Internet Draft
Document: draft-pate-pwe3-tdm-03.txt
Expires: July, 2002

Prayson Pate
Overture Networks
Ron Cohen
Lycium Networks
David Zelig
Corrigent Systems

TDM Service Specification for Pseudo-Wire Emulation Edge-to-Edge (PWE3)

draft-pate-pwe3-tdm-03.txt

Status of this Memo

Abstract

This document describes the service-specific implementation and requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3) of TDM circuits. It discusses the emulation of circuits (such as T1, E1, T3 and E3) over packet networks using IP or MPLS.

Copyright Notice

Table of Contents

1. Introduction

This document describes the service-specific implementation and requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3) of TDM circuits. It discusses the emulation of circuits (such as T1, E1, T3 and E3) over packet networks using IP or MPLS. It is structured as an extension to [MALIS].

See [PATE] and [XIAO] for background, motivation and requirements concerning circuit emulation over PSNs. [MARTINI] and [MALIS] provide information on the very similar emulation of SONET circuits.

1.1. Goals

- Definition of encapsulation for T1, E1, DS1C, DS2 and T3 as an extension to [MALIS].

- Definition of mapping to IP and MPLS PSNs.

- Compatibility with existing circuit networks.

- Compatibility with ongoing work in PWE3.

1.2. Non-Goals

- Replication of existing works.

1.3. Acronyms

ADM	Add Drop Multiplexer
AIS	Alarm Indication Signal
BIP	Interleaved Parity
BITS	Building Integrated Timing Supply
CEP	Circuit Emulation over Packet - see [MALIS]
CI	Customer Installation
DBA	Dynamic Bandwidth Allocation - see [MALIS]
EBM	Equipped Bit Mask
LOF	Loss of Frame
LOS	Loss of Signal
NI	Network Interface

NPRM	Network Performance Report Message
PSN	Packet Switched Network
POH	Path Overhead
PTE	Path Terminating Equipment
PWE3	Pseudo-Wire Emulation Edge-to-Edge
RAI	Remote Alarm Indication
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SPRM	Supplementary Performance Report Message
TDM	Time Division Multiplexing
TSA	Time Slot Assignment
VT	Virtual Tributary
VTG	Virtual Tributary Group

2. Example Network Diagrams

2.1. Example 1 - T1 Transport

Figure 1 below shows a T1 being used to connect Sites A and B via a TDM/SONET network. The node marked "M" is an M13 multiplexer, while the nodes marked "S" are SONET ADMs. The framing would be terminated at the M13 and SONET ADM in a structured application and carried transparently in an unstructured application.

Note that Figure 1 is also applicable for the transport of E1 and DS3.

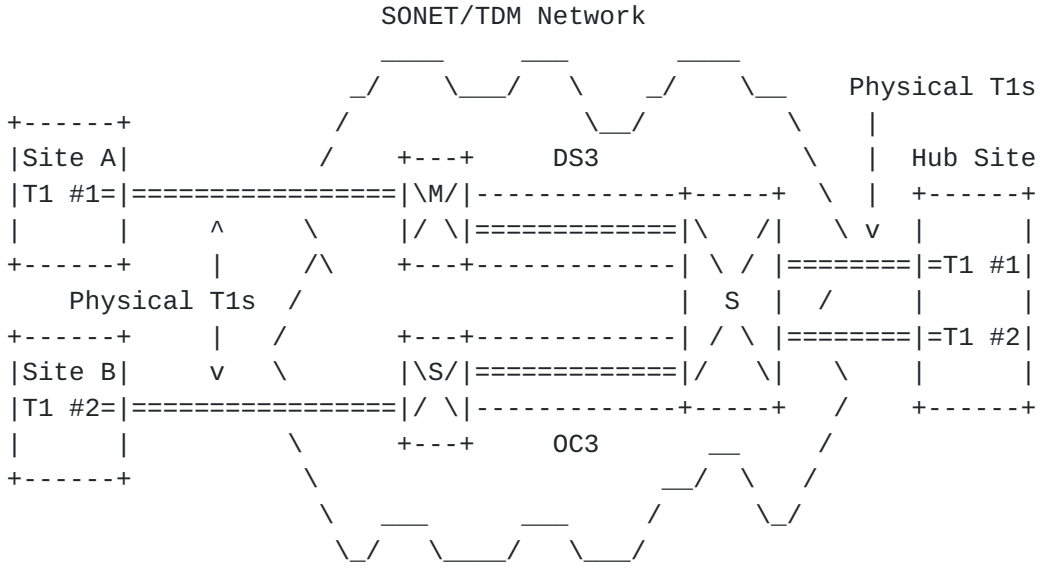

Figure 1: T1 Transport Example Diagram

Figure 2 below shows the same pair of T1s being carried over a packet network using the VT encapsulation defined in Section 4 of this draft. Here the emulation is performed by the boxes marked "E", and the routers marked "R" carry the resulting packets. Note that the emulation, routing and/or SONET functions could be combined in the same device.

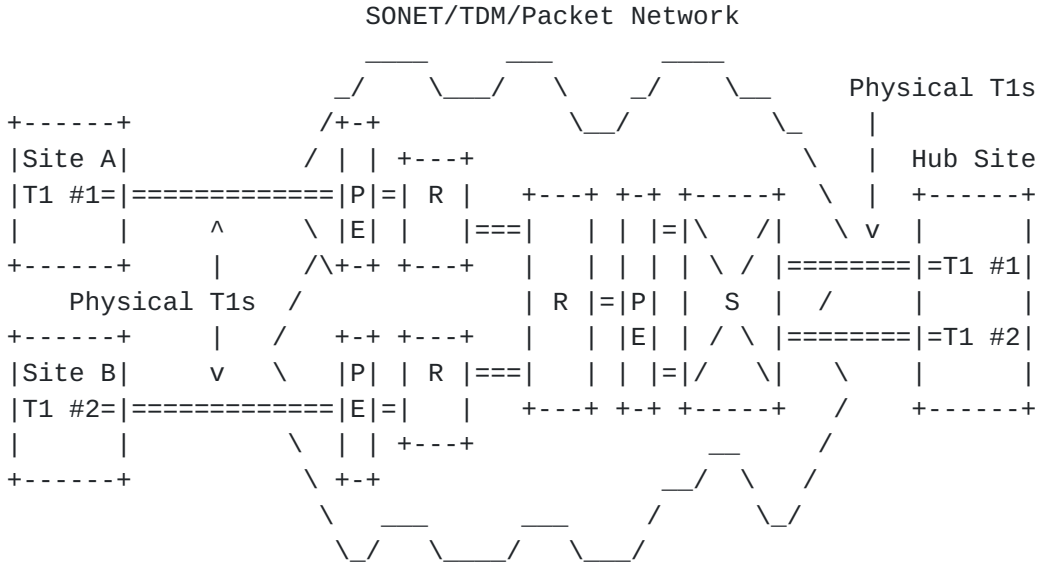

Figure 2: T1 Transport Emulation Example Diagram

2.2. Example 2 - T1 Access

Figure 3 below shows a pair of T1s being carried over a TDM/SONET network. This example differs from the previous one in that the

signal format differs between the ingress and egress nodes.

