

Network Working Group
Request for Comments: 4842
Category: Standards Track

A. Malis
Verizon Communications
P. Pate
Overture Networks
R. Cohen, Ed.
Resolute Networks
D. Zelig
Corrigent Systems
April 2007

Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH)
Circuit Emulation over Packet (CEP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document provides encapsulation formats and semantics for emulating Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) circuits and services over MPLS.

Table of Contents

1.	Introduction	4
2.	Scope	4
3.	Requirements Notation	4
4.	Acronyms	4
5.	CEP Encapsulation Format	5
5.1.	SONET/SDH Fragment	6
5.2.	CEP Header	7
5.3.	RTP Header	9
5.4.	PSN Encapsulation	11
6.	CEP Operation	11
6.1.	CEP Packetizer and De-Packetizer	11
6.2.	Packet Synchronization	12
6.2.1.	Acquisition of Packet Synchronization	13
6.2.2.	Loss of Packet Synchronization	13
7.	SONET/SDH Maintenance Signals	13
7.1.	SONET/SDH to PSN	13
7.1.1.	CEP-AIS: AIS-P/V Indication	13
7.1.2.	Unequipped Indication	14
7.1.3.	CEP-RDI: Remote Defect Indication	15
7.2.	PSN to SONET/SDH	15
7.2.1.	CEP-AIS: AIS-P/V Indication	15
7.2.2.	Unequipped Indication	16
8.	SONET/SDH Transport Timing	16
9.	SONET/SDH Pointer Management	17
9.1.	Explicit Pointer Adjustment Relay (EPAR)	17
9.2.	Adaptive Pointer Management (APM)	19
10.	CEP Performance Monitors	19
10.1.	Near-End Performance Monitors	19
10.2.	Far-End Performance Monitors	20
11.	Payload Compression Options	20
11.1.	Dynamic Bandwidth Allocation	21
11.2.	Service-Specific Payload Formats	21
11.2.1.	Fractional STS-1 (VC-3) Encapsulation	21
11.2.1.1.	Fractional STS-1 CEP Header	23
11.2.1.2.	B3 Compensation	24
11.2.1.3.	Actual Payload Size	24
11.2.2.	Asynchronous T3/E3 STS-1 (VC-3) Encapsulation	25
11.2.2.1.	T3 Payload Compression	25
11.2.2.2.	E3 Payload Compression	26
11.2.3.	Fractional VC-4 Encapsulation	26
11.2.3.1.	Fractional VC-4 Mapping	27
11.2.3.2.	Fractional VC-4 CEP Header	28
11.2.3.3.	B3 Compensation	29
11.2.3.4.	Actual Payload Sizes	30
12.	Signaling of CEP Pseudowires	30
12.1.	CEP/TDM Payload Bytes	31

12.2. CEP/TDM Bit Rate	31
12.3. CEP Options	32
13. Congestion Control	34
14. Security Considerations	34
15. IANA Considerations	35
16. Acknowledgments	35
17. Co-Authors	35
Appendix A. SONET/SDH Rates and Formats	36
Appendix B. Example Network Diagrams	37
18. References	40
18.1. Normative References	40
18.2. Informative References	41

1. Introduction

This document provides encapsulation formats and semantics for emulating SONET/SDH circuits and services over MPLS.

2. Scope

The generic requirements and architecture for Pseudowire Emulation Edge-to-Edge (PWE3) are described in [PWE3-REQ] and [PWE3-ARCH]. Additional requirements for emulation of SONET/SDH and lower-rate TDM circuits are described in [PWE3-TDM-REQ].

This document provides encapsulation formats and semantics for emulating SONET/SDH circuits and services over MPLS packet-switched networks (PSNs). Emulation of the following digital signals are defined:

1. Synchronous Payload Envelope (SPE)/Virtual Container (VC-n): STS-1/VC-3, STS-3c/VC-4, STS-12c/VC-4-4c, STS-48c/VC-4-16c, STS-192c/VC-4-64c, etc.
2. Virtual Tributary (VT)/Virtual Container (VC-n): VT1.5/VC-11, VT2/VC-12, VT3, VT6/VC-2

For the remainder of this document, these constructs are referred to as SONET/SDH channels.

3. Requirements Notation

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].

4. Acronyms

ADM	Add Drop Multiplexer
AIS	Alarm Indication Signal
APM	Adaptive Pointer Management
AU-n	Administrative Unit-n (SDH)
AUG	Administrative Unit Group (SDH)
BIP	Bit Interleaved Parity
BITS	Building Integrated Timing Supply
CEP	Circuit Emulation over Packet
DBA	Dynamic Bandwidth Allocation
EBM	Equipped Bit Mask
EPAR	Explicit Pointer Adjustment Relay

LOF	Loss of Frame
LOS	Loss of Signal
LTE	Line Terminating Equipment
POH	Path Overhead
PSN	Packet Switched Network
PTE	Path Terminating Equipment
PW	Pseudowire
RDI	Remote Defect Indication
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SPE	Synchronous Payload Envelope
STM-n	Synchronous Transport Module-n (SDH)
STS-n	Synchronous Transport Signal-n (SONET)
TDM	Time Division Multiplexing
TOH	Transport Overhead
TU-n	Tributary Unit-n (SDH)
TUG-n	Tributary Unit Group-n (SDH)
UNEQ	Unequipped
VC-n	Virtual Container-n (SDH)
VT	Virtual Tributary (SONET)
VTG	Virtual Tributary Group (SONET)

5. CEP Encapsulation Format

In order to transport SONET/SDH circuits through a packet-oriented network, the SPE (or VT) is broken into fragments, and a CEP header and optionally an RTP header are prepended to each fragment.

The basic CEP packet appears in Figure 1.




Figure 1: Basic CEP Packet

5.1. SONET/SDH Fragment

The SONET/SDH fragments MUST be byte aligned with the SONET/SDH SPE or VT. The first bit received from each byte of the SONET/SDH MUST be the Most Significant Bit of each byte in the SONET/SDH fragment.

SONET/SDH bytes are placed into the SONET/SDH fragment in the same order in which they are received.

SONET/SDH optical interfaces use binary coding and therefore are scrambled prior to transmission to ensure an adequate number of transitions. For clarity, this scrambling is referred to as physical layer scrambling/descrambling.

In addition, many payload formats (such as for Asynchronous Transfer Mode (ATM) and High-Level Data Link Control (HDLC)) include an additional layer of scrambling to provide protection against transition density violations within the SPEs. This function is referred to as payload scrambling/unscrambling.

CEP assumes that physical layer scrambling/unscrambling occurs as part of the SONET/SDH section/line termination Native Service Processing (NSP) functions.

However, CEP makes no assumption about payload scrambling. The SONET/SDH fragments MUST be constructed without knowledge or processing of any incidental payload scrambling.

CEP implementations MUST include the H3 (or V3) byte in the SONET/SDH fragment during negative pointer adjustment events, and MUST remove the stuff byte from the SONET/SDH fragment during positive pointer adjustment events.

VT emulation implementations MUST NOT carry the VT pointer bytes V1, V2, V3, and V4 as part of the CEP payload. The only exception is the carriage of V3 during negative pointer adjustment as described above.

All CEP SPE implementations MUST support a packet size of 783 bytes and MAY support other packet sizes.

VT emulation implementations MUST support payload size of full VT super-frame, and MAY support 1/2 and 1/4 VT super-frame payload sizes.

Table 1 below describes single super-frame payload size per VT type.

VT type	Size (Bytes)
VT1.5/VC-11	104
VT2/VC-12	140
VT3	212
VT6/VC-2	428

Table 1: VT Super-Frame Payload Sizes

OPTIONAL SONET/SDH Fragment formats have been defined to reduce the bandwidth requirements of the emulated SONET/SDH circuits under certain conditions. These OPTIONAL formats are included in Section 11.

5.2. CEP Header

The CEP header supports both a basic and extended mode. The Basic CEP header provides the minimum functionality necessary to accurately emulate a SONET/SDH circuit over a PSN. Extended headers are utilized for some of the OPTIONAL SONET/SDH fragment formats described in Section 11.

Enhanced functionality and commonality with other real-time Internet applications is provided by RTP encapsulation.

The CEP header has the following format:




Figure 2: CEP Header Format

L bit: CEP-AIS. This bit MUST be set to 1 to signal to the downstream PE that a failure condition has been detected on the attachment circuit. Failure conditions leading to generation of CEP-AIS and the mapping of CEP-AIS signal on the downstream attachment circuit are described in Section 7.

R bit: CEP-RDI. This bit MUST be set to 1 to signal to the upstream PE that a loss of packet synchronization has occurred. This bit MUST be set to 0 once packet synchronization is acquired. See Section 6.2 for details.

