

Internet Engineering Task Force (IETF)
Request for Comments: 9801
Category: Standards Track
ISSN: 2070-1721

S. Gringeri
J. Whittaker
Verizon
N. Leymann
Deutsche Telekom
C. Schmutzer, Ed.
Cisco Systems, Inc.
C. Brown
Ciena Corporation
July 2025

Private Line Emulation over Packet Switched Networks

Abstract

This document expands the applicability of Virtual Private Wire Service (VPWS) bit-stream payloads beyond Time Division Multiplexing (TDM) signals and provides pseudowire transport with complete signal transparency over Packet Switched Networks (PSNs).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9801>.

Copyright Notice

Copyright (c) 2025 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 Motivation
2. Requirements Notation
3. Terminology and Reference Models
 - 3.1. Abbreviations
 - 3.2. Reference Models
4. Emulated Services
 - 4.1. Generic PLE Service
 - 4.2. Ethernet Services
 - 4.2.1. 1000BASE-X
 - 4.2.2. 10GBASE-R and 25GBASE-R

- 4.2.3. 40GBASE-R, 50GBASE-R, and 100GBASE-R
- 4.2.4. 200GBASE-R and 400GBASE-R
- 4.2.5. Energy Efficient Ethernet (EEE)
- 4.3. SONET/SDH Services
- 4.4. Fibre Channel Services
 - 4.4.1. 1GFC, 2GFC, 4GFC, and 8GFC
 - 4.4.2. 16GFC
 - 4.4.3. 32GFC and 4-Lane 128GFC
 - 4.4.4. 64GFC
- 4.5. OTN Services
- 5. PLE Encapsulation Layer
 - 5.1. PSN and VPWS Demultiplexing Headers
 - 5.1.1. New SRv6 Behaviors
 - 5.2. PLE Header
 - 5.2.1. PLE Control Word
 - 5.2.2. RTP Header
- 6. PLE Payload Layer
 - 6.1. Basic Payload
 - 6.2. Byte-Aligned Payload
- 7. PLE Operation
 - 7.1. Common Considerations
 - 7.2. PLE IWF Operation
 - 7.2.1. PSN-Bound Encapsulation Behavior
 - 7.2.2. CE-Bound Decapsulation Behavior
 - 7.3. PLE Performance Monitoring
 - 7.4. PLE Fault Management
- 8. QoS and Congestion Control
- 9. Security Considerations
- 10. IANA Considerations
 - 10.1. Bit-Stream Next Header Type
 - 10.2. SRv6 Endpoint Behaviors
- 11. References
 - 11.1. Normative References
 - 11.2. Informative References
- Acknowledgements
- Contributors
- Authors' Addresses

1. Introduction and Motivation

This document describes a method called Private Line Emulation (PLE) for encapsulating not only Time Division Multiplexing (TDM) signals as bit-stream Virtual Private Wire Service (VPWS) over Packet Switched Networks (PSN). In this regard, it complements methods described in [RFC4553].

This emulation suits applications, where carrying Protocol Data Units (PDUs) as defined in [RFC4906] or [RFC4448] is not enough, physical layer signal transparency is required and data or framing structure interpretation of the Provider Edge (PE) would be counterproductive.

One example of such case is two Ethernet-connected Customer Edge (CE) devices and the need for Synchronous Ethernet operation (see [G.8261]) between them without the intermediate PE devices interfering or addressing concerns about Ethernet control protocol transparency for PDU-based carrier Ethernet services, beyond the behavior definitions of MEF Forum (MEF) specifications.

Another example would be a Storage Area Networking (SAN) extension between two data centers. Operating at a bit-stream level allows for a connection between Fibre Channel switches without interfering with any of the Fibre Channel protocol mechanisms defined by [T11].

Also, SONET/SDH (Synchronous Optical Network (SONET) / Synchronous Digital Hierarchy (SDH)) add/drop multiplexers or cross-connects can be interconnected without interfering with the multiplexing

structures and networks mechanisms. This is a key distinction to Circuit Emulation over Packet (CEP) defined in [RFC4842] where multiplexing and demultiplexing is desired in order to operate per SONET Synchronous Payload Envelope (SPE) and Virtual Tributary (VT) or SDH Virtual Container (VC). In other words, PLE provides an independent layer network underneath the SONET/SDH layer network, whereas CEP operates at the same level and peer with the SONET/SDH layer network.

The mechanisms described in this document follow principles similar to Structure-Agnostic TDM over Packet (SAToP) (defined in [RFC4553]). The applicability is expanded beyond the narrow set of Plesiochronous Digital Hierarchy (PDH) interfaces (T1, E1, T3, and E3) to allow the transport of signals from many different technologies such as Ethernet, Fibre Channel, SONET/SDH ([GR253] / [G.707]), and Optical Transport Network (OTN) [G.709] at gigabit speeds. The signals are treated as bit-stream payload, which was defined in the Pseudo Wire Emulation Edge-to-Edge (PWE3) architecture in Sections 3.3.3 and 3.3.4 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 Models

3.1. Abbreviations

ACH: Associated Channel Header [RFC7212]

AIS: Alarm Indication Signal

AIS-L: Line AIS

MS-AIS: Multiplex Section AIS

BITS: Building Integrated Timing Supply [ATIS-0900105.09.2013]

CBR: Constant Bit Rate

CE: Customer Edge

CEP: Circuit Emulation over Packet [RFC4842]

CSID: Compressed SID

CSRC: Contributing Source [RFC3550]

DEG: Degradation

ES: Errored Second

FEC: Forward Error Correction

ICMP: Internet Control Message Protocol [RFC4443]

IEEE: Institute of Electrical and Electronics Engineers

INCITS: International Committee for Information Technology Standards

IWF: Interworking Function

LDP: Label Distribution Protocol [RFC5036], [RFC8077]

LF: Local Fault

LOF: Loss Of Frame

LOM: Loss Of Multiframe

LOS: Loss Of Signal

LPI: Low Power Idle

LSP: Label Switched Path

MEF: MEF Forum

MPLS: Multiprotocol Label Switching [RFC3031]

NOS: Not Operational

NSP: Native Service Processing [RFC3985]

ODUk: Optical Data Unit k

OOF: Out Of Frame

OTN: Optical Transport Network

OTUk: Optical Transport Unit k

PCS: Physical Coding Sublayer

PDV: Packet Delay Variation

PE: Provider Edge

PLE: Private Line Emulation

PLOS: Packet Loss Of Signal

PLR: Packet Loss Rate

PMA: Physical Medium Attachment

PMD: Physical Medium Dependent

PSN: Packet Switched Network

PTP: Precision Time Protocol

PW: Pseudowire [RFC4664]

