

Workgroup: Internet Engineering Task Force
Internet-Draft: draft-schmutzer-pals-ple-00
Published: 22 February 2022
Intended Status: Standards Track
Expires: 26 August 2022
Authors: S. Gringeri J. Whittaker N. Leymann
 Verizon Verizon Deutsche Telekom
 C. Schmutzer, Ed. L. Della Chiesa
 Cisco Systems, Inc. Cisco Systems, Inc.
 N. Nainar, Ed. C. Pignataro
 Cisco Systems, Inc. Cisco Systems, Inc.
 G. Smallegange C. Brown F. Dada
 Ciena Corporation Ciena Corporation Xilinx
Private Line Emulation over Packet Switched Networks

Abstract

This document describes a method for encapsulating high-speed bit-streams as virtual private wire services (VPWS) over packet switched networks (PSN) providing complete signal transport transparency.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 26 August 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this

document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction and Motivations](#)
- [2. Requirements Notation](#)
- [3. Terminology and Reference Model](#)
 - [3.1. Terminology](#)
 - [3.2. Reference Models](#)
- [4. PLE Encapsulation Layer](#)
 - [4.1. PSN and VPWS Demultiplexing Headers](#)
 - [4.2. PLE Header](#)
 - [4.2.1. PLE Control Word](#)
 - [4.2.2. RTP Header](#)
- [5. PLE Payload Layer](#)
 - [5.1. Structure Agnostic Payload](#)
 - [5.2. Byte aligned Payload](#)
 - [5.3. 10280bit-block aligned Payload](#)
- [6. PLE Operation](#)
 - [6.1. Common Considerations](#)
 - [6.2. PLE IWF Operation](#)
 - [6.2.1. PSN-bound Encapsulation Behavior](#)
 - [6.2.2. CE-bound Decapsulation Behavior](#)
 - [6.3. PLE Performance Monitoring](#)
 - [6.4. QoS and Congestion Control](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
- [9. Acknowledgements](#)
- [10. References](#)
 - [10.1. Normative References](#)
 - [10.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction and Motivations

This document describes a method for encapsulating high-speed bit-streams as VPWS over packet switched networks (PSN). This emulation suits applications where complete signal transparency is required and data interpretation of the PE would be counter productive.

One example is two ethernet connected CEs and the need for synchronous ethernet operation between them without the intermediate PEs interfering. Another example is addressing common ethernet control protocol transparency concerns for carrier ethernet services, beyond the behavior definitions of MEF specifications.

The mechanisms described in this document follow principals similar to [\[RFC4553\]](#) but expanding the applicability beyond TDM and allow the transport of signals from many technologies such as ethernet, fibre channel, SONET/SDH [\[GR253\]](#)/[\[G.707\]](#) and OTN [\[G.709\]](#) at gigabit speeds by treating them as bit-stream payload defined in Section 3.3.3 of [\[RFC3985\]](#).

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

3. Terminology and Reference Model

3.1. Terminology

- *ACH - Associated Channel Header
- *AIS - Alarm Indication Signal
- *CBR - Constant Bit Rate
- *CE - Customer Edge
- *CSRC - Contributing SouRCe
- *ES - Errored Second
- *FEC - Forward Error Correction
- *IWF - InterWorking Function
- *LDP - Label Distribution Protocol
- *LF - Local Fault
- *MPLS - Multi Protocol Label Switching
- *NSP - Native Service Processor
- *ODUk - Optical Data Unit k
- *OTN - Optical Transport Network
- *OTUk - Optical Transport Unit k
- *PCS - Physical Coding Sublayer

- *PE - Provider Edge
- *PLE - Private Line Emulation
- *PLOS - Packet Loss Of Signal
- *PSN - Packet Switched Network
- *P2P - Point-to-Point
- *QOS - Quality Of Service
- *RSVP-TE - Resource Reservation Protocol Traffic Engineering
- *RTCP - RTP Control Protocol
- *RTP - Realtime Transport Protocol
- *SES - Severely Errored Seconds
- *SDH - Synchronous Digital Hierarchy
- *SRTP - Secure Realtime Transport Protocol
- *SRV6 - Segment Routing over IPv6 Dataplane
- *SSRC - Synchronization SouRCe
- *SONET - Synchronous Optical Network
- *TCP - Transmission Control Protocol
- *UAS - Unavailable Seconds
- *VPWS - Virtual Private Wire Service

Similarly to [[RFC4553](#)] and [[RFC5086](#)] the term Interworking Function (IWF) is used to describe the functional block that encapsulates bit streams into PLE packets and in the reverse direction decapsulates PLE packets and reconstructs bit streams.

