

Network Working Group
Request for Comments: 5143
Obsoleted by: 4842
Category: Historic

A. Malis
Verizon Communications
J. Brayley
J. Shirron
ECI Telecom Inc.
L. Martini
Cisco Systems, Inc.
S. Vogelsang
Alcatel-Lucent
February 2008

Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH)
Circuit Emulation Service over MPLS (CEM) Encapsulation

Status of This Memo

This memo defines a Historic Document for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

IESG Note

The IESG thinks that this work is related to IETF work done in WG PWE3, but this does not prevent publishing.

Abstract

This document describes a historical method for encapsulating Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Path signals for transport across packet-switched networks (PSNs). The PSNs explicitly supported by this document include MPLS and IP. Note that RFC 4842 describes the standards-track protocol for this functionality, and new implementations must use RFC 4842 rather than this document except when interoperability with older implementations is desired.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	3
3. Scope	3
4. CEM Encapsulation Format	4
4.1. Transport Encapsulation	6
4.1.1. MPLS Transport	6
4.1.2. IP Transport	7
5. CEM Operation	8
5.1. Introduction and Terminology	8
5.1.1. CEM Packetizer and De-Packetizer	9
5.1.2. CEM DBA	9
5.2. Description of Normal CEM Operation	10
5.3. Description of CEM Operation during DBA	10
5.4. Packet Synchronization	11
5.4.1. Acquisition of Packet Synchronization	11
5.4.2. Loss of Packet Synchronization	11
6. SONET/SDH Maintenance Signals	12
6.1. SONET/SDH to PSN	12
6.1.1. AIS-P Indication	13
6.1.2. STS SPE Unequipped Indication	14
6.1.3. CEM-RDI	14
6.2. PSN to SONET/SDH	15
6.2.1. AIS-P Indication	15
6.2.2. STS SPE Unequipped Indication	15
7. Clocking Modes	16
7.1. Synchronous	16
7.1.1. Synchronous Unstructured CEM	16
7.1.2. Synchronous Structured CEM	16
7.2. Asynchronous	17
8. CEM LSP Signaling	17
9. Security Considerations	18
10. IANA Considerations	18
11. References	18
11.1. Normative References	18
11.2. Informative References	19
Appendix A. SONET/SDH Rates and Formats	20
Appendix B. ECC-6 Definition	21

1. Introduction

This document describes a historical method for encapsulating SONET/SDH Path signals for transport across packet-switched networks (PSNs).

The native transmission system for circuit-oriented Time Division Multiplexing (TDM) signals is the Synchronous Optical Network (SONET) [T1.105], [GR-253]/Synchronous Digital Hierarchy (SDH) [G.707]. To support TDM traffic (which includes voice, data, and private leased line services), PSNs must emulate the circuit characteristics of SONET/SDH payloads. MPLS labels and a new circuit emulation header are used to encapsulate TDM signals and provide the Circuit Emulation Service over MPLS (CEM) function. The MPLS encapsulation may be further encapsulated in IP for carriage across IP PSNs [RFC4023].

This document also describes an optional extension to CEM called Dynamic Bandwidth Allocation (DBA). This is a method for dynamically reducing the bandwidth utilized by emulated SONET/SDH circuits in the packet network. This bandwidth reduction is accomplished by not sending the SONET/SDH payload through the packet network under certain conditions, such as Alarm Indication Signal - Path (AIS-P) or Synchronous Transport Signal Synchronous Payload Envelope (STS SPE) Unequipped.

2. Conventions Used in This Document

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].

3. Scope

This document describes how to provide CEM for the following digital signals:

1. SONET STS-1 synchronous payload envelope (SPE)/SDH VC-3
2. STS-Nc SPE (N = 3, 12, or 48)/SDH VC-4, VC-4-4c, VC-4-16c
3. Unstructured SONET Emulation, where the entire SONET bit-stream (including the transport overhead) is packetized and transported across the PSN.

