BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 4, 2021

S. Gringeri
J. Whittaker
Verizon
C. Schmutzer, Ed.
P. Brissette
Cisco Systems, Inc.

May 3, 2021

Private Line Emulation VPWS Signalling
draft-schmutzer-bess-ple-vpws-signalling-02

Abstract

This document specifies the mechanisms to allow for dynamic signalling of Virtual Private Wire Services (VPWS) carrying bit-stream signals over Packet Switched Networks (PSN).

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 November 4, 2021.

Copyright Notice

Table of Contents

1. Introduction

Virtual Private Wire Service (VPWS) is a widely deployed technology for providing point-to-point (P2P) services for various layer 2 and also layer 1 technologies. Initially VPWS were define in the Pseudowire Emulation Edge-to-Edge (PWE3) architecture [RFC3985] for Frame Relay, ATM, HDLC, PPP, Ethernet, TDM and SONET/SDH.

This document focuses on bit stream VPWS instance types which already got introduced in [RFC3985]. Possible bit stream VPWS instance types and their encapsulation specification documents are:

- TDM services using SAToP [RFC4553]
- TDM services using CESoP [RFC5086]
- SONET/SDH services using CEP [RFC4842]
- High-speed private line services using PLE [PLE]

Signalling mechanisms and extensions to [RFC8077] required to dynamically signal TDM bit-stream services ([RFC4553], [RFC5086]) and SONET/SDH bit-stream services ([RFC4842]) are already described in [RFC5287].

The scope of this document is to specify extensions to [RFC8077] required to dynamically signal PLE bit-stream services defined in [PLE] and to specify extensions required to use EVPN-VPWS [RFC8214] as a signalling protocol for all bit-stream services mentioned in this document.

A generic VPWS reference model similar to the one defined in [RFC3985] and [PLE] is shown in Figure 1. Data received from a CEs is encapsulated by PEs into the respective VPWS established between the attachment circuits of the local and remote PE and transmitted across the Packet Switched Network (PSN) using a PSN tunnel.

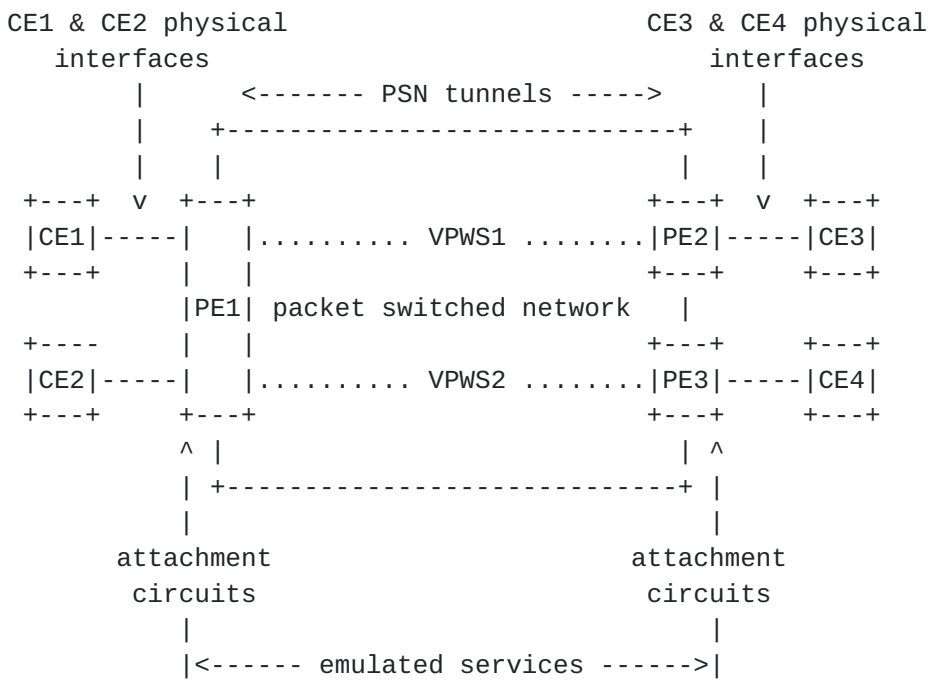

Figure 1: VPWS Reference Model

In the example shown in Figure 1 there are two CE nodes (CE1 and CE2) connected to the same PE node (PE1). CE3 is connected to PE2 and CE4 is connected to PE3. There are two VPWS instances established. VPWS1 between CE1 and CE3 and VPWS2 between CE2 and CE4. For traffic to be carried across the network PSN tunnels between PE1 and PE2 and between PE1 and PE3 are needed.