3.2. Reference Models

The generic models defined in [[RFC4664](#)] are applicable to PLE.

PLE embraces the minimum intervention principle outlined in section 3.3.5 of [[RFC3985](#)] whereas the data is flowing through the PLE encapsulation layer as received without modifications.

For some applications the NSP function is responsible for performing operations on the native data received from the CE. Examples are

terminating FEC in case of 100GE or terminating the OTUK layer for OTN. After the NSP the IWF is generating the payload of the VPWS which carried via a PSN tunnel.

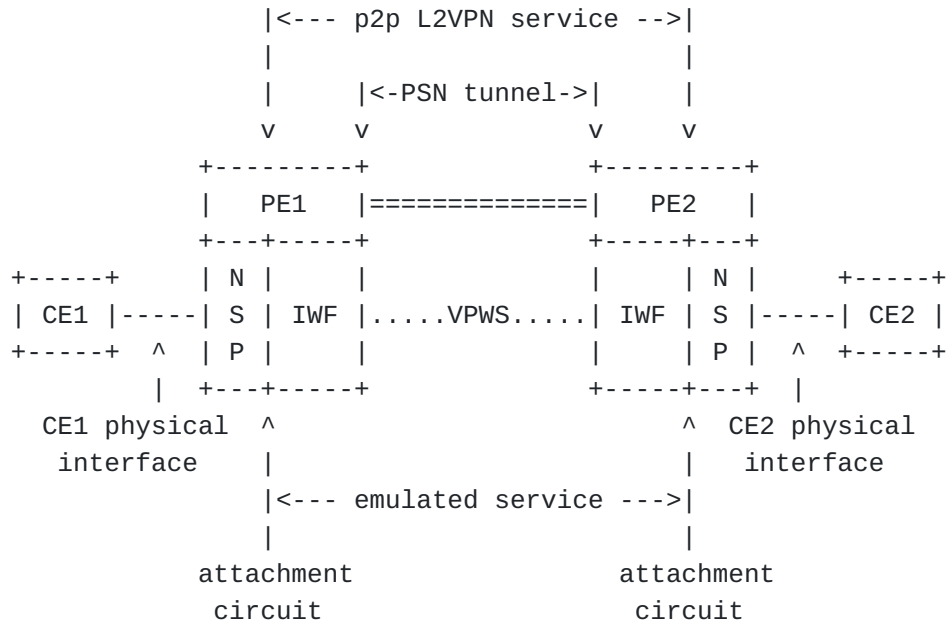


Figure 1: PLE Reference Model

To allow the clock of the transported signal to be carried across the PLE domain in a transparent way the network synchronization reference model and deployment scenario outlined in section 4.3.2 of [\[RFC4197\]](#) is applicable.




Figure 2: Relative Network Scenario Timing

The attachment circuit clock E is generated by PE2 via a differential clock recovery method in reference to a common clock I. For this to work the difference between clock I and clock A MUST be explicitly transferred between the PE1 and PE2 using the timestamp inside the RTP header.

For the reverse direction PE1 does generate the clock J in reference to clock I and the clock difference between I and G.

The way the common clock I is implemented is out of scope of this document. Well established concepts for achieving frequency synchronization in a PSN have already been defined in [[G.8261](#)] and can be applied here as well.

4. PLE Encapsulation Layer

The basic packet format used by PLE is shown in the [Figure 3](#).