Note that Figure 3 is also applicable for E1 and DS3 access.

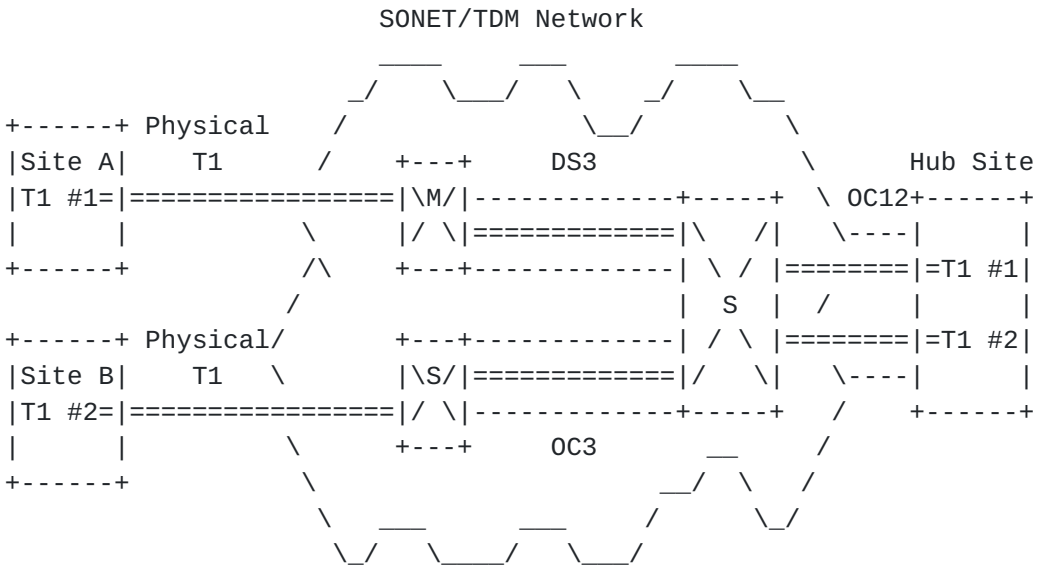

Figure 3: T1 Access Example Diagram

Figure 4 below shows the same pair of T1s being carried over a packet network, again using the VT encapsulation defined in Section 4 of this draft. As with the previous example, the emulation, routing and/or SONET functions could be combined in the same device. Such combinations are likely, so the VT encapsulation format defined in this draft facilitates implementation of such applications and devices.

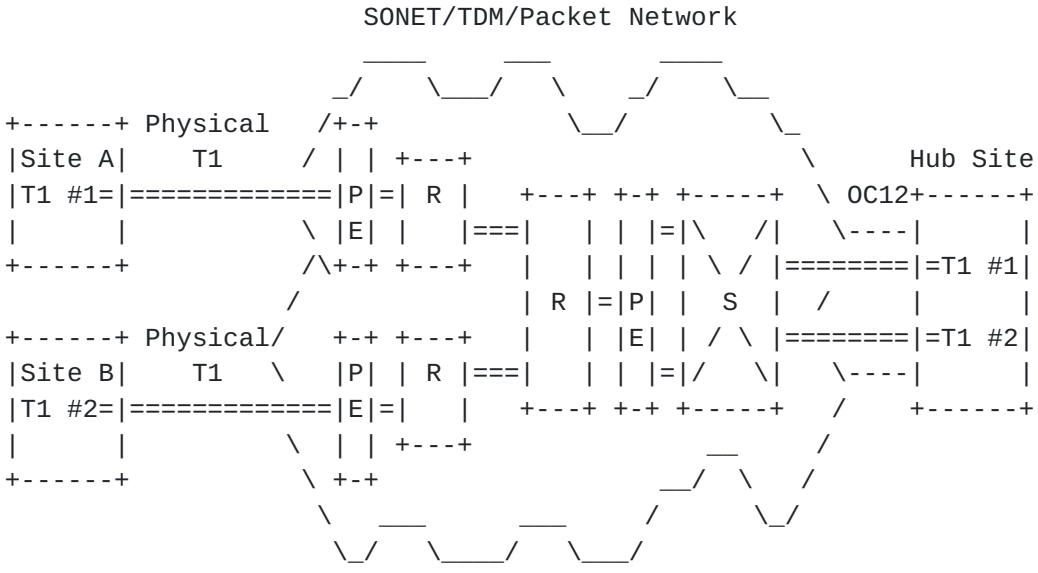

Figure 4: T1 Access Emulation Example Diagram

2.2.1. Example 3 - Fractional SONET Transport

Figure 5 below shows an example of fractional SONET transport. OC3 #1 and OC3 #2 are both channelized and are partially equipped.

