

Network Working Group
Request for Comments: 4197
Category: Informational

M. Riegel
Siemens AG
October 2005

Requirements for Edge-to-Edge Emulation of
Time Division Multiplexed (TDM) Circuits over
Packet Switching Networks

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.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document defines the specific requirements for edge-to-edge emulation of circuits carrying Time Division Multiplexed (TDM) digital signals of the Plesiochronous Digital Hierarchy as well as the Synchronous Optical Network/Synchronous Digital Hierarchy over packet-switched networks. It is aligned to the common architecture for Pseudo Wire Emulation Edge-to-Edge (PWE3). It makes references to the generic requirements for PWE3 where applicable and complements them by defining requirements originating from specifics of TDM circuits.

Table of Contents

1. Introduction	3
1.1. TDM Circuits Belonging to the PDH Hierarchy	3
1.1.1. TDM Structure and Transport Modes	4
1.2. SONET/SDH Circuits	4
2. Motivation	5
3. Terminology	6
4. Reference Models	7
4.1. Generic PWE3 Models	7
4.2. Clock Recovery	7
4.3. Network Synchronization Reference Model	8
4.3.1. Synchronous Network Scenarios	10
4.3.2. Relative Network Scenario	12
4.3.3. Adaptive Network Scenario	12
5. Emulated Services	13
5.1. Structure-Agnostic Transport of Signals out of the PDH Hierarchy	13
5.2. Structure-Aware Transport of Signals out of the PDH Hierarchy	14
5.3. Structure-Aware Transport of SONET/SDH Circuits	14
6. Generic Requirements	14
6.1. Relevant Common PW Requirements	14
6.2. Common Circuit Payload Requirements	15
6.3. General Design Issues	16
7. Service-Specific Requirements	16
7.1. Connectivity	16
7.2. Network Synchronization	16
7.3. Robustness	16
7.3.1. Packet loss	17
7.3.2. Out-of-order delivery	17
7.4. CE Signaling	17
7.5. PSN Bandwidth Utilization	18
7.6. Packet Delay Variation	19
7.7. Compatibility with the Existing PSN Infrastructure	19
7.8. Congestion Control	19
7.9. Fault Detection and Handling	20
7.10. Performance Monitoring	20
8. Security Considerations	20
9. References	20
9.1. Normative References	20
9.2. Informative References	21
10. Contributors Section	22

1. Introduction

This document defines the specific requirements for edge-to-edge emulation of circuits carrying Time Division Multiplexed (TDM) digital signals of the Plesiochronous Digital Hierarchy (PDH) as well as the Synchronous Optical NETWORK (SONET)/Synchronous Digital Hierarchy (SDH) over Packet-Switched Networks (PSN). It is aligned to the common architecture for Pseudo Wire Emulation Edge-to-Edge (PWE3) as defined in [RFC3985]. It makes references to requirements in [RFC3916] where applicable and complements [RFC3916] by defining requirements originating from specifics of TDM circuits.

The term "TDM" will be used in this documents as a general descriptor for the synchronous bit streams belonging to either the PDH or the SONET/SDH hierarchies.

1.1. TDM Circuits Belonging to the PDH Hierarchy

The bit rates traditionally used in various regions of the world are detailed in the normative reference [G.702]. For example, in North America, the T1 bit stream of 1.544 Mbps and the T3 bit stream of 44.736 Mbps are mandated, while in Europe, the E1 bit stream of 2.048 Mbps and the E3 bit stream of 34.368 Mbps are utilized.

Although TDM can be used to carry unstructured bit streams at the rates defined in [G.702], there is a standardized method of carrying bit streams in larger units called frames, each frame contains the same number of bits.

Related to the sampling frequency of voice traffic the bitrate is always a multiple of 8000, hence the T1 frame consists of 193 bits and the E1 frame of 256 bits. The number of bits in a frame is called the frame size.