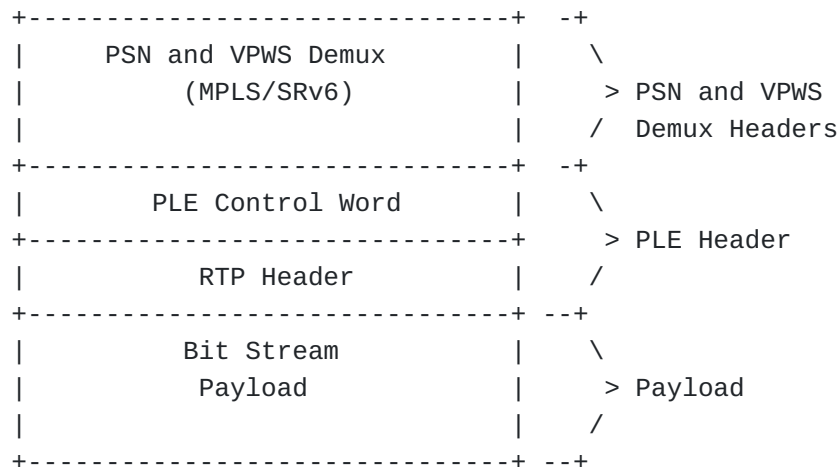


Figure 3: PLE Encapsulation Layer

4.1. PSN and VPWS Demultiplexing Headers

This document does not imply any specific technology to be used for implementing the VPWS demultiplexing and PSN layers.

When a MPLS PSN layer is used. A VPWS label provides the demultiplexing mechanism as described in section 5.4.2 of [[RFC3985](#)]. The PSN tunnel can be a simple best path Label Switched Path (LSP) established using LDP [[RFC5036](#)] or Segment Routing [[RFC8402](#)] or a traffic engineered LSP established using RSVP-TE [[RFC3209](#)] or SR-TE [[SRPOLICY](#)].

When PLE is applied to a SRv6 based PSN, the mechanisms defined in [RFC8402] and the End.DX2 endpoint behavior defined in [SRV6NETPROG] do apply.

4.2. PLE Header

The PLE header MUST contain the PLE control word (4 bytes) and MUST include a fixed size RTP header [RFC3550]. The RTP header MUST immediately follow the PLE control word.

4.2.1. PLE Control Word

The format of the PLE control word is inline with the guidance in [RFC4385] and as shown in Figure 4:




Figure 4: PLE Control Word

The first nibble is used to differentiate if it is a control word or Associated Channel Header (ACH). The first nibble MUST be set to 0000b to indicate that this header is a control word as defined in section 3 of [RFC4385].

The other fields in the control word are used as defined below:

L

Set by the PE to indicate that data carried in the payload is invalid due to an attachment circuit fault (client signal failure). The downstream PE MUST play out an appropriate replacement data. The NSP MAY inject an appropriate native fault propagation signal.

R

Set by the downstream PE to indicate that the IWF experiences packet loss from the PSN or a server layer backward fault indication is present in the NSP. The R bit MUST be cleared by the PE once the packet loss state or fault indication has cleared.

RSV

These bits are reserved for future use. This field MUST be set to zero by the sender and ignored by the receiver.

FRG

These bits MUST be set to zero by the sender and ignored by the receiver.

LEN

In accordance to [\[RFC4385\]](#) section 3 the length field MUST always be set to zero as there is no padding added to the PLE packet. To detect malformed packets the default, preconfigured or signaled payload size MUST be assumed.

Sequence Number

The sequence number field is used to provide a common PW sequencing function as well as detection of lost packets. It MUST be generated in accordance with the rules defined in Section 5.1 of [\[RFC3550\]](#) for the RTP sequence number and MUST be incremented with every PLE packet being sent.

4.2.2. RTP Header

The RTP header MUST be included and is used for explicit transfer of timing information. The RTP header is purely a formal reuse and RTP mechanisms, such as header extensions, contributing source (CSRC) list, padding, RTP Control Protocol (RTCP), RTP header compression, Secure Realtime Transport Protocol (SRTCP), etc., are not applicable to PLE VPWS.

The format of the RTP header is as shown in [Figure 5](#):

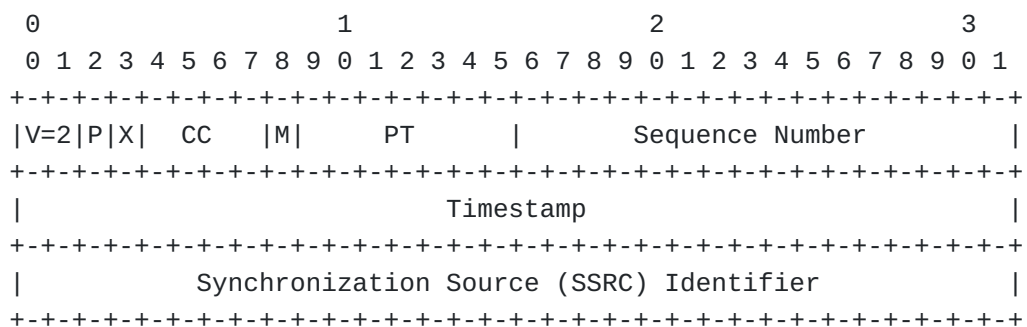


Figure 5: RTP Header

V: Version

The version field MUST be set to 2.