PWE3: Pseudo Wire Emulation Edge-to-Edge [RFC3985]

RDI: Remote Defect Indication

RSVP-TE: Resource Reservation Protocol Traffic Engineering [RFC4875]

RTCP: RTP Control Protocol [RFC3550]

RTP: Real-time Transport Protocol [RFC3550]

SD: Signal Degrade

SES: Severely Errored Seconds

SDH: Synchronous Digital Hierarchy

SID: Segment Identifier [RFC8402]

SR: Segment Routing [RFC8402]

SRH: Segment Routing Header [RFC8754]

S RTP: Secure Real-time Transport Protocol [RFC3711]

SRv6: Segment Routing over IPv6 [RFC8986]

SSRC: Synchronization Source [RFC3550]

SONET: Synchronous Optical Network

TCP: Transmission Control Protocol [RFC9293]

TDM: Time Division Multiplexing

TTS: Transmitter Training Signal

UAS: Unavailable Seconds

VPWS: Virtual Private Wire Service [RFC3985]

Note: 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 reference model for PLE is illustrated in Figure 1 and is inline with the reference model defined in Section 4.1 of [RFC3985]. PLE relies on PWE3 preprocessing, in particular the concept of an NSP function defined in Section 4.2.2 of [RFC3985].

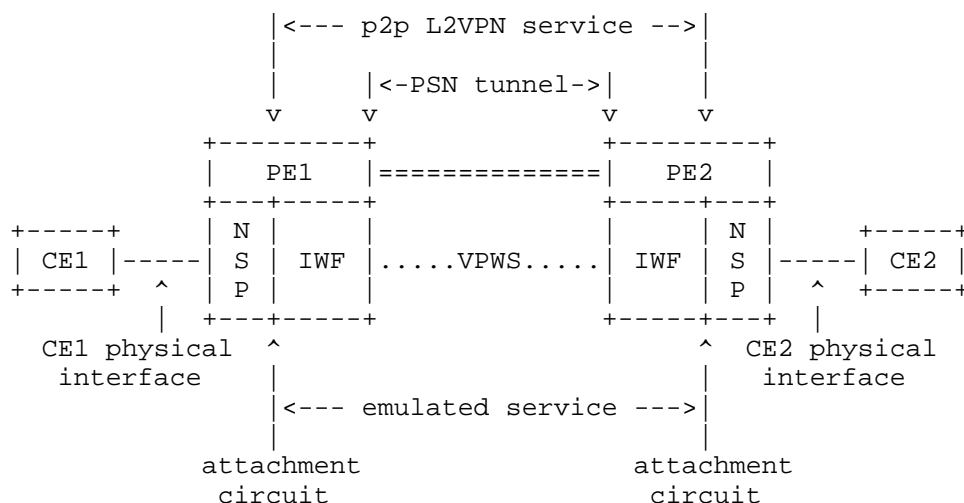


Figure 1: PLE Reference Model

PLE embraces the minimum intervention principle outlined in Section 3.3.5 of [RFC3985], which means the data is flowing through the PLE encapsulation layer as received without modifications.

For some service types, the NSP function is responsible for performing operations on the data received from the CE. Examples are terminating FEC, terminating the OTUk layer for OTN, or dealing with multi-lane processing. After the NSP, the IWF is generating the

payload of the VPWS, which is carried via a PSN tunnel.

To allow the clock of the transported signal to be carried across the PLE domain in a transparent way, the relative network synchronization reference model and deployment scenario outlined in Section 4.3.2 of [RFC4197] are applicable and are shown in Figure 2.

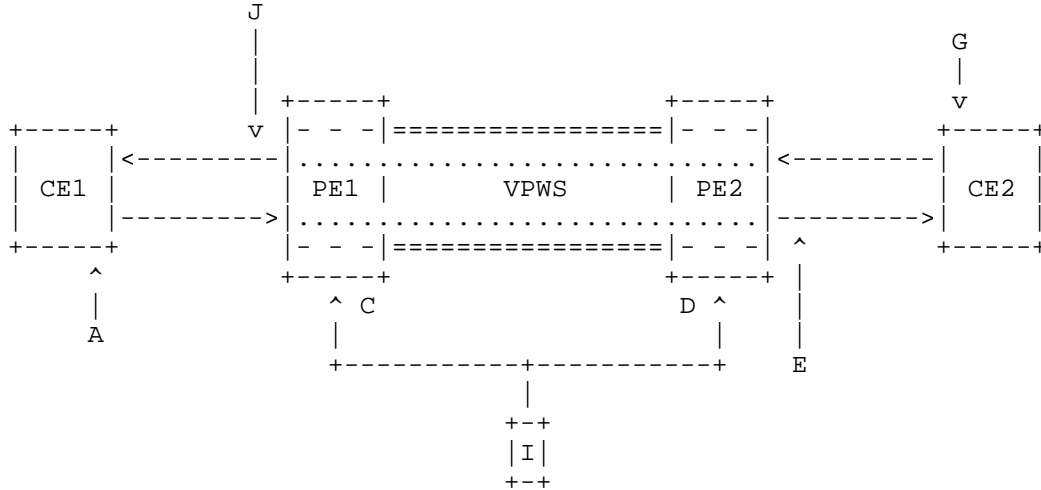


Figure 2: Relative Network Scenario Timing

The local oscillators C of PE1 and D of PE2 are locked to a common clock I.

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

For the reverse direction, PE1 generates the attachment circuit clock J and the clock difference between G and D (locked to I) is transferred from PE2 to PE1.

The method used to lock clocks C and D to the common clock I is out of scope of this document; however, there are already several well-established concepts for achieving clock synchronization (commonly also referred to as "frequency synchronization") available.

While using external timing inputs (aka BITS [ATIS-0900105.09.2013]) or synchronous Ethernet (as defined in [G.8261]), the characteristics and limits defined in [G.8262] have to be considered.

While relying on PTP (as defined in [G.8265.1]), the network limits defined in [G.8261.1] have to be considered.

4. Emulated Services

This specification describes the emulation of services from a wide range of technologies, such as TDM, Ethernet, Fibre Channel, or OTN, as bit-streams or structured bit-streams, as defined in Sections 3.3.3 and 3.3.4 of [RFC3985].

4.1. Generic PLE Service

The generic PLE service is an example of the bit-stream defined in Section 3.3.3 of [RFC3985].

Under the assumption that the CE-bound IWF is not responsible for any service-specific operation, a bit-stream of any rate can be carried

using the generic PLE payload.

There is no NSP function present for this service.