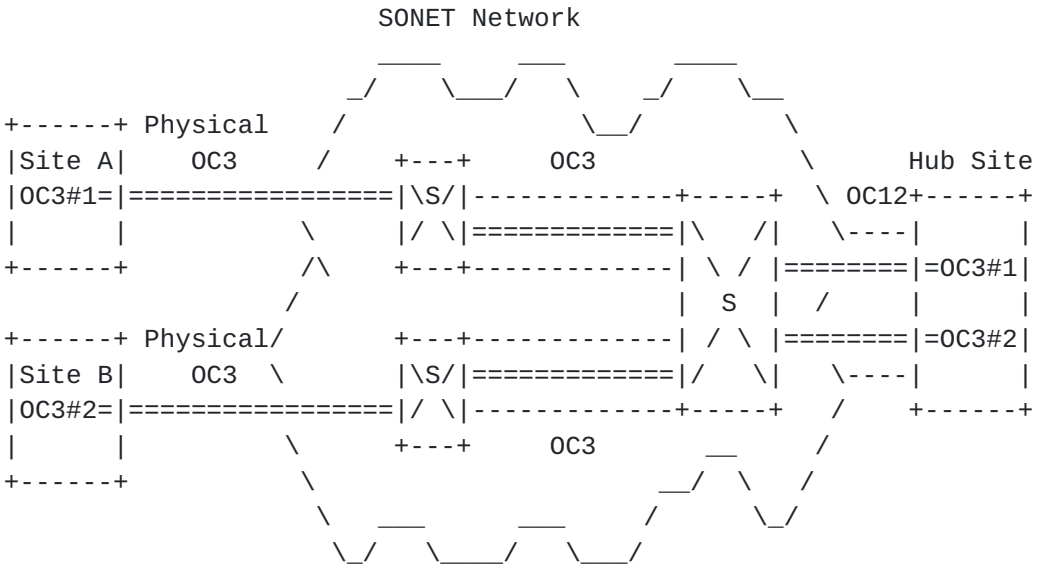

Figure 5: Fractional SONET Transport Example Diagram

Figure 6 below shows the same pair of OC3s being emulated over a PSN using the fractional STS mapping defined in Section 5 of this draft. Note that the PWE3 emulation device is not acting as Path Terminating Equipment (PTE) and should preserve the path BIP and OAM functions.

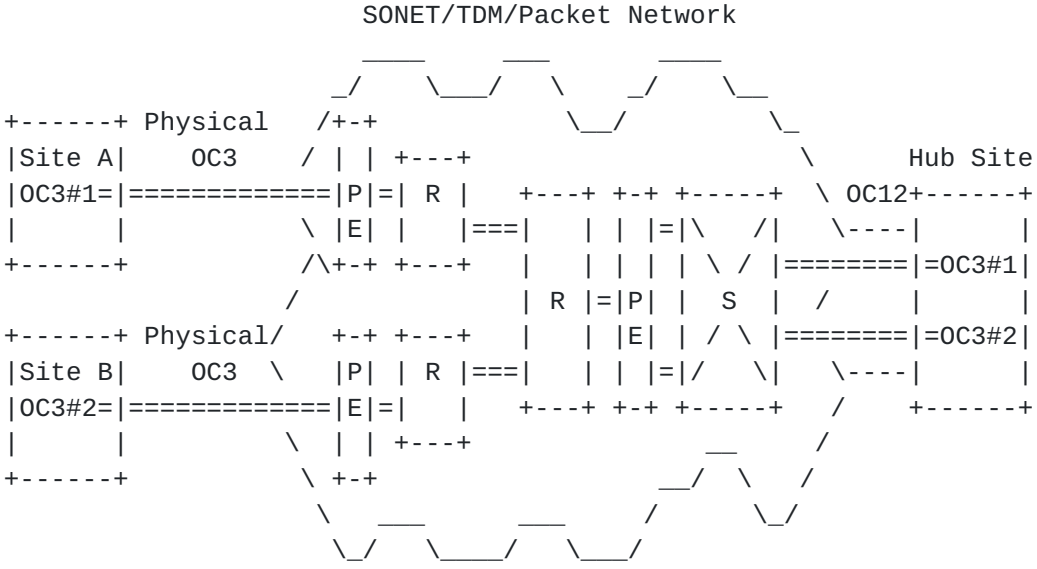

Figure 6: Fractional SONET Transport Emulation Example Diagram

2.2.2. Example 4 - SONET Interconnect

Figure 7 below shows an example of SONET interconnect. As with the previous example, OC3 #1 and OC3 #2 are both channelized and are partially equipped. However, the VTs in OC3 #1 and OC3 #2 are groomed into OC3 #3. The right hand SONET ADM is acting as PTE, and the path BIP and OAM functions are terminated.

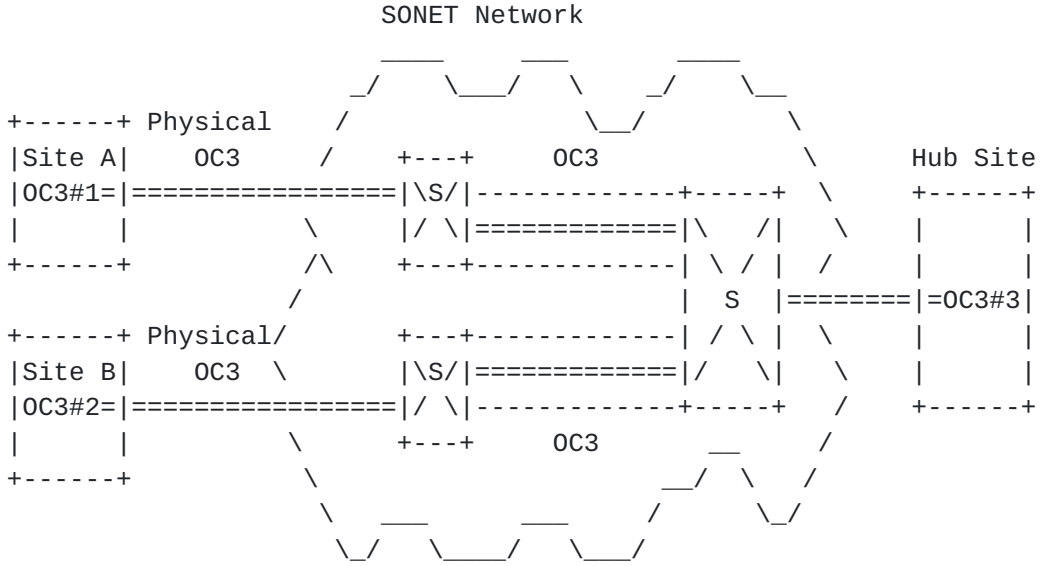

Figure 7: SONET Interconnect Example Diagram

Figure 8 below shows the same pair of OC3s being emulated over a PSN, again using the fractional STS mapping defined in Section 5 of this draft.