P: Padding

The padding flag MUST be set to zero by the sender and ignored by the receiver.

X: Header Extension

The X bit MUST be set to zero by sender and ignored by receiver.

CC: CSRC Count

The CC field MUST be set to zero by the sender and ignored by the receiver.

M: Marker

The M bit MUST be set to zero by sender and ignored by receiver.

PT: Payload Type

A PT value MUST be allocated from the range of dynamic values define by [[RFC3551](#)] for each direction of the VPWS. The same PT value MAY be reused both for direction and between different PLE VPWS.

Sequence Number

The packet sequence number MUST continuously cycle from 0 to 0xFFFF. It is generated and processed in accordance with the rules established in [[RFC3550](#)]. The PLE receiver MUST sequence packets according to the Sequence Number field of the PLE control word and MAY verify correct sequencing using RTP Sequence Number field.

Timestamp

Timestamp values are used in accordance with the rules established in [[RFC3550](#)]. For bit-streams up to 200 Gbps the frequency of the clock used for generating timestamps MUST be 125 MHz based on a the common clock I. For bit-streams above 200 Gbps the frequency MUST be 250 MHz.

SSRC: Synchronization Source

The SSRC field MAY be used for detection of misconnections.

5. PLE Payload Layer

A bit-stream is mapped into a PLE packet with a fixed payload size which MUST be defined during VPWS setup, MUST be the same in both directions of the VPWS and MUST remain unchanged for the lifetime of the VPWS.

All PLE implementations MUST be capable of supporting the default payload size of 1024 bytes.

5.1. Structure Agnostic Payload

The PLE payload is filled with incoming bits of the bit-stream starting from the most significant to the least significant bit without considering any structure of the bit-stream.

For PCS based attachment circuits supporting FEC the NSP function MUST terminate the FEC and pass the PCS encoded signal to the IWF function.

For PCS based attachment circuits supporting virtual lanes (i.e. 100GE) a PLE payload MUST carry information from all virtual lanes in a bit interleaved manner after the NSP function has performed PCS layer de-skew and re-ordering.

A PLE implementation MUST support the structure agnostic payload for all bit-streams except the following:

- *OTN

- *200GBASE-R ethernet

- *400GBASE-R ethernet

5.2. Byte aligned Payload

In case of OTN bit-streams, the NSP function MUST present to the IWF an extended ODUk including a valid frame alignment overhead. The IWF is performing byte-aligned mapping into PLE packets. The egress NSP function will recover the ODUk by searching for the frame alignment overhead.

For byte aligned payloads PLE uses the following order for packetization:

- *The order of the payload bytes corresponds to their order on the attachment circuit.

- *Consecutive bits coming from the attachment circuit fill each payload byte starting from most significant bit to least significant.

All PLE implementations MUST support the transport of OTN bit-streams using the byte aligned payload.

5.3. 10280bit-block aligned Payload

In IEEE 802.3BS the PCS layer for 200GBASE-R and 400GBASE-R is defined with the functions shown in [Figure 6](#).

Reconciliation Sublayer (RS)

