

Network Working Group
Internet Draft
Expiration Date: February 2002

Andrew G. Malis
Ken Hsu
Vivace Networks, Inc.

Jeremy Brayley
Steve Vogelsang
John Shirron
Laurel Networks, Inc.

Luca Martini
Level 3 Communications, LLC.

Tom Johnson
Marlene Drost
Ed Hallman
Litchfield Communications, Inc.

July 2001

**SONET/SDH Circuit Emulation Service Over MPLS (CEM) Encapsulation
draft-malis-sonet-ces-mpls-05.txt**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [section 10 of RFC 2026](#) [1].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

1. Abstract

This document describes a method for encapsulating SONET/SDH Path signals for transport across packet-switched networks (PSNs). The PSNs explicitly supported by this document include MPLS and IP.

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 [RFC 2119](#) [2].

3. Introduction

This document describes a method for encapsulating time division multiplexed (TDM) digital signals for transmission over packet-switched networks.

The transmission system for circuit-oriented TDM signals is the Synchronous Optical Network (SONET) [3], [6] / Synchronous Digital Hierarchy (SDH) [4]. 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 [8].

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 AIS-P or STS SPE Unequipped.

This document is closely related to references [5], which describes the control protocol methods used to signal the usage of CEM, [7], which describes a related method of encapsulating Layer 2 frames over MPLS and which shares the same signaling, and [11] which describes a MIB for controlling and observing CEM services.

4. 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.

5. 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: Signals DBA Mode. MUST be set to zero for Normal Operation. 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. 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 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 [3], [4] and [6] 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: 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 0.

+---+---+---+-----+			
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 [5], 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.

5.1 Transport Encapsulation

Malis, et al.

Expires February 2002

[Page 4]

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.

5.1.1 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 [5] and [7].




Figure 3. Typical MPLS Transport Encapsulation

5.1.1 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 [8].

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

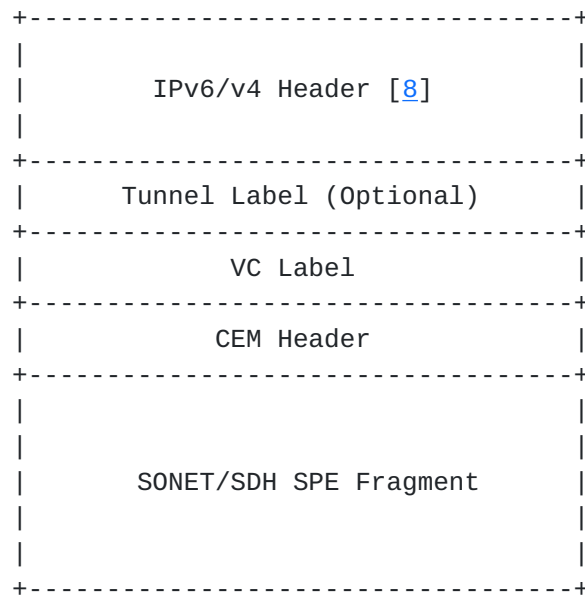


Figure 4. MPLS Transport Encapsulation

account for delay variation in the CEM packet stream. This buffering mechanism will be generically referred to as the CEM jitter buffer.

6.1.2 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 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 Unequipped, both, or neither on a per channel basis.

If DBA is supported, the determination of AIS-P and STS 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 Unequipped, the SONET/SDH payload pointer is valid, and therefore pointer adjustments MUST be supported even during DBA. See Table 1 for details.

6.2 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 packets 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 receive 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.

6.3 Description of CEM Operation during DBA

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

Malis, et al.

Expires February 2002

[Page 8]

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 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.

6.4 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.

6.4.1 Acquisition of Packet Synchronization

Malis, et al.

Expires February 2002

[Page 9]

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.

6.4.2 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 6.2](#), a CEM de-packetizer MAY or MAY NOT support 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, AAL1 [9] 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.

7. 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 described below.

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

7.1 SONET/SDH to PSN

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

7.1.1.1 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 [\[5\]](#), 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).

[7.1.2](#) 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 SPE is unequipped. Note: this is typically signaled by setting the entire SPE to zeros.

For normal structured CEM operation during 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.

[7.1.3](#) 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.

[7.2](#) PSN to SONET/SDH

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

[7.2.1](#) 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 [3], [4], and [6].

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 [3], [4], and [6].

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.

7.2.2 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 DBA active and the N and P bits set to 00 binary, 01 binary, or 10 binary to indicate 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.

8. 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.

8.1 Synchronous

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

8.1.1 Synchronous Unstructured CEM

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

8.1.2 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 justification events between structured CEM end-points.

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 N/P bit 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 [3],[4],and [6] 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.

8.2 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, 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 [10].

9. CEM LSP Signaling

For maximum network scaling in MPLS applications, CEM LSP signaling may be performed using the LDP Extended Discovery mechanism as augmented by the VC FEC Element defined in [5]. MPLS traffic tunnels may be dedicated to CEM, or shared with other MPLS-based services. The value 8008 is used for the VC Type in the VC FEC Element in order to signify that the LSP being signaled is to carry CEM. Note that the generic control word defined in [6] 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 [5], or is statically provisioned for each CEM service.

At this time, other aspects of the CEM service MUST be statically provisioned. The CEM-MIB [[11](#)] provides a method for statically provisioning CEM services.