The framing is imposed by introducing a periodic pattern into the bit stream to identify the boundaries of the frames (e.g., 1 framing bit per T1 frame, a sequence of 8 framing bits per E1 frame). The details of how these framing bits are generated and used are elucidated in [G.704], [G.706], and [G.751]. Unframed TDM has all bits available for payload.

Framed TDM is often used to multiplex multiple channels (e.g., voice channels each consisting of 8000 8-bit-samples per second) in a sequence of "timeslots" recurring in the same position in each frame. This multiplexing is called "channelized TDM" and introduces additional structure.

In some cases, framing also defines groups of consecutive frames called multiframes. Such grouping imposes an additional level of structure on the TDM bit-stream.

1.1.1. TDM Structure and Transport Modes

Unstructured TDM:

TDM that consists of a raw bit-stream of rate defined in [G.702], with all bits available for payload.

Structured TDM:

TDM with one or more levels of structure delineation, including frames, channelization, and multiframes (e.g., as defined in [G.704], [G.751], and [T1.107]).

Structure-Agnostic Transport:

Transport of unstructured TDM, or of structured TDM when the structure is deemed inconsequential from the transport point of view. In structure-agnostic transport, any structural overhead that may be present is transparently transported along with the payload data, and the encapsulation provides no mechanisms for its location or utilization.

Structure-Aware Transport:

Transport of structured TDM taking at least some level of the structure into account. In structure-aware transport, there is no guarantee that all bits of the TDM bit-stream will be transported over the PSN network (specifically, the synchronization bits and related overhead may be stripped at ingress and usually will be regenerated at egress) or that transported bits will be situated in the packet in their original order (but in this case, bit order is usually recovered at egress; one known exception is loss of multiframe synchronization between the TDM data and CAS bits introduced by a digital cross-connect acting as a Native Service Processing (NSP) block, see [TR-NWT-170]).

1.2. SONET/SDH Circuits

The term SONET refers to the North American Synchronous Optical NETWORK as specified by [T1.105]. It is based on the concept of a Nx783 byte payload container repeated every 125us. This payload is referred to as an STS-1 SPE and may be concatenated into higher bandwidth circuits (e.g., STS-Nc) or sub-divided into lower bandwidth circuits (Virtual Tributaries). The higher bandwidth concatenated circuits can be used to carry anything from IP Packets to ATM cells to Digital Video Signals. Individual STS-1 SPEs are frequently used

to carry individual DS3 or E3 TDM circuits. When the 783 byte containers are sub-divided for lower rate payloads, they are frequently used to carry individual T1 or E1 TDM circuits.

The Synchronous Digital Hierarchy (SDH) is the international equivalent and enhancement of SONET and is specified by [G.707].

Both SONET and SDH include a substantial amount of transport overhead that is used for performance monitoring, fault isolation, and other maintenance functions along different types of optical or electrical spans. This also includes a pointer-based mechanism for carrying payloads asynchronously. In addition, the payload area includes dedicated overhead for end-to-end performance monitoring, fault isolation, and maintenance for the service being carried. If the main payload area is sub-divided into lower rate circuits (such as T1/E1), additional overhead is included for end-to-end monitoring of the individual T1/E1 circuits.

This document discusses the requirements for emulation of SONET/SDH services. These services include end-to-end emulation of the SONET payload (STS-1 SPE), emulation of concatenated payloads (STS-Nc SPE), as well as emulation of a variety of sub-STS-1 rate circuits jointly referred to as Virtual Tributaries (VT) and their SDH analogs.

2. Motivation

[RFC3916] specifies common requirements for edge-to-edge emulation of circuits of various types. However, these requirements, as well as references in [RFC3985], do not cover specifics of PWs carrying TDM circuits.