4.2. Ethernet Services

Ethernet services are special cases of the structured bit-stream defined in Section 3.3.4 of [RFC3985].

The IEEE has defined several layers for Ethernet in [IEEE802.3]. Emulation is operating at the physical (PHY) layer, more precisely at the PCS.

Over time, many different Ethernet interface types have been specified in [IEEE802.3] with a varying set of characteristics, such as optional versus mandatory FEC and single-lane versus multi-lane transmission.

Ethernet interface types with backplane PMD variants and Ethernet interface types mandating auto-negotiation (except 1000Base-X) are out of scope for this document.

All Ethernet services are leveraging the basic PLE payload and interface-specific mechanisms are confined to the respective service specific NSP functions.

4.2.1. 1000BASE-X

The PCS layer of 1000BASE-X (defined in Section 36 of [IEEE802.3]) is based on 8B/10B code.

The PSN-bound NSP function does not modify the received data and is transparent to auto-negotiation; however, it is responsible for detecting attachment circuit faults specific to 1000BASE-X such as LOS and sync loss.

When the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set, the CE-bound NSP function MAY disable its transmitter as no appropriate maintenance signal was defined for 1000BASE-X by the IEEE.

4.2.2. 10GBASE-R and 25GBASE-R

The PCS layers of 10GBASE-R (defined in Section 49 and 25GBASE-R defined in Section 107 of [IEEE802.3]) are based on a 64B/66B code.

Sections 74 and 108 of [IEEE802.3] define an optional FEC layer; if present, the PSN-bound NSP function MUST terminate the FEC and the CE-bound NSP function MUST generate the FEC.

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to 10GBASE-R and 25GBASE-R such as LOS and sync loss.

The PSN-bound IWF maps the scrambled 64B/66B code stream into the basic PLE payload.

The CE-bound NSP function MUST perform:

- * PCS code sync (Section 49.2.9 of [IEEE802.3]) and
- * descrambling (Section 49.2.10 of [IEEE802.3])

in order to properly:

- * transform invalid 66B code blocks into proper error control

characters /E/ (Section 49.2.4.11 of [IEEE802.3]) and

- * insert LF ordered sets (Section 46.3.4 of [IEEE802.3]) when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid 66B code blocks typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets or the far-end PSN-bound NSP function setting sync
| headers to 11 due to uncorrectable FEC errors.

Before sending the bit-stream to the CE, the CE-bound NSP function MUST also scramble the 64B/66B code stream (Section 49.2.6 [IEEE802.3]).

4.2.3. 40GBASE-R, 50GBASE-R, and 100GBASE-R

The PCS layers of 40GBASE-R and 100GBASE-R (defined in Section 82 of [IEEE802.3]) and of 50GBASE-R (defined in Section 133 of [IEEE802.3]) are based on a 64B/66B code transmitted over multiple lanes.

Sections 74 and 91 of [IEEE802.3] define an optional FEC layer; if present, the PSN-bound NSP function MUST terminate the FEC and the CE-bound NSP function MUST generate the FEC.

To gain access to the scrambled 64B/66B code stream, the PSN-bound NSP further MUST perform:

- * block synchronization (Section 82.2.12 of [IEEE802.3]),
- * PCS lane de-skew (Section 82.2.13 of [IEEE802.3]), and
- * PCS lane reordering (Section 82.2.14 of [IEEE802.3]).

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to 40GBASE-R, 50GBASE-R, and 100GBASE-R such as LOS and loss of alignment.

The PSN-bound IWF maps the serialized and scrambled 64B/66B code stream including the alignment markers into the basic PLE payload.

The CE-bound NSP function MUST perform:

- * PCS code sync (Section 82.2.12 of [IEEE802.3]),
- * alignment-marker removal (Section 82.2.15 of [IEEE802.3]), and
- * descrambling (Section 49.2.10 of [IEEE802.3])

in order to properly:

- * transform invalid 66B code blocks into proper error control characters /E/ (Section 82.2.3.10 of [IEEE802.3]) and
- * insert LF ordered sets (Section 81.3.4 of [IEEE802.3]) when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid 66B code blocks typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets or the far-end PSN-bound NSP function not setting sync
| headers to 11 due to uncorrectable FEC errors.

When sending the bit-stream to the CE, the CE-bound NSP function MUST also perform:

- * scrambling of the 64B/66B code (Section 49.2.6 of [IEEE802.3]),
- * block distribution (Section 82.2.6 of [IEEE802.3]), and
- * alignment-marker insertion (Sections 82.2.7 and 133.2.2 of [IEEE802.3]).

4.2.4. 200GBASE-R and 400GBASE-R

The PCS layers of 200GBASE-R and 400GBASE-R (defined in Section 119 of [IEEE802.3]) are based on a 64B/66B code transcoded to a 256B/257B code to reduce the overhead and make room for a mandatory FEC.

To gain access to the 64B/66B code stream, the PSN-bound NSP further MUST perform:

- * alignment lock and de-skew (Section 119.2.5.1 of [IEEE802.3]),
- * PCS lane reordering and de-interleaving (Section 119.2.5.2 of [IEEE802.3]),
- * FEC decoding (Section 119.2.5.3 of [IEEE802.3]),
- * post-FEC interleaving (Section 119.2.5.4 of [IEEE802.3]),
- * alignment-marker removal (Section 119.2.5.5 of [IEEE802.3]),
- * descrambling (Section 119.2.5.6 of [IEEE802.3]), and
- * reverse transcoding from 256B/257B to 64B/66B (Section 119.2.5.7 of [IEEE802.3]).

Further, the PSN-bound NSP MUST perform rate compensation and scrambling (Section 49.2.6 of [IEEE802.3]) before the PSN-bound IWF maps the same into the basic PLE payload.

Rate compensation is applied so that the rate of the 66B encoded bit-stream carried by PLE is 528/544 times the nominal bitrate of the 200GBASE-R or 400GBASE-R at the PMA service interface. X number of 66-byte-long rate compensation blocks are inserted every X*20479 number of 66B client blocks. For 200GBASE-R, the value of X is 16; for 400GBASE-R, the value of X is 32. Rate compensation blocks are special 66B control characters of type 0x00 that can easily be searched for by the CE-bound IWF in order to remove them.

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to 200GBASE-R and 400GBASE-R such as LOS and loss of alignment.