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

Other SONET/SDH signals, such as virtual tributary (VT) structured sub-rate mapping, are not explicitly discussed in this document; however, it can be extended in the future to support VT and lower speed non-SONET/SDH services. OC-192c SPE/VC-4-64c are also not included at this point, since most PSNs use OC-192c or slower trunks, and thus would not have sufficient capacity. As trunk capacities increase in the future, the scope of this document can be accordingly extended.

4. CEM Encapsulation Format

In order to transport SONET/SDH SPEs through a packet-oriented network, the SPE is broken into fragments. A 32-bit CEM header is pre-pended to each fragment. The Basic CEM packet appears in Figure 1.




Figure 1. Basic CEM Packet

The 32-bit CEM header has the following format:




Figure 2. CEM Header Format

The above fields are defined as follows:

D-bit: This bit signals DBA Mode. It MUST be set to zero for normal operation, and it MUST be set to one if CEM is currently in DBA mode. DBA is an optional mode during which trivial SPEs are not transmitted into the packet network. See Table 1 and sections 7 and 8 for further details.

Note: for unstructured CEM, the D-bit MUST be set to zero.

R bit: CEM-RDI (Remote Defect Indicator). This bit is set to one to signal to the remote CEM function that a loss of packet synchronization has occurred.

Rvd: These bits are reserved for future use, and MUST be set to zero.

Sequence Number: This is a packet sequence number, which MUST continuously cycle from 0 to 1023. It SHOULD begin at zero when a CEM LSP (Label Switched Path) is created.

Structure Pointer: The Structure Pointer MUST contain the offset of the J1 byte within the CEM payload. The value is from 0 to 1,022, where 0 means the first byte after the CEM header. The Structure Pointer MUST be set to 0x3FF (1,023) if a packet does not carry the J1 byte. See [T1.105], [G.707], and [GR-253] for more information On the J1 byte and the SONET/SDH payload pointer.

Note: for unstructured CEM, the Structure Pointer field MUST be set to 0x3FF.

The N and P bits: These bits indicate negative and positive pointer adjustment events. They are also used to relay SONET/SDH maintenance signals, such as AIS-P. See Table 1 and sections 7 and 8 for more details.

Note: for unstructured CEM, the N and P bits MUST both be set to zero.

D	N	P	Interpretation	
0	0	0	Normal Mode	No Ptr Adjustment
0	0	1	Normal Mode	Positive Ptr Adjustment
0	1	0	Normal Mode	Negative Ptr Adjustment
0	1	1	Normal Mode	AIS-P
1	0	0	DBA Mode	STS SPE Unequipped
1	0	1	DBA Mode	STS SPE Unequipped Pos Ptr Adj
1	1	0	DBA Mode	STS SPE Unequipped Neg Ptr Adj
1	1	1	DBA Mode	AIS-P

Table 1. Interpretation of D, N, and P bits

ECC-6: An Error Correction Code to protect the CEM header. This offers the ability to correct single bit errors and detect up to two bit errors. The ECC algorithm is described in Appendix B. The ECC-6 can be optionally disabled at provisioning time. If the ECC-6 is not utilized, it MUST be set to zero.

Note: Normal CEM packets are fixed in length for all of the packets of a particular emulated TDM stream. This length is signaled using the CEM Payload Bytes parameter defined in [RFC4447], or is statically provisioned for each TDM stream. Therefore, the length of each CEM packet does not need to be carried in the CEM header.

Note: Setting the D-bit to 0 and the R bit to 1 violates the Best Current Practice defined in [RFC4928] when operating on MPLS networks. In this situation, MPLS networks could mistake a CEM payload as the first nibble of an IPv4 packet, potentially causing mis-ordering of packets on the pseudowire in the presence of IP Equal Cost Multi-Path (ECMP) in the MPLS network. The use of this CEM header preceded the use of MPLS ECMP. As stated earlier, [RFC4842] describes the standards-track protocol for this functionality, and it does not share this violation.

4.1. Transport Encapsulation

In principle, CEM packets can be transported over any packet-oriented network. The following sections describe specifically how CEM packets MUST be encapsulated for transport over MPLS or IP networks.