N and P bits: These bits are used to explicitly relay negative and positive pointer adjustments events across the PSN. The use of N and P bits is OPTIONAL. If not used, N and P bits MUST be set to 0. See Section 9 for details.

Table 2 describes the interpretation of N and P bits settings.

N	P	Interpretation
0	0	No Pointer Adjustments
0	1	Positive Pointer Adjustment
1	0	Negative Pointer Adjustment
1	1	Loss of Pointer Alarm

Table 2: Interpretation of N and P Bits

FRG bits: FRG bits MUST be set to 0 by sender and ignored by receiver.

SONET data is continuously fragmented into packets. The structure pointer field designates the offset between the SONET SPE or VT structure and the packet boundary.

Length [0:5]: The Length field, if other than zero, indicates the length of the CEP header, plus the length of the RTP header if used, plus the length of the payload. The Length field MUST be set if the length of CEP header plus RTP header if used, plus payload is less than or equal to 64 bytes and MUST be set to 0 otherwise. In particular, if the payload is suppressed (e.g., DBA) the Length field MUST be set to the CEP header length plus the RTP header length if used.

Sequence Number [0:15]: The packet sequence number MUST continuously cycle from 0 to 0xFFFF. It is generated and processed in accordance with the rules established in [RTP].

Structure Pointer [0:11]: The structure pointer MUST contain the offset of the first byte of the SONET structure within the packet payload. For SPE emulation, the structure pointer locates the J1 byte within the CEP packet. For VT emulation, the structure pointer locates the V5 byte within the packet. The structure pointer value ranges between 0 to 0xFFE, where 0 represents the first byte after the CEP header. The structure pointer MUST be set to 0xFFFF if a packet does not carry the J1 (or V5) byte. An arbitrary pointer change (New Data Flag (NDF) event) in the attachment circuit changes the offset of the SONET structure within the CEP packet and therefore changes the structure pointer value.

Reserved field: The reserved field MUST be set to 0 by the sender and ignored by receiver.

5.3. RTP Header

Usage of the RTP header is OPTIONAL. Although CEP MAY employ an RTP header when explicit transfer of timing information is required, this is purely a 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 Realtime Transport Protocol (SRTP), etc., are not applicable to pseudowires. CEP uses the RTP Header as shown below.

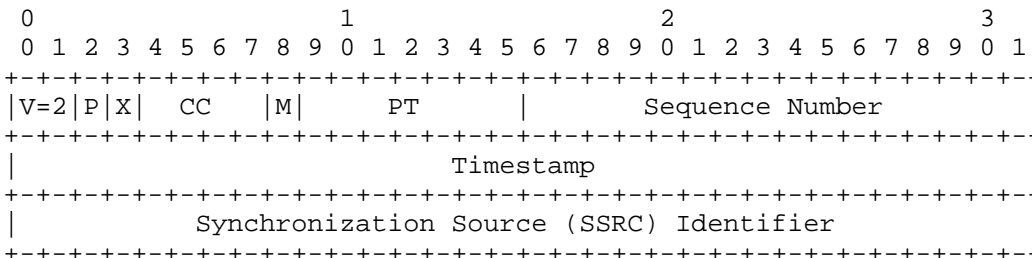


Figure 3: RTP Header

- V: Version. The Version field MUST be set to 2.
- P: Padding. No padding is required. The P bit MUST be set to 0 by sender and ignored by receiver.
- X: Header extension. No extensions are defined. The X bit MUST be set to 0 by sender and ignored by receiver.
- CC: CSRC count. The CC field MUST be set to 0 by sender and ignored by receiver.
- M: Marker. The M bit MUST be set to 0 by sender and ignored by receiver.
- PT [0:6]: Payload type. A PT value SHOULD be allocated from the range of dynamic values for each direction of the PW. The same PT value MAY be reused both for direction and between different CEP PWs.
- Sequence Number [0:15]: The packet sequence number MUST continuously cycle from 0 to 0xFFFF. It is generated and processed in accordance with the rules established in [RTP]. The CEP receiver MUST sequence packets according to the Sequence Number field of the CEP header and MAY verify correct sequencing using RTP Sequence Number field.
- Timestamp [0:31]: Timestamp values are used in accordance with the rules established in [RTP]. Frequency of the clock used for generating timestamps MUST be 19.44 MHz based on a local reference.
- SSRC [0:31]: Synchronization source. The SSRC field MAY be used for detection of misconnections.

5.4. PSN Encapsulation

This document defines the transport of CEP over MPLS PSNs. The bottom label in the MPLS label stack MUST be used to de-multiplex individual CEP channels. In keeping with the conventions used in [PWE3-CONTROL], this de-multiplexing label is referred to as the PW Label and the upper labels are referred to as Tunnel Labels. The CEP header follows the generic PWE3 Control Word format specified in [PWE3-MPLSCW] for use over an MPLS PSN.

Transport of CEP over other PSN technologies is out of scope of this document.

6. CEP Operation

A CEP implementation MUST support a normal mode of operation and MAY support additional bandwidth conserving modes as described in Section 11. During normal operation, SONET/SDH payloads are fragmented, prepended with the appropriate headers, and then transmitted into the packet network.

6.1. CEP Packetizer and De-Packetizer

As with all adaptation functions, CEP has two distinct components: adapting TDM SONET/SDH into a CEP packet stream, and converting the CEP packet stream back into a TDM SONET/SDH. The first function is referred to as CEP packetizer or sender and the second as CEP de-packetizer or receiver. This terminology is illustrated below.




Figure 4: CEP Terminology

The CEP de-packetizer requires a buffering mechanism to account for delay variation in the CEP packet stream. This buffering mechanism is generically referred to as the CEP jitter buffer.

During normal operation, the CEP packetizer receives a fixed-rate byte stream from a SONET/SDH interface. When a packet worth of data is received from a SONET/SDH channel, the necessary headers are prepended to the SPE fragment and the resulting CEP packet is transmitted into the packet network. Because all CEP packets

associated with a specific SONET/SDH channel have the same length, the transmission of CEP packets for that channel SHOULD occur at regular intervals.

At the far end of the packet network, the CEP de-packetizer receives packets into a jitter buffer and then plays out the received byte stream at a fixed rate onto the corresponding SONET/SDH channel. The jitter buffer SHOULD be adjustable in length to account for varying network delay behavior. On average, the receive packet rate from the packet network should be balanced by the transmission rate onto the SONET/SDH channel.

The CEP sequence numbers provide a mechanism to detect lost and/or misordered packets. The sequence number in the CEP header MUST be used when transmission of the RTP header is suppressed. The CEP de-packetizer MUST detect lost or misordered packets. The CEP de-packetizer SHOULD play out an all-ones pattern (AIS) in place of any dropped packets. The CEP de-packetizer SHOULD re-order packets received out of order. If the CEP de-packetizer does not support re-ordering, it MUST drop misordered packets.

6.2. Packet Synchronization

A key component in declaring the state of a CEP service is whether or not the CEP de-packetizer is in or out of packet synchronization. The following paragraphs describe how that determination is made.

As packets are received from the PSN, they are placed into a jitter buffer prior to play out on the SONET/SDH interface. If a CEP de-packetizer supports re-ordering, any packet received before its play out time will still be considered valid.

The determination of acquisition or loss of packet synchronization is always made at SONET/SDH play out time. During SONET/SDH play out, the CEP de-packetizer will play received CEP packets onto the SONET/SDH interface. However, if the jitter buffer is empty or the packet to be played out has not been received, the CEP de-packetizer will play out an 'empty packet' composed of an all-ones AIS pattern onto the SONET/SDH interface in place of the unavailable packet.

The acquisition of packet synchronization is based on the number of sequential CEP packets that are played onto the SONET/SDH interface. Loss of packet synchronization is based on the number of sequential 'empty' packets that are played onto the SONET/SDH interface. Specific details of these two cases are described below.

6.2.1. Acquisition of Packet Synchronization

At startup, a CEP de-packetizer will be out of packet synchronization by default. To declare packet synchronization at startup or after a loss of packet synchronization, the CEP de-packetizer must play out a configurable number of CEP packets with sequential sequence numbers towards the SONET/SDH interface.

6.2.2. Loss of Packet Synchronization

Once a CEP de-packetizer is in packet synchronization state, it may encounter a set of events that will cause it to lose packet synchronization.

If the CEP de-packetizer encounters more than a configurable number of sequential empty packets, the CEP de-packetizer MUST declare a Loss of Packet Synchronization (LOPS) defect.

LOPS failure is declared after 2.5 +/- 0.5 seconds of LOPS defect, and cleared after 10 seconds free of LOPS defect state. The circuit is considered down as long as LOPS failure is declared.