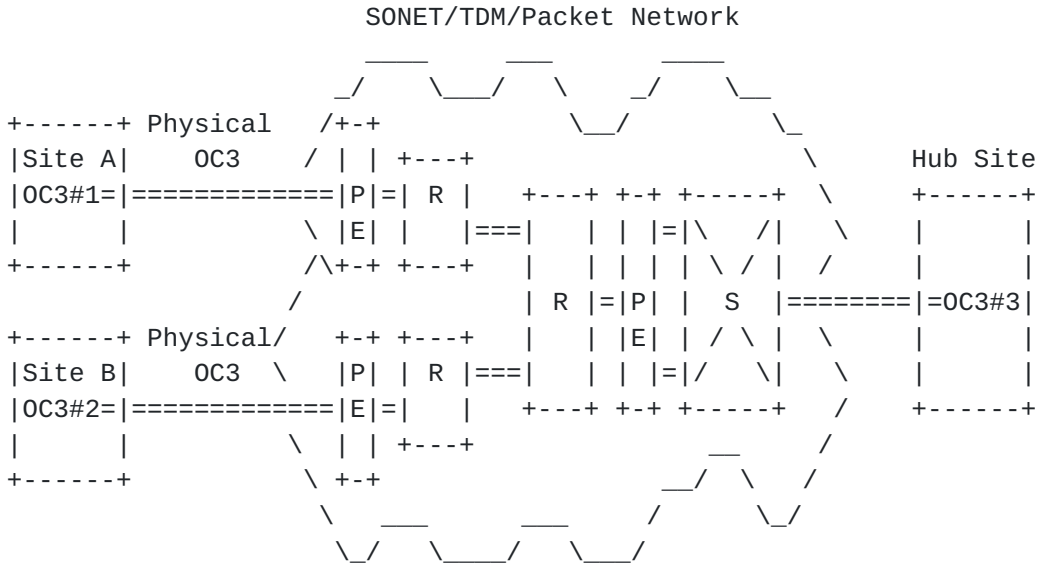

Figure 8: SONET Interconnect Emulation Example Diagram

3. Encapsulation Overview

3.1. Packet Format

Section 4 of [MALIS] defines a mapping for SONET SPEs into a format for transport over various Packet Switched Networks (PSNs). That format is extended here to sub-SPE rates using the standard VT (virtual tributary) mapping mechanism. The format for a TDM CEP (Circuit Emulation over Packet) packet is shown in Figure 9 below.

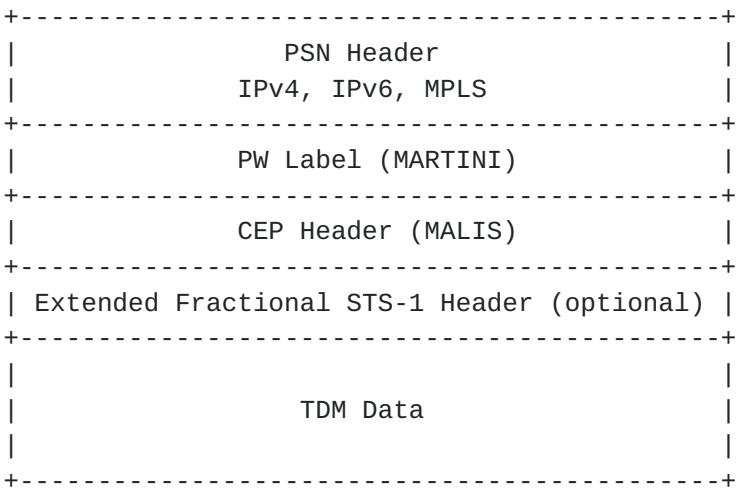

Figure 9: TDM CEP Packet Format

The "PSN Header" could be an IP or GRE header or MPLS label. See Section 4 of [MALIS] for a description of the overall structure, and for the definition of the CEP Header. The formats of the "Extended Fractional STS-1 Header" and "TDM Data" are described in the following sections.

3.2. TDM Encapsulation

This document builds upon the existing standards for its definitions of encapsulations. The SONET VT mapping for DS1/T1, E1, DS1C and DS2 into VT1.5, VT2, VT3 and VT6 is defined in [GR253]. [G.707] defines the mapping of T1s, E1s and DS2 into VC-11, VC-12 and VC-2 containers within the SDH hierarchy. Mapping of E3 and T3 into SDH VC-3 container and to the SONET STS-1 container are defined as well. Both synchronous and asynchronous mappings are defined for E1s and T1s. These mapping are well known and widely used for various applications. This draft extends SONET/SDH circuit emulation [MALIS] to carry the tributary TDM signals.

In order to save paper and wear and tear on the reader's eyeballs, SONET terminology is used throughout this document. All SONET discussions are applicable for SDH terminology as well.

An example of VT1.5 mapping into an STS-1 SPE is shown in Figure 10 below.

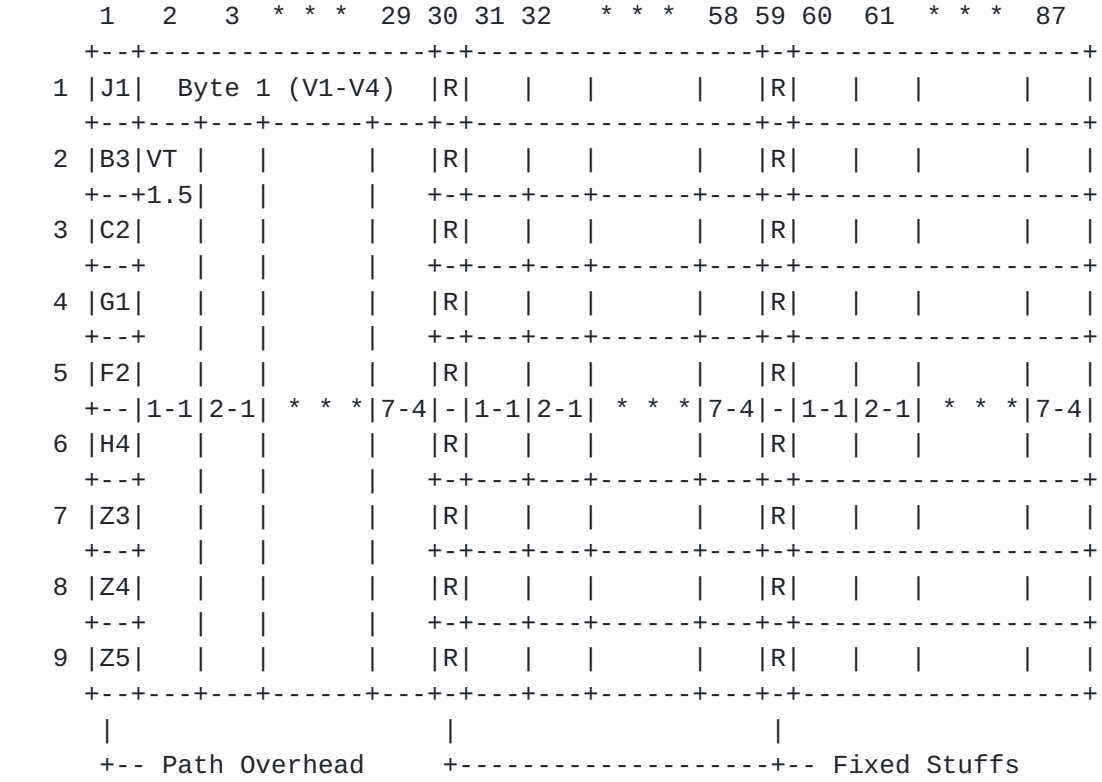

Figure 10: SONET SPE Mapping of VT1.5

The SPE always contains seven interleaved VT groups (VTGs). Each VTG contains a single type of VT, and each VTG occupies 12 columns (108 bytes) within each SPE. A VTG can contain 4 VT1.5s (T1s), 3 VT2s (E1s), 2 VT3s or a single VT6. Altogether, the SPE can carry 28 T1s or carry 21 E1s. SONET carries DS3 signals within a single STS-1,

The encapsulations described in this document use SONET containers to carry TDM signals. Three formats are defined in this document:

- The VT encapsulation maps a single T1, E1, DS1C or DS2 into a VT and then into packets.