In order for a bit stream VPWS instance to come up, the attachment circuit parameters must be identical on both endpoints. The control plane mechanisms described in this document are leveraged to meet this requirement. Mechanisms for misconnection detection and protection switch coordination are also described.

1.1. Requirements Language

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.

1.2. Terminology

- AIS - Alarm Indication Signal
- AFI - Address Family Identifier
- ATM - Asynchronous Transfer Mode
- BGP - Border Gateway Protocol
- CBR - Constant Bit Rate
- CE - Customer Edge
- CEP - SONET/SDH Circuit Emulation over Packet
- CESoP - Structure-aware TDM Circuit Emulation Service over Packet Switched Network
- DF - Designated Forwarder
- EAD - Ethernet Auto Discovery
- FC - Fibre Channel
- EBM - Equipped Bit Mask
- EVI - EVPN Instance
- EVPN - Ethernet Virtual Private Network
- HDLC - High-level Data Link Control
- LDP - Label Distribution Protocol
- MPLS - Multi Protocol Label Switching
- MTU - Maximum Transmission Unit
- NDF - Non-Designated Forwarder
- NLRI - Network Layer Reachability Information
- OC - Optical Carrier
- ODUk - Optical Data Unit k
- PDH - Plesynchronous Digital Hierarchy

- PE - Provider Edge
- PLE - Private Line Emulation
- PPP - Point-to-Point Protocol
- PSN - Packet Switched Network
- PW - Pseudo Wire
- PWE3 - Pseudowire Emulation Edge-to-Edge
- P2P - Point-to-Point
- RTP - Realtime Transport Protocol
- SAFI - Subsequent Address Family Identifier
- SAToP - Structure Agnostic TDM over Packet
- SDH - Synchronous Digital Hierarchy
- SONET - Synchronous Optical Network
- SRv6 - Segment Routing over IPv6 Dataplane
- STM - Synchronous Transport Module
- STS - Synchronous Transport Signal
- TDM - Time Division Multiplexing
- TLV - Type Length Value
- UNE - Unequipped
- VC - Virtual Circuit
- VPWS - Virtual Private Wire Service
- VT - Virtual Tributary
-

2. Solution Requirements

To avoid redefining PW types for [RFC8214] the notion of "PW type" from [RFC8077] is maintained and only a new PW type for [PLE] has been assigned by IANA.

- TBD1 - Private Line Emulation (PLE) over Packet

The concept of "CEP type" from [RFC5287] to distinguish different connection types that use the same PW type is adopted. In this document it is referred to as "PLE/CEP type". Two new connection types are defined (see Section 4.2.3).

To unambiguously identify the rate of an attachment circuit, also the concept of "CEP/TDM bit-rate" from [RFC5287] is adopted and called "PLE bit-rate" herein.

The VPWS signalling requirements are as follows:

- EVPN-VPWS [RFC8214] as signalling protocol MUST be supported
- LDP [RFC8077] MAY be supported as VPWS signalling protocol
- Implementations MUST support MPLS as underlay PSN
- The VPWS instance MAY be signalled as SRv6 overlay service per [srv6_overlay] leveraging on [srv6_netprog] using the End.DX2 function. In such case, the implementation MUST support SRv6 as underlay PSN.
- The use of control word MUST be signalled, as defined in [RFC4553], [RFC5086], [RFC4842] and [PLE].
- The PW type MUST be signalled and the PE nodes MUST validate that the PW type is identical on both endpoint.
- For CEP [RFC4842] and PLE [PLE] the PLE/CEP type MUST be signalled and the PE nodes MUST validate that the PLE/CEP type is identical on both endpoints.
- The PLE/CEP/TDM bit-rate MUST be signalled if the attachment circuit can not be unambiguously identified from the PW type alone and the PE nodes MUST validate that the attachment circuit is identical on both endpoints.
- A non-default payload size MAY be signalled. Both PE nodes MUST validate that the payload size is identical on both endpoints.