7. SONET/SDH Maintenance Signals

This section describes the mapping of maintenance and alarm signals between the SONET/SDH network and the CEP pseudowire. For clarity, the mappings are split into two groups: SONET/SDH to PSN, and PSN to SONET/SDH.

For coherent failure detection, isolation, monitoring, and troubleshooting, SONET/SDH failure indications MUST be transferred correctly over the CEP pseudowire, and CEP connection failures MUST be mapped to SONET/SDH SPE/VT failure indications. Failure detection capabilities and performance monitoring capabilities are dependent on the NSP functionality, e.g., LTE, PTE, Tandem Connection Monitoring [G.707], or Non-intrusive Monitoring (intermediate connection monitoring).

7.1. SONET/SDH to PSN

The following sections describe the mapping of SONET/SDH Maintenance Signals and Alarm conditions into CEP pseudowire indications.

7.1.1. CEP-AIS: AIS-P/V Indication

SONET/SDH Path outages are signaled by using maintenance alarms such as Path AIS (AIS-P). AIS-P, in particular, indicates that the SONET/SDH Path is not currently transmitting valid end-user data, and the

SPE contains all ones. Similarly, AIS-V indicates that the VT is not currently transmitting valid end-user data, and the VT contains all ones.

It should be noted that nearly every type of service-affecting section or line defect would result in an AIS-P/V condition.

The mapping of SONET/SDH failures into the SPE/VT layer is considered part of the NSP function and is based on the principles specified in [GR253], [SONET], [G.707], [G.806], and [G.783]. For example, should the SONET Section Layer detect a Loss of Signal (LOS) or Loss of Frame (LOF) or Section Trace Mismatch (TIM) conditions, an AIS-L is sent up to the Line Layer. If the Line Layer detects AIS-L or Loss of Pointer (LOP), an AIS-P is sent to the Path Layer. If the Path is terminated at the PE (i.e., a PTE) and the Path Layer detects AIS-P or UNEQ-P or TIM-P or PLM-P an AIS-V is sent to the VT Layer.

The AIS-P/V indication is transferred to the CEP packetizer. During AIS-P/V, CEP packets are generated as usual. The L bit in the CEP header MUST be set to 1 to signal AIS-P/V explicitly through the packet network. The N and P bits SHOULD be set to 1 to indicate loss of pointer indication.

If DBA has been enabled for AIS-P/V, only the necessary headers and optional padding are transmitted into the packet network. The Length field MUST be set to the size of the CEP header plus the size of the RTP header if used.

7.1.2. Unequipped Indication

Unequipped indication is used for bandwidth conserving modes defined in Section 11 as a trigger for payload removal.

The declaration of the SPE/VT channel as Unequipped MUST conform to [GR253], [SONET], [G.806], and [G.783]. Unequipped circuits do not carry valid end-user data, but MAY be used for transporting valid information in the SPE/VT overhead bytes. Supervisory Unequipped signals and Tandem Connection transport are two such applications:

Supervisory Unequipped signal (called TEST-P in [SONET]) is used to provide a test signal for pre-service testing or in-service supervision of a path connection to a remote PTE from a PTE or an intermediate non-terminating path network element. Both Unequipped and Supervisory Unequipped signals carry Unequipped Signal Label defined to be zero. To distinguish between Unequipped and Supervisory Unequipped signals, [G.806] recommends that the SPE/VT Trace bytes J1/J2 be set to a non-zero value in Supervisory Unequipped signals.

The SPE/VT overhead bytes N1/Z6 (SDH refers to Z6 as N2) optionally transport Tandem Connection signals between intermediate network elements. Unequipped signals transporting Tandem Connection would have non-zero N1 or N2/Z6 bytes.

Therefore, the CEP packetizer MUST declare a circuit unequipped only if the Signal Label, Trace (J1/J2), and Tandem Connection (N1/N2/Z6) bytes all have zero value.

During SPE/VT Unequipped, the N and P bits MUST be interpreted as usual. The SPE/VT MUST be transmitted into the packet network along with the appropriate headers.

If DBA has been enabled for Unequipped SPE/VT, only the necessary headers and optional padding are transmitted into the packet network. The Length field MUST be set to the size of the CEP header plus the size of the RTP header if used. The N and P bits MAY be used to signal pointer adjustments as normal.

7.1.3. CEP-RDI: Remote Defect Indication

The CEP function MUST send CEP-RDI indication towards the packet network during loss of packet synchronization by setting the R bit to one in the CEP header. The CEP function SHOULD clear the R bit once packet synchronization is restored.

7.2. PSN to SONET/SDH

The following sections describe the mapping of pseudowire indications to SONET/SDH Maintenance Signals and Alarm conditions.

7.2.1. CEP-AIS: AIS-P/V Indication

There are several conditions in the packet network that cause the CEP de-packetization function to play out an AIS-P/V indication towards a SONET/SDH channel. The CEP de-packetizer MUST play out one packet's worth of all ones for each packet received, and MUST set the SONET/SDH Overhead to signal AIS-P/V as defined in [SONET], [GR253], and [G.707].

The first of these is the receipt of CEP packets with the L bit set to one indicating that the far end has detected an error leading to declaration of AIS-P/V alarm. In addition to the play out of AIS-P/V, the CEP de-packetizer SHOULD set the pointer value to all-ones value.

The second case that will cause a CEP de-packetization function to play out an AIS-P/V indication onto a SONET/SDH channel is during loss of packet synchronization.

The third case is the receipt of CEP packets with both the N and P bits set to 1. This is an explicit indication of Loss of Pointer LOP-P/V received at the far-end of the packet network. In addition to play out of AIS-P/V, the CEP de-packetizer SHOULD set the pointer value to all-ones value.

7.2.2. Unequipped Indication

There are several conditions in the packet network that cause the CEP function to transmit Unequipped indications onto the SONET/SDH channel.

The first, which is transparent to CEP, is the receipt of regular CEP packets that happen to carry an SPE/VT containing the appropriate Path overhead or VT overhead set to Unequipped. This case does not require any special processing on the part of the CEP de-packetizer.

The second case is the receipt of CEP packets with the Length field indicating that the payload was removed by DBA, while the L bit is set to 0, indicating that the DBA was triggered by an Unequipped indication and not by an AIS-P/V indication. The CEP de-packetizer MUST use this information to transmit a packet of all zeros onto the SONET/SDH interface.

The third case when a CEP de-packetizer MUST play out an SPE/VT Unequipped indication towards the SONET/SDH interface is when the circuit has been de-provisioned.

8. SONET/SDH Transport Timing

It is assumed that the distribution of SONET/SDH transport timing information is addressed either through external mechanisms such as Building Integrated Timing Supply (BITS), Stand Alone Synchronization Equipment (SASE), Global Positioning System (GPS), or other such methods, or is obtained through an adaptive timing recovery mechanism.

Details about specific adaptive algorithms for recovery of SONET/SDH transport timing are not considered to be within scope for this document. The wander and jitter limits for networks based on the SDH hierarchy are defined in [G.825] and for the SONET hierarchy in [GR253]. The wander and jitter limits specified in these standards must be maintained when CEP PWs are used.

9. SONET/SDH Pointer Management

A pointer management system is defined as part of the definition of SONET/SDH. Details on SONET/SDH pointer management can be found in [SONET], [GR253], [G.707], and [G.783]. If there is a frequency offset between the frame rate of the transport overhead and that of the SONET/SDH SPE, the alignment of the SPE will periodically slip back or advance in time through positive or negative stuffing. Similarly, if there is a frequency offset between the SPE rate and the VT rate it carries, the alignment of the VT will periodically slip back or advance in time through positive or negative stuffing within the SPE.

The emulation of this aspect of SONET/SDH networks may be accomplished using a variety of techniques including Explicit Pointer Adjustment Relay (EPAR) and Adaptive Pointer Management (APM).

In any case, the handling of the SPE or VT data by the CEP packetizer is the same.

During a negative pointer adjustment event, the CEP packetizer MUST incorporate the H3 (or V3) byte from the SONET/SDH stream into the CEP packet payload in order with the rest of the SPE (or VT). During a positive pointer adjustment event, the CEP packetizer MUST strip the stuff byte from the CEP packet payload.

When playing out a negative pointer adjustment event, the appropriate byte of the CEP payload MUST be placed into the H3 (or V3) byte of the SONET/SDH stream. When playing out a positive pointer adjustment, the CEP de-packetizer MUST insert a stuff byte into the appropriate position within the SONET/SDH stream.

The details regarding the use of the H3 (and V3) byte and stuff byte during positive and negative pointer adjustments can be found in [SONET], [GR253], and [G.707].

9.1. Explicit Pointer Adjustment Relay (EPAR)

CEP provides an OPTIONAL mechanism to explicitly relay pointer adjustment events from one side of the PSN to the other. This technique is referred to as Explicit Pointer Adjustment Relay (EPAR). EPAR is only effective when both ends of the PW have access to a common timing reference.