- The fractional STS-1 encapsulation defined herein can carry any number of VTs up to the maximal allowed within a single STS-1.

- A DS3 is encapsulated within an STS-1 container and sent over an STS-1 emulated circuit.

These encapsulations are described in more detail in the following sections.

4. VT Encapsulation

The VT encapsulation carries a single VT1.5 (T1), VT2 (E2), VT3 or VT6 circuit. The E1 or T1 signal may be carried either in byte-synchronous mapping or asynchronous mapping. This format is suitable for applications that carry only one TDM signal between sites, such as those shown in Figures 2 and 4. Byte-synchronous mode is typically used for applications that require DS0 access for voice switching, CCS and fractional data transfer (e.g. for Frame Relay). Byte-synchronous mapping may require that a framer is available at the T1 or E1 connection point. A framer is needed for performance monitoring on the incoming signal, loopback commands detection and extracting and insertion of the CCS signaling for T1 within the overheads bytes. In this case the T1 or E1 signals MAY be mapped using the byte-synchronous mode as defined in section 3.4.1.1 in [GR253]. If transparent transmission of the data and clock signals is required, without changing the framing patterns, bit-asynchronous mapping is used as defined in section 3.4.1.2 in [GR253]. Bit-asynchronous mapping does not require a framer at the T1 connection point.

4.1. Multi-frame Format

VTs are organized in SONET multi-frames, where a SONET multi-frame is a sequence of four SONET SPEs. The SPE path overhead byte H4 indicates the SPE number within the multi-frame. The VT overhead bytes (V1, V2, V3 and V4) of each VT occupy the same SPE byte at a fixed position in SPEs 1, 2, 3 and 4 of the multi-frame, respectively. The VT data can float relative to the SPE position. The VT overhead bytes V1, V2 and V3 are used as pointer and stuffing byte similar to the use of the H1, H2 and H3 TOH bytes. V4 is currently unused.

The VT encapsulation does not carry the overhead bytes V1-V4 within the payload, but rather maps the relevant information into the CEP pointer and N/P indications. The CEP pointer indicates the position of the V5 byte within the payload.

Figure 11 below indicates the number of bytes occupied by a VT within a multi-frame.

Mapping	Bytes per Multi-frame	Reference
VT1.5	108 bytes	[GR253] Section 3.4.1.1
VT2	144 bytes	[G.707] Section 10.1.4.1
VT3	216 bytes	[GR253] Section 3.4.1.3
VT6	432 bytes	[GR253] Section 3.4.1.4

Figure 11: Number of Bytes in a Multi-Frame

Each CEP packet carries a fixed payload size that can go up to a single SONET multi-frame. This limitation is due to the restriction

of carrying only one pointer within each CEP header. In particular, a VT1.5 emulation packet can carry up to 104 bytes of payload (leaving out V1-V4).

4.2. VT Header

The basic VT CEP header is defined in Figure 12 per Section 4 of [MALIS]:

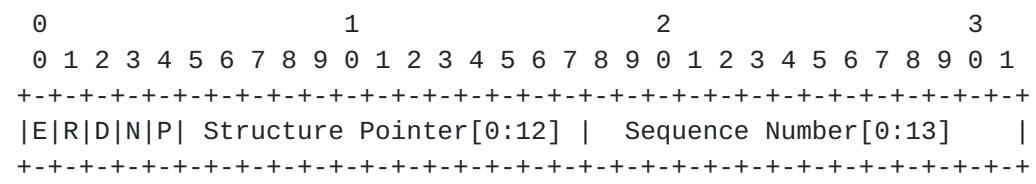

Figure 12: Basic VT CEP Header Format

The following fields are used within the header:

- E: Extension bit. The E bit indicates whether the extended header (to be defined in future revision of [MALIS]) is used.

  E=0: indicates that extended header is not used.

  E=1: indicates that extended header is carried within the packet.

- R bit: RDI indication. The RDI indication is sent whenever a remote defect indication needs to be sent to the PW far side. See [MALIS] for more details.

- D bit: Support for DBA mode for unequipped and AIS indication payload. See Section 7.7.2 below for more details.

- N/P bits: Indicate negative and positive pointer adjustment events. They are also used to relay SONET/SDH maintenance signals such as AIS-V. N indicates a negative pointer event, and P indicates a positive pointer event. Both N and P are set to 1 to indicate the AIS-V signal.

- Structure pointer: The Structure Pointer MUST contain the offset of the V5 byte within the VT Fragment. A value 0 means the first byte after the CEP header. The maximal structure pointer value corresponds to the maximal number of VT bytes contained within a multi-frame, minus the 4 overhead bytes. The Structure Pointer MUST be set to 0x1FF if a packet does not carry the V5 byte.

5. Fractional STS-1 Encapsulation

The fractional STS-1 encapsulation carries VTs within an STS-1

container. This format is suitable for applications such as those shown in Figures 6 and 8.

The STS-1 container includes the path overhead bytes, and the normal SONET encapsulation is used. The additional benefit in using the fractional STS-1 encapsulation is that it does not require sending any unused VTs, giving the ability to be used as a compressed version of the STS-1 encapsulation defined in [MALIS].

The fractional STS-1 encapsulation can optionally carry a bit mask that specifies which VTs are carried within the STS-1 payload and which VTs have been removed. This optional bit mask attribute allows the ingress circuit emulation node to remove an unequipped VT on the fly, providing the egress circuit emulation node enough information for reconstructing the VTs in the right order. The use of bit mask enables "on the fly" compression, whereby only equipped VTs (carrying actual data) are sent. This compression saves bandwidth in the PSN.

