

**Applicability Statement for PPP over SONET/SDH
draft-ietf-pppext-sonet-as-00.txt**

Status of this Memo

This document is an Internet-Draft. 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 not appropriate to use Internet Drafts as reference material, or to cite them other than as a ``working draft'' or ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the internet-drafts Shadow Directories on:

ftp.is.co.za (Africa)
nic.nordu.net (Northern Europe)
ftp.nis.garr.it (Southern Europe)
ftp.ietf.org (Eastern USA)
ftp.isi.edu (Western USA)
munnari.oz.au (Pacific Rim)

Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) William Allen Simpson (1994, 1997-1998). All Rights Reserved.

Abstract

The Point-to-Point Protocol (PPP) [[RFC-1661](#)] provides a standard method for transporting multi-protocol datagrams over point-to-point links. This document describes the use of PPP over Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) circuits.

1. Applicability Statement

PPP over SONET/SDH is intended for those implementations that desire the PPP encapsulation over high speed private point-to-point links, such as intra-campus single-mode fiber that may already be installed and unused.

Some installations have attempted to use commercial long-distance carriers. Because these services were installed and tested for voice, and carriers have bid least-cost components that do not meet the needs of high speed data services, numerous interoperability problems have arisen. The profile identifies many of these pitfalls, and recommends practices that avoid them.

1.1. Terminology

The various committees changing the SONET/SDH specifications have been inconsistent in their terminology. This specification uses a few simplified terms:

block	A fixed-size time-based multiplexing unit, carried from interface to interface; also known as the SONET Synchronous Transport Signal (STS-N) Frame, or the SDH Synchronous Transport Module (STM-N) Frame Structure. This use of "frame" conflicts with link-layer use of the same term. The format has more in common with fixed-length unit record equipment, such as a magnetic tape, than with a variable-length packet.
byte	An 8-bit quantity; also known as "octet" in standardese.
envelope	A fixed-size time-based multiplexing unit, carried within successive blocks along a path; also known as the SONET Synchronous Payload Envelope (SPE), or the SDH Administrative Unit Group (AUG).
frame	A variable-length unit of transmission, which is passed across the interface between the link-layer and the physical-layer. A single packet is usually wrapped in a frame, unless link-layer packet fragmentation is performed, or multiple packets are aggregated into a single frame.
packet	The basic unit of link encapsulation, which is passed across the interface between the network-

Simpson

expires in six months

[Page 1]

layer and the link-layer. The packet information is variable-length.

In this document, the key words "MAY", "MUST", "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [[RFC-2119](#)].

To remain consistent with standard Internet practice, and avoid confusion for people used to reading RFCs, all binary numbers in the following descriptions are in Most Significant Bit to Least Significant Bit order, from Most Significant Byte to Least Significant Byte, reading from left to right, unless otherwise indicated. Note that this is contrary to ISO and ITU practice, which orders bits as transmitted (network bit order). Keep this in mind when comparing this document with the other documents.

2. SONET/SDH Profile

"The nice thing about standards is that you have so many to choose from; furthermore, if you do not like any of them, you can just wait for next year's model." [Tanenbaum]

"The problem with standards is the last 's'." [Kleinrock]

SONET and SDH were "standardized" by at least 4 different organizations.

- Bellcore developed the original SONET specification, and has issued significantly reorganized and revised documents in 1991 and 1995. According the [[SONET](#)] Revision History: "Since TR-TSY-000253, Issue 1, there have been so many revisions to this material that the use of change bars is impractical."
- ANSI (T1X1 and T1M3) has a separate series of documents, that are continuously revised, and strongly influenced by a single vendor.
- ETSI also has a separate series of documents, but has greatly benefited by close attention to detail, and coordination with CCITT/ITU-T.
- CCITT (since renamed ITU-T) published the SDH specifications in 1988. Minor revisions occurred in 1993 and 1994, with a major reorganization in 1996.