- A locally configured connection identifier as defined in Section 4.2.6 SHOULD be sent to the remote PE node. A locally configured expected identifier MAY be used to identify a misconnection by comparing it with the identifier received from the remote PE node.

- When using EVPN-VPWS [RFC8214] as a signalling protocol, multi-homed PE scenarios per [RFC7432] and [RFC8214] SHOULD be supported where the load-balancing mode single-active MUST be supported. Port-active load-balancing mode MAY also be supported.

- For EVPN-VPWS [RFC8214] multi-homed PE scenarios non-revertive mode MUST and revertive mode SHOULD be supported in compliance to [pref_df].

3. Service Types

The following sections list all possible service types that are supported by the proposed signalling mechanisms.

3.1. Ethernet Service Types

Service Type	Encapsulation Standard	PW Type	PLE/CEP Type	PLE/CEP/TDM Bit-rate
1000BASE-X	[PLE]	TBD1	0x3	1,250,000
10GBASE-R	[PLE]	TBD1	0x3	10,312,500
25GBASE-R	[PLE]	TBD1	0x3	25,791,300
40GBASE-R	[PLE]	TBD1	0x3	41,250,000
100GBASE-R	[PLE]	TBD1	0x3	103,125,000

3.2. Fibre Channel Service Types

Service Type	Encapsulation Standard	PW Type	PLE/CEP Type	PLE/CEP/TDM Bit-rate
1GFC	[PLE]	TBD1	0x3	1,062,500
2GFC	[PLE]	TBD1	0x3	2,125,000
4GFC	[PLE]	TBD1	0x3	4,250,000
8GFC	[PLE]	TBD1	0x3	8,500,000
10GFC	[PLE]	TBD1	0x3	19,518,750
16GFC	[PLE]	TBD1	0x3	14,025,000
32GFC	[PLE]	TBD1	0x3	28,050,000
128GFC	[PLE]	TBD1	0x3	112,200,000

3.3. OTN Service Types

Service Type	Encapsulation Standard	PW Type	PLE/CEP Type	PLE/CEP/TDM Bit-rate
ODU0	[PLE]	TBD1	0x4	1,244,160
ODU1	[PLE]	TBD1	0x4	2,498,775
ODU2	[PLE]	TBD1	0x4	10,037,273
ODU2e	[PLE]	TBD1	0x4	10,399,525
ODU3	[PLE]	TBD1	0x4	40,319,218
ODU4	[PLE]	TBD1	0x4	104,794,445

3.4. TDM Service Types

Service Type	Encapsulation Standard	PW Type	PLE/CEP Type	PLE/CEP/TDM Bit-rate
CESoPSN basic mode	[RFC5086]	0x0015	N/A	N
CESoPSN with CAS	[RFC5086]	0x0017	N/A	N
E1	[RFC4553]	0x0011	N/A	32
DS1	[RFC4553]	0x0012	N/A	24
DS1 octet-aligned	[RFC4553]	0x0012	N/A	25
E3	[RFC4553]	0x0013	N/A	535
T3	[RFC4553]	0x0014	N/A	699

N is the number of DS0 channels in the attachment circuit

3.5. SONET/SDH Service Types

Service Type	Encapsulation Standard	PW Type	PLE/CEP Type	PLE/CEP/TDM Bit-rate
VT1.5/VC-11	[RFC4842]	0x0010	0x1	26
VT2/VC-12	[RFC4842]	0x0010	0x1	35
VT3	[RFC4842]	0x0010	0x1	53
VT6/VC-2	[RFC4842]	0x0010	0x1	107
STS-Nc	[RFC4842]	0x0010	0x0	783*N
VC-4-Mc	[RFC4842]	0x0010	0x0	783*3*M
Fract. STS1/VC-3	[RFC4842]	0x0010	0x2	783
Fract. VC-4	[RFC4842]	0x0010	0x2	783*4
Async STS1/VC-3	[RFC4842]	0x0010	0x2	783
OC3/STM1	[PLE]	TBD1	0x3	155,520
OC12/STM4	[PLE]	TBD1	0x3	622,080
OC48/STM16	[PLE]	TBD1	0x3	2,488,320
OC192/STM64	[PLE]	TBD1	0x3	9,953,280
OC768/STM256	[PLE]	TBD1	0x3	39,813,120

N=1,3,12,48,192,768 and M=1,4,16,64,256

4. EVPN-VPWS signalling

4.1. Reuse of existing BGP EVPN-VPWS capabilities

A PLE VPWS instance is identified by a pair of per-EVI ethernet A-D routes advertised by two PE nodes establishing the VPWS in accordance to [RFC8214].