The need for a specific document to complement [RFC3916] addressing of edge-to-edge emulation of TDM circuits arises from the following:

- o Specifics of the TDM circuits. For example,
 - * the need for balance between the clock of ingress and egress attachment circuits in each direction of the Pseudo Wire (PW),
 - * the need to maintain jitter and wander of the clock of the egress end service, within the limits imposed by the appropriate normative documents, in the presence of the packet delay variation produced by the PSN.
- o Specifics of applications using TDM circuits. For example, voice applications,
 - * put special emphasis on minimization of one-way delay, and

- * are relatively tolerant to errors in data.
- o Other applications might have different specifics. For example, transport of signaling information
 - * is relatively tolerant to one-way delay, and
 - * is sensitive to errors in transmitted data.
- o Specifics of the customers' expectations regarding end-to-end behavior of services that contain emulated TDM circuits. For example, experience with carrying such services over SONET/SDH networks increases the need for
 - * isolation of problems introduced by the PSN from those occurring beyond the PSN bounds,
 - * sensitivity to misconnection,
 - * sensitivity to unexpected connection termination, etc.

3. 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 are used consistently. However some terms and acronyms are used in conjunction with the TDM services. In particular:

TDM networks employ Channel-Associated Signaling (CAS) or Common Channel Signaling (CCS) to supervise and advertise status of telephony applications, provide alerts to these applications (as to requests to connect or disconnect), and to transfer routing and addressing information. These signals must be reliably transported over the PSNs for the telephony end-systems to function properly.

CAS (Channel-Associated Signaling)

CAS is carried in the same T1 or E1 frame as the voice signals, but not in the speech band. Since CAS signaling may be transferred at a rate slower than the TDM traffic in a timeslot, one need not update all the CAS bits in every TDM frame. Hence, CAS systems cycle through all the signaling bits only after some number of TDM frames, which defines a new structure known as a multiframe or superframe. Common multiframes are 12, 16, or 24 frames in length, corresponding to 1.5, 2, and 3 milliseconds in duration.

CCS (Common Channel Signaling)

CCS signaling uses a separate digital channel to carry asynchronous messages pertaining to the state of telephony applications over related TDM timeslots of a TDM trunk. This channel may be physically situated in one or more adjacent timeslots of the same TDM trunk (trunk associated CCS) or may be transported over an entirely separate network.

CCS is typically HDLC-based, with idle codes or keep-alive messages being sent until a signaling event (e.g., on-hook or off-hook) occurs. Examples of HDLC-based CCS systems are SS7 [Q.700] and ISDN PRI signaling [Q.931].

Note: For the TDM network, we use the terms "jitter" and "wander" as defined in [G.810] to describe short- and long-term variance of the significant instants of the digital signal, while for the PSN we use the term packet delay variation (PDV) (see [RFC3393]).

4. Reference Models

4.1. Generic PWE3 Models

Generic models that have been defined in [RFC3985] in sections

- 4.1 (Network Reference Model),
- 4.2 (PWE3 Pre-processing),
- 4.3 (Maintenance Reference Model),
- 4.4 (Protocol Stack Reference Model) and
- 4.5 (Pre-processing Extension to Protocol Stack Reference Model).

They are fully applicable for the purposes of this document without modification.

All the services considered in this document represent special cases of the Bit-stream and Structured bit-stream payload type defined in Section 3.3 of [RFC3985].

4.2. Clock Recovery

Clock recovery is extraction of the transmission bit timing information from the delivered packet stream. Extraction of this information from a highly jittered source, such as a packet stream, may be a complex task.

4.3. Network Synchronization Reference Model

Figure 1 shows a generic network synchronization reference model.