The following text only applies to CEP implementations that choose to implement EPAR. Any CEP implementation that does not support EPAR MUST set the N and P bits to 0.

The pointer adjustment event MUST be transmitted in three consecutive packets by the packetizer. The de-packetizer MUST play out the pointer adjustment event when any one packet with N/P bit set is received. The CEP de-packetizer MUST utilize the CEP sequence numbers to ensure that SONET/SDH pointer adjustment events are not played any more frequently than once per every three CEP packets transmitted by the remote CEP packetizer.

The VT EPAR packetizer MUST relay pointer adjustment indications received at the SPE level in addition to relaying VT pointer adjustment indications. Because of the rate differences between VT and SPE, the magnitude of a VT pointer adjustment is much larger than that of an SPE adjustment. Therefore, the VT EPAR packetizer has to convert multiple SPE pointer adjustments to fewer VT pointer adjustment indications signaled over the PSN using the N and P CEP header bits. A simple algorithm can be used for this purpose using an accumulator (counter):

The accumulator value is reset to 0 when the circuit is in Loss of Packet Synchronization (LOPS) state.

A positive pointer adjustment indication increases the accumulator value by a fixed quota, while negative pointer adjustment subtracts the same quota from the accumulator. A VT pointer adjustment changes the accumulator value by 783 units (one STS-1 SPE size). An SPE pointer adjustment changes the accumulator value by quota that depends on the VT emulation type. The quota is 1/4 of the VT size as defined in Table 1, e.g., 26 bytes for VT1.5 emulation and 35 bytes for VT2 emulation.

When the accumulator value is larger than or equal to 783 units, a positive pointer adjustment is signaled towards the PSN using the CEP header P bit and 783 units are subtracted from the accumulator.

When the accumulator value is smaller than or equal to minus 783 units, a negative pointer adjustment is signaled towards the PSN using the CEP header N bit and 783 units are added to the accumulator.

The same algorithm is applicable for SDH Virtual Container carried in VC-4, i.e., positive VC-4 pointer adjustment adds 35 units to a VC-12 accumulator, while positive VC-12 pointer adjustment adds 783 units to the accumulator.

If both N and P bits are set, then a Loss of Pointer event has been detected at the PW ingress, making the pointer invalid. The de-packetizer MUST play out an AIS-P/V indication and SHOULD set the pointer value to all-ones value.

9.2. Adaptive Pointer Management (APM)

Another OPTIONAL method that may be used to emulate SONET/SDH pointer management is Adaptive Pointer Management (APM). In basic terms, APM uses information about the depth of the CEP jitter buffers to introduce pointer adjustments in the reassembled SONET/SDH SPE.

Details about specific APM algorithms are not considered to be within scope for this document.

10. CEP Performance Monitors

SONET/SDH as defined in [SONET], [GR253], [G.707], and [G.784] includes a definition of several counters used to monitor the performance of SONET/SDH services. These counters are referred to as Performance Monitors.

In order for CEP to be utilized by traditional SONET/SDH network operators, CEP SHOULD provide similar functionality. The following sections describe a number of counters that are collectively referred to as CEP Performance Monitors.

10.1. Near-End Performance Monitors

These performance monitors are maintained by the CEP de-packetizer during reassembly of the SONET/SDH stream.

The performance monitors are based on two types of defects.

Type 1: missing or dropped packet.

Type 2: buffer underrun, buffer overrun, LOPS, missing packets above predefined configurable threshold.

The specific performance monitors defined for CEP are as follows:

ES-CEP	- CEP Errored Seconds
SES-CEP	- CEP Severely Errored Seconds
UAS-CEP	- CEP Unavailable Seconds

Each second that contains at least one type 1 defect SHALL be declared as ES-CEP. Each second that contains at least one type 2 defect SHALL be declared as SES-CEP.

UAS-CEP SHALL be declared after configurable number of consecutive SES-CEP, and cleared after a configurable number of consecutive seconds without SES-CEP. Default value for each is 10 seconds.

Once unavailability is declared, ES and SES counts SHALL be inhibited up to the point where the unavailability was started. Once unavailability is removed, ES and SES that occurred along the clearing period SHALL be added to the ES and SES counts.

CEP-NE failure is declared after 2.5 +/- 0.5 seconds of CEP-NE type 2 defect, and cleared after 10 seconds free of CEP-NE defect state. Sending notification to the OS for CEP-NE failure is local policy.

10.2. Far-End Performance Monitors

Far-end performance monitors provide insight into the CEP de-packetizer at the far end of the PSN.

Far-end statistics are based on the CEP-RDI indication carried in the CEP header R bit. CEP-FE defect is declared when CEP-RDI is set in the incoming CEP packets.

CEP-FE failure is declared after 2.5 +/- 0.5 seconds of CEP-FE defect, and cleared after 10 seconds free of CEP-FE defect state. Sending notification to the OS for CEP-FE failure is local policy.

11. Payload Compression Options

In addition to pure emulation, CEP also offers a number of options for reducing the total bandwidth utilized by the emulated circuit. These options fall into two categories: Dynamic Bandwidth Allocation (DBA) and Service-Specific Payload Formats.

DBA suppresses transmission of the CEP payload altogether under certain circumstances such as AIS-P/V and SPE/VT Unequipped. The use of DBA is dependent on network architectures, e.g., support of Tandem Connection Monitoring, test signals (TEST-P) [SONET], or Supervisory Unequipped [G.806] signals.

Service-Specific Payload Formats reduce bandwidth by suppressing transmission of portions of the SPE based on specific knowledge of the SPE payload.

Details on these payload compression options are described in the following subsections.

11.1. Dynamic Bandwidth Allocation

Dynamic Bandwidth Allocation (DBA) is an OPTIONAL mechanism for suppressing the transmission of the SPE (or VT) fragment when one of two trigger conditions are met, AIS-P/V or SPE/VT Unequipped.

Implementations that support DBA MUST include a mechanism for disabling DBA on a channel-by-channel basis to allow for interoperability with implementations that do not support DBA.

When a DBA trigger is recognized at PW ingress, the CEP payload will be suppressed. The CEP Length field MUST be set to the CEP header length plus the RTP header length if used, and padding bytes SHOULD be added if the intervening packet network has a minimum packet size that is larger than the payload-suppressed DBA packet.

Other than the suppression of the CEP payload, the CEP behavior during DBA should be equivalent to normal CEP behavior. In particular, the packet transmission rate during DBA should be equivalent to the normal packet transmission rate.

11.2. Service-Specific Payload Formats

In addition to the standard payload encapsulations for SPE and VT transport, several OPTIONAL payload formats have been defined to provide varying amounts of payload compression depending on the type and amount of user traffic present within the SPE. These are described below.

11.2.1. Fractional STS-1 (VC-3) Encapsulation

Fractional STS-1 (VC-3) encapsulation carries only a selected set of VTs within an STS-1 container. This mode is applicable for STS-1 with POH signal label byte C2=2 (VT-structured SPE) and for C2=3 (Locked VT mode).

Implementations of fractional STS-1 (VC-3) encapsulation MUST support payload length of one SPE and MAY support payload length of 1/3 SPE.

The mapping of VTs into an STS-1 container is described in Section 3.2.4 of [GR253], and the mapping into VC-3 is defined in Section 7.2.4 in [G.707]. The CEP packetizer removes all fixed column bytes (columns 30 and 59) and all bytes belonging to the removed VTs. Only

STS-1 POH bytes and bytes that belong to the selected VTs are carried within the payload. The CEP de-packetizer adds the fixed stuff bytes and generates unequipped VT data replacing the removed VT bytes.

The figure below illustrates VT1.5 mapping into an STS-1 SPE.




Figure 5: 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.

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.

11.2.1.1. Fractional STS-1 CEP Header

The fractional STS-1 CEP header uses the STS-1 CEP header encapsulation as defined in this document. Optionally, an additional 4-byte header extension word is added.

The extended header has the following format:

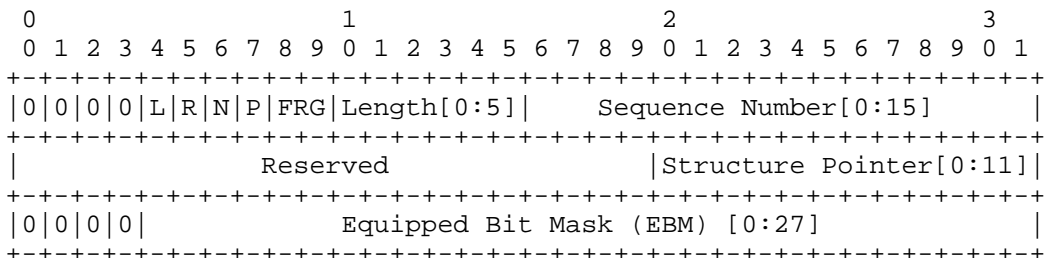


Figure 6: Extended Fractional STS-1 Header

The L, R, N, P, FRG, Length, Sequence Number, and Structured Pointer fields are used as defined in this document for STS-1 encapsulation.