The EVPN layer 2 attribute extended community defined in [RFC8214] MUST be supported and added to the per-EVI ethernet A-D route.

- C bit set to 1 to indicate Control Word MUST be present.
- P and B bits are set by dual-homing PEs as per [RFC8214] and [pref_df]
- L2 MTU MUST be set to zero and ignored by the receiver

4.2. BGP PLE Attribute

To exchange and validate bit-stream specific attachment circuit parameters during the VPWS setup, a new BGP path attribute called "BGP PLE attribute" is defined.

The BGP PLE attribute defined in this document can be attached to EVPN VPWS routes [RFC8214]. The usage for other Address Family

Identifier (AFI) / Subsequent Address Family Identifier (SAFI) combinations is not defined herein but may be specified in future specifications.

The BGP PLE attribute is an optional and transitive BGP path attribute. The attribute type code TBD2 has been assigned by IANA (see section Section 6)

The format is defined as a set of Type/Length/Value (TLV) triplets, described in the following sections and listed in Table 1. This attribute SHOULD only be included with EVPN Network Layer Reachability Information (NLRI).

TLV Type	Name	Length	Mandatory
1	PW Type TLV	3	Y
2	PLE/CEP/TDM Bit-rate TLV	5	Y
3	PLE/CEP Options TLV	3	Y 1*
4	TDM Options TLV	13	Y 2*
5	PLE/CEP/TDM Payload Bytes TLV	3	N
6	Endpoint-ID TLV	0..80	N

1* PLE/CEP only, 2* TDM only

Table 1: BGP PLE attribute TLVs

For a particular PSN it is expected that the network operator will choose a common set of parameters per VPWS type, hence efficient BGP update packing as discussed in section 12 of [RFC4277] is expected to happen.

4.2.1. PW Type TLV

The PW Type TLV MUST be present in the BGP PLE attribute to signal what type of VPWS instance has to be established. Valid PW types for the mechanisms described in this document can be found in Section 3.

The PW Type TLV format is shown in Figure 2.

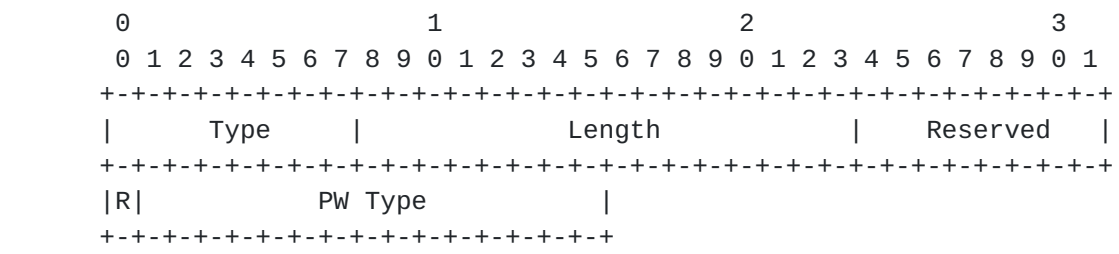

Figure 2: PW type TLV

Type : 1

Length : 3

The total length in octets of the value portion of the TLV.

Reserved / R :

For future use. MUST be set to ZERO and ignored by receiver.

PW Type :

A 15-bit quantity containing a value that represents the type of VPWS. Assigned Values are specified in "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)" [RFC4446].

4.2.2. PLE/CEP/TDM Bit-rate TLV

The PLE/CEP/TDM Bit-rate TLV is MANDATORY but MAY be omitted if the attachment circuit type can be unambiguously derived from the PW Type carried in the PW Type TLV. The PLE/CEP/TDM Bit-rate TLV format is shown in Figure 3.