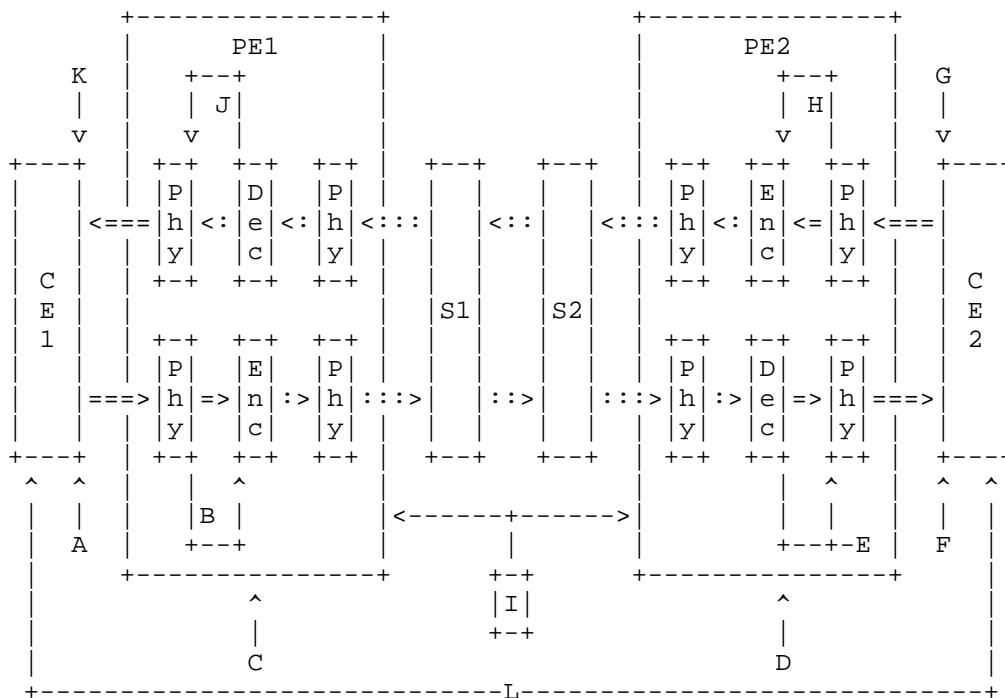


Figure 1: The Network Synchronization Reference Model

The following notation is used in Figure 1:

CE1, CE2

Customer edge devices terminating TDM circuits to be emulated.

PE1, PE2

Provider edge devices adapting these end services to PW.

S1, S2

Provider core routers.

Phy

Physical interface terminating the TDM circuit.

Enc

PSN-bound interface of the PW, where the encapsulation takes place.

Dec

CE-bound interface of the PW, where the decapsulation takes place. It contains a compensation buffer (also known as the "jitter buffer") of limited size.

"==>"

TDM attachment circuits.

"::>"

PW providing edge-to-edge emulation for the TDM circuit.

The characters "A" - "L" denote various clocks:

"A"

The clock used by CE1 for transmission of the TDM attachment circuit towards CE1.

"B"

The clock recovered by PE1 from the incoming TDM attachment circuit. "A" and "B" always have the same frequency.

"G"

The clock used by CE2 for transmission of the TDM attachment circuit towards CE2.

"H"

The clock recovered by PE2 from the incoming TDM attachment circuit. "G" and "H" always have the same frequency.

"C", "D"

Local oscillators available to PE1 and PE2, respectively.

"E"

Clock used by PE2 to transmit the TDM attachment service circuit to CE2 (the recovered clock).

"F"

Clock recovered by CE2 from the incoming TDM attachment service ("E" and "F" have the same frequency).

"I"

If the clock exists, it is the common network reference clock available to PE1 and PE2.

"J"

Clock used by PE1 to transmit the TDM attachment service circuit to CE1 (the recovered clock).

"K"

Clock recovered by CE1 from the incoming TDM attachment service ("J" and "K" have the same frequency).

"L"

If it exists, it is the common reference clock of CE1 and CE2. Note that different pairs of CE devices may use different common reference clocks.

A requirement of edge-to-edge emulation of a TDM circuit is that clock "B" and "E", as well as clock "H" and "J", are of the same frequency. The most appropriate method will depend on the network synchronization scheme.

