multiframe.

CESoPSN, multiplexing layer and PSN headers are prepended to the packetized service data.

The resulting packets are transmitted over the PSN.

6.2.2. CE-bound Direction

The CE-bound CESoPSN IWF SHOULD include a jitter buffer where payload of the received CESoPSN packets is stored prior to play-out to the local TDM attachment circuit. The size of this buffer SHOULD be locally configurable to allow accommodation to the PSN-specific packet delay variation.

The CE-bound CESoPSN IWF SHOULD use the sequence number in the control word for detection of lost and mis-ordered packets. If the RTP header is used, the RTP sequence numbers MAY be used for the same purposes.

The CE-bound CESoPSN IWF MAY re-order mis-ordered packets. Mis-ordered packets that cannot be reordered MUST be discarded and treated as lost.

The payload of the received CESoPSN data packets marked with the L bit set SHOULD be replaced by the equivalent amount of some locally configured "idle" bit pattern even if it has not been omitted. In addition, the CE-bound CESoPSN IWF will be locally configured to command its local NSP to perform one of the following actions:

- o None (MUST be supported by all the implementations)
- o Transmit the AIS pattern towards the local CE on the E1 or T1 trunk carrying the local attachment circuit (support of this action is RECOMMENDED)
- o Send the "Channel Idle" signal to the local CE for all the DS0 channels comprising the local attachment circuit (support of this action is OPTIONAL).

If the data packets received are marked with L bit cleared and M bits set to '10' or with R bit set, the CE-bound CESoPSN IWF will be locally configured to command its local NSP to perform one of the following actions:

- o None (MUST be supported by all the implementations)
- o Transmit the RAI pattern towards the local CE on the E1 or T1 trunk carrying the local attachment circuit (support of this action is RECOMMENDED)
- o Send the "Channel Idle" signal to the local CE for all the DS0 channels comprising the local attachment circuit (support of this action is OPTIONAL and requires also that the CE-bound CES IWF replaces the actually received payload with the equivalent amount of the locally configured "idle" bit pattern).

Notes:

1. If the pair of IWFs at the two ends of the PW have been configured to force the TDM trunks carrying their ACs to transmit AIS upon reception of data packets with the L bit set and to transmit RAI upon reception of data packets with the R bit set or with the L bit cleared and M bits set to '10', this PW provides a BW-saving emulation of a fractional E1 or T1 service between the pair of CE devices
2. If the pair of IWFs at the two ends of the PW have been configured to signal "Channel Idle" CE application state to its local CE upon reception of packets marked with L bit set, R bit set or (L,M) set to '010' and to replace the actually received payload with the locally configured "idle" bit pattern, the resulting PW will comply with the requirements for Downstream Trunk conditioning as defined in [\[TR-NWT-170\]](#).
3. Usage of bits R,L and M described above additionally provides the tools for "single-ended" management of the CESoPSN pseudo-wires with ability to distinguish between the problems in the PSN and in the TDM attachment circuits.

The payload of each lost CESoPSN data packet MUST be replaced with the equivalent amount of the replacement data. The contents of the replacement data are implementation-specific and MAY be locally configurable. By default, all CESoPSN implementations MUST support generation of the locally configurable "idle" pattern as the replacement data.

Before a PW has been set up and after a PW has been torn down, the IWF MUST play out the locally configurable "idle" pattern to its TDM attachment circuit.

Once the PW has been set up, the CE-bound IWF begins to receive CESoPSN packets and to store their payload in the jitter buffer but continues to play out the locally configurable "idle" pattern to its TDM attachment circuit. This intermediate state persists until a preconfigured amount of TDM data (usually half of the jitter buffer)

has been received in consecutive CESoPSN packets or until a preconfigured intermediate state timer expires.

Once the preconfigured amount of the TDM data has been received, the CE-bound CESoPSN IWF enters its normal operation state where it continues to receive CESoPSN packets and to store their payload in the jitter buffer while playing out the contents of the jitter buffer in accordance with the required clock. In this state the CE-bound IWF performs clock recovery, MAY monitor PW defects, and MAY collect PW performance monitoring data.

If the CE-bound CESoPSN IWF detects loss of a preconfigured number of consecutive packets or if the intermediate state timer expires before the required amount of TDM data has been received, it enters its packet loss state. While in this state:

- o The locally configurable "idle" pattern SHOULD be played out to the TDM attachment circuit
- o The local PSN-bound CESoPSN IWF SHOULD mark every packet it transmits with the R bit set.