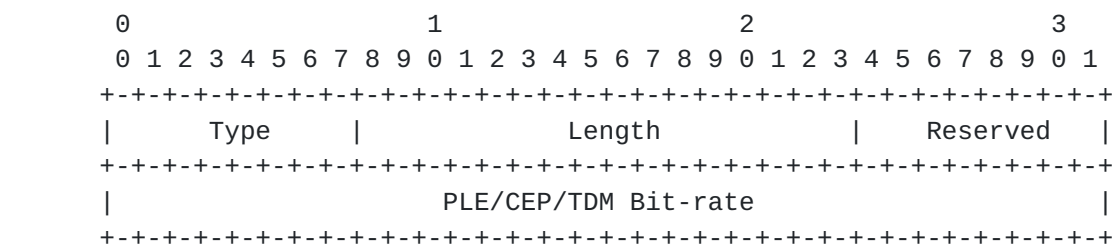

Figure 3: PLE/CEP/TDM Bit-rate TLV

Type : 2

Length : 5

The total length in octets of the value portion of the TLV.

Reserved :

8-bit field for future use. MUST be set to ZERO and ignored by receiver.

PLE/CEP/TDM Bit-rate :

A four byte field denoting the desired payload size to be used. Rules defined in [RFC5287] do apply for signalling TDM VPWS. Rules for CEP VPWS are defined in [RFC4842].

* For PLE [PLE] the bit rate MUST be set to the data rate in units of 1-kbps of the PLE payload.

* Guidelines for setting the bit rate for SAToP VPWS and CESoP VPWS can be found in [RFC5287]. And for CEP VPWS in [RFC4842].

4.2.3. PLE/CEP Options TLV

The PLE/CEP Options TLV MUST be present when signalling CEP and PLE VPWS instances. The PLE/CPE Options TLV format is shown in Figure 4.

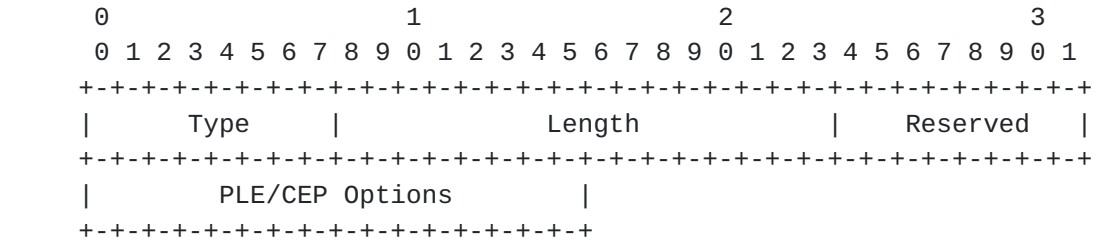

Figure 4: PLE/CEP Options TLV

Type : 3

Length : 3

The total length in octets of the value portion of the TLV.

Reserved :

8-bit field for future use. MUST be set to ZERO and ignored by receiver.

PLE/CEP Options :

A two byte field with the format as shown in Figure 5

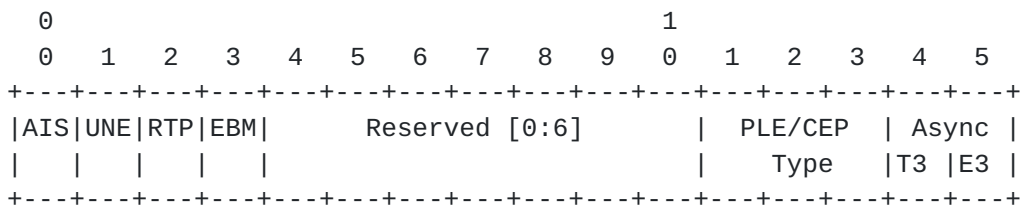

Figure 5: PLE/CEP Options

AIS, UNE, RTP, EBM :

These bits MUST be set to zero and ignored by the receiver except for CEP VPWS. Guidelines for CEP are defined in [RFC4842]

Reserved :

7-bit field for future use. MUST be set to ZERO and ignored by receiver.

CEP/PLE Type :

Indicates the connection type for CEP and PLE. CEP connection types are defined in [RFC4842]. Two new values for PLE are defined in this document:

0x3 - Constand Bit Rate (CBR) PLE payload

0x4 - Byte aligned PLE payload

Async :

These bits MUST be set to zero and ignored by the receiver except for CEP VPWS. Guidelines for CEP are defined in [RFC4842]

4.2.4. TDM options TLV

Whether when signalling TDM VPWS the TDM Options TLV MUST be present or MAY be omitted when signalling TDM VPWS instances is defined in [RFC5287]. The TDM Options TLV format is shown in Figure 6.