The latter organizations have not maintained consistent naming of fields and payloads, have made conflicting changes and extensions, and have been careless about conformity and backward compatibility

Simpson

expires in six months

[Page 2]

with existing deployment.

This specification was developed utilizing carefully chosen documents that reflect a particular point in time, and that correspond to extant fielded implementations. Where naming conventions are inconsistent, or conflict with other network usage of the same terms, a simpler taxonomy is chosen.

To promote interoperability, this document provides a condensed description of basic SONET/SDH functionality, together with some of the recent changes and their relevance, and a profile of recommended practices.

2.1. Bit Transmission

All bytes are transmitted Most Significant Bit first.

Discussion

Care must be taken when converting from other media, such as serial links and ethernet, that are transmitted Least Significant Bit first.

2.1.1. Electrical Bits

SONET specifies an optional electrical transmission capability using a pair of 75 Ohm coaxial cables, one for each direction.

At 156 Mbps, Coded Mark Inversion (CMI) is specified.

single bits

-	-	-	---
-	-	---	
"0"	"0"	"1"	"1"

Simpson

expires in six months

[Page 3]

A1A2 pattern

```

      ---      ---      -      ---      -      -      -      -      ---      -      -      -
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
      ---      ---      -      ---      -      -      -      -      ---      -      -      -
      1      1      1      1      0      1      1      0      0      0      1      0      1      0      0      0

```

or its inverse.

Lower signal rates (52 Mbps) use Bipolar with Three-Zero Substitution (B3ZS). These rates are not relevant to PPP over SONET/SDH (see "Physical Layer Requirements").

Discussion

There are no recent changes that might affect interoperability.

At the time of writing, there are no reported router implementations using CMI electrical transmission to directly feed electro-optical section or line terminating equipment.

Typically, pseudo-ECL signals over very short inter-component distances are used to drive the optics with the same encoding pattern as the optical bits.

Recommendations

As the cost of optical interfaces and cabling has rapidly decreased, the use of optical transmission is preferred, even for moderately short intra-office distances.

For electrical intra-office router interconnection, use of 100baseT (or future gigabit ethernet) is preferable to PPP over SONET/SDH.

2.1.2. Optical Bits

SONET specifies an optical transmission capability using a pair of fibers, one for each direction. The use of bi-directional fiber transmission is "for further study".

The optical line coding used for all system interfaces is binary Non-Return-to-Zero (NRZ). The convention adopted for optical logic level is

Simpson

expires in six months

[Page 4]

logical 1 -- emission of light
 logical 0 -- no emission of light

single bits

```

      ---
    |   |
    |   |
    ---
  "0"  "1"
  
```

A1A2 pattern

```

-----
          | | | |
          | | | |
          ---
1  1  1  1  0  1  1  0  0  0  1  0  1  0  0  0
  
```

To simplify the development of compatible multi-vendor systems, it is desirable to define a relatively small set of categories and corresponding optical interface specifications.

Short Reach (intra-office)

SONET optical sections having loss budgets of 0 dB to 7 dB; alternatively, SDH interconnect distances less than 2 kilometers.

Intermediate Reach (short haul inter-office)

SONET optical sections having loss budgets of 0 dB to 12 dB; alternatively, SDH interconnect distances up to 15 kilometers.

Long Reach (long haul inter-office)

SONET optical sections having loss budgets of 10 dB to 28 dB; alternatively, SDH interconnect distances up to 40 kilometers at 1310 nanometers, and up to 60 kilometers at 1550 nanometers.

Unfortunately, the current proliferation of specifications falls far short of this goal, having different category definitions (as seen above) and more than 20 total pages of parameter tables.

Discussion

There are too many recent changes that impede interoperability to detail in this short overview.

Care must be taken when comparing signals from other media, such

Simpson

expires in six months

[Page 5]

as high speed serial links, where "no signal" is "logical 1".

At the time of writing, the majority of reported implementations use "intermediate" level 1310 nanometer optics with single mode fiber. This can accommodate multi-mode fiber and short reach (intra-office) interconnection with the addition of transmitter attenuation.

