

Network Working Group
Request for Comments: 5086
Category: Informational

A. Vainshtein, Ed.
I. Sasson
Axerra Networks
E. Metz
KPN
T. Frost
Zarlink Semiconductor
P. Pate
Overture Networks
December 2007

Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation
Service over Packet Switched Network (CESoPSN)

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document describes a method for encapsulating structured (NxDS0) Time Division Multiplexed (TDM) signals as pseudowires over packet-switching networks (PSNs). In this regard, it complements similar work for structure-agnostic emulation of TDM bit-streams (see RFC 4553).

Table of Contents

1. Introduction	3
2. Terminology and Reference Models	3
2.1. Terminology	3
2.2. Reference Models	4
2.3. Requirements and Design Constraint	4
3. Emulated Services	5
4. CESoPSN Encapsulation Layer	6
4.1. CESoPSN Packet Format	6
4.2. PSN and Multiplexing Layer Headers	8
4.3. CESoPSN Control Word	9
4.4. Usage of the RTP Header	11
5. CESoPSN Payload Layer	12
5.1. Common Payload Format Considerations	12
5.2. Basic NxDS0 Services	13
5.3. Extending Basic NxDS0 Services with CE Application Signaling	15
5.4. Trunk-Specific NxDS0 Services with CAS	18
6. CESoPSN Operation	20
6.1. Common Considerations	20
6.2. IWF Operation	20
6.2.1. PSN-Bound Direction	20
6.2.2. CE-Bound Direction	20
6.3. CESoPSN Defects	23
6.4. CESoPSN PW Performance Monitoring	24
7. QoS Issues	25
8. Congestion Control	25
9. Security Considerations	27
10. IANA Considerations	27
11. Applicability Statement	27
12. Acknowledgements	29
13. Normative References	30
14. Informative References	31
Appendix A. A Common CE Application State Signaling Mechanism	33
Appendix B. Reference PE Architecture for Emulation of NxDS0 Services	34
Appendix C. Old Mode of CESoPSN Encapsulation Over L2TPV3	36

1. Introduction

This document describes a method for encapsulating structured (NxDS0) Time Division Multiplexed (TDM) signals as pseudowires over packet-switching networks (PSN). In this regard, it complements similar work for structure-agnostic emulation of TDM bit-streams [RFC4553].

Emulation of NxDS0 circuits provides for saving PSN bandwidth, and 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 Pseudowire Emulation Edge-to-Edge (PWE3) architecture described in [RFC3985], satisfies the general requirements put forth in [RFC3916], and specific requirements for structured TDM emulation put forth in [RFC4197].

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 [RFC3985], Section 1.4, and in [RFC4197], Section 3, are consistently used without additional explanations.

This document uses some terms and acronyms that are commonly used in conjunction with 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 and T1 circuits. These patterns are described in [G.704].
- 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 an LOS, OOF, or AIS condition is detected. Specific rules for encoding this pattern in the TDM framing are discussed in [G.775].

We also use the term Interworking Function (IWF) to describe the functional block that segments and encapsulates TDM into CESoPSN packets and, in the reverse direction, decapsulates CESoPSN packets and reconstitutes TDM.

2.2. Reference Models

Generic models that have been defined in Sections 4.1, 4.2, and 4.4 of [RFC3985] 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 [RFC4197], Section 4.3.

Structured services considered in this document represent special cases of the "Structured bit stream" payload type defined in Section 3.3.4 of [RFC3985]. 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 ([RFC3985], Section 3.3.5), the TDM payload is encapsulated without any changes.

2.3. Requirements and Design Constraints

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 a given pair of CEs 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 (e.g., Channel-Associated Signaling (CAS)-- see [RFC4197]): 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 remain unchanged.

3. Emulated Services

In accordance with [RFC4197], structured services considered in this specification are NxDS0 services, with and without CAS.

NxDS0 services are usually carried within appropriate physical trunks, and Provider Edges (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 Appendix 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.

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 follow the CESoPSN control word in all cases except UDP demultiplexing, where it MUST precede it (see Figures 1a, 1b, and 1c below).

Note: The difference in the CESoPSN packet formats for IP PSN using UDP-based demultiplexing and the rest of the PSN and demultiplexing combinations, is based on the following considerations:

1. Compliance with the existing header compression mechanisms for IPv4/IPv6 PSNs with UDP demultiplexing requires placing the RTP header immediately after the UDP header.
2. Compliance with the common PWE3 mechanisms for keeping PWs Equal Cost Multipath (ECMP)-safe for the MPLS PSN by providing for PW-IP packet discrimination (see [RFC3985], Section 5.4.3). This requires placing the PWE3 control word immediately after the PW label.
3. Commonality of the CESoPSN packet formats for MPLS networks and IPv4/IPv6 networks with Layer 2 Tunneling Protocol Version 3 (L2TPv3) demultiplexing facilitates smooth stitching of L2TPv3-based and MPLS-based segments of CESoPSN PWs (see [PWE3-MS]).