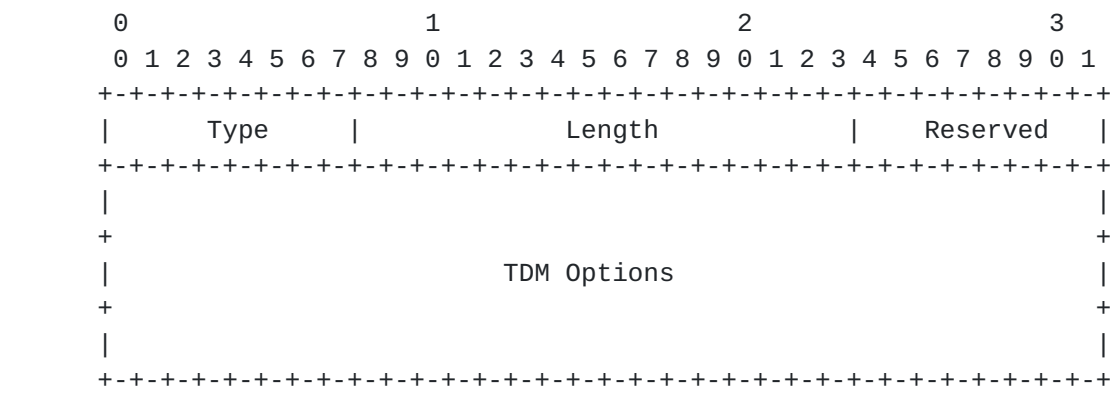

Figure 6: TDM Options TLV

Type : 4

Length : 13

The total length in octets of the value portion of the TLV.

Reserved :

8-bit field for future use. MUST be set to ZERO and ignored by receiver.

TDM Options :

A twelve byte field with the format as defined in section 3.8 of [RFC5287]

4.2.5. PLE/CEP/TDM Payload Bytes TLV

The PLE/CEP/TDM Payload Bytes TLV MAY be included if a non-default payload size is to be used. If this TLV is omitted then the default payload sizes defined in [RFC4553], [RFC5086], [RFC4842] and [PLE] MUST be assumed. The format of the PLE/CEP/TDM Payload Bytes TLV is shown in Figure 7.

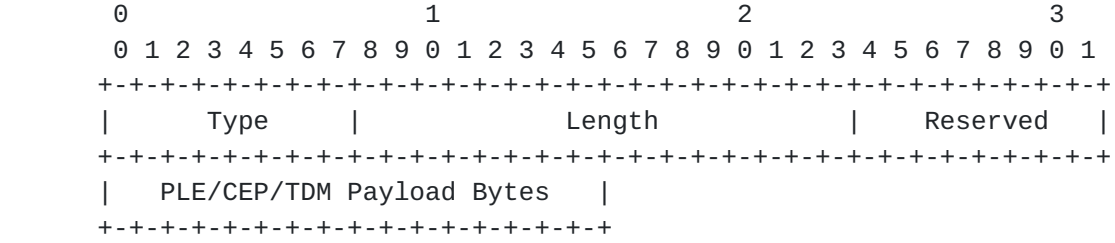

Figure 7: PLE/CEP/TDM Payload Bytes TLV

Type : 5

Length : 3

The total length in octets of the value portion of the TLV.

Reserved :

8-bit field for future use. MUST be set to ZERO and ignored by receiver.

PLE/CEP/TDM Payload Bytes :

A two byte field denoting the desired payload size to be used. Rules defined in [RFC5287] do apply for signalling TDM VPWS. Rules for CEP VPWS are defined in [RFC4842].

4.2.6. Endpoint-ID TLV

The Endpoint-ID TLV MAY be included to allow for misconnection detection. The Endpoint-ID TLV format is shown in Figure 8.