4.1.1. MPLS Transport

To transport a CEM packet over an MPLS network, an MPLS label stack MUST be pushed on top of the CEM packet.

The last two labels prior to the CEM header are referred to as the Tunnel and Virtual Circuit (VC) labels.

The VC label is required, and is the last label prior to the CEM Header. The VC label MUST be used to identify the CEM connection within the MPLS tunnel.

The optional tunnel label is immediately above the VC label on the label stack. If present, the tunnel label MUST be used to identify the MPLS LSP used to tunnel the TDM packets through the MPLS network (the tunnel LSP).

This is similar to the label stack usage defined in [RFC4447].




Figure 3. Typical MPLS Transport Encapsulation

4.1.2. IP Transport

It is highly desirable to define a single encapsulation format that will work for both IP and MPLS. Furthermore, it is desirable that the encapsulation mechanism be as efficient as possible.

One way to achieve these goals is to map CEM directly onto IP by mapping the previously described MPLS packets onto IP.

A mechanism for carrying MPLS over IP is described in [RFC4023].

Using this encapsulation scheme would result in the packet format illustrated in Figure 4.




Figure 4. MPLS Transport Encapsulation

5. CEM Operation

The following sections describe CEM operation.

5.1. Introduction and Terminology

There are two types of CEM: structured and unstructured.

Unstructured CEM packetizes the entire SONET/SDH bit-stream (including transport overhead).

Structured CEM terminates the transport overhead and packetizes individual channels (STS-1/Nc) within the SONET/SDH frame. Because structured CEM terminates the transport overhead, structured CEM implementations SHOULD meet the generic requirements for SONET/SDH Line Terminating Equipment as defined in [T1.105], [G.707], and [GR-253].

Implementations MUST support structured CEM and MAY support unstructured CEM.

Structured CEM MUST support a normal mode of operation and MAY support an optional extension called Dynamic Bandwidth Allocation (DBA). During normal operation, SONET/SDH payloads are fragmented, pre-pended with the CEM header, the VC label, and the PSN header, and

then transmitted into the packet network. During DBA mode, only the CEM header, the VC label, and PSN header are transmitted. This is done to conserve bandwidth when meaningful user data is not present in the SPE, such as during AIS-P or STS SPE Unequipped.

5.1.1. CEM Packetizer and De-Packetizer

As with all adaptation functions, CEM has two distinct components: adapting TDM SONET/SDH into a CEM packet stream, and converting the CEM packet stream back into a TDM SONET/SDH. The first function will be referred to as CEM packetizer and the second as CEM de-packetizer. This terminology is illustrated in Figure 5.




Figure 5. CEM Terminology

Note: the CEM de-packetizer requires a buffering mechanism to account for delay variation in the CEM packet stream. This buffering mechanism will be generically referred to as the CEM jitter buffer.

5.1.2. CEM DBA

DBA is an optional mode of operation for structured CEM that only transmits the CEM header, the VC label, and PSN header into the packet network under certain circumstances, such as AIS-P or STS SPE Unequipped.

If DBA is supported by a CEM implementation, the user SHOULD be able to configure if DBA will be triggered by AIS-P, STS SPE Unequipped, both, or neither on a per channel basis.

If DBA is supported, the determination of AIS-P and STS SPE Unequipped MUST be based on the state of SONET/SDH Section, Line, and Path Overhead bytes. DBA based on pattern detection within the SPE (i.e., all zeros, 7Es, or ATM idle cells) is for further study.

During AIS-P, there is no valid payload pointer, so pointer adjustments cannot occur. During STS SPE Unequipped, the SONET/SDH payload pointer is valid, and therefore pointer adjustments MUST be supported even during DBA. See Table 1 for details.

5.2. Description of Normal CEM Operation

During normal operation, the CEM packetizer will receive a fixed rate byte stream from a SONET/SDH interface. When a packet's worth of data has been received from a SONET/SDH channel, the CEM header, the VC Label, and PSN header are pre-pended to the SPE fragment and the resulting CEM packet is transmitted into the packet network. Because all normal CEM packets associated with a specific SONET/SDH channel will have the same length, the transmission of CEM packets for that channel SHOULD occur at regular intervals.