The CE-bound CESoPSN IWF leaves this state and transits to the normal one once a preconfigured number of consecutive CESoPSN packets have been received.

6.3. CESoPSN Defects

In addition to the packet loss state of the CE-bound CESoPSN IWF defined above, it MAY detect the following defects:

- o Stray packets
- o Malformed packets
- o Excessive packet loss rate
- o Buffer overrun
- o Remote packet loss.

Corresponding to each defect is a defect state of the IWF, a detection criterion that triggers transition from the normal operation state to the appropriate defect state, and an alarm that MAY be reported to the management system and thereafter cleared. Alarms are only reported when the defect state persists for a preconfigured amount of time (typically [2.5](#) seconds) and MUST be cleared after the corresponding defect is undetected for a second preconfigured amount of time (typically 10 seconds). The trigger and release times for the various alarms may be independent.

Stray packets MAY be detected by the PSN and multiplexing layers. When RTP is used, the SSRC field in the RTP header MAY be used for this

purpose as well. Stray packets MUST be discarded by the CE-bound IWF and their detection MUST NOT affect mechanisms for detection of packet loss.

Structured TDM Circuit Emulation Service over PSN

July 2004

Malformed packets are detected by mismatch between the expected packet size (taking the value of the L bit into account) and the actual packet size inferred from the PSN and multiplexing layers. When RTP is used, lack of correspondence between the PT value and that allocated for this direction of the PW MAY also be used for this purpose. Malformed in-order packets MUST be discarded by the CE-bound IWF and replacement data generated as for lost packets.

Excessive packet loss rate is detected by computing the average packet loss rate over a configurable amount of times and comparing it with a preconfigured threshold.

Buffer overrun is detected in the normal operation state when the jitter buffer of the CE-bound IWF cannot accommodate newly arrived CESoPSN packets.

Remote packet loss is indicated by reception of packets with their R bit set.

6.4. CESoPSN PW Performance Monitoring

Performance monitoring (PM) parameters are routinely collected for TDM services and provide an important maintenance mechanism in TDM networks. Ability to collect compatible PM parameters for CESoPSN PWs enhances their maintenance capabilities.

Collection of the CESoPSN PW performance monitoring parameters is OPTIONAL, and if implemented, is only performed after the CE-bound IWF has exited its intermediate state.

CESoPSN defines error events, errored blocks and defects as follows:

- o A CESoPSN error event is defined as insertion of a single replacement packet into the jitter buffer (replacement of payload of CESoPSN packets with the L bit set is not considered as insertion of a replacement packet)
- o A CESoPSN errored data block is defined as a block of data played out to the TDM attachment circuit and of size defined in accordance with the [\[G.826\]](#) rules for the

- corresponding TDM service that has experienced at least one CESoPSN error event
- o A CESoPSN defect is defined as the packet loss state of the CE-bound CESoPSN IWF.

The CESoPSN PW PM parameters (Errored, Severely Errored and Unavailable Seconds) are derived from these definitions in accordance with [G.826].

7. 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).

8. 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.

9. 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.

10. Applicability Statement

CESoPSN is an encapsulation layer intended for carrying circuits NxDS0 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 NxDS0 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. These applications can be described like following:

- o The pair of CE devices operates as if they were connected by an emulated E1 or T1 circuit. In particular they react to AIS and RAI states of their local ACs in the standard way
- o The PSN carries only an NxDS0 service where N is the number of actually used timeslots in the circuit connecting the pair of CE devices thus saving the BW.

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 provides for detection of lost packets and allows using various 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.

11. Disclaimer of Validity

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use

of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

ACKNOWLEDGEMENTS

Akiva Sadovski has been an active participant of the team that co-authored early versions of this document.

We express deep gratitude to Stephen Casner who reviewed an early version of 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

[RFC1122] R. Braden (ed.), Requirements for Internet Hosts -- Communication Layers, [RFC 1122](#), IETF, 1989

[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

[RFC3550] H. Schulzrinne et al, RTP: A Transport Protocol for Real-Time Applications, [RFC 1889](#), IETF, 2003

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

[G.702] ITU-T Recommendation G.702 (11/88) - Digital Hierarchy Bit Rates

[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.826] ITU-T Recommendation G.826 (02/99) - Error performance parameters and objectives for international, constant bit rate digital paths at or above the primary rate