Upgrading to future higher speeds would be facilitated by installing single mode fiber instead of multi-mode fiber.

Recommendations

The greatest interoperability and economies of scale will be achieved with deployment of a few general interface choices, all for single mode fiber:

Intermediate Reach equipment at 1310 nanometers for each STM line speed.

Long Reach equipment at 1550 nanometers for 2 Gbps and above.

2.1.3. Bit Scrambling

Optical interface signals use NRZ binary line coding (described above). A series of repeated bits will not feature any signal level transitions. Such transitions (zeros to ones and ones to zeros) are needed at the receiver to dynamically correct bit timing variance for line rate clock recovery. Therefore, although it has no effect on random data, the signal is scrambled against the possibility that a very long series of ones or zeros might naturally occur in the transmitted bit stream.

Electrical interface signals use line encoding (B3ZS and CMI) that assures adequate transitions. Scrambling is not needed. However, they are also scrambled for consistency between the electrical and optical component interfaces.

In both cases, a 7-bit linear feedback shift register (LFSR) is identically applied at the transmitter and receiver. The sequence of bits from the x^7 stage of the LFSR is exclusive-or'd (in most significant bit order) with the output bits before transmission, and with the corresponding input bits after signal recovery.

The first row of the STM-N transmission block overhead ($N * 9$ bytes described later) is not scrambled. The 7-bit LFSR state is reset to bit value 1111111 (127 decimal) on the most significant bit of the

Simpson

expires in six months

[Page 6]

byte following the block overhead. The scrambler runs continuously from that bit onward, throughout the remainder of the transmission block.

The generating polynomial for the LFSR is $x^{**7} + x^{**6} + 1$. The resulting 127-bit sequence repeats in a 127 byte stream (shown in hex):

```
fe 04 18 51 e4 59 d4 fa 1c 49 b5 bd 8d 2e e6 55
fc 08 30 a3 c8 b3 a9 f4 38 93 6b 7b 1a 5d cc ab
f8 10 61 47 91 67 53 e8 71 26 d6 f6 34 bb 99 57
f0 20 c2 8f 22 ce a7 d0 e2 4d ad ec 69 77 32 af
e0 41 85 1e 45 9d 4f a1 c4 9b 5b d8 d2 ee 65 5f
c0 83 0a 3c 8b 3a 9f 43 89 36 b7 b1 a5 dc ca bf
81 06 14 79 16 75 3e 87 12 6d 6f 63 4b b9 95 7f
02 0c 28 f2 2c ea 7d 0e 24 da de c6 97 73 2a
```

Discussion

There are no recent changes that might affect interoperability.

The scrambling mechanism is defective in a packet transmission context, where adjacent bytes are sequentially related. Apparently, the short polynomial assumed unrelated data that is obtained from DS0 through DS3 circuits. This is characteristic for legacy technology such as Asynchronous Transfer Mode (ATM), interleaved virtual circuits, or voice traffic.

Since the feedback is based on the internal state of the LFSR, and is not dependent on its interaction with previous bytes, any subset of the bit sequence that exactly matches the phase of the LFSR will generate a series of zeros. The bitwise inverse of the bit sequence will generate a series of ones.

This effect is not limited to direct PPP over SONET/SDH. Any packet oriented transmissions carried by SONET/SDH without interleaving would have the same result.

Moreover, although the 127 byte stream may have been appropriate for the original STS-1 (90 bytes per row), this length is manifestly insufficient for STM-N ($N * 270$ bytes per row).

Finally, the most egregious flaw in this part of the SONET/SDH specification, the number of bits permitted without a transition is not explicitly defined. Instead, the value is implicitly derived from the fixed length scrambling sequence versus the variable length Loss of Signal (LOS) condition based on the link speed (described later).

Simpson

expires in six months

[Page 7]

Note that the complete scrambling byte stream cannot be represented in a valid PPP packet. The pair "fd 03" is not conformant, limiting the maximum sequence without a transition to 127 bytes.