At the far-end of the packet network, the CEM de-packetizer will receive packets into a jitter buffer and then play 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. The received packet rate from the packet network should be exactly balanced by the transmission rate onto the SONET/SDH channel, on average. The time over which this average is taken corresponds to the depth of the jitter buffer for a specific CEM channel.

The CEM sequence numbers provide a mechanism to detect lost and/or mis-ordered packets. The CEM de-packetizer MUST detect lost or mis-ordered packets. The CEM de-packetizer MUST play out a programmable byte pattern in place of any dropped packets. The CEM de-packetizer MAY re-order packets received out of order. If the CEM de-packetizer does not support re-ordering, it MUST drop mis-ordered packets.

5.3. Description of CEM Operation during DBA

(Note: DBA is only applicable to structured CEM.)

There are several issues that should be addressed by a workable CEM DBA mechanism. First, when DBA is invoked, there should be a substantial savings in bandwidth utilization in the packet network. The second issue is that the transition in and out of DBA should be tightly coordinated between the local CEM packetizer and CEM de-packetizer at the far side of the packet network. A third is that the transition in and out of DBA should be accomplished with minimal disruption to the adapted data stream.

Another goal is that the reduction of CEM traffic due to DBA should not be mistaken for a fault in the packet network or vice-versa. Finally, the implementation of DBA should require minimal modifications beyond what is necessary for the nominal CEM case. The mechanism described below is a reasonable balance of these goals.

During DBA, packets MUST be emitted at exactly the same rate as they would be during normal operation. This SHOULD be accomplished by transmitting each DBA packet after a complete packet of data has been received from the SONET/SDH channel. The only change from normal operation is that the CEM packets during DBA MUST only carry the CEM header, the VC label, and the PSN header.

Because some links have a minimum supported packet size, the CEM packetizer MAY append a configurable number of bytes immediately after the CEM header to pad out the CEM packet to reach the minimum supported packet size. The value of the padding bytes is implementation specific. The D-bit MUST be set to one, to indicate that DBA is active.

The CEM de-packetizer MUST assume that each packet received with the D-bit set represents a normal-sized packet containing an AIS-P or STS SPE Unequipped payload as noted by N and P, (see Table 1). The CEM de-packetizer MUST accept DBA packets with or without padding.

This allows the CEM packetization and de-packetization logic during DBA to be similar to the nominal case. It insures that the correct SONET/SDH indication is reliably transmitted between CEM adaptation points. It minimizes the risk of under or over running the jitter buffer during the transition in and out of DBA. And, it guarantees that faults in the packet network are recognized as distinctly different from line conditioning on the SONET/SDH interfaces.

5.4. Packet Synchronization

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

5.4.1. Acquisition of Packet Synchronization

At startup, a CEM de-packetizer will be out of packet synchronization by default. To declare packet synchronization at startup or after a loss of packet synchronization, the CEM de-packetizer must receive a configurable number of CEM packets with sequential sequence numbers.

5.4.2. Loss of Packet Synchronization

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

As discussed in section 5.2, a CEM de-packetizer MAY support the re-ordering of mis-ordered packets.

If a CEM de-packetizer supports re-ordering, then the determination that packet synchronization has been lost cannot be made at the time the packets are received from the PSN. Instead, the determination MUST be made as the packets are being played out onto the SONET/SDH interface. This is because it is only at play-out time that the determination can be made as to whether the entire emulated SONET/SDH stream was received from the PSN.

If a CEM de-packetizer does not support re-ordering, a number of approaches may be used to minimize the impact of mis-ordered or lost packets on the final re-assembled SONET/SDH stream. For example, ATM Adaptation Layer 1 (AAL1) [I.363.1] uses a simple state-machine to re-order packets in a subset of possible cases. The algorithm for these state-machines is outside of the scope of CEM. However, the final determination as to whether or not to declare loss of packet synchronization MUST be based on the same criteria as for implementations that do support re-ordering.