Figure 1a. CESoPSN Packet Format for an IPv4/IPv6 PSN with UDP demultiplexing




Figure 1b. CESoPSN Packet Format for an MPLS PSN




Figure 1c. CESoPSN Packet Format for an IPv4/IPv6 PSN with L2TPv3 Demultiplexing

4.2. 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.

Usage of MPLS and L2TPv3 as demultiplexing layers is explained in [RFC3985] and [RFC3931], respectively.

Setup and maintenance of CESoPSN PWs over MPLS PSN is described in [PWE3-TDM-CONTROL].

Setup and maintenance of CESoPSN PWs over IPv4/IPv6 using L2TPv3 demultiplexing is defined in [L2TPEXT-TDM].

The destination UDP port MUST be used to multiplex and demultiplex individual PWs between nodes. Architecturally (see [RFC3985]) this makes the destination UDP port act as the PW Label.

UDP ports MUST be manually configured by both endpoints of the PW. The configured destination port together with both the source and destination IP addresses uniquely identifies the PW for the receiver. All UDP port values that function as PW labels SHOULD be in the range of dynamically allocated UDP port numbers (49152 through 65535).

While many UDP-based protocols are able to traverse middleboxes without dire consequences, the use of UDP ports as PW labels makes middlebox traversal more difficult. Hence, it is NOT RECOMMENDED to use UDP-based PWs where port-translating middleboxes are present between PW endpoints.

4.3. CESoPSN Control Word

The structure of the CESoPSN Control Word that **MUST** be used with all combinations of the PSN and demultiplexing mechanisms described in the previous section is shown in Figure 2 below.




Figure 2. Structure of the CESoPSN Control Word

The use of Bits 0 to 3 is described in [RFC4385]. These bits MUST be set to zero unless they are being used to indicate the start of an Associated Channel Header (ACH). An ACH is needed if the state of the CESoPSN PW is being monitored using Virtual Circuit Connectivity Verification [RFC5085].

L - if set, indicates some abnormal condition of the attachment circuit.

M - a 2-bit modifier field. In case of L cleared, this field allows discrimination of signaling packets and carrying RDI of the attachment circuit across the PSN. 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 pre-configured 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 pre-configured 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: 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

Notes:

1. Bits in the M field are shown in the same order as in Figure 2 (i.e., bit 6 of the CW followed by bit 7 of the CW).
2. 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 [RFC4623]; see Section 5.4 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. Note: If fixed RTP header is used in the encapsulation, it is considered part of the CESoPSN header.

The sequence number is used to provide the common PW sequencing function, as well as detection of lost packets. It MUST be generated in accordance with the rules defined in Section 5.1 of [RFC3550] for the RTP sequence number, i.e.:

- o Its space is a 16-bit unsigned circular space
- o Its initial value SHOULD be random (unpredictable)
- o It MUST be incremented with each CESoPSN data packet sent in the specific PW.

4.4. Usage of the RTP Header

Although CESoPSN MAY employ an RTP header when explicit transfer of timing information is required, this is purely formal reuse of the header format. RTP mechanisms, such as header extensions, contributing source (CSRC) list, padding, RTP Control Protocol (RTCP), RTP header compression, Secure RTP (SRTP), etc., are not applicable to CESoPSN pseudowires.

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 that preserve 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 bandwidth. 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 NxDS0 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 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 Figure 3 below.




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

This mode of operation complies with the recommendation in [RFC3985] 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:

$$L = 8 * N * D$$

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 $2 \leq N \leq 4$: 4 millisecond (with the corresponding packet payload size of $32 \cdot N$ bytes)
- o For $N \geq 5$: 1 millisecond (with the corresponding packet payload size of $8 \cdot N$ 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 described 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.

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

Signaling packets SHOULD be carried in 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 an 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 Appendix A.

Encoding of CE application state for telephony applications using CAS follows [RFC2833] (which has since been obsoleted by [RFC4733] and [RFC4734], but they do not affect the relevant text).

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




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 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) $16*N + \text{floor}((N+1)/2)$ for E1-NxDS0
 - b) $24*N + \text{floor}((N+1)/2)$ 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 [RFC4623].
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.

Format of CESoPSN packets carrying trunk-specific NxDS0 service with CAS that do and do not contain signaling substructures is shown in Figures 5 (a) and (b), respectively. In these figures, the number of the trunk frames per multiframe fragment ("m") MUST be an integer divisor of the number of frames per trunk 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

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 [RFC4447].