The CE-bound NSP function MUST perform:

- * PCS code sync (Section 49.2.13 of [IEEE802.3]),
- * descrambling (Section 49.2.10 of [IEEE802.3]), and
- * rate compensation block removal

in order to properly:

- * transform invalid 66B code blocks into proper error control characters /E/ (Section 119.2.3.9 of [IEEE802.3]) and
- * insert LF ordered sets (Section 81.3.4 of [IEEE802.3]) when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

Note: Invalid 66B code blocks typically are a consequence of the CE-bound IWF inserting replacement data in case of lost PLE packets or the far-end PSN-bound NSP function not setting sync headers to 11 due to uncorrectable FEC errors.

When sending the bit-stream to the CE, the CE-bound NSP function MUST also perform:

- * transcoding from 64B/66B to 256B/257B (Section 119.2.4.2 of [IEEE802.3]),
- * scrambling (Section 119.2.4.3 of [IEEE802.3]),
- * alignment-marker insertion (Section 119.2.4.4 of [IEEE802.3]),
- * pre-FEC distribution (Section 119.2.4.5 of [IEEE802.3]),
- * FEC encoding (Section 119.2.4.6 of [IEEE802.3]), and
- * PCS lane distribution (Section 119.2.4.8 of [IEEE802.3]).

4.2.5. Energy Efficient Ethernet (EEE)

Section 78 of [IEEE802.3] defines the optional LPI capability for Ethernet. Two modes are defined:

- * deep sleep
- * fast wake

Deep sleep mode is not compatible with PLE due to the CE ceasing transmission. Hence, there is no support for LPI for 10GBASE-R services across PLE.

In fast wake mode, the CE transmits /LI/ control code blocks instead of /I/ control code blocks and, therefore, PLE is agnostic to it. For 25GBASE-R and higher services across PLE, LPI is supported as only fast wake mode is applicable.

4.3. SONET/SDH Services

SONET/SDH services are special cases of the structured bit-stream defined in Section 3.3.4 of [RFC3985].

SDH interfaces are defined in [G.707]; SONET interfaces are defined in [GR253].

The PSN-bound NSP function does not modify the received data but is responsible for detecting attachment circuit faults specific to SONET/SDH such as LOS, LOF, and OOF.

Data received by the PSN-bound IWF is mapped into the basic PLE payload without any awareness of SONET/SDH frames.

When the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set, the CE-bound NSP function is responsible for generating the:

- * MS-AIS maintenance signal (defined in Section 6.2.4.1.1 of [G.707]) for SDH services and
- * AIS-L maintenance signal (defined in Section 6.2.1.2 of [GR253]) for SONET services

at client-frame boundaries.

4.4. Fibre Channel Services

Fibre Channel services are special cases of the structured bit-stream defined in Section 3.3.4 of [RFC3985].

The T11 technical committee of INCITS has defined several layers for Fibre Channel. PLE operates at the FC-1 layer that leverages mechanisms defined by [IEEE802.3].

Over time, many different Fibre Channel interface types have been specified with a varying set of characteristics such as optional versus mandatory FEC and single-lane versus multi-lane transmission.

Speed negotiation is not supported by PLE.

All Fibre Channel services leverage the basic PLE payload, and interface-specific mechanisms are confined to the respective service-specific NSP functions.

4.4.1. 1GFC, 2GFC, 4GFC, and 8GFC

[FC-PI-2] specifies 1GFC and 2GFC. [FC-PI-5] and [FC-PI-5aml] define 4GFC and 8GFC.

The PSN-bound NSP function is responsible for detecting attachment circuit faults specific to the Fibre Channel such as LOS and sync loss.

The PSN-bound IWF maps the received 8B/10B code stream as is directly into the basic PLE payload.

The CE-bound NSP function MUST perform transmission word sync in order to properly:

- * replace invalid transmission words with the special character K30.7 and
- * insert NOS ordered sets when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid transmission words typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets.

[FC-PI-5aml] defines the use of scrambling for 8GFC; in this case, the CE-bound NSP MUST also perform descrambling before replacing invalid transmission words or inserting NOS ordered sets. Before sending the bit-stream to the CE, the CE-bound NSP function MUST scramble the 8B/10B code stream.

4.4.2. 16GFC

[FC-PI-5] and [FC-PI-5aml] specify 16GFC and define an optional FEC layer.

If FEC is present, it must be indicated via TTS when the attachment circuit is brought up. Further, the PSN-bound NSP function MUST terminate the FEC and the CE-bound NSP function must generate the FEC.

The PSN-bound NSP function is responsible for detecting attachment circuit faults specific to the Fibre Channel such as LOS and sync loss.

The PSN-bound IWF maps the received scrambled 64B/66B code stream as is into the basic PLE payload.

The CE-bound NSP function MUST perform:

- * transmission word sync (Section 49.2.13 of [IEEE802.3]) and
- * descrambling (Section 49.2.10 of [IEEE802.3])

in order to properly:

- * replace invalid transmission words with the error transmission word 1Eh and
- * insert NOS ordered sets when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid transmission words typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets or the far-end PSN-bound NSP function not setting sync
| headers to 11 due to uncorrectable FEC errors.

Before sending the bit-stream to the CE, the CE-bound NSP function MUST also scramble the 64B/66B code stream (Section 49.2.6 of [IEEE802.3]).

4.4.3. 32GFC and 4-Lane 128GFC

[FC-PI-6] specifies 32GFC and [FC-PI-6P] specifies 4-lane 128GFC, both with FEC layer and TTS support being mandatory.

To gain access to the 64B/66B code stream the PSN-bound NSP further MUST perform:

- * descrambling (Section 49.2.10 of [IEEE802.3]),
- * FEC decoding (Section 91.5.3.3 of [IEEE802.3]), and
- * reverse transcoding from 256B/257B to 64B/66B (Section 119.2.5.7 of [IEEE802.3]).

Further, the PSN-bound NSP MUST perform scrambling (Section 49.2.6 of [IEEE802.3]) before the PSN-bound IWF maps the same into the basic PLE payload.

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to the Fibre Channel such as LOS and sync loss.

The CE-bound NSP function MUST perform:

- * transmission word sync (Section 119.2.6.3 of [IEEE802.3]) and
- * descrambling (Section 49.2.10 of [IEEE802.3])

in order to properly:

- * replace invalid transmission words with the error transmission word 1Eh and
- * insert NOS ordered sets when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid transmission words typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets or the far-end PSN-bound NSP function not setting sync
| headers to 11 due to uncorrectable FEC errors.

When sending the bit-stream to the CE, the CE-bound NSP function MUST also perform:

- * transcoding from 64B/66B to 256B/257B (Section 119.2.4.2 of [IEEE802.3]),
- * FEC encoding (Section 91.5.2.7 of [IEEE802.3]), and
- * scrambling (Section 49.2.6 of [IEEE802.3]).