Recommendations

Packet transmissions carried at rates less than 156 Mbps MUST be interleaved under the VT/TU scheme (see "Physical Layer Requirements"). This is consistent with prior use of bit-synchronous HDLC (and Frame Relay) over DS0 through DS3 circuits.

Line rate clock recovery SHOULD be sufficiently stable to withstand periods of 1016 repeated one or zero bits (127 bytes).

2.1.4. Loss of Signal (LOS)

To detect a failure that occurs at the transmitter (such as laser failure) or within the transmission facility (such as a fiber cut), all incoming signals (optical and electrical) are monitored for Loss of Signal (LOS) before descrambling, by detecting a lack of sufficient signal level transitions to ensure accurate line rate clock recovery.

No transitions on the incoming signal is not considered LOS when the condition lasts 2.3 microseconds or less, and is considered LOS when the condition lasts 100 microseconds or more. The treatment of no transitions lasting between 2.3 microseconds and 100 microseconds for the purpose of LOS detection is not specified (left to the choice of the equipment designer).

For electrical (B3ZS and CMI) interfaces, the no voltage transitions condition is easily distinguished from an "all-zeros pattern".

For SONET optical (and component) interfaces, the no light pulses condition corresponds to an "all-zeros pattern". For testing SONET compliance with the LOS detection requirement, it suffices to apply an "all-zeros pattern" lasting at most 2.3 microseconds, and to apply an "all-zeros pattern" lasting at least 100 microseconds. Optical equipment shall enter the LOS state within 100 microseconds of the onset of the incoming "all-zeros pattern."

For SDH optical (and component) interfaces, equipment may also indicate LOS within 3 milliseconds of the time the received signal level drops below an implementation-determined threshold (based on receiver sensitivity). The threshold should be selected such that LOS will not be indicated while the Bit Error Rate (BER) is still acceptable. For SDH, the BER is defined to be 1 in $10^{(-3)}$.

Simpson

expires in six months

[Page 8]

Note that the secondary method is in addition to the primary method, not in lieu of it, and is not required to be implemented in SONET installations.

The equipment shall terminate the LOS condition when the incoming signal has two consecutive valid block alignment patterns (described later), and during the intervening time (125 microseconds) sufficient transitions prevent another LOS defect.

For trouble isolation purposes, especially for intermittent failures, it is important to distinguish between LOS and other signal failure conditions, such as Loss of Frame (LOF).

Discussion

There are significant differences between SONET, ANSI, and SDH; SDH does not describe the primary method of detection. Circa 1994, this resulted in the addition of the secondary method to SONET.

There is no explicit requirement given for the "all-ones pattern" that would indicate a stuck clock or laser transmitter. The same difficulty with lack of transitions will apply.

There is no rationale given for the selection of 2.3 us and 100 us. Presumably, 100 us was chosen to allow sufficient variance between 125 us block alignment patterns. However, 100 us will cause multiple errors in the block overhead (described later), while 2.3 us is too short to cause any significant damage.

The wide disparity between 2.3 us and 100 us, and leaving the error treatment to the vendor, has caused many interoperability problems. Incorrect LOS values as low as 1 ± 0.750 us have been discovered. Typical values for valid designs are 20 ± 3 us, and 26 ± 1 us. No values higher than 30 us have been found.

Line rate clock stability has been a related interoperability problem. Early 156 Mbps implementations were very robust, and could handle periods of 2,000 bits with no transitions. Some sample and/or low quality implementations have not been as robust; failure has been noted as low as 72 to 80 bits without transitions.

Simpson

expires in six months

[Page 9]

Recommendations

Using the minimum and maximum specification of LOS, line rate clock recovery **MUST** be sufficiently stable to withstand periods of repeated one or zero bits:

Mbps	min bits	max bits
156	358	15,552
622	1,431	62,208
2,488	5,724	248,832

Implementations compatible with this specification **SHOULD** observe a stricter range of LOS detection, between 13.89 us and 27.26 us:

Mbps	min bits	max bits
156	2,161	4,240
622	8,641	16,960
2,488	34,561	67,840

The lower limit affects no more than one row of overhead, while the upper limit includes at least one row of overhead.