Whether or not a CEM implementation supports re-ordering, the declaration of loss of packet synchronization MUST be based on the following criteria.

As packets are played out towards the SONET/SDH interface, the CEM de-packetizer will encounter empty packets in the place of packets that were dropped by the PSN, or effectively dropped due to limitations of the CEM implementation. If the CEM de-packetizer encounters more than a configurable number of sequential dropped packets, the CEM de-packetizer MUST declare loss of packet synchronization.

6. SONET/SDH Maintenance Signals

There are several issues that must be considered in the mapping of maintenance signals between SONET/SDH and a PSN. A description of how these signals and conditions are mapped between the two domains is given below.

For clarity, the mappings are split into two groups: SONET/SDH to PSN and PSN to SONET/SDH.

6.1. SONET/SDH to PSN

The following sections describe how SONET/SDH Maintenance Signals and Alarm conditions are mapped into a Packet-Switched Network.

6.1.1. AIS-P Indication

In a SONET/SDH network, SONET/SDH Path outages are signaled using maintenance alarms, such as Path AIS (AIS-P). In particular, AIS-P indicates that the SONET/SDH Path is not currently transmitting valid end-user data, and the SPE contains all ones.

It should be noted that for structured CEM, nearly every type of service-affecting section or line defect will result in an AIS-P condition.

The SONET/SDH hierarchy is illustrated below.




Figure 6. SONET/SDH Fault Hierarchy

Should the Section Layer detect a Loss of Signal (LOS) or Loss of Frame (LOF) condition, it sends AIS-L up to the Line Layer. If the Line Layer detects AIS-L or Loss of Path (LOP), it sends AIS-P to the Path Layer.

In normal mode during AIS-P, structured CEM packets are generated as usual. The N and P bits MUST be set to 11 binary to signal AIS-P explicitly through the packet network. The D-bit MUST be set to zero

to indicate that the SPE is being carried through the packet network. Normal CEM packets with the SPE fragment, CEM header, the VC label, and PSN header MUST be transmitted into the packet network.

However, to conserve network bandwidth during AIS-P, DBA MAY be employed. If DBA has been enabled for AIS-P and AIS-P is currently occurring, the N and P bits MUST be set to 11 binary to signal AIS, and the D-bit MUST be set to one to indicate that the SPE is not being carried through the packet network. Only the CEM header, the VC label, and the PSN header MUST be transmitted into the packet network.

Also note that this differs from the outage mechanism in [RFC4447], which withdraws the VC label as a result of an endpoint outage. TDM circuit emulation requires the ability to distinguish between the de-provisioning of a circuit (which causes the VC label to be withdrawn), and temporary outages (which are signaled using AIS-P).

6.1.2. STS SPE Unequipped Indication

The STS SPE Unequipped Indication is a slightly different case than AIS-P. When byte C2 of the Path Overhead (STS path signal label) is 00h and Byte B3 (STS Path BIP-8) is valid, it indicates that the STS SPE is unequipped. Note: this is typically signaled by setting the entire SPE to zeros.

For normal structured CEM operation during STS SPE Unequipped, the N and P bits MUST be interpreted as usual. The SPE MUST be transmitted into the packet network along with the CEM header, the VC label, and PSN header, and the D-Bit MUST be set to zero.

If DBA has been enabled for STS SPE Unequipped and the Unequipped condition is occurring on the SONET/SDH channel, the D-bit MUST be set to one to indicate DBA is active. Only the CEM header, the VC Label, and PSN header MUST be transmitted into the packet network. The N and P bits MUST be used to signal pointer adjustments as normal. See Table 1 and section 8 for details.

6.1.3. CEM-RDI

The CEM function MUST send CEM-RDI towards the packet network during loss of packet synchronization. This MUST be accomplished by setting the R bit to one in the CEM header. This applies for both structured and unstructured CEM.