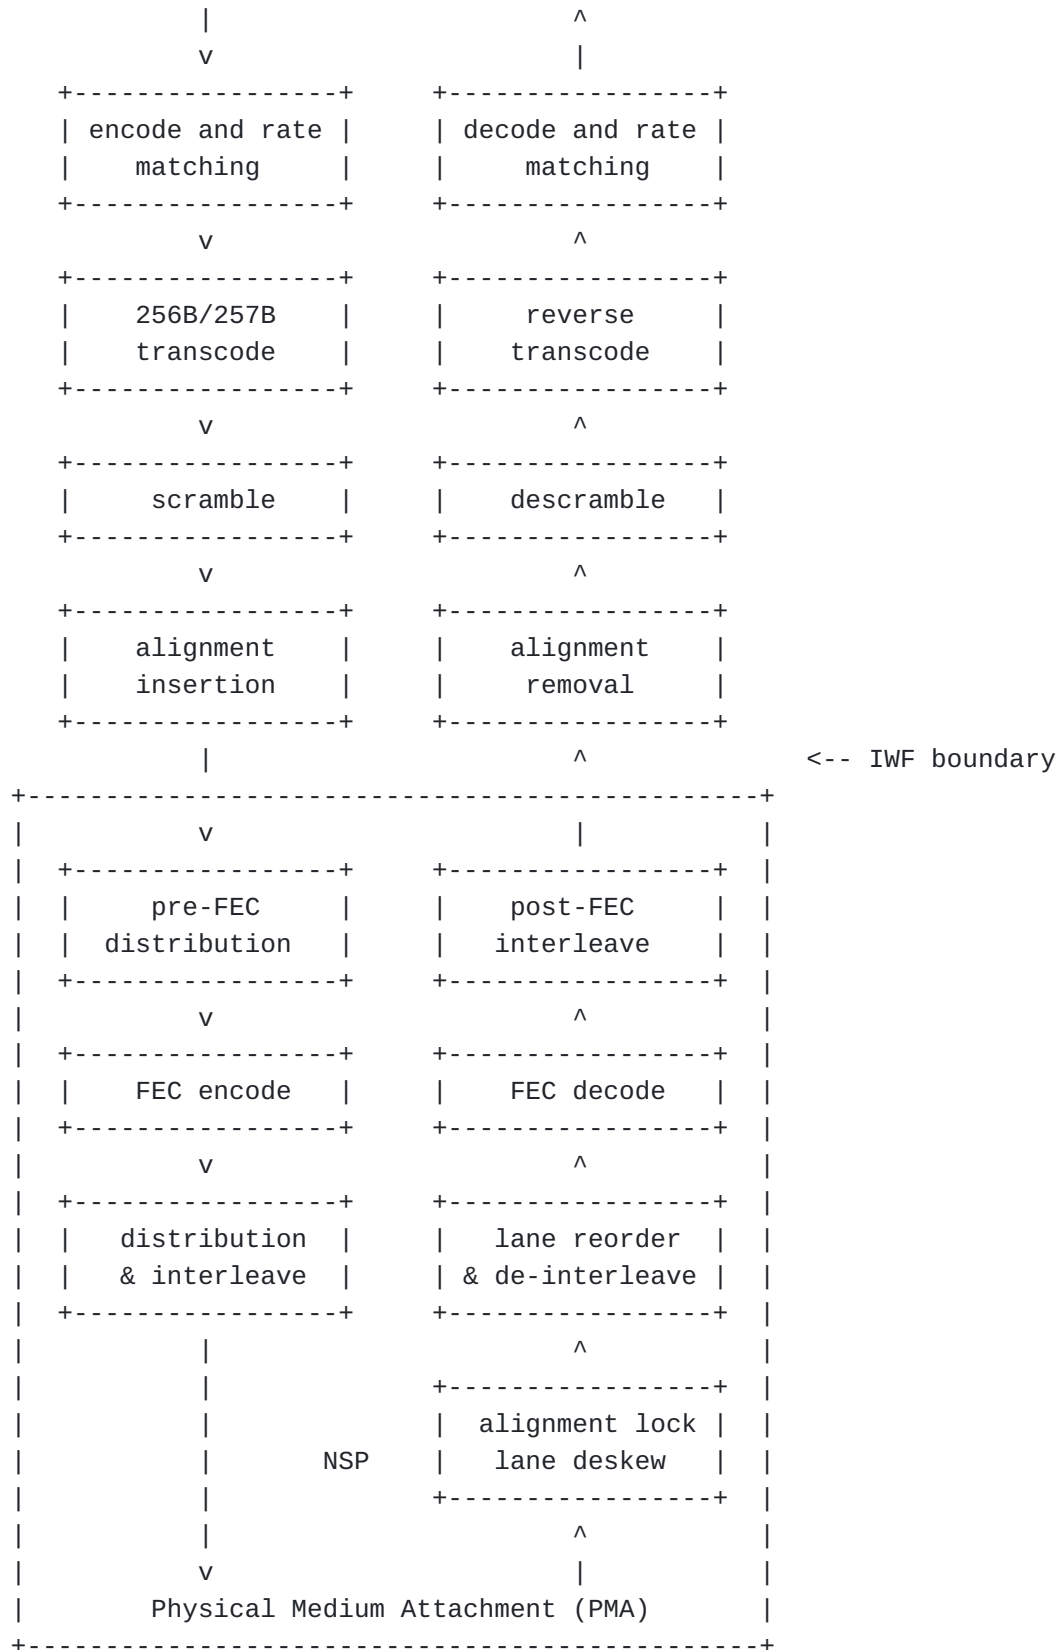


Figure 6: 200GBASE-R and 400GBASE-R Functional Block Diagram

For 200GBASE-R and 400GBASE-R bit-streams, on ingress the NSP function will perform alignment lock and lane de-skew, lane order and de-interleave, FEC decode and post-FEC interleave as shown in [Figure 6](#). After the post-FEC interleave the NSP function will create a stream of 10280 bit blocks (comprising of two 5140 code blocks).

On the egress the IWF sends a stream of 10280 bit blocks to the NSP function and which performs pre-FEC distribution, FEC encode and distribute and interleave functions as shown in [Figure 6](#).

In the 10280 bit block stream, alignment markers exist every 4096, 10280 bit blocks (8192 code blocks) for 400GBASE-R and every 2048, 10280 bit blocks (4096 code blocks) for 200GBASE-R.

On ingress the NSP must indicate to the IWF when a code word carries an alignment marker (or every n-th alignment marker where n is a multiple of 2). The IWF will create a PLE packet with the alignment marker bits at the beginning of the PLE payload. Considering the default PLE payload size of 1024 bytes, the PLE payload will contain the first 8096 bits (1024 bytes) of the 10280 bit block in the first packet. The following PLE packets will contain the remaining bits followed by the next 10280 bits.

The egress NSP will recover the 10280 bit block by searching for the alignment markers at the beginning of PLE packets and recover the 10280 bit block stream.

For the 10280 bit data streams the NSP will use the following order of packetization.

- *The first alignment bit of a 10280 bit block is always mapped to the first bit of a PLE payload

- *The order of the bits corresponds to their order in the attached circuit

- *Consecutive bits from the attached circuit are mapped directly into the PLE packet

With the default payload size of 1024 bytes the alignment markers will be present at the start of every 5140-th PLE packet for 400GBASE-R and every 2570-th PLE packet for 200GBASE-R.

Non-default payload sizes must be chosen so that alignment markers will always be at the start of every N-th packet.

Alignment of the signal may use the alignment marker state machine defined in IEEE802.3BS.

6. PLE Operation