Each bit within the Equipped Bit Mask (EBM) field 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 STS-1 EBM has the following format:




Figure 7: Equipped Bit Mask (EBM) for Fractional STS-1

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 read from right to left are used to indicate whether VT2 (E1) tributaries are carried within a VTG. The first 2 bits are used to indicate whether VT3 (DS1C) tributaries are carried within a VTG. The rightmost bit is used to indicate whether a VT6 (DS2) is carried within the VTG. The VTs within the VTG are numbered from right to left, starting from the first VT as the first bit on the right. For example, the EBM of a fully occupied STS-1

with VT1.5 tributaries is all ones, while that of an STS-1 fully occupied with VT2 (E1) tributaries has the binary value 0111011101110111011101110111.

11.2.1.2. B3 Compensation

Fractional STS-1 encapsulation can be implemented in Line Terminating Equipment (LTE) or in Path Terminating Equipment (PTE). PTE implementations terminate the path layer at the ingress PE and generate a new path layer at the egress PE.

LTE implementations do not terminate the path layer, and therefore need to keep the content and integrity of the POH bytes across the PSN. In LTE implementations, special care must be taken to maintain the B3 bit-wise parity POH byte. The B3 compensation algorithm is defined below.

Since the BIP-8 value in a given frame reflects the parity check over the previous frame, the changes made to BIP-8 bits in the previous frame shall also be considered in the compensation of BIP-8 in the current frame. Therefore, the following equation shall be used for compensation of the individual bits of the BIP-8:

$$B3[i]'(t) = B3[i](t-1) \text{ || } B3[i]'(t-1) \text{ || } B3[i](t) \text{ || } B*3[i](t-1)$$

Where:

B3[i] = the existing B3[i] value in the incoming signal
 B3[i]' = the new (compensated) B3[i] value
 B3*[i] = the B3[i] value of the unequipped VTs in the incoming signal
 || = exclusive OR operator
 t = the time of the current frame
 t-1 = the time of the previous frame

The egress PE MUST reconstruct the unequipped VTs and modify the B3 parity value in the same manner to accommodate the additional VTs added. In this way, the end-to-end BIP is preserved.

11.2.1.3. Actual Payload Size

The actual CEP payload size depends on the number of virtual tributaries carried within the fractional SPE. The contributions of each tributary to the fractional STS-1 payload size as well as the path overhead contribution are described below.

Each VT1.5 contributes 27 bytes

Each VT2 contributes 36 bytes

Each VT3 contributes 54 bytes

Each VT6 contributes 108 bytes

STS-1 POH contributes 9 bytes

For example, a fractional STS-1 carrying 7 VT2 circuit in full-SPE encapsulation would have an actual size of $261=36*7+9$ bytes. Divide by 3 to calculate the third-SPE encapsulation actual payload sizes.

11.2.2. Asynchronous T3/E3 STS-1 (VC-3) Encapsulation

Asynchronous T3/E3 STS-1 (VC-3) encapsulation is applicable for signals with POH signal label byte C2=4, carrying asynchronously mapped T3 or E3 signals.

Implementations of asynchronous T3/E3 STS-1 (VC-3) encapsulation MUST support payload length of one SPE and MAY support payload length of 1/3 SPE.

11.2.2.1. T3 Payload Compression

A T3 is encapsulated asynchronously into an STS-1 SPE using the format defined in Section 3.4.2.1 of [GR253]. The STS-1 SPE is then encapsulated as defined in this document.

Optionally, the STS-1 SPE can be compressed by removing the fixed columns leaving only data columns. STS-1 columns are numbered 1 to 87, starting from the POH column numbered 1. The following columns have fixed values and are removed: 2, 3, 30, 31, 59, and 60.

Bandwidth saving is approximately 7% (6 columns out of 87). The B3 parity byte need not be modified as the parity of the fixed columns amounts to 0. The actual payload size for full-SPE encapsulation would be 729 bytes and 243 bytes for third-SPE encapsulation.

A T3 is encapsulated asynchronously into a VC-3 container as described in Section 10.1.2.1 of [G.707]. VC-3 container has only 85 data columns, which is identical to the STS-1 container with the two fixed stuff columns 30 and 59 removed. Other than that, the mapping is identical.

11.2.2.2. E3 Payload Compression

An E3 is encapsulated asynchronously into a VC-3 SPE using the format defined in Section 10.1.2.2 of [G.707]. The VC-3 SPE is then encapsulated as defined in this document.

Optionally, the VC-3 SPE can be compressed by removing the fixed columns leaving only data columns. VC-3 columns are numbered 1 to 85 (and not 87), starting from the POH column numbered 1. The following columns have fixed values and are removed: 2, 6, 10, 14, 18, 19, 23, 27, 31, 35, 39, 44, 48, 52, 56, 60, 61, 65, 69, 73, 77, and 81.

Bandwidth saving is approximately 28% (24 columns out of 85). The B3 parity byte need not be modified as the parity of the fixed columns amounts to 0. The actual payload size for full-SPE encapsulation would be 567 bytes and 189 bytes for third-SPE encapsulation.

11.2.3. Fractional VC-4 Encapsulation

SDH defines a mapping of VC-11, VC-12, VC-2, and VC-3 directly into VC-4. This mapping does not have an equivalent within the SONET hierarchy. The VC-4 mapping is common in SDH implementations.




Figure 8: Mapping of VCs into VC-4

Figure 8 describes the mapping options of VCs into VC-4. A VC-4 contains three TUG-3s. Each TUG-3 is composed of either a single TU-3 or 7 TUG-2s. A TU-3 contains a single VC-3. A TUG-2 contains either 4 VC-11s (T1s), 3 VC-12s (E1s), or one VC-2. Therefore, a VC-4 may contain 3 VC-3s, 1 VC-3 and 42 VC-12s, 63 VC-12s, etc.

Fractional VC-4 encapsulation carries only a selected set of VCs within a VC-4 container. This mode is applicable for VC-4 with POH signal label byte C2=2 (TUG structure) and for C2=3 (Locked TU-n). The mapping of VCs into a VC-4 container is described in Section 7.2 of [G.707]. The CEP packetizer removes all fixed column bytes and all bytes that belong to the removed VCs. Only VC-4 POH bytes and bytes that belong to the selected VCs are carried within the payload. The CEP de-packetizer adds the fixed stuff bytes and generates unequipped VC data replacing the removed VC bytes.

The fractional VC-4 encapsulation can optionally carry a bit mask that specifies which VCs are carried within the VC-4 payload and which VCs have been removed. This optional bit mask attribute allows the ingress circuit emulation node to remove unequipped VCs on the fly, providing the egress circuit emulation node enough information for reconstructing the VCs in the right order. The use of bit mask enables on-the-fly compression, whereby only equipped VCs (carrying actual data) are sent.

VC-3 carrying asynchronous T3/E3 signals within the VC-4 container can optionally be compressed by removing the fixed column bytes as described in Section 11.2.2, providing additional bandwidth saving.

Implementations of fractional VC-4 encapsulation MUST support payload length of 1/3 SPE and MAY support payload lengths of 4/9, 5/9, 6/9, 7/9, 8/9, and full SPE. The actual payload size of fractional VC-4 encapsulation depends on the number of VCs carried within the payload.

11.2.3.1. Fractional VC-4 Mapping

[G.707] defines the mapping of TUG-3 to a VC-4 in Section 7.2.1. Each TUG-3 includes 86 columns. TUG-3#1, TUG-3#2, and TUG-3#3 are byte multiplexed, starting from column 4. Column 1 is the VC-4 POH, while columns 2 and 3 are fixed and therefore removed in the fractional VC-4 encapsulation.

The mapping of TU-3 into TUG-3 is defined in Section 7.2.2 of [G.707]. The TU-3 consists of the VC-3 with a 9-byte VC-3 POH and the TU-3 pointer. The first column of the 9-row-by-86-column TUG-3 is allocated to the TU-3 pointer (bytes H1, H2, H3) and fixed stuff. The phase of the VC-3 with respect to the TUG-3 is indicated by the TU-3 pointer.

The mapping of TUG-2 into TUG-3 is defined in Section 7.2.3 of [G.707]. The first two columns of the TUG-3 are fixed and therefore removed in the fractional VC-4 encapsulation. The 7 TUG-2s, each 12 columns wide, are byte multiplexed starting from column 3 of the TUG-3. This is the equivalent of multiplexing 7 VTGs within STS-1 container in SONET terminology, except for the location of the fixed columns.