6.2. PSN to SONET/SDH

The following sections discuss how the various conditions on the packet network are converted into SONET/SDH indications.

6.2.1. AIS-P Indication

There are several conditions in the packet network that will cause the structured CEM de-packetization function to send an AIS-P indication onto a SONET/SDH channel.

The first of these is the receipt of structured CEM packets with the N and P bits set to one, and the D-bit set to zero. This is an explicit indication of AIS-P being received at the far-end of the packet network, with DBA disabled for AIS-P. The CEM de-packetizer MUST play out the received SPE fragment (which will incidentally be carrying all ones), and MUST configure the SONET/SDH Overhead to signal AIS-P as defined in [T1.105], [G.707], and [GR-253].

The second case is the receipt of structured CEM packets with the N and P bits set to one, and the D-bit set to one. This is an explicit indication of AIS-P being received at the far-end of the packet network, with DBA enabled for AIS-P. The CEM de-packetizer MUST play out one packet's worth of all ones for each packet received, and MUST configure the SONET/SDH Overhead to signal AIS-P as defined in [T1.105], [G.707], and [GR-253].

A third case that will cause a structured CEM de-packetization function to send an AIS-P indication onto a SONET/SDH channel is loss of packet synchronization.

6.2.2. STS SPE Unequipped Indication

There are three conditions in the packet network that will cause the CEM function to transmit STS SPE Unequipped Indications onto the SONET/SDH channel.

The first, which is transparent to CEM, is the receipt of regular CEM packets that happen to be carrying an SPE that contains the appropriate Path Overhead to signal STS SPE Unequipped. This case does not require any special processing on the part of the CEM de-packetizer.

The second case is the receipt of structured CEM packets that have the D-bit set to one to indicate that DBA is active and the N and P bits set to 00 binary, 01 binary, or 10 binary to indicate STS SPE

Unequipped with or without pointer adjustments. The CEM de-packetizer MUST use this information to transmit a packet of all zeros onto the SONET/SDH interface, and adjust the payload pointer as necessary.

The third case when a structured CEM de-packetizer MUST send an STS SPE Unequipped Indication towards the SONET/SDH interface is when the VC-label has been withdrawn due to de-provisioning of the circuit.

7. Clocking Modes

It is necessary to be able to regenerate the input service clock at the output interface. Two clocking modes are supported: synchronous and asynchronous. Selection of the clocking mode is made as part of service provisioning. Both ends of the emulated circuit must be configured with the same clocking mode.

7.1. Synchronous

When synchronous SONET/SDH timing is available at both ends of the circuit, the issue of clock recovery becomes much simpler.

7.1.1. Synchronous Unstructured CEM

For unstructured CEM, the external clock is used to clock each bit onto the optical carrier.

7.1.2. Synchronous Structured CEM

For structured CEM, the external clock is used to clock the SONET/SDH carrier. The N and P bits are used to signal negative or positive pointer adjustment events between structured CEM endpoints.

If there is a frequency offset between the frame rate of the transport overhead and that of the SONET/SDH SPE, then the alignment of the SPE shall periodically slip back or advance in time through positive or negative stuffing. The N and P bits are used to replay the pointer adjustment events and eliminate transport jitter.

During a negative pointer adjustment event, the H3 byte from the SONET/SDH stream is incorporated into the CEM packet payload in order with the rest of the SPE. During a positive pointer adjustment event, the stuff byte is not included in the CEM packet payload.

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 the first packet of the three with the N/P bits set is received.

The CEM de-packetizer MUST utilize the CEM sequence numbers to insure that SONET/SDH pointer adjustment events are not played any more frequently than once per every three CEM packets transmitted by the remote CEM packetizer.

References [T1.105], [G.707], and [GR-253] specify that pointer adjustment events MUST be separated by three SONET/SDH frames without a pointer adjustment event. In order to relay all legal pointer adjustment events, the packet size for a specific circuit MUST be no larger than $(783 * 4 * N)/3$, where N is the STS-Nc multiplier.