4.4.4. 64GFC

[FC-PI-7] specifies 64GFC with a mandatory FEC layer.

To gain access to the 64B/66B code stream, the PSN-bound NSP further MUST perform:

- * alignment lock (Section 134.5.4 of [IEEE802.3] modified to single FEC lane operation),
- * FEC decoding (Section 134.5.3.3 of [IEEE802.3]),
- * alignment-marker removal (Section 134.5.3.4 of [IEEE802.3]), and
- * reverse transcoding from 256B/257B to 64B/66B (Section 91.5.3.5 of [IEEE802.3]).

Further, the PSN-bound NSP MUST perform scrambling (Section 49.2.6 of [IEEE802.3]) before the PSN-bound IWF maps the same into the basic PLE payload.

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to the Fibre Channel such as LOS and sync loss.

The CE-bound NSP function MUST perform:

- * transmission word sync (Section 49.2.13 of [IEEE802.3]) and
- * descrambling (Section 49.2.10 of [IEEE802.3])

in order to properly:

- * replace invalid transmission words with the error transmission word 1Eh and
- * insert NOS ordered sets when the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set.

| Note: Invalid transmission words typically are a consequence of
| the CE-bound IWF inserting replacement data in case of lost PLE
| packets or the far-end PSN-bound NSP function not setting sync
| headers to 11 due to uncorrectable FEC errors.

When sending the bit-stream to the CE, the CE-bound NSP function MUST also perform:

- * transcoding from 64B/66B to 256B/257B (Section 91.5.2.5 of [IEEE802.3]),
- * alignment-marker insertion (Section 134.5.2.6 of [IEEE802.3]), and
- * FEC encoding (Section 134.5.2.7 of [IEEE802.3]).

4.5. OTN Services

OTN services are special cases of the structured bit-stream defined in Section 3.3.4 of [RFC3985].

OTN interfaces are defined in [G.709].

The PSN-bound NSP function MUST terminate the FEC and replace the OTUk overhead in row 1, columns 8-14 with an all-zeros pattern; this results in an extended ODUk frame as illustrated in Figure 3. The frame alignment overhead (FA OH) in row 1, columns 1-7 is kept as it is.




Figure 3: Extended ODUk Frame

The PSN-bound NSP function is also responsible for detecting attachment circuit faults specific to OTUk such as LOS, LOF, LOM, and AIS.

The PSN-bound IWF maps the extended ODUk frame into the byte-aligned PLE payload.

The CE-bound NSP function will recover the ODUk by searching for the frame alignment overhead in the extended ODUk received from the CE-bound IWF and generating the FEC.

When the CE-bound IWF is in PLOS state or when PLE packets are received with the L bit set, the CE-bound NSP function is responsible for generating the ODUk-AIS maintenance signal defined in Section 16.5.1 of [G.709] at client-frame boundaries.

5. PLE Encapsulation Layer

The basic packet format used by PLE is shown in Figure 4.

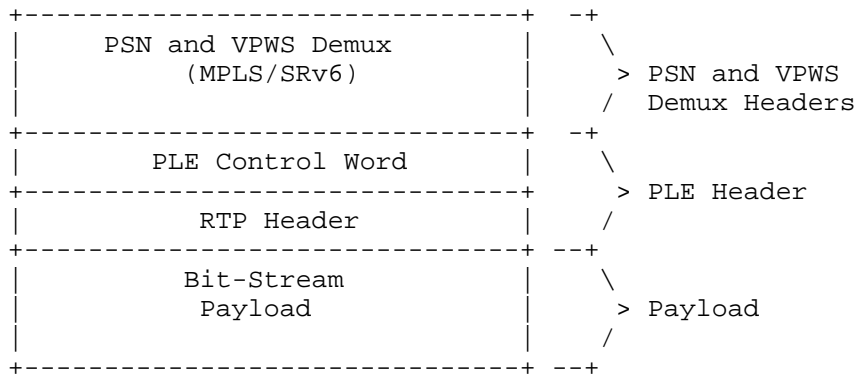


Figure 4: PLE Encapsulation Layer

5.1. PSN and VPWS Demultiplexing Headers

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

The total size of a PLE packet for a specific PW MUST NOT exceed the path MTU between the pair of PEs terminating this PW.

When an 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 LSP established using LDP (see [RFC5036]) or Segment Routing (SR) (see [RFC8402]); or it can be a traffic-engineered LSP established using RSVP-TE (see [RFC3209]) or SR policies (see [RFC9256]).

When an SRv6 PSN layer is used, an SRv6 service SID (as defined in [RFC8402]) provides the demultiplexing mechanism and definitions of Section 6 of [RFC9252] apply. Both SRv6 service SIDs with the full IPv6 address format defined in [RFC8986] and CSIDs with the format defined in [RFC9800] can be used.

5.1.1. New SRv6 Behaviors

Two new encapsulation behaviors, H.Encaps.L1 and H.Encaps.L1.Red, are defined in this document. The behavior procedures are applicable to both SIDs and CSIDs.

The H.Encaps.L1 behavior encapsulates a frame received from an IWF in an IPv6 packet with a segment routing header (SRH). The received frame becomes the payload of the new IPv6 packet.

- * The next header field of the SRH or the last extension header present MUST be set to 147.
- * The insertion of the SRH MAY be omitted per [RFC8986] when the SRv6 policy only contains one segment and there is no need to use any flag, tag, or TLV.

The H.Encaps.L1.Red behavior is an optimization of the H.Encaps.L1 behavior.

- * H.Encaps.L1.Red reduces the length of the SRH by excluding the first SID in the SRH. The first SID is only placed in the Destination Address field of the IPv6 header.
- * The insertion of the SRH MAY be omitted per [RFC8986] when the SRv6 policy only contains one segment and there is no need to use any flag, tag, or TLV.

Three new "Endpoint with decapsulation and bit-stream cross-connect" behaviors called "End.DX1", "End.DX1 with NEXT-CSID", and "End.DX1 with REPLACE-CSID" are defined in this document. These new behaviors are variants of End.DX2 defined in [RFC8986], and they all have the following procedures in common:

The End.DX1 SID MUST be the last segment in an SR Policy, and it is associated with a CE-bound IWF I. When N receives a packet destined to S and S is a local End.DX1 SID, N does the following:

```
S01. When an SRH is processed {
S02.   If (Segments Left != 0) {
S03.     Send an ICMP Parameter Problem to the Source Address
        with Code 0 (Erroneous header field encountered)
        and Pointer set to the Segments Left field,
        interrupt packet processing, and discard the packet.
S04.   }
S05.   Proceed to process the next header in the packet
S06. }
```

When processing the next (Upper-Layer) header of a packet matching a FIB entry locally instantiated as an End.DX1 SID, N does the


```

S01. If (Upper-Layer header type == 147 (bit-stream) ) {
S02.     Remove the outer IPv6 header with all its extension headers
S03.     Forward the remaining frame to the IWF I
S04. } Else {
S05.     Process as per {{Section 4.1.1 of RFC 8986}}
S06. }

```

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.