The following groups of synchronization scenarios can be considered:

4.3.1. Synchronous Network Scenarios

Depending on which part of the network is synchronized by a common clock, there are two scenarios:

- o PE Synchronized Network:

Figure 2 is an adapted version of the generic network reference model, and presents the PE synchronized network scenario.

The common network reference clock "I" is available to all the PE devices, and local oscillators "C" and "D" are locked to "I":

- * Clocks "E" and "J" are the same as "D" and "C", respectively.
- * Clocks "A" and "G" are the same as "K" and "F", respectively (i.e., CE1 and CE2 use loop timing).




Figure 2: PE Synchronized Scenario

o CE Synchronized Network:

Figure 3 is an adapted version of the generic network reference model, and presents the CE synchronized network scenario.

The common network reference clock "L" is available to all the CE devices, and local oscillators "A" and "G" are locked to "L":

* Clocks "E" and "J" are the same as "G" and "A", respectively (i.e., PE1 and PE2 use loop timing).




Figure 3: CE Synchronized Scenario

No timing information has to be transferred in these cases.

4.3.2. Relative Network Scenario

In this case, each CE uses its own transmission clock source that must be carried across the PSN and recovered by the remote PE, respectively. The common PE clock "I" can be used as reference for this purpose.

Figure 4 shows the relative network scenario.

The common network reference clock "I" is available to all the PE devices, and local oscillators "C" and "D" are locked to "I":

- o Clocks "A" and "G" are generated locally without reference to a common clock.
- o Clocks "E" and "J" are generated in reference to a common clock available at all PE devices.

In a slight modification of this scenario, one (but not both!) of the CE devices may use its receive clock as its transmission clock (i.e., use loop timing).




Figure 4: Relative Network Scenario Timing

In this case, timing information (the difference between the common reference clock "I" and the incoming clock "A") MUST be explicitly transferred from the ingress PE to the egress PE.

4.3.3. Adaptive Network Scenario

The adaptive scenario is characterized by:

- o No common network reference clock "I" is available to PE1 and PE2.
- o No common reference clock "L" is available to CE1 and CE2.

Figure 5 presents the adaptive network scenario.




Figure 5: Adaptive Scenario

Synchronizing clocks "A" and "E" in this scenario is more difficult than it is in the other scenarios.

Note that the tolerance between clocks "A" and "E" must be small enough to ensure that the jitter buffer does not overflow or underflow.

In this case, timing information MAY be explicitly transferred from the ingress PE to the egress PE, e.g., by RTP.

5. Emulated Services

This section defines requirements for the payload and encapsulation layers for edge-to-edge emulation of TDM services with bit-stream payload as well as structured bit-stream payload.

Wherever possible, the requirements specified in this document SHOULD be satisfied by appropriate arrangements of the encapsulation layer only. The (rare) cases when the requirements apply to both the encapsulation and payload layers (or even to the payload layer only) will be explicitly noted.

The service-specific encapsulation layer for edge-to-edge emulation comprises the following services over a PSN.

5.1. Structure-Agnostic Transport of Signals out of the PDH Hierarchy

Structure-agnostic transport is considered for the following signals:

- o E1 as described in [G.704].
- o T1 (DS1) as described in [G.704].

- o E3 as defined in [G.751].
- o T3 (DS3) as described in [T1.107].

5.2. Structure-Aware Transport of Signals out of the PDH Hierarchy

Structure-aware transport is considered for the following signals:

- o E1/T1 with one of the structures imposed by framing as described in [G.704].
- o NxDS0 with or without CAS.

5.3. Structure-Aware Transport of SONET/SDH Circuits

Structure-aware transport is considered for the following SONET/SDH circuits:

- o SONET STS-1 synchronous payload envelope (SPE)/SDH VC-3.
- o SONET STS-Nc SPE (N = 3, 12, 48, 192) / SDH VC-4, VC-4-4c, VC-4-16c, VC-4-64c.
- o SONET VT-N (N = 1.5, 2, 3, 6) / SDH VC-11, VC-12, VC-2.
- o SONET Nx VT-N / SDH Nx VC-11/VC-12/VC-2/VC-3.