However, some SONET/SDH equipment allows pointer adjustments to occur in back-to-back SONET/SDH frames. In order to support this possibility, the packet size for a particular circuit SHOULD be no larger than $(783*N)/3$, where N is the STS-Nc multiplier.

Since the minimum value of N is one, CEM implementations SHOULD support a minimum payload length of 783/3 or 261 bytes. Smaller payload lengths MAY be supported as an option.

7.2. Asynchronous

If synchronous timing is not available, other methods MAY be employed to regenerate the circuit timing.

For structured CEM, the CEM packetizer SHOULD generate the N and P bits as usual. However, without external synchronization, this information is not sufficient to reliably justify the SPE within the SONET/SDH transport framing at the CEM de-packetizer. The de-packetizer MAY employ an adaptive algorithm to introduce pointer adjustment events to map the CEM SPE to the SONET/SDH transport framing. Regardless of whether the N and P bits are used by the de-packetizer as part of the adaptive clock recovery algorithm, they MUST still be used in conjunction with the D-bit to signal AIS-P, STS SPE Unequipped, and DBA.

For unstructured CEM, the CEM de-packetizer MAY use an adaptive clock recovery technique to regenerate the SONET/SDH transport clock.

An example adaptive clock recovery method can be found in section 3.4.2 of [VTOA].

8. CEM LSP Signaling

For maximum network scaling in MPLS applications, CEM LSP signaling may be performed using the Label Distribution Protocol (LDP) Extended Discovery mechanism as augmented by the Pseudo-Wire id Forward Error Correction (PWid FEC) Element defined in [RFC4447]. MPLS traffic

tunnels may be dedicated to CEM, or shared with other MPLS-based services. The value 0x8008 is used for the PWE3 PW Type in the PWid FEC Element in order to signify that the LSP being signaled is to carry CEM. Note that the generic control word defined in [GR-253] is not used, as its functionality is included in the CEM encapsulation header.

Alternatively, static label assignment may be used, or a dedicated traffic engineered LSP may be used for each CEM service.

Normal CEM packets are fixed in length for all of the packets of a particular emulated TDM stream. This length is signaled using the CEM Payload Bytes parameter defined in [RFC4447], or it is statically provisioned for each CEM service.

At this time, other aspects of the CEM service MUST be statically provisioned.

9. Security Considerations

The CEM encapsulation is subject to all of the general security considerations discussed in [RFC3985] and [RFC4447]. 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 CEM 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.

10. IANA Considerations

IANA has already allocated the PWE3 PW Type value 0x0008 for this specification. No further actions are required.

11. References

11.1. Normative References

- [G.707] ITU Recommendation G.707, "Network Node Interface For The Synchronous Digital Hierarchy", 1996.
- [GR-253] Telcordia Technologies, "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria", GR-253-CORE, Issue 3, September 2000.

- [I.363.1] ITU-T, "Recommendation I.363.1, B-ISDN Adaptation Layer Specification: Type AAL1", Appendix III, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4842] Malis, A., Pate, P., Cohen, R., Ed., and D. Zelig, "Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Circuit Emulation over Packet (CEP)", RFC 4842, April 2007.
- [RFC4928] Swallow, G., Bryant, S., and L. Andersson, "Avoiding Equal Cost Multipath Treatment in MPLS Networks", BCP 128, RFC 4928, June 2007.
- [T1.105] American National Standards Institute, "Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates and Formats," ANSI T1.105-1995.
- [VTOA] ATM Forum, "Circuit Emulation Service Interoperability Specification Version 2.0", af-vtoa-0078.000, January 1997.

11.2. Informative References

- [RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

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 2 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 2. Standard SONET Line Rates

Each SONET frame is 125 us 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 nine 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 exactly 64 times larger than the 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 3 lists the SPE and payload rates supported.

SONET STS Level	STS-1	STS-3c	STS-12c	STS-48c	STS-192c
SDH VC Level	-	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 3. 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.

Appendix B. ECC-6 Definition

ECC-6 is an Error Correction Code to protect the CEM header. This provides single bit correction and the ability to detect up to two bit errors.