6.1. Common Considerations

A PLE VPWS can be established using manual configuration or leveraging mechanisms of a signalling protocol

Furthermore emulation of bit-stream signals using PLE is only possible when the two attachment circuits of the VPWS are of the same type (OC192, 10GBASE-R, ODU2, etc) and are using the same PLE payload type and payload size. This can be ensured via manual configuration or via a signalling protocol

Extensions to the PWE3 [[RFC4447](#)] and EVPN-VPWS [[RFC8214](#)] control protocols are described in a separate document [[PLESIG](#)].

6.2. PLE IWF Operation

6.2.1. PSN-bound Encapsulation Behavior

After the VPWS is set up, the PSN-bound IWF does perform the following steps:

- *Packetise the data received from the CE is into a fixed size PLE payloads
- *Add PLE control word and RTP header with sequence numbers, flags and timestamps properly set
- *Add the VPWS demultiplexer and PSN headers
- *Transmit the resulting packets over the PSN
- *Set L bit in the PLE control word whenever attachment circuit detects a fault
- *Set R bit in the PLE control word whenever the local CE-bound IWF is in packet loss state

6.2.2. CE-bound Decapsulation Behavior

The CE-bound IWF is responsible for removing the PSN and VPWS demultiplexing headers, PLE control word and RTP header from the received packet stream and play-out of the bit-stream to the local attachment circuit.

A de-jitter buffer MUST be implemented where the PLE packets are stored upon arrival. The size of this buffer SHOULD be locally configurable to allow accommodation of specific PSN packet delay variation expected.

The CE-bound IWF SHOULD use the sequence number in the control word to detect lost and mis-ordered packets. It MAY use the sequence number in the RTP header for the same purposes.

The payload of a lost packet MUST be replaced with equivalent amount of replacement data. The contents of the replacement data MAY be locally configurable. All PLE implementations MUST support generation of "0xAA" as replacement data. The alternating sequence of 0s and 1s of the "0xAA" pattern does ensure clock synchronization is maintained. While playing out the replacement data, the IWF will apply a holdover mechanism to maintain the clock.