In [5], Parameter ID 0x05 is allocated for CEM Options. At this time, the CEM Options parameter is reserved and MUST be set to zero. The use of the CEM Options parameter is for future consideration.

10. Open Issues

Future revisions of this document may discuss the following items.

Underlying MPLS QoS requirements are not covered by this revision of the draft. Future revisions may discuss underlying QoS requirements.

Support for VT and lower speed non-SONET/SDH services are not covered in this revision of the draft. Future revisions may address VT and non-SONET/SDH TDM services.

Use of the CEM Options parameter in [5] is currently undefined. Future revision of this draft will determine how the CEM Options word will be used.

The CEM-MIB[11] calls out a number of CEM performance parameters such as Errored Seconds, Severely Errored Seconds, and Unavailable Seconds. These terms are currently not defined in this draft. The definition of these terms is for future consideration.

The current draft only considers DBA based on SONET/SDH Overhead. It may be desirable to extending DBA to include pattern-based suppression such as long runs of HDLC flags (i.e. 0x7E).

11. Security Considerations

As with [5], this document does not affect the underlying security issues of MPLS.

12. Intellectual Property Disclaimer

This document is being submitted for use in IETF standards discussions. Vivace Networks, Inc. has filed one or more patent applications relating to the CEM technology outlined in this document. Vivace Networks, Inc. will grant free unlimited licenses for use of this technology.

13. References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997
- [3] American National Standards Institute, "Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates and Formats," ANSI T1.105-1995.
- [4] ITU Recommendation G.707, "Network Node Interface For The Synchronous Digital Hierarchy", 1996.
- [5] Martini et al, "Transport of Layer 2 Frames Over MPLS", [draft-martini-l2circuit-trans-mpls-06.txt](#), work in progress, July 2001.
- [6] Telcordia Technologies, Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria÷, GR-253-CORE, Issue 3, September 2000.
- [7] Martini et al, "Encapsulation Methods for Transport of Layer 2 Frames Over MPLS", [draft-martini-l2circuit-encap-mpls-02.txt](#), work in progress, July 2001.
- [8] Worster, MPLS Label Stack Encapsulation in IP÷, [draft-worster-mpls-in-ip-05](#), work in progress, July 2001.
- [9] ITU-T, Recommendation I.363.1, B-ISDN Adaptation Layer Specification: Type AAL1÷, [Appendix III](#), August 1996.
- [10] ATM Forum, "Circuit Emulation Service Interoperability Specification Version 2.0", af-vtoa-0078.000, January 1997.
- [11] Danenberg et al, "SONET/SDH Circuit Emulation Service Over MPLS (CEM) Management Information Base Using SMIV2", [draft-danenberg-pw-cem-mib-01.txt](#), work in progress, July 2001.

14. Acknowledgments

Malis, et al.

Expires February 2002

[Page 18]

The authors would like to thank Mitri Halabi and Bob Colvin, both of Vivace Networks, for their comments and suggestions.

15. Authors' Addresses

Andrew G. Malis
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Email: Andy.Malis@vivacenetworks.com

Ken Hsu
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Email: Ken.Hsu@vivacenetworks.com

Jeremy Brayley
Laurel Networks, Inc.
2706 Nicholson Rd.
Sewickley, PA 15143
Email: jbrayley@laurelnetworks.com

Steve Vogelsang
Laurel Networks, Inc.
2706 Nicholson Rd.
Sewickley, PA 15143
Email: sjv@laurelnetworks.com

John Shirron
Laurel Networks, Inc.
2607 Nicholson Rd.
Sewickley, PA 15143
Email: jshirron@laurelnetworks.com

Luca Martini
Level 3 Communications, LLC.
1025 Eldorado Blvd.
Broomfield, CO 80021
Email: luca@level3.net

Thomas K. Johnson
Litchfield Communications, Inc.
76 Westbury Park Rd.
Watertown, CT 06795
Email: tom_johnson@litchfieldcomm.com

Ed Hallman

Litchfield Communications, Inc.
76 Westbury Park Rd.
Watertown, CT 06795
Email: ed_hallman@litchfieldcomm.com

Malis, et al.

Expires February 2002

[Page 19]

Marlene Drost
Litchfield Communications, Inc.
76 Westbury Park Rd.
Watertown, CT 06795
Email: marlene_drost@litchfieldcomm.com

[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 $125 \mu\text{s}$ and consists of nine rows. An STS-N frame has nine rows and $N \times 90$ columns. Of the $N \times 90$ columns, the first $N \times 3$ columns are transport overhead and the other $N \times 87$ 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 $(N \times 87) - (N/3 - 1)$ columns per frame. Thus, the payload capacity of an STS-3c is $(3 \times 87 - 1) \times 9 = 2,340$ 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 $783 \times 8 \times 8,000 = 50.112$ 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 2 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 2. 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 2
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1 1 1 1 1 0 0 0 1 0 0 0 1 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 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 1, 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. $ECC[k] = CEM[25+k] = \text{even parity of row } k$).

The receiver also uses matrix X to calculate an intermediate matrix $Y\check{E}$ based on all 32 bits of the CEM header. Therefore $Y\check{E}$ is 32 columns wide and includes the ECC-6 code.

```

for i = 0 to 31 {
    if CEM[i] = 0 {
         $Y\check{E}[i] = 0$ ;
    } else {
         $Y\check{E}[i] = X[i]$ ;
    }
}

```

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

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this

document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for

copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