When a circuit path is divided into multiple line sections, the LOS is not propagated end to end. A hidden section can fail independently. Therefore, the LOS indication is only advisory, and the (possibly improperly) decoded signal bits **MUST** be passed to the next section. The framed PPP packet FCS is used to detect the failure.

Some SONET framers provide a test feature to send an "all-zeroes pattern" (post scrambling) that forces the LOS alarm downstream. Alternatively, the pattern can be simulated with software. This **SHOULD** be used to test for compatible equipment during the installation of the circuit, by generating progressively longer zero patterns, to verify that the circuit meets the requirements of this profile.

2.2. Block Transmission

SONET/SDH transmits each block at precise 125 microsecond intervals. This corresponds to the traditional 4 kHz voice digital sampling rate used for DS0.

Each block is divided into 9 logical "rows".

For "concatenated" STS/STM, each row consists of some multiple of 270 byte "columns".

Simpson

expires in six months

[Page 10]

Block transmission overhead is arranged into layers, which correspond to different equipment interfaces that might occur in the path of a single circuit. Although SONET and SDH use inconsistent terms, the layers have a general equivalence.

SONET divides its Network Element (NE) interface overhead into Section (SOH), Line (LOH), and Path (POH).

SDH has a corresponding Network Node Interface (NNI) overhead called Regenerator-Section (RSOH), Multiplex-Section (MSOH), and Path (POH).

There are 3 rows of (Regenerator) Section Overhead (SOH or RSOH), and 6 rows of Line or Multiplex-Section Overhead (LOH or MSOH).

2.2.1. Block Overhead

The STS-3c/STM-1 transmission block overhead consists of 9 byte columns at the beginning of each 270 byte row, with the remaining 261 bytes used for the data envelope(s).

```

+---+---+---+---+---+---+---+---+---+---+
| A1  A1  A1  A2  A2  A2  J0  Z0  Z0  envelope
+---+---+---+---+---+---+---+---+---+---+
| B1  ++  ++  E1  ++  ??  F1  **  **  envelope
+---+---+---+---+---+---+---+---+---+---+
| D1  ++  ++  D2  ++  ??  D3  ??  ??  envelope
+---+---+---+---+---+---+---+---+---+---+
| H1  H1# H1# H2  H2# H2# H3  H3  H3  envelope
+---+---+---+---+---+---+---+---+---+---+
| B2  B2  B2  K1  ??  ??  K2  ??  ??  envelope
+---+---+---+---+---+---+---+---+---+---+
| D4  ??  ??  D5  ??  ??  D6  ??  ??  envelope
+---+---+---+---+---+---+---+---+---+---+
| D7  ??  ??  D8  ??  ??  D9  ??  ??  envelope
+---+---+---+---+---+---+---+---+---+---+
| D10 ??  ??  D11 ??  ??  D12 ??  ??  envelope
+---+---+---+---+---+---+---+---+---+---+
| S1  Z1  Z1  Z2  Z2  M1  E2  **  **  envelope
+---+---+---+---+---+---+---+---+---+---+

```

++ media dependent.

** reserved for national use.

?? reserved for future use, SHOULD be zero.

Values relevant to this specification:

Simpson

expires in six months

[Page 11]

- (A1) bit value 1111 0110 (f6 hex)
- (A2) bit value 0010 1000 (28 hex)
- (H1) and (H2) Each of the 3 consecutive pairs of H1H2 is taken as a 16-bit mask, with various combinations of bits determining the offset (in multiples of 3*N) to the start of the data envelope. An offset of 0 indicates that the envelope immediately follows the last H3 overhead byte. An offset of 522 indicates that the envelope immediately follows the last Z0 overhead byte in the next transmission block.
- When used for this specification, the second and third H1#H2# pairs are always bit value 1001 xx11 1111 1111 (93ff hex), indicating a "concatenated" envelope, where xx is ignored on receipt.