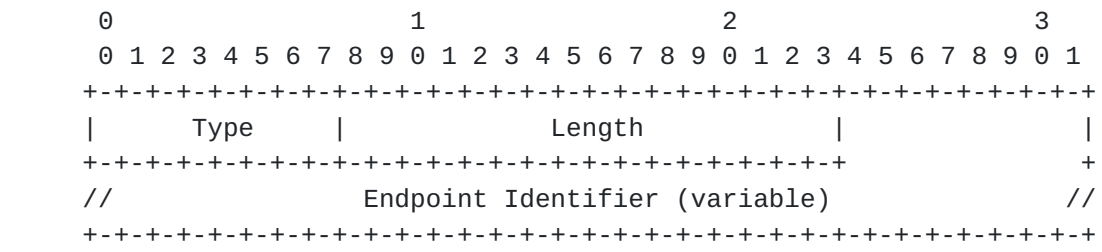

Figure 8: Endpoint-ID TLV

Type : 6

Length : 0-80

The total length in octets of the value portion of the TLV.

Endpoint Identifier :

Arbitrary string of variable length from 0 to 80 octets used to describe the attachment circuit to the remote PE node.

4.3. Control Plane Operations

The deployment model shown in figure 3 of [RFC8214] does equally apply to the operations defined in this document.

4.3.1. VPWS Setup and Teardown

After an attachment circuit has been configured to be part of a VPWS instance and has not declared any local defect, the PE node announces his endpoint using a per-EVI ethernet A-D route to other PEs in the PSN via BGP. The Ethernet Tag ID is set to the VPWS instance identifier and the BGP PLE attribute is included to carry mandatory and optional bit-stream specific attachment circuit parameters.

Both endpoints receiving the EVPN per-EVI A-D route, validate the end to end connectivity by comparing BGP PLE attributes. Upon successful validation, the VPWS instance comes up and traffic can flow through the PSN. In the scenario where the validation phase fails, the remote PE reachability information is simply ignored and dismissed as a destination candidate. The VPWS instance validation is performed as follow:

- The mandatory PW type parameter MUST be identical
- The mandatory PLE/CEP/TDM Bit-rate parameter MUST be identical. This MAY be skipped if this parameter was not signaled because the attachment circuit rate can be unambiguously derived from the PW type [RFC5287].
- For CEP and PLE, the mandatory CEP/PLE Type parameter signalled via the CEP/PLE Options TLV MUST be identical
- If the payload size was signalled via the optional PLE/CEP/TDM Payload Bytes TLV it MUST be identical and supported by the PE node. Else the default payload size MUST be assumed.
- If any of the previous statements is no true or any of the signal CEP/PLE or TDM options is not supported by the PE node, the VPWS instance must stay down and a appropriate defect MUST be declared.

PLE is structure agnostic for SONET/SDH service types and hence can not validate whether a mix of SONET and SDH attachment circuits are connected (by incident) via VPWS. The detection of such misconfiguration is the responsibility of the operator managing the CE nodes.

In case of multi-homed CEs the mechanisms defined in [RFC8214] apply but are limited to the single-active and port-active scenarios.

Whenever the VPWS instance configuration is removed, the PE node MUST widthdraw its associated per-EVI ethernet A-D route.

4.3.2. Misconnection Handling

In circuit switched networks it is a common requirement to have the ability to check if the correct two endpoints got connected via a circuit. To confirm that the established bit-stream VPWS service is connecting the appropriate pair of attachment circuits, a Endpoint-ID string MAY be configured on each attachment circuit and communicated to the peer PE node using the Endpoint-ID TLV defined in Section 4.2.6.

Each endpoint MAY be configured to compare the Endpoint-ID received from the peer PE node to a locally configured expected Endpoint-ID and raise a fault (defect) when the IDs don't match. When a fault is raised, the R bit in the control word must bet set to 1 (backward defect indication) for the VPWS packets sent to the peer PE node. Each endpoint MAY be configured to only compare and report mismatches, but not to raise a fault.

4.3.3. Failure Scenarios

4.3.3.1. Single-homed CEs

Whenever a attachment circuit does declare a local fault the following operations MUST happen:

- Operations defined in [RFC4553], [RFC5086], [RFC4842] and [PLE] MUST happen
- The per-EVI ethernet A-D route MAY be withdrawn

Whenever the CE-bound IWF does enter packet loss state the operations defined in [RFC4553], [RFC5086], [RFC4842] and [PLE] MUST happen.

4.3.3.2. Multi-homed CEs

Figure 9 demonstrates a multi-homing scenario. CE1 is connected to PE1 and PE2 where PE1 is the designated forwarder while PE2 is the non designated forwarder.