The format of the PLE control word is in line with the guidance in [RFC4385] and is shown in Figure 5.



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

Set by the PE to indicate that data carried in the payload is invalid due to an attachment circuit fault. The downstream PE MUST send appropriate replacement data. The NSP MAY inject an appropriate downstream fault-indication signal.

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.

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

These bits MUST be set to zero by the sender and ignored by the receiver as PLE does not use payload fragmentation.

In accordance with Section 3 of [RFC4385], the length field MUST always be set to zero as there is no padding added to the PLE packet. The size of the PLE payload MUST be assumed to be as described in Section 6; if the actual packet size is inconsistent with this, the packet MUST be considered malformed.

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] and MUST be incremented with every PLE packet being sent.

5.2.2. RTP Header

The RTP header MUST be included to explicitly convey timing information.

The RTP header (as defined in [RFC3550]) is reused to align with other bit-stream emulation pseudowires defined by [RFC4553], [RFC5086], and [RFC4842] and to allow PLE implementations to reuse preexisting work.

There is no intention to support full RTP topologies and protocol mechanisms, such as header extensions, contributing source (CSRC) list, padding, RTCP, RTP header compression, SRTP, etc., as these are not applicable to PLE VPWS.

The format of the RTP header is as shown in Figure 6.




Figure 6: 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 the sender and ignored by the receiver.

PT:
Payload type

A PT value MUST be allocated from the range of dynamic values defined in Section 6 of [RFC3551] for each direction of the VPWS. The same PT value MAY be reused for both for directions and

between different PLE VPWSs.

The PT field MAY be used for detection of misconnections.

Sequence number:

When using a 16-bit sequence number space, the sequence number in the RTP header MUST be equal to the sequence number in the PLE control word. When using a sequence number space of 32 bits, the initial value of the RTP sequence number MUST be 0 and incremented whenever the PLE control word sequence number cycles through from 0xFFFF to 0x0000.

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.

6. 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. The payload size SHOULD be configurable to be able to address specific packetization delay and overhead expectations. The smallest supported payload size is 64 bytes.

6.1. Basic 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.

6.2. Byte-Aligned Payload

The PLE payload is filled in a byte-aligned manner, where 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. The PLE payload size MUST be an integer number of bytes.

7. PLE Operation

7.1. Common Considerations

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

Furthermore, emulation of bit-stream signals using PLE is only possible when the two attachment circuits of the VPWS are of the same service 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 the mechanisms of a signaling protocol.

PLE-related control protocol extensions to LDP [RFC8077] or EVPN-VPWS [RFC8214] are out of scope for this document.

Extensions for EVPN-VPWS are proposed in [EVPN-VPWS] and for LDP in [LDP-PLE].

7.2. PLE IWF Operation

7.2.1. PSN-Bound Encapsulation Behavior

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

- * Packetize the data received from the CE into PLE payloads, all of the same configured size,
- * 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 the L bit in the PLE control word whenever the attachment circuit detects a fault, and
- * Set the R bit in the PLE control word whenever the local CE-bound IWF is in packet loss state.

7.2.2. 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 sending the bit-stream out via 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 PDV expected.

The CE-bound IWF SHOULD use the sequence number in the control word to detect lost and misordered packets. It MAY use the sequence number in the RTP header for the same purpose. The CE-bound IWF MAY support reordering of packets received out of order. If the CE-bound IWF does not support reordering, it MUST drop the misordered packets.

The payload of a lost or dropped packet MUST be replaced with an equivalent amount of replacement data. The contents of the replacement data MAY be locally configurable. By default, all PLE implementations MUST support generation of "0xAA" as replacement data. The alternating sequence of 0s and 1s of the "0xAA" pattern ensures clock synchronization is maintained and, for 64B/66B code-based services, ensures no invalid sync headers are generated. While sending 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 downstream fault-indication signal.

Whenever a VPWS comes up, the CE-bound IWF will enter 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 downstream fault-indication signal until a preconfigured number of payload s stored in the jitter buffer.

After the preconfigured 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 streamed out to the CE in

accordance with the required clock. In this state, the CE-bound IWF MUST perform egress clock recovery.

Considerations for choosing the preconfigured amount of payload required to be present for transitioning into the normal state:

- * Typically set to 50% of the de-jitter buffer size to equally allow compensating for increasing and decreasing delay
- * A compromise between the maximum amount of tolerable PDV and delay introduced to the emulated service

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

- * [G.825], [G.783], and [G.823] for SDH
- * [GR253] and [GR499] 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 downstream fault-indication signal instead of streaming out the payload.

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

If the CE-bound IWF detects a PLR above a configurable SD threshold for a configurable amount of consecutive 1-second intervals, it enters the DEG state and a corresponding defect is declared. The SD-PLR threshold can be defined as a percentage with the default being 15% or absolute packet count for finer granularity for higher rate interfaces. Possible values for consecutive intervals are 2..10 with the default 7.

While the PLOS defect is declared, the CE-bound NSP function MUST inject the appropriate downstream fault-indication signal. If the emulated service does not have an appropriate maintenance signal defined, the CE-bound NSP function MAY disable its transmitter instead. Also, the PSN-bound IWF SHOULD set the R bit in the PLE control word of every packet transmitted.

The CE-bound IWF changes from the PLOS to normal state after the preconfigured amount of payload has been received similar 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.

7.3. PLE Performance Monitoring

Attachment circuit performance monitoring SHOULD be provided by the NSP. The performance monitors are service specific, documented in related specifications, and beyond the scope of this document.

The PLE IWF SHOULD provide 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 or DEG defect SHALL be counted as an ES-PLE. Each second with a PLR greater than 15% or a PLOS or DEG defect SHALL be counted as an SES-PLE.

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

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

A PLE far-end performance monitor provides 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].

Performance monitoring data MUST be provided by the management interface and SHOULD be provided by a YANG data model. The YANG data model specification is out of scope for this document.

7.4. PLE Fault Management

Attachment circuit faults applicable to PLE are detected by the NSP, are service specific, and are documented in Section 4.

The two PLE faults, PLOS and DEG, are detected by the IWF.