Since the block is transmitted by row, the overhead columns are interspersed with envelope columns, complicating generation and recovery.

An additional complication is added by the section-level bit scrambling. Overhead bytes of the first row (A1, A2, J0, Z0) are not scrambled.

Discussion

J0 was formerly called C1. Z0 was formerly additional copies of C1, but these bytes are now reserved for national use.

The xx bit values in H1# were originally 00 for SONET and 10 for SDH.

These changes might affect interoperability between international facilities.

The scrambler needs a reset signal to indicate the position of the beginning of the transmission block and length of the overhead. Failure of the reset signal timing could prevent recognition of the block.

Teleopolists seem inordinately fond of groups of 3. This was evident in early digital efforts (24 DS0 in T1), and continues into SONET/SDH (3 T3/E3 in STM-1). Note that all of the values, including 261 and 270, are divisible by 3. This does not scale well in a digital processing environment.

Simpson

expires in six months

[Page 12]

Recommendations

Envelopes SHOULD be sent with the H1H2 offset of 522. However, every implementation MUST be prepared to receive a variable offset.

2.2.2. Block Synchronization

Timing alignment is detected by searching for (a subset of) the A1 and A2 overhead pattern. There are two levels of synchronization recovery.

When synchronization has been lost for 3 milliseconds or more (or has never been achieved), long-term synchronization recovery requires detection of a minimum of eight (8) successive error-free A1 and A2 overhead patterns. The maximum recovery time is 3 milliseconds (24 error-free patterns).

Once alignment has been achieved, short-term synchronization loss is declared when four (4) or more consecutive A1 and A2 pattern errors have been detected. The maximum detection time is 625 microseconds. The algorithm used shall be such that, under normal operation, a $10^{(-3)}$ Poisson distribution of bit errors will not cause a false loss indication more than once every 6 minutes.

Recovery from a short-term synchronization loss requires detection of only two (2) successive error-free A1 and A2 overhead patterns. Any recovery circuitry is acceptable that achieves resynchronization within the 250 microsecond interval implied by this definition.

Failure to obtain resynchronization within 24 blocks (3 milliseconds) results in a long-term synchronization loss declaration, called Loss of Frame (LOF).

Discussion

The SDH specification is incomplete, and is missing some time intervals.

Conversely, SDH has a somewhat stricter requirement for resynchronization after a short-term loss; the algorithm used shall be such that, with a random signal, the probability for false recovery is no more than $10^{(-5)}$ per 250 microsecond interval.

Simpson

expires in six months

[Page 13]

Recommendations

Because of the interaction between block synchronization detection, termination of the Loss of Signal (LOS) condition, and the first row reset of the section-level bit scrambling, the detection circuitry **MUST** be independent (upstream) of the scrambler, and trigger the subsequent scrambler reset timing.

[2.2.3.](#) Envelope Overhead

A series of envelopes appear within the transmission block. These envelopes float to allow precise timing alignment and jitter correction as the payload is carried over the circuit path. The position of the envelope(s) head within the block is specified by the H1 and H2 pair(s). The envelope(s) tail can carry over into the next transmission block.

The STS-3c/STM-1 Path Overhead (POH for both SONET and SDH) consists of a 1 byte column, with the remaining bytes used for the payload data.

```

+----+-----
| J1  payload
+----+-----
| B3  payload
+----+-----
| C2  payload
+----+-----
| G1  payload
+----+-----
| F2  payload
+----+-----
| H4  payload
+----+-----
| Z3  payload
+----+-----
| Z4  payload
+----+-----
| Z5  payload
+----+-----

```

Values relevant to this specification:

- (J1) Path Trace. Carries one byte at a time of a repeating 64-byte fixed-length string (SONET) or a 16-byte E.164 number (SDH), so that a path receiving terminal can verify its continued connection to the

Simpson

expires in six months

[Page 14]

intended transmitter.