Error Correction Code:

-----Header bits 0-25 -----																									ECC-6 code								
0										1										2					3								
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1		
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		
	1	1	1	1	1	0	0	0	1	0	0	0	1	1	1	1	0	1	0	0	0	1	0	1	1		1	0	0	0	0	0	
	1	1	1	1	0	1	0	0	0	1	0	0	1	0	0	0	0	1	0	1	1	1	1	1	1		0	1	0	0	0	0	
	1	0	0	0	1	1	1	1	0	0	1	0	1	1	1	0	0	0	1	1	1	1	0	0	1	1		0	0	1	0	0	0
	0	1	0	0	1	1	1	1	0	0	0	1	1	0	0	1	1	1	1	1	0	0	1	1	0	1		0	0	0	1	0	0
	0	0	1	0	0	0	1	0	1	1	1	1	1	1	0	0	1	1	1	1	1	0	1	0	1	0		0	0	0	0	1	0
	0	0	0	1	0	0	0	1	1	1	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1		0	0	0	0	0	1	
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+		

Figure 7. ECC-6 Check Matrix X

The ECC-6 code protects the 32-bit CEM header as follows:

The encoder generates the 6-bit ECC using the matrix shown in Figure 7. In brief, the encoder builds another 26 column by 6 row matrix and calculates even parity over the rows. The matrix columns represent CEM header bits 0 through 25.

Denote each column of the ECC-6 check matrix by $X[]$, and each column of the intermediate encoder matrix as $Y[]$. $CEM[]$ denotes the CEM header and $ECC[]$ is the error correction code that is inserted into CEM header bits 26 through 31.

```
for i = 0 to 25 {  
    if CEM[i] = 0 {  
        Y[i] = 0;  
    } else {  
        Y[i] = X[i];  
    }  
}
```

In other words, for each CEM header bit (i) set to one, set the resulting matrix column Y[i] according to Figure 7.

The final ECC-6 code is calculated as even parity of each row in Y (i.e., ECC[k]=CEM[25+k]=even parity of row k).

The receiver also uses matrix X to calculate an intermediate matrix Y' based on all 32 bits of the CEM header. Therefore, Y' is 32 columns wide and includes the ECC-6 code.

```
for i = 0 to 31 {  
    if CEM[i] = 0 {  
        Y'[i] = 0;  
    } else {  
        Y'[i] = X[i];  
    }  
}
```

The receiver then appends the incoming ECC-6 code to Y as column 32 (ECC[0] should align with row 0) and calculates even parity for each row. The result is a single 6-bit column Z. If all 6 bits are 0, there are no bit errors (or at least no detectable errors). Otherwise, it uses Z to perform a reverse lookup on X[] from Figure 7. If Z matches column X[i], then there is a single bit error. The receiver should invert bit CEM[i] to correct the header. If Z fails to match any column of X, then the CEM header contains more than one bit error and the CEM packet MUST be discarded.

Note that the ECC-6 code provides single-bit correction and 2-bit detection of errors within the received ECC-6 code itself.

Acknowledgments

The authors would like to thank Mitri Halabi, Bob Colvin, and Ken Hsu, all formerly of Vivace Networks and Tellabs; Tom Johnson, Marlene Drost, Ed Hallman, and Dave Danenberg, all formerly of Litchfield Communications, for their contributions to the document.

Authors' Addresses

Andrew G. Malis
Verizon Communications
40 Sylvan Road
Waltham, MA 02451
EMail: andrew.g.malis@verizon.com

Jeremy Brayley
ECI Telecom Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: jeremy.brayley@ecitele.com

John Shirron
ECI Telecom Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: john.shirron@ecitele.com

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO, 80112
EMail: lmartini@cisco.com

Steve Vogelsang
Alcatel-Lucent
600 March Road
Kanata, ON K2K 2T6
Canada
EMail: steve.vogelsang@alcatel-lucent.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at www.rfc-editor.org/copyright.html, 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.