5.1. Fractional STS-1 Mapping

Figure 13 below shows a mapping of 3 VT1.5s, designated 1-1, 2-1 and 3-1.

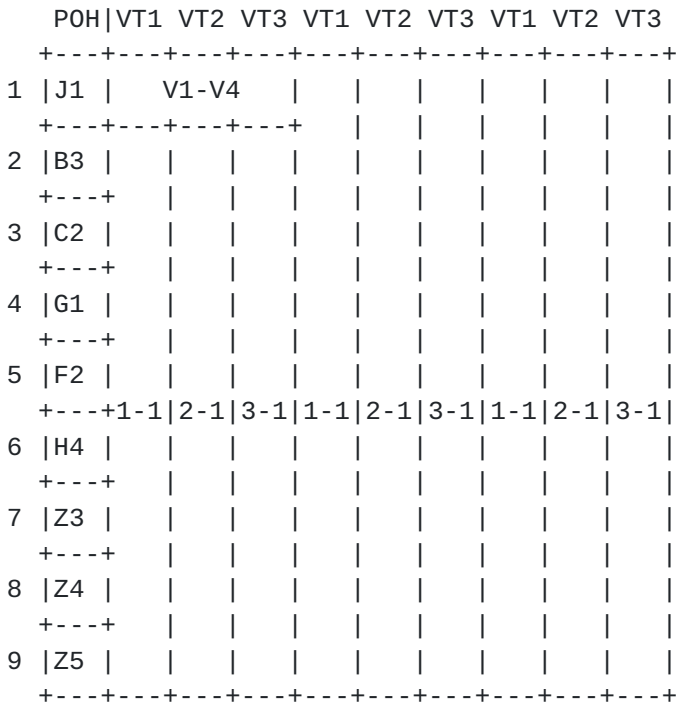

Figure 13: Fractional SPE Mapping of VT1.5

Note that the fixed stuffs shown in Figure 10 are not sent when using this mode. Also note that Figure 13 shows the bytes from the VTs interleaved, as with the SONET SPE shown in Figure 10. This interleaving reduces the buffering required at the ingress and egress PEs. It also helps simplify the construction of combined PW/ADM PEs to operate in networks such as that shown in Figure 1. The

"fractional" SPE in Figure 13 could be expanded out to a full SPE by the addition of "dummy" VTs, Path Overhead and fixed stuffs.

Section 3.3.3 of [GR253] states that "Four bytes (V5, J2, Z6 and Z7) are allocated for VT POH." The same section also defines how these bits are set.

5.2. Fractional STS-1 CEP header

The fractional STS-1 CEP header uses the STS-1 CEP header encapsulation as defined in [MALIS]. Optionally, an additional 4 byte header extension word is added. The extended header is described in Figure 14.

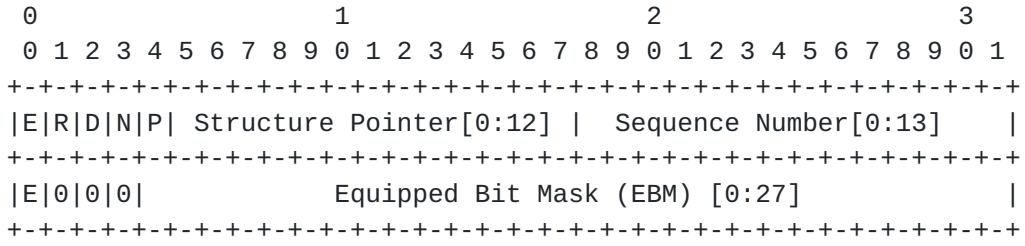

Figure 14: Extended Fractional STS-1 Header

The following fields are used within the extended header:

- R, D, N, P, Structured Pointer and Sequence Number: All fields are used as defined in [MALIS] for STS-1 encapsulation.

- E: Extension bit.

  E=0: indicates that extended header is not used.

  E=1: indicates that extended header is carried within the packet.

  The E bit in the first word is set to 1 to indicate use of the Equipped Bit Mask (EBM). The E bit in the second word indicates whether the extended header (to be defined in future revision of [MALIS]) is used.

- EBM: Each bit within the bit mask refers to a different VT within the STS-1 container. A bit set to 1 indicates that the corresponding VT is equipped, hence carried within the fractional STS-1 payload.

The format of the EBM is defined in Figure 15.

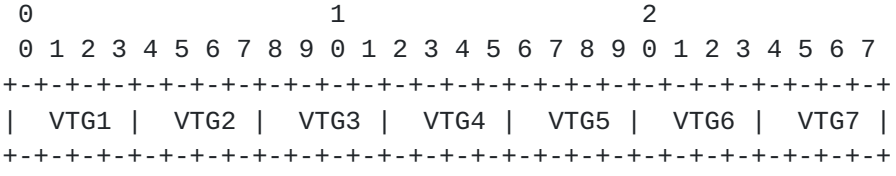

Figure 15: Equipped Bit Mask (EBM)

The 28 bits of the EBM are divided into groups of 4 bits, each corresponding to a different VTG within the STS container. All 4 bits are used to indicate whether VT1.5 (T1) tributaries are carried within a VTG. The first 3 bits are used to indicate whether VT2 (E1) tributaries are carried within a VTG. For example, the EBM of a fully occupied STS-1 with VT1.5 is all ones, while that of an STS-1 fully occupied with VT2 (E1) tributaries has the binary value 11101110111011101110111011101110.

5.3. BIP

The B3 byte within the POH indicates the bit-wise parity of the payload.

5.3.1. Case 1: Path Terminating Equipment (PTE)

In some applications the Path is terminated at the PE, and some PEs may integrate many fractional STS-1's into one STS-1. Figure 8 shows an example of PTE, where a PE is using a fractional STS-1 to carry TDM signals between the ingress and egress emulation edges. In these cases B3 is re-generated at the concentrating PE toward the SONET equipment, and B3 is calculated as the payload is sent, including VTs and POH bytes.

5.3.2. Case 2: Non-PTE

In some applications the Path is not terminated at the PE. For example, the PEs in Figure 6 are using a fractional STS-1 to carry a partially-equipped STS-1, and are not acting as PTEs.

In this case the B3 value should be modified to reflect the removal of the unequipped VTs. Let B3* be the bit-wise parity of the removed unequipped VTs, and let B3p be the value carried within the fractional STS-1 payload. Then B3p = B3 || B3*, where || indicates a bit-wise OR operation. The egress PE can reconstruct the unequipped VTs and modify the B3p value in the same manner to accommodate for the additional VTs added. In this way the end-to-end BIP is preserved.