- (C2) Path Signal Label. Indicates the composition, construction and content of the envelope payload.
- (H4) Multipurpose Position Indicator. Used by specific payload mappings, such as Floating Virtual Tributary.

Since the envelope is transmitted by row, the path overhead columns are interspersed with payload columns, and the combination are interspersed with block overhead columns, complicating generation and recovery.

Discussion

The envelope payload provides the physical-layer interface for insertion of PPP encapsulated packets.

The J1 16-byte E.164 number format is rarely applicable, as there is no voice circuit involved in packets. Neither have fielded implementations observed strict adherence to the 64-byte limitation. A widely used implementation has generated a string consisting of the following pattern, each line CRLF terminated:

```
Remote hostname : sl-bb10
Remote interface: POS9
Remote IP addr  : 127.255.255.255
Remote Rx(K1/K2): 00/00 Tx(K1/K2): 00/00
```

The C2 value is primarily informational. It is difficult to use for applying alternate processing for adjacent envelopes, as it occurs after two rows of the envelope have already been processed. Although there is no need for intermediate transit equipment to interpret the value, some equipment fails upon encountering an unrecognized value.

At first glance this floating envelope may appear to be a highly efficient technique to avoid interface buffering of synchronized incoming data. Unfortunately, due to interspersing the overhead with the data, it is impossible to determine where the envelope begins until at least a third of the transmission block has been received.

Also, the envelope head can begin much later in the same block, or even in the next block. An envelope tail can carry over into a block that is two blocks later than the H1H2 overhead that indicates its head.

Simpson

expires in six months

[Page 15]

Thus, in a packet switched network, no matter how fast the link speed, a minimum of 125 microseconds is added at every hop, and cut-through routing is difficult or impossible.

Recommendations

See [[RFC-xxxx](#)] "Physical Layer Requirements".

2.2.4. Envelope Ordering

SONET and SDH differ in the ordering of envelope rows. Several schemes are described, including:

J1 ...
B3 ...
C2 ...
J1 ...
B3 ...
C2 ...
J1 ...
B3 ...
C2 ...

J1 ... J1 ... J1 ...
B3 ... B3 ... B3 ...
C2 ... C2 ... C2 ...

J1 J1 J1 ...
B3 B3 B3 ...
C2 C2 C2 ...

Another difficulty has been the interleaving of lower rate channels within envelopes. These are time division multiplexed.

In the SONET and SDH specifications, much of the complexity is due to these multi-stage multiplexing schemes.

Discussion

After optical specifications, this has been the worst impediment to interoperability.

As carriers have upgraded equipment to STS-3c/STM-1 and STS-12c/STM-4, some have merely reconnected older equipment with

Simpson

expires in six months

[Page 16]

coax (see "Electrical Bits"). This is often haphazard. Although this has no effect on voice, where each byte is a separate channel, this is unacceptable for packets.

Recommendations

Higher data rate envelopes, and the bytes within envelopes, MUST be sequentially ordered by row.

SDH VC-3/VC-4/VC-4-Xc ordering MUST be supported.

Envelope data MUST be carried from source to sink in sequential order. That is, bytes inserted in the order "1234" MUST arrive in the same order "1234".

At the current setup and monthly charges of hundreds of thousands of US Dollars, installations can insist that equipment be carefully configured, or will need to find an alternative carrier.

2.2.5. Payload Insertion

Various ordering schemes differ only in the location of duplicate POH. These redundant POH bytes are specified as "fixed stuff", and the values are not examined.

For example, in the previous section, the latter schemes interleave the redundant POH bytes throughout the row (VC-3/VC-4), or concatenate the redundant POH bytes at the beginning of the row (VC-4-Xc).

Recommendations

See [[RFC-xxxx](#)] "The Data Link Layer". The specification finesses the issue by filling the "fixed stuff" with packet data. This make VC-4 and VC-4-Xc identical with respect to packet processing.