[T1.107] American National Standard for Telecommunications - Digital Hierarchy - Format Specifications, ANSI T1.107-1988

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

[TR-NWT-170] Digital Cross Connect Systems - Generic Requirements and Objectives, Bellcore, TR-NWT-170, January 1993

INFORMATIONAL REFERENCES

[PWE3-REQ] XiPeng Xiao et al, Requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3), Work in Progress, December 2003, [draft-ietf-pwe3-requirements-08.txt](#)

Structured TDM Circuit Emulation Service over PSN

July 2004

[PWE3-TDM-REQ] Maximilian Riegel et al, Requirements for Edge-to-Edge Emulation of TDM Circuits over Packet Switching Networks (PSN), Work in Progress, April 2004, [draft-ietf-pwe3-tdm-requirements-05.txt](#)

[PWE3-ARCH] S. Bryant, P. Pate et al, PWE3 Architecture, Work in progress, March 20034, [draft-ietf-pwe3-arch-07.txt](#)

[PWE3-IANA] L. Martini, M. Townsley, IANA Allocations for pseudo Wire Edge to Edge Emulation (PWE3), Work in progress, April 2004, [draft-ietf-pwe3-iana-allocation-04.txt](#)

[PWE3-FRAG] A. Malis, M. Townsley, PWE3 Fragmentation and Reassembly, Work in Progress, February 2004, [draft-ietf-pwe3-fragmentation-05.txt](#)

[PWE3-SATOP] A. Vainshtein, Y. Stein, Structure-Agnostic TDM over Packet (SAToP), Work in Progress, December 2004, [draft-ietf-pwe3-satop-01.txt](#)

[L2TPv3] J. Lau, M. Townsley, I Goyret (editors), Layer Two Tunneling Protocol (Version 3), Work in Progress, June 2004, [draft-ietf-l2tpext-l2tp-base-14.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

Eduard Metz
Thrupoint
Paasheuvelweg 16,
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

Vainshtein et al.

Expires January 2005

[Page 24]

Structured TDM Circuit Emulation Service over PSN

July 2004

Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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

Format of the CESoPSN signaling packets is discussed in [Section 5.3](#) above.

The sequence number in the CESoPSN control word for the signaling packets is generated according to the same rules as for the TDM data packets.

If the RTP header is used in the CESoPSN signaling packets, the timestamp in this header represents the time when the CE application state has been collected.

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 (including the timestamp in the case RTP header is used) is sent 3 times at an interval of 5 ms under one of the following conditions:
 - a) The CESoPSN PW has been set up
 - b) 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
 - c) Loss of packets defect has been cleared. These packets SHOULD be marked as "urgent"
 - d) Remote Loss of Packets indication has been cleared (after previously being set)
2. Otherwise, the CESoPSN signaling packet with the current CE application state information is 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 for telephony applications using CAS follows [[RFC2833](#)].

Encoding of CE application state for telephony application using CCS will be considered in a separate document.

ANNEX B. REFERENCE PE ARCHITECTURE FOR EMULATION OF NXDS0 SERVICES

Structured TDM services do not exist as physical circuits. They are always carried within appropriate physical attachment circuits (AC), and the PE providing their emulation always includes a Native Processing Block (NSP) commonly referred to as Framer. As a consequence, the architecture of a PE device providing edge-to-edge emulation for these services includes the Framer and Forwarder blocks.

In case of NxDS0 services (the only type of structured services considered in this document), the AC is either an E1 or a T1 trunk, and bundles of NxDS0 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 "trunk" AC, E1 and T1 framers commonly support some additional

functionality including:

1. Detection of special states of the incoming AC (e.g., AIS, OOF 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 NxDS0 services is shown in Fig. B.1 below. In this diagram:

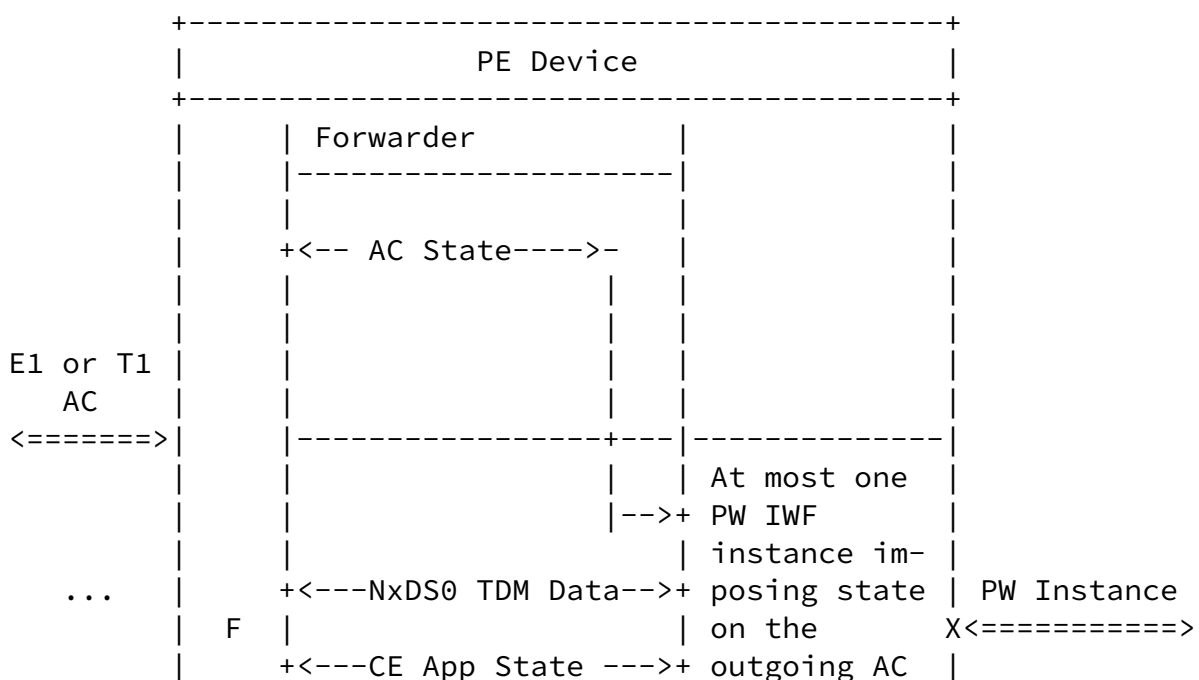
1. In the PSN-bound direction:
 - a) The Framer:
 - i) Detects frame alignment signal (FAS) and splits the incoming ACs into separate DS0 channels
 - ii) Detects special AC states
 - iii) If necessary, extracts CE application signals accompanying each of the separate DS0 services
 - b) The Forwarder:
 - i) Creates one or more NxDS0 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
2. In the CE-bound direction:
 - i) 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 NxDS0 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: This model is asymmetric:

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



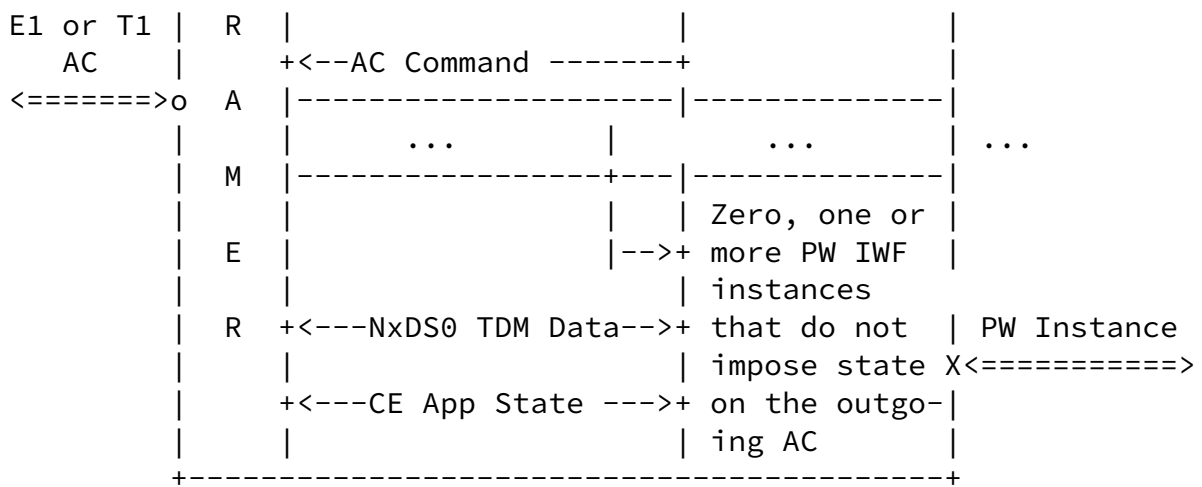


Figure B.1. Reference PE Architecture for NxDS0 Services