The static fractional VC-4 mapping assumes that both the ingress and egress nodes are preconfigured with the set of equipped VCs carried within the fractional VC-4 encapsulation. The ingress emulation edge removes the fixed columns as well as columns of the VCs agreed upon by the two edges, and updates the B3 VC-4 byte. The egress side adds the fixed columns and the unequipped VCs and updates B3.

11.2.3.2. Fractional VC-4 CEP Header

The fractional VC-4 CEP header uses the VC-4 CEP header defined in this document. Optionally, an additional 12-byte header extension word is added.

The extended header has the following format:

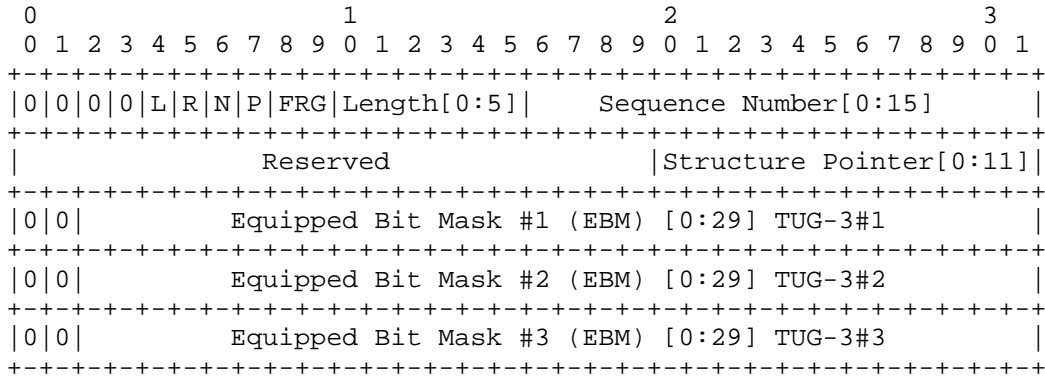


Figure 9: Extended Fractional VC-4 Header

The L, R, N, P, FRG, Length, Sequence Number, and Structured Pointer fields are used as defined in this document for STS-1 encapsulation.

Each bit within the Equipped Bit Mask (EBM) field refers to a different tributary within the VC-4 container. A bit set to 1 indicates that the corresponding tributary is equipped, hence carried within the fractional VC-4 payload.

Three EBM fields are used. Each EBM field corresponds to a different TUG-3 within the VC-4. The EBM includes 7 groups of 4 bits per TUG-2. A bit set to 1 indicates that the corresponding VC is equipped, hence carried within the fractional VC-4 payload. An additional 2 bits within the EBM indicate whether VC-3 carried within the TUG-3 is equipped and whether it is in AIS mode.

The VC-4 EBM has the following format:




Figure 10: Equipped Bit Mask (EBM) for Fractional VC-4

The 30 bits of the EBM are divided into 2 bits that control the VC-3 within the TUG-3 and 7 groups of 4 bits, each corresponding to a different TUG-2 within the TUG-3 container.

For a TUG-3 containing TUG-2, the first two A and T bits MUST be set to 0. The TUG-2 bits indicate whether the VCs within the TUG-2 are equipped. All 4 bits are used to indicate whether VC-11 (T1) tributaries are carried within a TUG-2. The first 3 bits read right to left are used to indicate whether VC-12 (E1) tributaries are carried within a TUG-2. The first bit is used to indicate that a VC-2 is carried within a TUG-2. The VCs within the TUG-2 are numbered from right to left, starting from the first VC as the first bit on the right. For example, 28 bits of the EBM of a fully occupied TUG-3 with VC-11 tributaries are all ones, while that of a TUG-3 fully occupied with VC-12 tributaries has the binary value 000111011101110111011101110111.

For a TUG-3 containing VC-3, all TUG-2 bits MUST be set to 0. The A and T bits are defined as follows:

T: TUG-3 carried bit. If set to 1, the VC-3 payload is carried within the TUG-3 container. If set to 0, all the TUG-3 columns are not carried within the fractional VC-4 encapsulation. The TUG-3 columns are removed either because the VC-3 is unequipped or in AIS mode.

A: VC-3 AIS bit. The A bit MUST be set to 0 when the T bit is 1 (i.e., when the TUG-3 columns are carried within the fractional VC-4 encapsulation). The A bit indicate the reason for removal of the entire TUG-3 columns. If set to 0, the TUG-3 columns were removed because the VC-3 is unequipped. If set to 1, the TUG-3 columns were removed because the VC-3 is in AIS mode.

11.2.3.3. B3 Compensation

Fractional VC-4 encapsulation can be implemented in Line Terminating Equipment (LTE) or in Path Terminating Equipment (PTE). PTE implementations terminate the path layer at the ingress PE and

generate a new path layer at the egress PE. LTE implementations do not terminate the path layer, and therefore need to keep the content and integrity of the POH bytes across the PSN. In LTE implementations, special care must be taken to maintain the B3 bit-wise parity POH byte. The same procedures for B3 compensation as described in Section 11.2.1.2 for fractional STS-1 encapsulation are used.

11.2.3.4. Actual Payload Sizes

The actual CEP payload size depends on the number of virtual tributaries carried within the fractional SPE. The contributions of each tributary to the fractional VC-4 payload length as well as the path overhead contribution are described below.

Each VC-11 contributes 27 bytes

Each VC-12 contributes 36 bytes

Each VC-2 contributes 108 bytes

Each VC-3(T3) contributes 738 bytes

Each VC-3(E3) contributes 576 bytes

Each VC-3(uncompressed) contributes 774 bytes

VC-4 POH contributes 9 bytes

The VC-3 contribution includes the AU-3 pointer. For example, the payload size of a fractional VC-4 configured to third-SPE encapsulation that carries a single compressed T3 VC-3 and 6 VC-12s would be: $321 = (9 + 6 * 36 + 738) / 3$ bytes payload per each packet.

12. Signaling of CEP Pseudowires

[PWE3-CONTROL] specifies the use of the MPLS Label Distribution Protocol, LDP, as a protocol for setting up and maintaining pseudowires. In particular, it provides a way to bind a de-multiplexer field value to a pseudo-wire, specifying procedures for reporting pseudowire status changes and for releasing the bindings. [PWE3-CONTROL] assumes that the pseudowire de-multiplexer field is an MPLS label; however, the PSN tunnel itself can be either an IP or MPLS PSN.

The use of LDP for setting up and maintaining CEP pseudowires is OPTIONAL. This section describes the use of the CEP-specific fields and error codes.

The PW Type field in PwID Forwarding Equivalence Class (FEC) and PW generalized ID FEC elements MUST be set to SONET/SDH Circuit Emulation over Packet (CEP) [PWE3-IANA].

The control word is REQUIRED for CEP pseudowires. Therefore, the C bit in PwID FEC and PW generalized ID FEC elements MUST be set. If the C bit is not set, the pseudowire MUST not be established and a Label Release MUST be sent with an Illegal C bit status code [PWE3-IANA].

The PwID FEC and PW generalized ID FEC elements can include one or more Interface Parameters fields. The Interface Parameters fields are used to validate that the two ends of the pseudowire have the necessary capabilities to interoperate with each other. The CEP-specific Interface Parameters fields are the CEP/TDM Payload Bytes, the CEP/TDM Bit Rate, and the CEP Options parameters.

12.1. CEP/TDM Payload Bytes

This parameter MUST contain the expected CEP payload size in bytes. The payload size does not include network headers, CEP header or padding. If payload compression is used, the CEP/TDM Payload Bytes parameter MUST be set to the uncompressed payload size as if payload compression was disabled. In particular, when Fractional SPE (STS-1/VC-3 or VC-4) payload compression is used, the Payload Bytes parameter MUST be set to the payload size before removal of the unequipped VT containers and fixed value columns. Therefore, when fractional SPE mode is used, the actual (i.e., on the wire) packet length would normally be less than advertised, and in dynamic fractional SPE, even change while the connection is active. Similarly, when DBA payload compression is used, the CEP/TDM Payload Bytes parameter MUST be set to the payload size prior to compression.

The CEP/TDM Payload Bytes parameter is OPTIONAL. Default payload sizes are assumed if this parameter is not included as part of the Interface Parameters fields. The default payload size for VT is a single super frame. The default payload size for SPE is 783 bytes.

A PE that receives a label-mapping request with request for a CEP/TDM Payload Bytes value that is not locally supported MUST return CEP/TDM misconfiguration status error code [PWE3-IANA], and the pseudowire MUST not be established.