2.2.6. Protection Switching

SONET provides Automatic Protection Switching (APS), also known as the SDH Multiplex Section Protection Function (MSP), to switch from one circuit to another spare circuit whenever a problem is detected. This switch can take place based on Signal Fail (SF), Signal Degrade (SD), or by manual intervention. Once switching is initiated, it can take up to 50 milliseconds to complete the switch.

Signal Fail (SF) is defined as LOS, LOF, BER exceeding $10^{(-3)}$, and various other error conditions.

Simpson

expires in six months

[Page 17]

Signal Degrade (SD) is defined as BER exceeding a user selectable range from $10^{(-5)}$ to $10^{(-9)}$.

The protection switch applies to all circuits in one direction simultaneously. Bidirectional switching is optional.

Discussion

This time domain switching facility wastes entire backup links against the rare possibility that a link might fail. Worse, these links are frequently in the same cable, and backhoe fade kills all the links.

Conversely, Internet routing utilizes all available circuits concurrently, and gracefully degrades while providing best-effort packet switching around points of failure.

Since this specification is intended for carrying PPP packets over point-to-point links in a routed network, APS/MSP is superfluous.

During intermittent failures, due to the many orders of magnitude time difference between activation (microseconds) and completion of switching (milliseconds), and the severe extended packet loss, APS/MSP link flapping can interact badly with routing protocols.

It usually costs more, with no added benefit.

Recommendations

Where SONET/SDH circuit path configuration is under control of the user, APS/MSP SHOULD be disabled.

Security Considerations

This specification introduces no known security vulnerabilities.

Section-level bit scrambling consists of a simple repeated XOR with a short, easily computed, LFSR bit stream (listed earlier). This is not a pseudo-random number generator. There is no practical cryptographic security.

Simpson

expires in six months

[Page 18]

Acknowledgements

PPP over SONET was first proposed by Craig Partridge (BBN). Some information was obtained from the good folks at Bellcore.

Technical assistance and information was also provided by Victor Demjanenko (SUNY Buffalo).

Considerable explanatory text was contributed by C. M. Heard (VVNet).

Anonymous, Karl Fox (Ascend), Peter Lothberg (Sprint), Subhash Roy (TranSwitch), Stuart Venters (Adtran), and Eric Verwillow (Juniper) provided useful critiques of earlier versions of this document.

Special thanks to Ascend Communications (nee Morning Star Technologies) for providing computing resources and network access support for writing this specification.

References

- [RFC-1661] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", STD-51, DayDreamer, July 1994.
- [RFC-1662] Simpson, W., Editor, "PPP in HDLC-like Framing", STD-51, DayDreamer, July 1994.
- [RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP-14](#), Harvard University, March 1997.
- [RFC-xxxx] Simpson, W., Editor, "PPP over SONET/SDH", DayDreamer, work in progress.
- [SONET] "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria", Bellcore, TR-NWT-000253 Issue 2, December 1991.
- [G.707] CCITT Recommendation G.707, "Synchronous Digital Hierarchy Bit Rates", June 1992.

Simpson

expires in six months

[Page 19]

Contacts

Comments about this document should be discussed on the
pppsdh@greendragon.com or ietf-ppp@merit.edu mailing lists.

This document was reviewed by the Point-to-Point Protocol Working
Group of the Internet Engineering Task Force (IETF). The working
group can be contacted via the current chair:

Karl Fox
Ascend Communications
655 Metro Place South, Suite 370
Dublin, Ohio 43017

karl@Ascend.com

Questions about this document can also be directed to:

William Allen Simpson
DayDreamer
Computer Systems Consulting Services
1384 Fontaine
Madison Heights, Michigan 48071

wsimpson@UMich.edu
wsimpson@GreenDragon.com (preferred)

Simpson

expires in six months

[Page 20]

Full Copyright Statement

Copyright (C) William Allen Simpson (1994, 1997-1998). 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, except as required to translate it into languages other than English.

This document and the information contained herein is provided on an "AS IS" basis and the author(s) 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.

Simpson

expires in six months

[Page 21]