Faults MUST be timestamped as they are declared and cleared; fault-related information MUST be provided by the management interface and SHOULD be provided by a YANG data model. The YANG data model specification is out of scope for this document.

8. QoS and Congestion Control

The PSN carrying PLE VPWS may be subject to congestion. Congestion considerations for PWs are described in Section 6.5 of [RFC3985].

PLE VPWS represent inelastic CBR flows that cannot respond to congestion in a TCP-friendly manner (as described in [RFC2914]) and are sensitive to jitter, packet loss, and packets received out of order.

The PSN providing connectivity between PE devices of a PLE VPWS has to ensure low jitter and low loss. The exact mechanisms used are beyond the scope of this document and may evolve over time. Possible options, but not exhaustively, are as follows:

- * a Diffserv-enabled [RFC2475] PSN with a per-domain behavior (see [RFC3086]) supporting Expedited Forwarding (see [RFC3246]),
- * traffic-engineered paths through the PSN with bandwidth reservation and admission control applied, or
- * capacity over-provisioning.

9. Security Considerations

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

PLE does not enhance or detract from the security performance of the underlying PSN. It relies upon the PSN mechanisms for encryption, integrity, and authentication whenever required.

The PSN (MPLS or SRv6) is assumed to be trusted and secure. Attackers who manage to send spoofed packets into the PSN could easily disrupt the PLE service. This MUST be prevented by following best practices for the isolation of the PSN. These protections are described in Section 3.4 of [RFC4381], Section 4.2 of [RFC5920], Section 8 of [RFC8402], and Section 9.3 of [RFC9252].

PLE PWs share susceptibility to a number of pseudowire-layer attacks and will use whatever mechanisms for confidentiality, integrity, and authentication that are developed for general PWs. These methods are beyond the scope of this document.

Random initialization of sequence numbers, in both the control word and the RTP header, makes known-plaintext attacks more difficult.

Misconnection detection using the SSRC and/or PT field of the RTP header can increase the resilience to misconfiguration and some types of denial-of-service (DoS) attacks. Randomly chosen expected values decrease the chance of a spoofing attack being successful.

A data plane attack may force PLE packets to be dropped, reordered, or delayed beyond the limit of the CE-bound IWF's dejitter buffer leading to either degradation or service disruption. Considerations outlined in [RFC9055] are a good reference.

Clock synchronization leveraging PTP is sensitive to PDV and vulnerable to various threats and attack vectors. Considerations outlined in [RFC7384] should be taken into account.

10. IANA Considerations

10.1. Bit-Stream Next Header Type

This document introduces a new value to be used in the next header field of an IPv6 header or any extension header indicating that the payload is an emulated bit-stream. IANA has assigned the following from the "Assigned Internet Protocol Numbers" registry [IANA-Proto].

Decimal	Keyword	Protocol	IPv6 Extension Header	Reference
147	BIT-EMU	Bit-stream Emulation	Y	RFC 9801

Table 1

10.2. SRv6 Endpoint Behaviors

This document introduces three new SRv6 Endpoint behaviors. IANA has assigned identifier values in the "SRv6 Endpoint Behaviors" registry under the "Segment Routing" registry group [IANA-SRv6-End].

Value	Hex	Endpoint Behavior	Reference
-------	-----	-------------------	-----------

158	0x009E	End.DX1	RFC 9801	
159	0x009F	End.DX1 with NEXT-CSID	RFC 9801	
160	0x00A0	End.DX1 with REPLACE-CSID	RFC 9801	

Table 2

11. References

11.1. Normative References

- [G.707] ITU-T, "Network node interface for the synchronous digital hierarchy (SDH)", ITU-T Recommendation G.707, January 2007, <<https://www.itu.int/rec/T-REC-G.707>>.
- [G.709] ITU-T, "Interfaces for the optical transport network", ITU-T Recommendation G.709, June 2020, <<https://www.itu.int/rec/T-REC-G.709>>.
- [G.783] ITU-T, "Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks", ITU-T Recommendation G.783, March 2006, <<https://www.itu.int/rec/T-REC-G.783>>.
- [G.823] ITU-T, "The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy", ITU-T Recommendation G.823, March 2000, <<https://www.itu.int/rec/T-REC-G.823>>.
- [G.825] ITU-T, "The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)", ITU-T Recommendation G.825, March 2000, <<https://www.itu.int/rec/T-REC-G.825>>.
- [G.8251] ITU-T, "The control of jitter and wander within the optical transport network (OTN)", ITU-T Recommendation G.8251, November 2022, <<https://www.itu.int/rec/T-REC-G.8251>>.
- [G.8261] ITU-T, "Timing and synchronization aspects in packet networks", ITU-T Recommendation G.8261, August 2019, <<https://www.itu.int/rec/T-REC-G.8261>>.
- [G.8261.1] ITU-T, "Packet delay variation network limits applicable to packet-based methods (Frequency synchronization)", ITU-T Recommendation G.8261.1, February 2012, <<https://www.itu.int/rec/T-REC-G.8261.1>>.
- [G.8262] ITU-T, "Timing characteristics of synchronous equipment clocks", ITU-T Recommendation G.8262, October 2024, <<https://www.itu.int/rec/T-REC-G.8262>>.
- [G.8265.1] ITU-T, "Precision time protocol telecom profile for frequency synchronization", ITU-T Recommendation G.8265.1, November 2022, <<https://www.itu.int/rec/T-REC-G.8265.1>>.
- [GR253] Telcordia, "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria", GR-253, October 2009, <<https://telecom-info.njdepot.ericsson.net/site-cgi/ido/docs.cgi?ID=2111701336SEARCH&DOCUMENT=GR-253>>.
- [GR499] Telcordia, "Transport Systems Generic Requirements (TSGR) - Common Requirements", GR-499, November 2009,

<<https://telecom-info.njdepot.ericsson.net/site-cgi/ido/docs.cgi?ID=2111701336SEARCH&DOCUMENT=GR-499>>.

[IANA-Proto]

IANA, "Assigned Internet Protocol Numbers",
<<https://www.iana.org/assignments/protocol-numbers>>.

[IANA-SRv6-End]

IANA, "SRv6 Endpoint Behaviors",
<<https://www.iana.org/assignments/segment-routing>>.

[IEEE802.3]

IEEE, "IEEE Standard for Ethernet", IEEE Std 802.3-2022,
DOI 10.1109/IEEESTD.2022.9844436, July 2022,
<<https://ieeexplore.ieee.org/document/9844436>>.

[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>>.

[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>>.

[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>>.