Note: There is no requirement for the structure-agnostic transport of SONET/SDH. For this case, it would seem that structure must be taken into account.

6. Generic Requirements

6.1. Relevant Common PW Requirements

The encapsulation and payload layers MUST conform to the common PW requirements defined in [RFC3916]:

1. Conveyance of Necessary Header Information:

- A. For structure-agnostic transport, this functionality MAY be provided by the payload layer.
- B. For structure-aware transport, the necessary information MUST be provided by the encapsulation layer.

- C. Structure-aware transport of SONET/SDH circuits MUST preserve path overhead information as part of the payload. Relevant components of the transport overhead MAY be carried in the encapsulation layer.
2. Support of Multiplexing and Demultiplexing if supported by the native services. This is relevant for Nx DS0 circuits (with or without signaling) and Nx VT-x in a single STS-1 SPE or VC-4.:
 - A. For these circuits, the combination of encapsulation and payload layers MUST provide for separate treatment of every sub-circuit.
 - B. Enough information SHOULD be provided by the pseudo wire to allow multiplexing and demultiplexing by the NSP. Reduction of the complexity of the PW emulation by using NSP circuitry for multiplexing and demultiplexing MAY be the preferred solution.
 3. Intervention or transparent transfer of Maintenance Messages of the Native Services, depending on the particular scenario.
 4. Consideration of Per-PSN Packet Overhead (see also Section 7.5 below).
 5. Detection and handling of PW faults. The list of faults is given in Section 7.9 below.

Fragmentation indications MAY be used for structure-aware transport when the structures in question either exceed desired packetization delay or exceed Path MTU between the pair of PEs.

The following requirement listed in [RFC3916] is not applicable to emulation of TDM services:

- o Support of variable length PDUs.

6.2. Common Circuit Payload Requirements

Structure-agnostic transport treats TDM circuits as belonging to the 'Bit-stream' payload type defined in [RFC3985].

Structure-aware transport treats these circuits as belonging to the "Structured bit-stream" payload type defined in [RFC3985].

Accordingly, the encapsulation layer MUST provide the common Sequencing service and SHOULD provide Timing information (Synchronization services) when required (see Section 4.3 above).

Note: Length service MAY be provided by the encapsulation layer, but is not required.

6.3. General Design Issues

The combination of payload and encapsulation layers SHOULD comply with the general design principles of the Internet protocols as presented in Section 3 of [RFC1958] and [RFC3985].

If necessary, the payload layer MAY use some forms of adaptation of the native TDM payload in order to achieve specific, well-documented design objectives. In these cases, standard adaptation techniques SHOULD be used.

7. Service-Specific Requirements

7.1. Connectivity

1. The emulation MUST support the transport of signals between Attachment Circuits (ACs) of the same type (see Section 5) and, wherever appropriate, bit-rate.
2. The encapsulation layer SHOULD remain unaffected by specific characteristics of connection between the ACs and PE devices at the two ends of the PW.

7.2. Network Synchronization

1. The encapsulation layer MUST provide synchronization services that are sufficient to:
 - A. match the ingress and egress end service clocks regardless of the specific network synchronization scenario, and
 - B. keep the jitter and wander of the egress service clock within the service-specific limits defined by the appropriate normative references.
2. If the same high-quality synchronization source is available to all the PE devices in the given domain, the encapsulation layer SHOULD be able to make use of it (e.g., for better reconstruction of the native service clock).

7.3. Robustness

The robustness of the emulated service depends not only upon the edge-to-edge emulation protocol, but also upon proper implementation of the following procedures.