12.2. CEP/TDM Bit Rate

The CEP/TDM Bit Rate parameter MUST be set to the data rate in 64-Kbps units of the CEP payload. If payload compression is used, the CEP/TDM Bit Rate parameter MUST be set to the uncompressed payload

data rate as if payload compression was disabled. Table 3 specifies the CEP/TDM Bit Rate parameters that MUST be set for each of the pseudowire circuits.

Circuit	Bit Rate Parameter
VT1.5/VC-11	26
VT2/VC-12	35
VT3	53
VT6/VC-2	107
STS-Nc	783*N N=1,3,12,48,192

Table 3: CEP/TDM Bit Rates

The CEP/TDM Bit Rate parameter is REQUIRED. Attempts to establish a pseudowire between two peers with different bit rates MUST be rejected with incompatible bit rate status error code [PWE3-IANA], and the pseudowire MUST not be established.

12.3. CEP Options

The CEP Options parameter is REQUIRED. The format of the CEP Options parameter is described below:




Figure 11: CEP Options

AIS: When set, indicates that the PE sending the label-mapping request is configured to send DBA packets when AIS indication is detected.

UNE: When set, indicates that the PE sending the label-mapping request is configured to send DBA packets when unequipped circuit indication is detected.

RTP: When set, indicates that the PE sending the label-mapping request is configured to send packets with RTP header.

EBM: When set, indicates that the PE sending the label-mapping request is configured to send packets with EBM extension header.

CEP Type: indicates the CEP connection type:

0x0 SPE mode (STS-1/STS-Mc)

0x1 VT mode (VT1.5/VT2/VT3/VT6)

0x2 Fractional SPE (STS-1/VC-3/VC-4)

Async Type: indicates the Async E3/T3 bandwidth reduction configuration. Relevant only when CEP type is set to fractional SPE, and fractional SPE is expected to carry Asynchronous T3/E3 payload:

T3: When set, indicates that the PE sending the label-mapping request is configured to send Fractional SPE packets with T3 bandwidth reduction.

E3: When set, indicates that the PE sending the label-mapping request is configured to send Fractional SPE packets with E3 bandwidth reduction.

Reserved field: MUST be set to 0 by the PE sending the label-mapping request and ignored by the receiver.

A PE that does not support one of the CEP options set in the label-mapping request MUST send a label-release message with status code of CEP/TDM misconfiguration [PWE3-IANA], report to the operator, and wait for a new consistent label-mapping. A PE MUST send a new label-mapping request once it is reconfigured or when it receives a label-mapping request from its peer with consistent configuration.

A pseudowire can be configured asymmetrically. One PE can be configured to use bandwidth reduction modes, while the other PE can be configured to send the entire circuit unmodified. A PE can compare the CEP Options settings received in the label-mapping request with its own configuration and detect an asymmetric pseudowire configuration. A PE that identifies an asymmetric configuration MAY report it to the operator.

13. Congestion Control

The PSN carrying the CEP PW may be subject to congestion. Congestion considerations for PWs are described in Section 6.5 of [PWE3-ARCH]. CEP PWs represent inelastic constant bit rate (CBR) flows and cannot respond to congestion in a TCP-friendly manner prescribed by [CONG]. CEP 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 [INTSERV] supporting Guaranteed Service [GS] and Diffserv-enabled domains [DIFFSERV] supporting Expedited Forwarding [EF] provide examples of such PSNs. It is expected that PWs emulating high-rate SONET STS-Nc or SDH virtual circuits will be tunneled over traffic-engineered MPLS PSN.

CEP PWs SHOULD monitor packet loss in order to detect "severe congestion". If such a condition is detected, a CEP PW SHOULD shut down bi-directionally. This specification does not define the exact criteria for detecting "severe congestion" using the CEP packet loss rate and the consequent restart criteria after a suitable delay. This is left for further study.

If the CEP PW has been set up using the PWE3 control protocol [PWE3-CONTROL], the regular PW teardown procedures SHOULD be used upon detection of "severe congestion".

The SONET/SDH services emulated by CEP PWs have high availability objectives that MUST be taken into account when deciding on temporary shutdown of CEP PWs. CEP performance monitoring provides entry and exit criteria for the CEP PW unavailable state (UAS-CEP). Detection of "severe congestion" MAY be based on unavailability criteria of the CEP PW.

14. Security Considerations

The CEP encapsulation is subject to all of the general security considerations discussed in [PWE3-ARCH]. In addition, this document specifies only encapsulations, and not the protocols used to carry the encapsulated packets across the PSN. Each such protocol may have its own set of security issues, but those issues are not affected by the encapsulations specified herein. Note that the security of the transported CEP service will only be as good as the security of the PSN. This level of security may be less rigorous than that available from a native TDM service due to the inherent differences between circuit-switched and packet-switched public networks.

Although CEP MAY employ an RTP header when explicit transfer of timing information is required, SRTP [RFC3711] mechanisms are not a substitute for securing the PW and underlying MPLS network.

15. IANA Considerations

IANA considerations for pseudowires are covered in [PWE3-IANA]. CEP does not introduce additional requirements from IANA.

16. Acknowledgments

The authors would like to thank the members of the PWE3 Working Group for their assistance on this document. We thank Sasha Vainshtein, Deborah Brungard, Juergen Heiles, and Nick Weeds for their review and valuable feedback.

17. Co-Authors

The individuals listed below are co-authors of this document. Tom Johnson from Litchfield Communications was the editor of this document from the pre-WG versions of the SONET SPE work through version 01 of this document.

Craig White	Level3 Communications
Ed Hallman	Litchfield Communications
Jeremy Brayley	Laurel Networks
Jim Boyle	Juniper Networks
John Shirron	Laurel Networks
Luca Martini	Cisco Systems
Marlene Drost	Litchfield Communications
Steve Vogelsang	Laurel Networks
Tom Johnson	Litchfield Communications
Ken Hsu	Tellabs

Appendix A. SONET/SDH Rates and Formats

For simplicity, the discussion in this section uses SONET terminology, but it applies equally to SDH as well. SDH-equivalent terminology is shown in the tables.

The basic SONET modular signal is the synchronous transport signal-level 1 (STS-1). A number of STS-1s may be multiplexed into higher-level signals denoted as STS-N, with N synchronous payload envelopes (SPEs). The optical counterpart of the STS-N is the Optical Carrier-level N, or OC-N. Table 4 lists standard SONET line rates discussed in this document.

OC Level	OC-1	OC-3	OC-12	OC-48	OC-192
SDH Term	-	STM-1	STM-4	STM-16	STM-64
Line Rate (Mb/s)	51.840	155.520	622.080	2,488.320	9,953.280

Table 4: Standard SONET Line Rates

Each SONET frame is 125us and consists of nine rows. An STS-N frame has nine rows and N*90 columns. Of the N*90 columns, the first N*3 columns are transport overhead and the other N*87 columns are SPEs. A number of STS-1s may also be linked together to form a super-rate signal with only one SPE. The optical super-rate signal is denoted as OC-Nc, which has a higher payload capacity than OC-N.

The first 9-byte column of each SPE is the path overhead (POH) and the remaining columns form the payload capacity with fixed stuff (STS-Nc only). The fixed stuff, which is purely overhead, is N/3-1 columns for STS-Nc. Thus, STS-1 and STS-3c do not have any fixed stuff, STS-12c has three columns of fixed stuff, and so on.

The POH of an STS-1 or STS-Nc is always 9 bytes in nine rows. The payload capacity of an STS-1 is 86 columns (774 bytes) per frame. The payload capacity of an STS-Nc is (N*87)-(N/3) columns per frame. Thus, the payload capacity of an STS-3c is $(3*87 - 1)*9 = 2,340$ bytes per frame. As another example, the payload capacity of an STS-192c is 149,760 bytes, which is 64 times the capacity of an STS-3c.

There are 8,000 SONET frames per second. Therefore, the SPE size, (POH plus payload capacity) of an STS-1 is $783*8*8,000 = 50.112$ Mb/s. The SPE size of a concatenated STS-3c is 2,349 bytes per frame or

150.336 Mb/s. The payload capacity of an STS-192c is 149,760 bytes per frame, which is equivalent to 9,584.640 Mb/s. Table 5 lists the SPE and payload rates supported.

SONET STS Level	STS-1	STS-3c	OC-12c	OC-48c	OC-192c
SDH VC Level	VC-3	VC-4	VC-4-4c	VC-4-16c	VC-4-64c
Payload Size(Bytes)	774	2,340	9,360	37,440	149,760
Payload Rate(Mb/s)	49.536	149.760	599.040	2,396.160	9,584.640
SPE Size(Bytes)	783	2,349	9,396	37,584	150,336
SPE Rate(Mb/s)	50.112	150.336	601.344	2,405.376	9,621.504

Table 5: Payload Size and Rate

To support circuit emulation, the entire SPE of a SONET STS or SDH VC level is encapsulated into packets, using the encapsulation defined in Section 5, for carriage across packet-switched networks.