[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>>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986,
DOI 10.17487/RFC8986, February 2021,
<<https://www.rfc-editor.org/info/rfc8986>>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252,
DOI 10.17487/RFC9252, July 2022,
<<https://www.rfc-editor.org/info/rfc9252>>.

[RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025,
<<https://www.rfc-editor.org/info/rfc9800>>.

11.2. Informative References

[ATIS-0900105.09.2013]

ATIS, "Synchronous Optical Network (SONET) - Network

Element Timing and Synchronization", ATIS-0900105.09.2013(S2023), 2023,
<<https://webstore.ansi.org/standards/atis/atis0900105092013s2023>>.

[EVPN-VPWS]

Gringeri, S., Whittaker, J., Schmutzer, C., Ed., Vasudevan, B., and P. Brissette, "Ethernet VPN Signalling Extensions for Bit-stream VPWS", Work in Progress, Internet-Draft, draft-schmutzer-bess-bitstream-vpws-signalling-02, 18 October 2024,
<<https://datatracker.ietf.org/doc/html/draft-schmutzer-bess-bitstream-vpws-signalling-02>>.

[FC-PI-2] INCITS, "Information Technology - Fibre Channel Physical Interfaces - 2 (FC-PI-2)", INCITS 404-2006 (S2016), 2016,
<<https://webstore.ansi.org/standards/incits/incits4042006s2016>>.

[FC-PI-5] INCITS, "Information Technology - Fibre Channel - Physical Interface-5 (FC-PI-5)", INCITS 479-2011 (S2021), 2021,
<<https://webstore.ansi.org/standards/incits/incits4792011s2021>>.

[FC-PI-5aml]

INCITS, "Information Technology - Fibre Channel - Physical Interface - 5/Amendment 1 (FC-PI-5/AM1)", INCITS 479-2011/AM1-2016 (R2021), 2021,
<<https://webstore.ansi.org/standards/incits/incits4792011am2016r2021>>.

[FC-PI-6] INCITS, "Information Technology - Fibre Channel - Physical Interface - 6 (FC-PI-6)", INCITS 512-2015 (R2020), 2020,
<<https://webstore.ansi.org/standards/incits/incits5122015r2020>>.

[FC-PI-6P] INCITS, "Information Technology - Fibre Channel - Physical Interface - 6P (FC-PI-6P)", INCITS 533-2016 (R2021), 2021,
<<https://webstore.ansi.org/standards/incits/incits5332016r2021>>.

[FC-PI-7] ISO/IEC, "Information technology - Fibre channel - Part 147: Physical interfaces - 7 (FC-PI-7)", ISO/IEC 14165-147:2021, 2021,
<<https://www.iso.org/standard/80933.html>>.

[G.826] ITU-T, "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections", ITU-T Recommendation G.826, December 2002, <<https://www.itu.int/rec/T-REC-G.826>>.

[LDP-PLE] Schmutzer, C., Ed., "LDP Extensions to Support Private Line Emulation (PLE)", Work in Progress, Internet-Draft, draft-schmutzer-pals-ple-signaling-02, 20 October 2024,
<<https://datatracker.ietf.org/doc/html/draft-schmutzer-pals-ple-signaling-02>>.

[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>>.

[RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000,
<<https://www.rfc-editor.org/info/rfc2914>>.

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [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>>.
- [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>>.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C.R., Benson, K., Le Boudec, J.Y., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, DOI 10.17487/RFC3246, March 2002, <<https://www.rfc-editor.org/info/rfc3246>>.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <<https://www.rfc-editor.org/info/rfc3711>>.
- [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>>.
- [RFC4381] Behringer, M., "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4381, DOI 10.17487/RFC4381, February 2006, <<https://www.rfc-editor.org/info/rfc4381>>.
- [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>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, DOI 10.17487/RFC4448, April 2006, <<https://www.rfc-editor.org/info/rfc4448>>.
- [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>>.
- [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>>.
- [RFC4842] Malis, A., Pate, P., Cohen, R., Ed., and D. Zelig, "Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Circuit Emulation over Packet (CEP)", RFC 4842, DOI 10.17487/RFC4842, April 2007,

<<https://www.rfc-editor.org/info/rfc4842>>.

- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.
- [RFC4906] Martini, L., Ed., Rosen, E., Ed., and N. El-Aawar, Ed., "Transport of Layer 2 Frames Over MPLS", RFC 4906, DOI 10.17487/RFC4906, June 2007, <<https://www.rfc-editor.org/info/rfc4906>>.
- [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>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC7212] Frost, D., Bryant, S., and M. Bocci, "MPLS Generic Associated Channel (G-ACh) Advertisement Protocol", RFC 7212, DOI 10.17487/RFC7212, June 2014, <<https://www.rfc-editor.org/info/rfc7212>>.
- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <<https://www.rfc-editor.org/info/rfc7384>>.
- [RFC8077] Martini, L., Ed. and G. Heron, Ed., "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", STD 84, RFC 8077, DOI 10.17487/RFC8077, February 2017, <<https://www.rfc-editor.org/info/rfc8077>>.
- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <<https://www.rfc-editor.org/info/rfc9055>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <<https://www.rfc-editor.org/info/rfc9293>>.

[T11] INCITS, "T11 - Fibre Channel",
<<https://www.incits.org/committees/t11>>.

Acknowledgements

The authors would like to thank Alexander Vainshtein, Yaakov Stein, Erik van Veelen, Faisal Dada, Giles Heron, Luca Della Chiesa, and Ashwin Gumaste for their early contributions, review, comments, and suggestions.

Special thank you to:

- * Carlos Pignataro and Nagendra Kumar Nainar for giving the authors new-to-the-IETF guidance on how to get started
- * Stewart Bryant for being our shepherd
- * Tal Mizahi, Joel Halpern, Christian Huitema, Tony Li, and Tommy Pauly for their reviews and suggestions during Last Call
- * Andrew Malis and Gunter van de Velde for their guidance through the process

Contributors

Andreas Burk
l&l Versatel
Email: andreas.burk@magenta.de

Faisal Dada
AMD
Email: faisal.dada@amd.com

Gerald Smallegange
Ciena Corporation
Email: gsmalleg@ciena.com

Erik van Veelen
Aimvalley
Email: erik.vanveelen@aimvalley.com

Luca Della Chiesa
Cisco Systems, Inc.
Email: ldellach@cisco.com

Nagendra Kumar Nainar
Cisco Systems, Inc.
Email: naikumar@cisco.com

Carlos Pignataro
Blue Fern Consulting
Email: Carlos@Bluefern.consulting

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

Chris Brown
Ciena Corporation
Email: cbrown@ciena.com