6. DS3 Encapsulation

A DS3 is encapsulated asynchronously into an STS-1 SPE using the format defined in section 3.4.2.2 of [GR253]. The STS-1 SPE is then encapsulated as defined in Section 4 of [MALIS].

7. Operational Considerations

7.1. Payload Size

Payload sizes are statically provisioned for each TDM stream as

described in [MALIS], using the same management information base [MIB]. CEP packets are normally fixed in length for all of the

packets of a particular emulated TDM stream. The exceptions are DBA CEP packets and on the fly compression within the fractional STS-1 mode. When the fractional STS-1 encapsulation does not carry the equipped flag indications, the VTs to be transmitted MUST be statically provisioned at both ends. The static EBM provisioned at the egress must equal in the number of VTs equipped at the ingress, but the actual VT positions could vary. The length of each CEP packet does not need to be carried in the CEP header because it can be uniquely determined for each CEP packet as a function of the provisioned payload size, the type of VTs carried within the STS-1 signal, the DBA indication and the equipped flags (either dynamically or statically provisioned).

Only the following payload lengths can be statically provisioned for fractional STS-1 encapsulations:

1. Full SPE length (783 bytes)
2. Third of SPE length (261 bytes)

The actual payload sizes would be smaller, depending on the number of virtual tributaries carried within the fractional SPE. Figure 16 provides the actual payload length as a function of N, the number of tributaries carried within the fractional STS-1.

In particular, when all VTs are equipped, the fractional STS-1 full SPE payload size is 765 bytes.

VT Type	Full SPE	SPE/3	
VT1.5 (T1)	27*N+9	9*N+3	N=0..28
VT2 (E1)	36*N+9	12*N+3	N=0..21

Figure 16: Fractional STS-1 Actual Payload Size

7.2. Operational Modes

Figure 17 defines the various options for PW encapsulations and user interfaces physical connections.

Mode #	User Interface	PW Encapsulation
1 *	T1	VT
2	T1	Fractional STS-1
3	STS-1	VT
4 *	STS-1	Fractional STS-1

* Most common uses.

Figure 17: PW Encapsulations Related to User Interfaces

7.3. Time Slot Assignment (TSA)

For an application like that shown in Figure 8, it may be desirable to change the TSA for a given VT. For example, an operator may desire to take an E1 appearing in the first VT on the ingress side and place it in a different VT on the egress side. The PE MAY allow the operator to configure the assignment of Time Slots at each end of the PW.

7.4. Timing

This section will describe some clarifications about the EPAR operation and its applicability to this document in the various modes defined above. The term REFCLK is used to indicate the local PE node system clock that is synchronized via the network timing distribution to the source clock feeding the peer PE system clock.

Mode #1: The VT framing timing is based on the REFCLK. If asynchronous mode is used, there is no use of pointer adjustments on the CEP VT header, and timing differences between the incoming T1 signal and REFCLCK are accommodated by the use of stuffing bits as defined in [GR253]. If byte-synchronous mode is used, the timing difference is accommodated by the use of pointer adjustments indication on the VT CEP header.

Mode #2: The signal is first mapped to a VT as defined in [GR253], which than maps for into fractional STS-1. The T1 to VT timing differences are accommodated as defined in [GR-253] (based on the relevant mode - byte synchronize or asynchronous) and the pointers at the CEP level are fixed.

Mode #3: Relative pointer adjustments of the incoming STS to REFCLCK are added to the VT to STS pointer adjustments and played on the VT CEP header as defined in [MALIS].

Mode #4: Same as in [MALIS].

Notes:

1) In mode #1 the originator of the PW is not known to be operating in mode #1 or mode #3, so the receiver may need to accommodate both stuffing bits and CEP VT pointer adjustments.

2) If one STS-1 is created from multiple sources, the timing of the generated STS-1 toward the user is typically based on the local REFCLCK. In this case each VT's pointer bytes should be played to reflect the pointer adjustments on the incoming CEP header + the VT itself V1-V4 pointer adjustments (if they exist on the incoming encapsulation).

Each pointer justification in [MALIS] is indicated by 3 consecutive CEP header marking. The same procedure is used for the fractional STS-1 encapsulation. For VT encapsulation, pointer justification events are indicated only on the packet(s) that carry the justified multi-frame.

7.5. Loopbacks

Figure 18 below shows a T1 being delivered to a customer site. The nomenclature at the top of the figure is that used in Figure 7 of [T1.403].

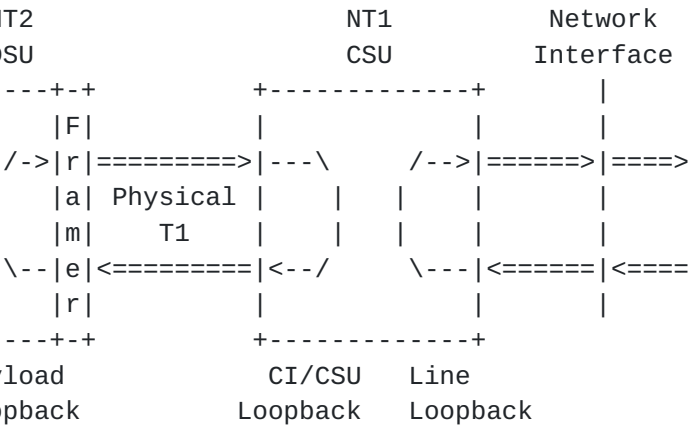

Figure 18: T1 Loopback Example

Note there are two types of loopback commands:

- Inband patterns as defined in Section 9.4.1 of [T1.403]. This type of loopback command was originally defined for voice equipment, and it often causes problems when used with data systems. Processing of inband loopbacks is usually disabled in newer equipment and applications and SHOULD NOT be implemented in a PE.

- FDL messages that carry loopback commands as defined in Section 9.4.2 of [T1.403]. Support of these messages requires the use of a framer.

Depending on the application, it may be desirable to process loopback messages and inband codes. The relevant considerations are:

- Whether the PE implements a CSU functionality. If not, the appropriate hardware may not be present to implement the loopbacks.

- What administrative domains are involved. A carrier will not want a customer to send commands that can interfere with network operation.

- Whether the T1 service is structured or unstructured.
  Processing of FDL loopback commands requires the use of a

framer.

Whenever loopback processing is supported, there MUST be a means to disable such processing.

The following examples point out when it may be appropriate to process loopback commands.

7.5.1. Neither End of T1 is in Carrier's Administrative Domain

A network like that shown in Figure 2 may be providing T1 transport. In this case, the carrier may not wish to generate or process any loopback messages.