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 MUST detect lost and misordered packets. It SHOULD use the sequence number in the control word for these purposes but, if the RTP header is used, the RTP sequence number MAY be used instead.

The CE-bound CESoPSN IWF MAY reorder misordered packets. Misordered 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 bandwidth-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 pseudowires 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 pre-configured amount of TDM data (usually half of the jitter buffer) has been received in consecutive CESoPSN packets, or until a pre-configured intermediate state timer expires.

Once the pre-configured 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 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 pre-configured 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 pre-configured 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 pre-configured amount of time (typically 2.5 seconds) and MUST be cleared after the corresponding defect is undetected for a second pre-configured 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.

Malformed packets MAY be 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. Other methods of detecting malformed packets are implementation-specific. 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 pre-configured 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 per-domain behavior (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

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion. CESoPSN PWs represent inelastic, constant bit rate (CBR) flows and cannot respond to congestion in a TCP-friendly manner prescribed by [RFC2914], although the percentage of total bandwidth they consume remains constant.

Unless appropriate precautions are taken, undiminished demand of bandwidth by CESoPSN PWs can contribute to network congestion that may impact network control protocols.

Whenever possible, CESoPSN PWs SHOULD be carried across traffic-engineered PSNs that provide either bandwidth reservation and admission control or forwarding prioritization and boundary traffic conditioning mechanisms. IntServ-enabled domains supporting Guaranteed Service (GS) [RFC2212] and Diffserv-enabled domains [RFC2475] supporting Expedited Forwarding (EF) [RFC3246] provide examples of such PSNs. Such mechanisms will negate, to some degree, the effect of the CESoPSN PWs on the neighboring streams. In order to facilitate boundary traffic conditioning of CESoPSN traffic over IP PSNs, the CESoPSN IP packets SHOULD NOT use the Diffserv Code Point (DSCP) value reserved for the Default PHB [RFC2474].

If CESoPSN PWs run over a PSN providing best-effort service, they SHOULD monitor packet loss in order to detect "severe congestion". If such a condition is detected, a CESoPSN PW SHOULD shut down bidirectionally for some period of time as described in Section 6.5 of [RFC3985].

Note that:

1. The CESoPSN IWF can inherently provide packet loss measurement, since the expected rate of arrival of CESoPSN packets is fixed and known

2. The results of the CESoPSN packet loss measurement may not be a reliable indication of presence or absence of severe congestion if the PSN provides enhanced delivery, e.g.,:
 - a) If CESoPSN traffic takes precedence over non-CESoPSN traffic, severe congestion can develop without significant CESoPSN packet loss.
 - b) If non-CESoPSN traffic takes precedence over CESoPSN traffic, CESoPSN may experience substantial packet loss due to a short-term burst of high-priority traffic.
3. The TDM services emulated by the CESoPSN PWs have high availability objectives (see [G.826]) that MUST be taken into account when deciding on temporary shutdown of CESoPSN PWs.

This specification does not define the exact criteria for detecting "severe congestion" using the CESoPSN packet loss rate, or the specific methods for bidirectional shutdown that the CESoPSN PWs (when such severe congestion has been detected) and their consequent restart after a suitable delay. This is left for further study. However, the following considerations may be used as guidelines for implementing the CESoPSN severe congestion shutdown mechanism:

1. CESoPSN Performance Monitoring techniques (see Section 6.4) provide entry and exit criteria for the CESoPSN PW "Unavailable" state that make it closely correlated with the "Unavailable" state of the emulated TDM circuit as specified in [G.826]. Using the same criteria for "severe congestion" detection may decrease the risk of shutting down the CESoPSN PW while the emulated TDM circuit is still considered available by the CE.
2. If the CESoPSN PW has been set up using either PWE3 control protocol [RFC4447] or L2TPv3 [RFC3931], the regular PW teardown procedures of these protocols SHOULD be used.
3. If one of the CESoPSN PW end points stops transmission of packets for a sufficiently long period, its peer (observing 100% packet loss) will necessarily detect "severe congestion" and also stop transmission, thus achieving bidirectional PW shutdown.

9. Security Considerations

CESoPSN does not enhance or detract from the security performance of the underlying PSN; rather, it relies upon the PSN mechanisms for encryption, integrity, and authentication whenever required.

CESoPSN PWs share susceptibility to a number of pseudowire-layer attacks, and will use whatever mechanisms for confidentiality, integrity, and authentication that are developed for general PWs. These methods are beyond the scope of this document.

Although CESoPSN PWs MAY employ an RTP header when explicit transfer of timing information is required, it is not possible to use SRTP (see [RFC3711]) mechanisms as a substitute for PW layer security.

Misconnection detection capabilities of CESoPSN increase its resilience to misconfiguration and some types of DoS attacks.