VTs are organized in SONET super-frames, where a SONET super-frame is a sequence of four SONET SPEs. The SPE path overhead byte H4 indicates the SPE number within the super-frame. The VT data can float relative to the SPE position. The 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.

Appendix B. Example Network Diagrams

Figure 12 below illustrates a SONET interconnect example. Site A and Site B are connected back to a Hub Site, Site C by means of a SONET infrastructure. The OC-12 from Site A and the OC-12 from Site B are partially equipped. Each of them is transported through a SONET network back to a hub site C. Equipped SPEs (or VTs) are then groomed onto the OC-12 towards site C.

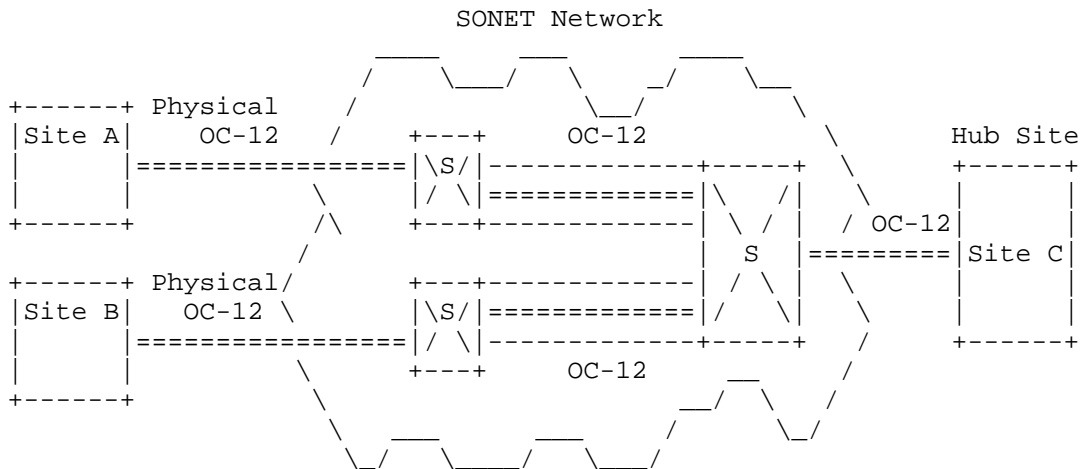


Figure 12: SONET Interconnect Example Diagram

Figure 13 below illustrates the same pair of OC-12s being emulated over a PSN. This configuration frees up bandwidth in the grooming network, since only equipped SPEs (or VTs) are sent through the PSN. Additional bandwidth savings can be realized by taking advantage of the various payload compression options described in Section 11.

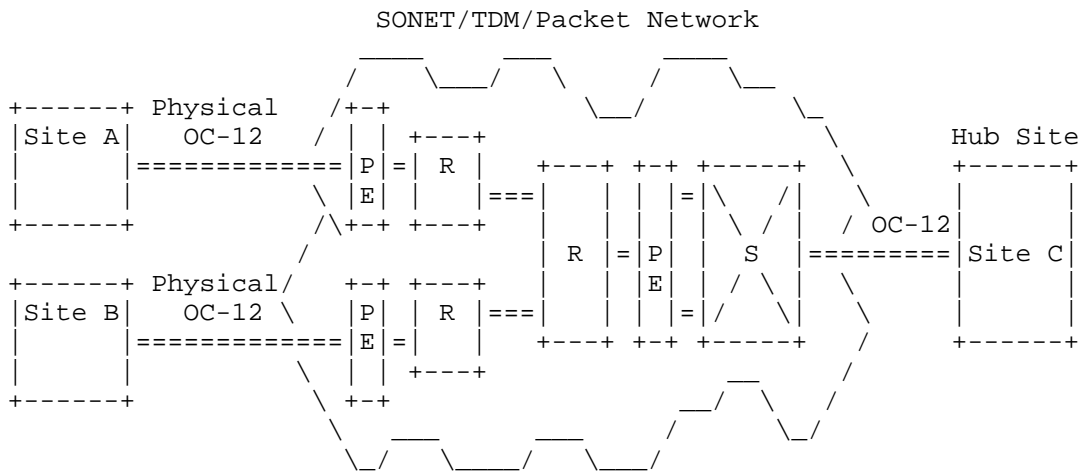


Figure 13: SONET Interconnect Emulation Example Diagram

Figure 14 below shows an example of T1 grooming into OC-12 in access networks. The VT encapsulation is used to transport the T1s from the Hub site to customer sites, maintaining SONET/SDH Operations and Management (OAM).

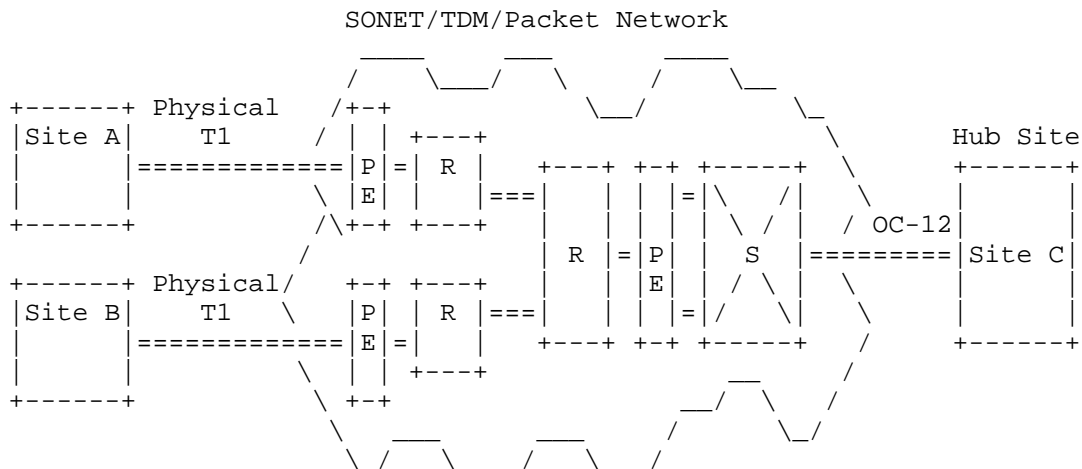


Figure 14: T1 to OC-12 Grooming Emulation Example Diagram

18. References

18.1. Normative References

- [G.707] "Network Node Interface For The Synchronous Digital Hierarchy", ITU-T Recommendation G.707, December 2003.
- [G.783] "Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks", ITU-T Recommendation G.783, February 2004.
- [G.784] "Synchronous Digital Hierarchy (SDH) management", ITU-T Recommendation G.784, July 1999.
- [G.806] "Characteristics of transport equipment-Description methodology and generic functionality", ITU-T Recommendation G.806, February 2004.
- [G.825] "The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)", ITU-T Recommendation G.825, March 2000.
- [GR253] "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria", Telcordia GR-253-CORE Issue 3, September 2000.
- [MPLS] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [PWE3-CONTROL] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [PWE3-IANA] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RTP] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3005, July 2003.

[SONET] "Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates and Formats", ANSI T1.105-2001, October 2001.

18.2. Informative References

- [CONG] Floyd, S., "Congestion Control Principles", RFC 2914, September 2000.
- [DIFFSERV] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [EF] Davie, B., Charny, A., Bennett, J., 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.
- [GS] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, September 1997.
- [INTSERV] Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [PWE3-ARCH] Bryant, S. and P. Pate, "PWE3 Architecture", RFC 3985, March 2005.
- [PWE3-MPLSCW] Bryant, S., Swallow, G., and D. McPherson, "Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [PWE3-REQ] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.
- [PWE3-TDM-REQ] Riegel, M., "Requirements for Edge-to-Edge Emulation of TDM Circuits over Packet Switching Networks (PSN)", RFC 4197, October 2005.
- [RFC3711] Baugher, M., McGrew, D., Naslund, N., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

Authors' Addresses

Andrew G. Malis
Verizon Communications
40 Sylvan Road
Waltham, MA 02451
USA

EEmail: andrew.g.malis@verizon.com

Prayson Pate
Overture Networks
507 Airport Blvd, Suite 111
Morrisville, NC 27560
USA

EEmail: prayson.pate@overturenetworks.com

Ron Cohen (editor)
Resolute Networks
15 Central Avenue
Modiin, 71700
Israel

EEmail: ronc@resolutenetworks.com

David Zelig
Corrigent Systems
126 Yigal Alon st.
Tel Aviv,
Israel

EEmail: davidz@corrigent.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.

Acknowledgement

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