7.5.2. One End of T1 is in Carrier's Administrative Domain

Consider the left-hand PEs in Figure 4. These PEs are within the administrative domain of the carrier. If these PEs are implementing the CSU functionality, then it would be desirable for these PEs to process loopback messages originating from within the carrier's network e.g. from the ADM on the right side.

7.6. Performance Processing and Reports

[T1.403] defines a Network Performance Report Message (NPRM) and a Supplementary Performance Report Message (SPRM). These messages are periodic reports on the performance of the link. A PE operating in a structured mode SHOULD generate these messages, as they are frequently used for surveillance and trouble-shooting. Many framers automatically generate and send these reports.

7.7. Alarms and Failure Propagation

7.7.1. Performance Monitoring

[MALIS] defines CEP level Errored Seconds (ES), Severely Errored Seconds (SES) and handling and reporting of CEP defects and failures. The same functionality is applicable here for both the fractional STS encapsulation and the VT encapsulation.

7.7.2. DBA Operation

A TDM multiplexer, switch or other path-terminating device generates AIS in the downstream direction in response to a LOS or LOF condition. This is done by sending a certain pattern in the data stream. Bandwidth can be saved by suppressing the AIS signal in the emulated stream and sending instead an indication in the control overhead. [MALIS] discusses the propagation of AIS using the pointer bits in the TDM control word.

7.7.2.1. Mandatory Procedures

These procedures MUST be implemented in equipment that does not support VT level DBA. The term AIS-P and UNEQ-P are related to the STS PW encapsulation indication, AIS-V and UNEQ-V are either for the VT encapsulation indications or the VT inside a compressed STS, depending on the context. For all modes: If a PW was set up but an association to a user interface is not yet available, or the PW is in an administrative down state, then the relevant UNEQ indication MUST be sent toward the PSN.

Mode #1:

PW Ingress: T1 is propagated transparently inside the VT. Physical layer defects are propagated as T1 AIS inside the VT.

PW Egress: AIS-V or UNEQ-V indications are played out as T1 AIS on the user interface.

Mode #2:

PW Ingress: T1 AIS is propagated transparently inside the VT containing the T1 on the STS. Physical layer defects are propagated as T1 AIS inside the VT.

PW Egress: AIS-P, UNEQ-P, TIM-P (if activated and supported), PLM-P (if supported) , LOP-V, AIS-V and UNEQ-V indication are played out as T1 AIS on the user interface.

Mode #3:

PW Ingress: Section/line layer defects, AIS-P, UNEQ-P, LOP-P, TIM-P (if activated), PLM-P, or LOP-V defects generate AIS-V.

PW Egress: AIS-V, UNEQ-V and LOP-V will be propagated inside the relevant VT inside the STS-1.

Mode #4: Same as in [MALIS].

7.7.2.2. Optional Procedures

These procedures MAY be implemented when a VT encapsulation is used on the PW to conserve BW for AIS and UNEQ states.

Mode #1:

PW Ingress: T1 AIS or T1 physical layer defects are propagated as VT AIS DBA.

Mode #2: No special procedures defined.

Mode #3:

PW Ingress: Section/line layer defects, AIS-P, UNEQ-P, LOP-P, TIM-P (if activated), PLM-P, or LOP-V defects generate AIS-V DBA. The signaling of DBA capability is the same as defined for CEP in [MALIS].

7.7.3. Missing Packet

In the case of a missing packet, all ones should be inserted instead of the missing bytes at the output signal. The only exception is where a single STS-1 is fed by multiple PWs. In this case the output POH is generated normally as this PE is PTE, and the all ones should be generated for the affected VTs only.

7.8. Session Multiplexing

Session multiplexing is accomplished by use of the PW label shown in Figure 9.

7.9. Packet Length Considerations

While the MTU introduced in this document is not an issue for packet networks (783 bytes for fully equipped STS-1 + PW + PSN header), some networks may have minimum packet length requirements. The following is defined to handle these cases:

a. DBA operation: As specified in [MALIS], the DBA packet may be an arbitrary length. This length may be configured at the PW ingress side to handle minimum packet length requirements. The PW egress ignores the DBA packet length (i.e. does not consider it as length error), eliminating the need to negotiate the DBA length. PW edges SHOULD support the ability to set the DBA packet length to accommodate minimum packet length requirements.

b. VT encapsulation: Selection of the packet length should be done based on the PSN minimum packet length requirement.

c. Fractional STS-1: PW ingress SHOULD have the option to be configured to add padding to the PW packet if the packet is less than the minimum packet requirement. At egress, the PE can calculate the number of payload byte using the procedures defined in Section 7.1. The PE MUST ignore the additional padding bytes and should not consider padding bytes as length errors.

8. Security Considerations

See Section 11 of [MALIS].

9. References

[MARTINI] Martini et al, "Transport of Layer 2 Frames Over MPLS", draft-martini-l2circuit-trans-mpls-08.txt, work in progress, November 2001.

[GR253] Bellcore, "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria" (GR253CORE), Issue 3, September 2000.

[G.707] ITU, ITU Recommendation G.707, "Network Node Interface For The Synchronous Digital Hierarchy", 1996.

[T1.403] ANSI, "Network and Customer Installation Interfaces - DS1 Electrical Interfaces", T1.403-1999, May 24, 1999.

[XIAO] Xiao et al, "Requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3)" (draft-pwe3-requirements-02.txt), work in progress, November 2001.

[PATE] Pate et al, "Framework for Pseudo-Wire Emulation Edge-to-Edge (PWE3)" (draft-pate-pwe3-framework-02.txt), work in progress, September 2001.

[MALIS] Malis et al, "SONET/SDH Circuit Emulation over Packet (CEP)" (draft-malis-pwe3-sonet-01.txt), work in progress, November 2001.

[MIB] Danenberg et al, "SONET/SDH Circuit Emulation Service Over PSN (CEP) Management Information Base Using SMIv2", draft-danenberg-pw-cem-mib-01.txt, work in progress, July 2001.

10. Authors' Addresses

Prayson Pate
Overture Networks, Inc.
P. O. Box 14864
RTP, NC, USA 27709
Email: prayson.pate@overturenetworks.com

Ron Cohen
Lycium Networks
P. O. Box 12256
Herzeliya, Israel 46733
Email: ronc@lyciumnetworks.com

David Zelig
Corrigent Systems LTD.
126, Yigal Alon st.

Tel Aviv, Israel
EMail: Davidz@corrigent.com

11. Full Copyright Section