Whenever the VPWS is not operationally up, the CE-bound NSP function MUST inject the appropriate native downstream fault indication signal (for example ODUK-AIS or ethernet LF).

Whenever a VPWS comes up, the CE-bound IWF enters the intermediate state, will start receiving PLE packets and will store them in the jitter buffer. The CE-bound NSP function will continue to inject the appropriate native downstream fault indication signal until a pre-configured amount of payloads is stored in the jitter buffer.

After the pre-configured amount of payload is present in the jitter buffer the CE-bound IWF transitions to the normal operation state and the content of the jitter buffer is played out to the CE in accordance with the required clock. In this state the CE-bound IWF MUST perform egress clock recovery.

The recovered clock MUST comply with the jitter and wander requirements applicable to the type of attachment circuit, specified in:

- *[\[G.825\]](#) and [\[G.823\]](#) for SDH

- *[\[GR253\]](#) for SONET

- *[\[G.8261\]](#) for synchronous ethernet

- *[\[G.8251\]](#) for OTN

Whenever the L bit is set in the PLE control word of a received PLE packet the CE-bound NSP function SHOULD inject the appropriate native downstream fault indication signal instead of playing out the payload.

If the CE-bound IWF detects loss of consecutive packets for a pre-configured amount of time (default is 1 millisecond), it enters packet loss (PLOS) state and a corresponding defect is declared.

If the CE-bound IWF detects a packet loss ratio (PLR) above a configurable signal-degrade (SD) threshold for a configurable amount of consecutive 1-second intervals, it enters the degradation (DEG) state and a corresponding defect is declared. Possible values for the SD-PLR threshold are between 1..100% with the default being 15%. Possible values for consecutive intervals are 2..10 with the default 7.

While either a PLOS or DEG defect is declared the CE-bound NSP function SHOULD inject the appropriate native downstream fault indication signal. Also the PSN-bound IWF SHOULD set the R bit in the PLE control word of every packet transmitted.

The CE-bound IWF does change from the PLOS to normal state after the pre-configured amount of payload has been received similarly to the transition from intermediate to normal state.

Whenever the R bit is set in the PLE control word of a received PLE packet the PLE performance monitoring statistics SHOULD get updated.

6.3. PLE Performance Monitoring

PLE SHOULD provide the following functions to monitor the network performance to be inline with expectations of transport network operators.

The near-end performance monitors defined for PLE are as follows:

ES-PLE : PLE Errored Seconds

SES-PLE : PLE Severely Errored Seconds

UAS-PLE : PLE Unavailable Seconds

Each second with at least one packet lost or a PLOS/DEG defect SHALL be counted as ES-PLE. Each second with a PLR greater than 15% or a PLOS/DEG defect SHALL be counted as SES-PLE.

UAS-PLE SHALL be counted after configurable number of consecutive SES-PLE have been observed, and no longer counted after a configurable number of consecutive seconds without SES-PLE have been observed. Default value for each is 10 seconds.

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

A PLE far-end performance monitor is providing insight into the CE-bound IWF at the far end of the PSN. The statistics are based on the PLE-RDI indication carried in the PLE control word via the R bit.

The PLE VPWS performance monitors are derived from the definitions in accordance with [G.826]

6.4. QoS and Congestion Control

The PSN carrying PLE VPWS may be subject to congestion, but PLE VPWS representing constant bit-rate (CBR) flows cannot respond to congestion in a TCP-friendly manner as described in [RFC2913].

Hence the PSN providing connectivity for the PLE VPWS between PE devices MUST be Diffserv [RFC2475] enabled and MUST provide a per domain behavior [RFC3086] that guarantees low jitter and low loss.

To achieve the desired per domain behavior PLE VPWS SHOULD be carried over traffic-engineering paths through the PSN with bandwidth reservation and admission control applied.

7. Security Considerations

As PLE is leveraging VPWS as transport mechanism the security considerations described in [RFC7432] and [RFC3985] are applicable.

8. IANA Considerations

Applicable signalling extensions are out of the scope of this document.

PLE does not introduce additional requirements from IANA.

9. Acknowledgements

The authors would like to thank Andreas Burk for reviewing this document and providing useful comments and suggestions.

10. References

10.1. Normative References

[G.823] International Telecommunication Union (ITU), "G.823: The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy", <<https://www.itu.int/rec/T-REC-G.823>>.