7.3.1. Packet loss

Edge-to-edge emulation of TDM circuits MAY assume very low probability of packet loss between ingress and egress PE. In particular, no retransmission mechanisms are required.

In order to minimize the effect of lost packets on the egress service, the encapsulation layer SHOULD:

1. Enable independent interpretation of TDM data in each packet by the egress PE (see [RFC2736]). This requirement MAY be disregarded if the egress PE needs to interpret structures that exceed the path MTU between the ingress and egress PEs.
2. Allow reliable detection of lost packets (see next section). In particular, it SHOULD allow estimation of the arrival time of the next packet and detection of lost packets based on this estimate.
3. Minimize possible effect of lost packets on recovery of the circuit clock by the egress PE.
4. Increase the resilience of the CE TDM interface to packet loss by allowing the egress PE to substitute appropriate data.

7.3.2. Out-of-order delivery

The encapsulation layer MUST provide the necessary mechanisms to guarantee ordered delivery of packets carrying the TDM data over the PSN. Packets that have arrived out-of-order:

1. MUST be detected, and
2. SHOULD be reordered if not judged to be too late or too early for playout.

Out-of-order packets that cannot be reordered MUST be treated as lost.

7.4. CE Signaling

Unstructured TDM circuits would not usually require any special mechanism for carrying CE signaling as this would be carried as part of the emulated service.

Some CE applications using structured TDM circuits (e.g., telephony) require specific signaling that conveys the changes of state of these applications relative to the TDM data.

The encapsulation layer SHOULD support signaling of state of CE applications for the relevant circuits providing for:

1. Ability to support different signaling schemes with minimal impact on encapsulation of TDM data,
2. Multiplexing of application-specific CE signals and data of the emulated service in the same PW,
3. Synchronization (within the application-specific tolerance limits) between CE signals and data at the PW egress,
4. Probabilistic recovery against possible, occasional loss of packets in the PSN, and
5. Deterministic recovery of the CE application state after PW setup and network outages.

CE signaling that is used for maintenance purposes (loopback commands, performance monitoring data retrieval, etc.) SHOULD use the generic PWE3 maintenance protocol.

7.5. PSN Bandwidth Utilization

1. The encapsulation layer SHOULD allow for an effective trade-off between the following requirements:
 - A. Effective PSN bandwidth utilization. Assuming that the size of the encapsulation layer header does not depend on the size of its payload, an increase in the packet payload size results in increased efficiency.
 - B. Low edge-to-edge latency. Low end-to-end latency is the common requirement for Voice applications over TDM services. Packetization latency is one of the components comprising edge-to-edge latency, and it decreases with the packet payload size.

The compensation buffer used by the CE-bound IWF increases latency to the emulated circuit. Additional delays introduced by this buffer SHOULD NOT exceed the packet delay variation observed in the PSN.

2. The encapsulation layer MAY provide for saving PSN bandwidth by not sending corrupted TDM data across the PSN.

3. The encapsulation layer MAY provide the ability to save the PSN bandwidth for the structure-aware case by not sending channels that are permanently inactive.
4. The encapsulation layer MAY enable the dynamic suppression of temporarily unused channels from transmission for the structure-aware case.

If used, dynamic suppression of temporarily unused channels MUST NOT violate the integrity of the structures delivered over the PW.

5. For NxDS0, the encapsulation layer MUST provide the ability to keep the edge-to-edge delay independent of the service rate.

7.6. Packet Delay Variation

The encapsulation layer SHOULD provide for the ability to compensate for packet delay variation, while maintaining jitter and wander of the egress end service clock with tolerances specified in the normative references.

The encapsulation layer MAY provide for run-time adaptation of delay introduced by the jitter buffer if the packet delay variation varies with time. Such an adaptation MAY introduce a low level of errors (within the limits tolerated by the application) but SHOULD NOT introduce additional wander of the egress end service clock.