Random initialization of sequence numbers, in both the control word and the optional RTP header, makes known-plaintext attacks on encrypted CESoPSN PWs more difficult. Encryption of PWs is beyond the scope of this document.

10. IANA Considerations

Allocation of PW Types for the corresponding CESoPSN PWs is defined in [RFC4446].

11. Applicability Statement

CESoPSN is an encapsulation layer intended for carrying NxDS0 services with or without CAS over PSN.

CESoPSN allows emulation of certain end-to-end delay properties of TDM networks. In particular, the end-to-end delay of a TDM circuit emulated by a CESoPSN PW does not depend upon the bit rate of the 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 bandwidth-saving Fractional point-to-point E1/T1 applications. These applications can be described as the following:

- o The pair of CE devices operates as if it was connected by an emulated E1 or T1 circuit. In particular, it reacts to AIS and RAI states of its 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 bandwidth.

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 re-convergence 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.

Security of TDM services provided by CESoPSN across a shared PSN may be below the level of security traditionally associated with TDM services carried across TDM networks.

12. 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 feedback.

We thank Alik Shimelmits for many fruitful discussions.

13. Normative References

- [ATM-CES] The ATM Forum Technical Committee. Circuit Emulation Service Interoperability Specification version 2.0 af-vtoa-0078.000, January 1997.
- [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.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
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2833] Schulzrinne, H. and S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, May 2000.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.
- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC 3086, April 2001.
- [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.
- [RFC4197] Riegel, M., "Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks", RFC 4197, October 2005.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [RFC4447] Martini L. et al, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP), RFC 4447, April 2006
- [RFC4623] Malis, A. and M. Townsley, "Pseudowire Emulation Edge-to-Edge (PWE3) Fragmentation and Reassembly", RFC 4623, August 2006.
- [RTP-TYPES] RTP PARAMETERS, <<http://www.iana.org/assignments/rtp-parameters>>.
- [TR-NWT-170] Digital Cross Connect Systems - Generic Requirements and Objectives, Bellcore, TR-NWT-170, January 1993

14. Informative References

- [L2TPEXT-TDM] Vainshtein, A. and S. Galtsur, "Layer Two Tunneling Protocol - Setup of TDM Pseudowires", Work in Progress, February 2007.
- [PWE3-MS] Martini, L., Metz, C., Nadeau, T., and M. Duckett, "Segmented Pseudo Wire", Work in Progress, November 2007.
- [PWE3-TDM-CONTROL] Vainshtein, A. and Y. Stein, "Control Protocol Extensions for Setup of TDM Pseudowires in MPLS Networks", Work in Progress, November 2007.
- [RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, September 1997.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC4553] Vainshtein, A. and YJ. Stein, "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, June 2006.
- [RFC4733] Schulzrinne, H. and T. Taylor, "RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals", RFC 4733, December 2006.
- [RFC4734] Schulzrinne, H. and T. Taylor, "Definition of Events for Modem, Fax, and Text Telephony Signals", RFC 4734, December 2006.
- [RFC5085] Nadeau, T., Ed., and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", Work in Progress, RFC 5085, December 2007.

Appendix 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 10 ms period (i.e., before all 3 signaling packets reporting the previous change have been sent), this process is re-started, i.e.:
 - i) The unsent signaling packet(s) with the previous CE application state are discarded
 - ii) Triple send of packets with the new CE application state begins.
 - c) Loss of packets defect has been cleared
 - 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.

Appendix 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 Service 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 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 Figure 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:
- 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 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

Appendix C. Old Mode of CESoPSN Encapsulation Over L2TPV3

Previous versions of this specification defined a CESoPSN PW encapsulation over L2TPv3, which differs from one described in Section 4.1 and Figure 1c. In these versions, the RTP header, if used, precedes the CESoPSN control word.

Existing implementations of the old encapsulation mode MUST be distinguished from the encapsulations conforming to this specification via the CESoPSN PW setup.

Authors' Addresses

Alexander ("Sasha") Vainshtein
Axerra Networks
24 Raoul Wallenberg St.,
Tel Aviv 69719, Israel
EMail: sasha@axerra.com, vainshtein.alex@gmail.com

Israel Sasson
Axerra Networks
24 Raoul Wallenberg St.,
Tel Aviv 69719, Israel
EMail: israel@axerra.com

Eduard Metz
KPN
Regulusweg 1
2316 AC The Hague
Netherlands
EMail: eduard.metz@kpn.com

Tim Frost
Symmetricom, Inc.
Tamerton Road
Roborough, Plymouth
PL6 7BQ, UK
EMail: tfrost@symmetricom.com

Prayson Pate
Overture Networks
507 Airport Boulevard
Building 111
Morrisville, North Carolina 27560 USA
EMail: prayson.pate@overturenetworks.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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, THE IETF TRUST 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.

Intellectual Property

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.