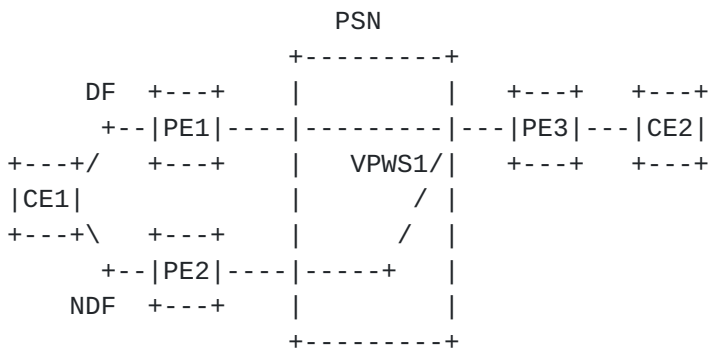

Figure 9: EVPN-VPWS multi-homing redundancy

In Figure 9 PE1 and PE2 are configured for single-active load-balancing mode. Both PEs are advertising a per-ES ethernet A-D route with the same non-zero Ethernet Segment (ES) value and the single-active bit set. This non-zero ES value is called Ethernet Segment Identifier (ESI).

In this example PE1 is elected as Designated Forwarder (DF) for the shared ESI where as PE2 is the Non-Designated Forwarder (NDF) for that segment. The signalling of primary / backup follows exactly the procedure defined in [RFC8214] where P and B bits of the layer 2 attribute extended community are used to settle proper connectivity.

Upon link failure between CE1 and PE1, PE1 and PE2 follows EVPN Ethernet Segment DF Election procedures described in [RFC8214] and [pref_df] for EVPN-VPWS. PE1 leverage mass-withdraw mechanism to tell PE3 to steer traffic over backup connectivity. The per-EVI ethernet A-D route advertisement remains intact. The main purpose is to keep reachability information available for fast convergence purpose. Therefore, the per-EVI ethernet A-D route MAY be withdrawn only under local fault and MUST be withdraw when the circuit is un-configured.

Port-active operation happens in the same way as single-active load-balancing mode described before but at the port level instead of being at the sub-interface level.

5. VPWS signalling using LDP

This section is already under construction and will be soon be publicly announced

6. IANA Considerations

This document defines a new BGP path attribute known as the BGP PLE attribute. IANA is requested to assign attribute code type TBD2 to the BGP PLE attribute from the "BGP Path Attributes" registry.

This document defines a new PW Type for PLE VPWS. IANA is requested to assign a PW type value TBD1 from the "MPLS Pseudowire Types" registry.

7. Security Considerations

The same Security Considerations described in [RFC8214] and [RFC5287] are valid for this document.

8. Acknowledgements

9. References

9.1. Normative References

[PLE] Schmutzer, C., "draft-cisco-bess-mpls-ple-00", 2020.

[pref_df] IETF, "Preference-based EVPN DF Election", <https://tools.ietf.org/html/draft-ietf-bess-evpn-pref-df>.

[RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, DOI 10.17487/RFC4446, April 2006, <https://www.rfc-editor.org/info/rfc4446>.

[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>.

[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>.

[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>.

[RFC5287] Vainshtein, A. and Y(J). Stein, "Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks", RFC 5287, DOI 10.17487/RFC5287, August 2008, <https://www.rfc-editor.org/info/rfc5287>.

[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>.

[srv6_netprog] IETF, "SRv6 Network Programming", <https://tools.ietf.org/html/draft-ietf-spring-srv6-network-programming>.

[srv6_overlay] IETF, "SRv6 BGP based Overlay services", <https://tools.ietf.org/html/draft-ietf-bess-srv6-services>.

9.2. Informative References

[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>.

[RFC4277] McPherson, D. and K. Patel, "Experience with the BGP-4 Protocol", RFC 4277, DOI 10.17487/RFC4277, January 2006, <https://www.rfc-editor.org/info/rfc4277>.

[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>.

Authors' Addresses

Steven Gringeri
Verizon

Email: steven.gringeri@verizon.com

Jeremy Whittaker
Verizon

Email: jeremy.whittaker@verizon.com

Christian Schmutzer (editor)
Cisco Systems, Inc.
Vienna
Austria

Email: cschmutz@cisco.com

Patrice Brissette
Cisco Systems, Inc.
Ottawa, ON
Canada

Email: pbrisset@cisco.com