7.7. Compatibility with the Existing PSN Infrastructure

The combination of encapsulation and PSN tunnel layers used for edge-to-edge emulation of TDM circuits SHOULD be compatible with existing PSN infrastructures. In particular, compatibility with the mechanisms of header compression over links where capacity is at a premium SHOULD be provided.

7.8. Congestion Control

TDM circuits run at a constant rate, and hence offer constant traffic loads to the PSN. The rate varying mechanism that TCP uses to match the demand to the network congestion state is, therefore, not applicable.

The ability to shut down a TDM PW when congestion has been detected MUST be provided.

Precautions should be taken to avoid situations wherein multiple TDM PWs are simultaneously shut down or re-established, because this leads to PSN instability.

Further congestion considerations are discussed in chapter 6.5 of [RFC3985].

7.9. Fault Detection and Handling

The encapsulation layer for edge-to-edge emulation of TDM services SHOULD, separately or in conjunction with the lower layers of the PWE3 stack, provide for detection, handling, and reporting of the following defects:

1. Misconnection, or Stray Packets. The importance of this requirement stems from customer expectation due to reliable misconnection detection in SONET/SDH networks.
2. Packet Loss. Packet loss detection is required to maintain clock integrity, as discussed in Section 7.3.1 above. In addition, packet loss detection mechanisms SHOULD provide for localization of the outage in the end-to-end emulated service.
3. Malformed packets.

7.10. Performance Monitoring

The encapsulation layer for edge-to-edge emulation of TDM services SHOULD provide for collection of performance monitoring (PM) data that is compatible with the parameters defined for 'classic', TDM-based carriers of these services. The applicability of [G.826] is left for further study.

8. Security Considerations

The security considerations in [RFC3916] are fully applicable to the emulation of TDM services. In addition, TDM services are sensitive to packet delay variation [Section 7.6], and need to be protected from this method of attack.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

- [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [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 structures defined in Recommendation G.704
- [G.707] ITU-T Recommendation G.707 (10/00) - Network node interface for the synchronous digital hierarchy (SDH)
- [G.751] ITU-T Recommendation G.751 (11/88) - Digital multiplex equipments operating at the third order bit rate of 34 368 Kbit/s and the fourth order bit rate of 139 264 Kbit/s and using positive justification
- [G.810] ITU-T Recommendation G.810 (08/96) - Definitions and terminology for synchronization networks
- [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
- [Q.700] ITU-T Recommendation Q.700 (03/93) - Introduction to CCITT Signalling System No. 7
- [Q.931] ITU-T Recommendation Q.931 (05/98) - ISDN user-network interface layer 3 specification for basic call control
- [RFC1958] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, June 1996.
- [RFC2736] Handley, M. and C. Perkins, "Guidelines for Writers of RTP Payload Format Specifications", BCP 36, RFC 2736, December 1999.

- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [T1.105] ANSI T1.105 - 2001 Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates, and Formats, May 2001
- [T1.107] ANSI T1.107 - 1995. Digital Hierarchy - Format Specification
- [TR-NWT-170] Digital Cross Connect Systems - Generic Requirements and Objectives, Bellcore, TR-NWT-170, January 1993

10. Contributors Section

The following have contributed to this document:

Sasha Vainshtein
Axerra Networks

EEmail: sasha@axerra.com

Yaakov Stein
RAD Data Communication

EEmail: yaakov_s@rad.com

Prayson Pate
Overture Networks, Inc.

EEmail: prayson.pate@overturenetworks.com

Ron Cohen
Lycium Networks

EEmail: ronc@lyciumnetworks.com

Tim Frost
Zarlink Semiconductor

EEmail: tim.frost@zarlink.com

Author's Address

Maximilian Riegel
Siemens AG
St-Martin-Str 76
Munich 81541
Germany

Phone: +49-89-636-75194
EMail: maximilian.riegel@siemens.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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.

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.

Acknowledgement

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