[G.825] International Telecommunication Union (ITU), "G.825: The control of jitter and wander within digital networks

which are based on the synchronous digital hierarchy (SDH)", <<https://www.itu.int/rec/T-REC-G.825>>.

- [G.8251] International Telecommunication Union (ITU), "G.8251: The control of jitter and wander within the optical transport network (OTN)", <<https://www.itu.int/rec/T-REC-G.8251>>.
- [G.8261] International Telecommunication Union (ITU), "G.8261: Timing and synchronization aspects in packet networks", <<https://www.itu.int/rec/T-REC-G.8251>>.
- [PLESIG] IETF, "Private Line Emulation VPWS Signalling", <<https://tools.ietf.org/html/draft-schmutzer-bess-ple-vpws-signalling>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <<https://www.rfc-editor.org/info/rfc2475>>.
- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC 3086, DOI 10.17487/RFC3086, April 2001, <<https://www.rfc-editor.org/info/rfc3086>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<https://www.rfc-editor.org/info/rfc3550>>.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, DOI 10.17487/RFC3551, July 2003, <<https://www.rfc-editor.org/info/rfc3551>>.
- [RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <<https://www.rfc-editor.org/info/rfc3985>>.
- [RFC4197] Riegel, M., Ed., "Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks", RFC 4197, DOI 10.17487/RFC4197, October 2005, <<https://www.rfc-editor.org/info/rfc4197>>.

[RFC4385]

Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.

[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, DOI 10.17487/RFC4447, April 2006, <<https://www.rfc-editor.org/info/rfc4447>>.

[RFC4664]

Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <<https://www.rfc-editor.org/info/rfc4664>>.

[RFC7432]

Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.

[RFC8174]

Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8214]

Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <<https://www.rfc-editor.org/info/rfc8214>>.

10.2. Informative References

[G.707]

ITU-T, "Network node interface for the synchronous digital hierarchy (SDH)", <<https://www.itu.int/rec/T-REC-G.707>>.

[G.709]

International Telecommunication Union (ITU), "G.709: Interfaces for the optical transport network", <<https://www.itu.int/rec/T-REC-G.709>>.

[G.826]

ITU-T, "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections", <<https://www.itu.int/rec/T-REC-G.826>>.

[GR253]

Telcordia, "SONET Transport Systems : Common Generic Criteria", <<https://telecom-info.telcordia.com>>.

[RFC2913]

Klyne, G., "MIME Content Types in Media Feature Expressions", RFC 2913, DOI 10.17487/RFC2913, September 2000, <<https://www.rfc-editor.org/info/rfc2913>>.

[RFC3209]

Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.

[RFC4553]

Vainshtein, A., Ed. and YJ. Stein, Ed., "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, DOI 10.17487/RFC4553, June 2006, <<https://www.rfc-editor.org/info/rfc4553>>.

[RFC5036]

Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.

[RFC5086]

Vainshtein, A., Ed., Sasson, I., Metz, E., Frost, T., and P. Pate, "Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)", RFC 5086, DOI 10.17487/RFC5086, December 2007, <<https://www.rfc-editor.org/info/rfc5086>>.

[RFC8402]

Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[SRPOLICY]

IETF, "Segment Routing Policy Architecture", <<https://tools.ietf.org/html/draft-ietf-spring-segment-routing-policy>>.

[SRV6NETPROG]

IETF, "SRv6 Network Programming", <<https://tools.ietf.org/html/draft-ietf-spring-srv6-network-programming>>.

Authors' Addresses

Steven Gringeri
Verizon

Email: steven.gringeri@verizon.com

Jeremy Whittaker
Verizon

Email: jeremy.whittaker@verizon.com

Nicolai Leymann
Deutsche Telekom

Email: N.Leymann@telekom.de

Christian Schmutzer (editor)
Cisco Systems, Inc.

Email: cschmutz@cisco.com

Luca Della Chiesa
Cisco Systems, Inc.

Email: ldellach@cisco.com

Nagendra Kumar Nainar (editor)
Cisco Systems, Inc.

Email: naikumar@cisco.com

Carlos Pignataro
Cisco Systems, Inc.

Email: cpignata@cisco.com

Gerald Smallegange
Ciena Corporation

Email: gsmalleg@ciena.com

Chris Brown
Ciena Corporation

Email: cbrown@ciena.com

Faisal Dada
Xilinx

Email: faisald@xilinx.com