

Internet Engineering Task Force (IETF)
Request for Comments: 9744
Category: Standards Track
ISSN: 2070-1721

A. Sajassi, Ed.
P. Brissette
Cisco Systems
J. Uttaro
Individual
J. Drake
Independent
S. Boutros
Ciena
J. Rabadan
Nokia
March 2025

EVPN Virtual Private Wire Service (VPWS) Flexible Cross-Connect (FXC) Service

Abstract

This document describes a new EVPN Virtual Private Wire Service (VPWS) service type specifically for multiplexing multiple attachment circuits across different Ethernet Segments (ESs) and physical interfaces into a single EVPN-VPWS service tunnel and still providing Single-Active and All-Active multi-homing. This new service is referred to as the EVPN-VPWS Flexible Cross-Connect (FXC) service. This document specifies a solution based on extensions to EVPN-VPWS (RFC 8214), which in turn is based on extensions to EVPN (RFC 7432). Therefore, a thorough understanding of RFCs 7432 and 8214 are prerequisites for this document.

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/rfc9744>.

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

- 1.1. Terminology
- 1.2. Requirements Language
- 2. Requirements
- 3. Solution
 - 3.1. VPWS Service Identifiers
 - 3.2. Default Flexible Cross-Connect
 - 3.2.1. Multi-homing
 - 3.3. VLAN-Signaled FXC
 - 3.3.1. Local Switching
 - 3.4. Service Instantiation
- 4. BGP Extensions
- 5. Failure Scenarios
 - 5.1. EVPN-VPWS Service Failure
 - 5.2. Attachment Circuit Failure
 - 5.3. PE Port Failure
 - 5.4. PE Node Failure
- 6. Security Considerations
- 7. IANA Considerations
- 8. References
 - 8.1. Normative References
 - 8.2. Informative References
- Contributors
- Authors' Addresses

1. Introduction

[RFC8214] describes a solution to deliver point-to-point (P2P) services using BGP constructs defined in [RFC7432]. It delivers this P2P service between a pair of Attachment Circuits (ACs), where an AC on a PE can represent a port, a VLAN on a port, or a group of VLANs on a port. It also leverages multi-homing and fast convergence capabilities of [RFC7432] in delivering these VPWS services. Multi-homing capabilities include the support of Single-Active and All-Active redundancy mode, and fast convergence is provided using a "mass withdrawal" message in control plane and fast protection switching using prefix-independent convergence in a data plane upon node or link failure [BGP-PIC]. Furthermore, the use of EVPN BGP constructs eliminates the need for multi-segment pseudowire auto-discovery and signaling if the VPWS service need to span across multiple Autonomous Systems (ASes) [RFC5659].

Service providers have a very large number of ACs (in millions) that need to be backhauled across their MPLS/IP network. These ACs may or may not require tag manipulation (e.g., VLAN translation). These service providers want to multiplex a large number of ACs across several physical interfaces spread across one or more PEs (e.g., several ESs) onto a single VPWS service tunnel in order to a) reduce the number of EVPN service labels associated with EVPN-VPWS service tunnels and thus the associated Operations, Administration, and Maintenance (OAM) monitoring and b) reduce EVPN BGP signaling (e.g., not to signal each AC as it is the case in [RFC8214]).

Service providers want the above functionality without sacrificing any of the capabilities of [RFC8214] including Single-Active and All-Active multi-homing and fast convergence.

This document specifies a solution based on extensions to EVPN-VPWS [RFC8214] to meet the above requirements. Furthermore, [RFC8214] is itself based on extensions to EVPN [RFC7432]. Therefore, a thorough understanding of [RFC7432] and [RFC8214] are prerequisites for this document.

1.1. Terminology

AC: Attachment Circuit

ES: Ethernet Segment

ESI: Ethernet Segment Identifier

EVI: EVPN Instance Identifier

EVPN: Ethernet Virtual Private Network

Ethernet A-D: Ethernet Auto-Discovery (per EVI or per Ethernet A-D per ESI routes, as defined in [RFC7432] and [RFC8214])

FXC: Flexible Cross-Connect

MAC: Media Access Control

MPLS: Multiprotocol Label Switching

OAM: Operations, Administration, and Maintenance

PE: Provider Edge

VCCV: Virtual Circuit Connectivity Verification

VID: VLAN ID

VPWS: Virtual Private Wire Service

VRF: VPN Routing and Forwarding

IP-VRF: VPN Routing and Forwarding for IP lookup

MAC-VRF: VPN Routing and Forwarding for MAC lookup

VID-VRF: VPN Routing and Forwarding for normalized VID lookup

VPWS Service Tunnel: It is represented by a pair of EVPN service labels associated with a pair of endpoints. Each label is downstream assigned and advertised by the disposition PE through an Ethernet A-D per EVI route. The downstream label identifies the endpoint on the disposition PE. A VPWS service tunnel can be associated with many VPWS service identifiers where each identifier is a normalized VID.

Single-Active Redundancy Mode: When a device or a network is multi-homed to two or more PEs and when only a single PE in such redundancy group can forward traffic to/from the multi-homed device or network for a given VLAN, then such multi-homing or redundancy is referred to as "Single-Active".

All-Active Redundancy Mode: When a device or a network is multi-homed to two or more PEs and when all PEs in such redundancy group can forward traffic to/from the multi-homed device or network for a given VLAN, then such multi-homing or redundancy is referred to as "All-Active".

1.2. 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.

2. Requirements

Two of the main motivations for service providers seeking a new

solution are: 1) to reduce the number of VPWS service tunnels by multiplexing a large number of ACs across different physical interfaces instead of having one VPWS service tunnel per AC and 2) to reduce the signaling of ACs as much as possible. Besides these two requirements, they also want the multi-homing and fast convergence capabilities of [RFC8214].

In [RFC8214], a PE signals an AC indirectly by first associating that AC to a VPWS service tunnel (e.g., a VPWS service instance) and then signaling the VPWS service tunnel via an Ethernet A-D per EVI route with the Ethernet Tag field set to a 24-bit VPWS service instance identifier (which is unique within the EVI) and the ESI field set to a 10-octet identifier of the ES corresponding to that AC.

Therefore, a PE device that receives such EVPN routes can associate the VPWS service tunnel to the remote ES using the ESI field. Additionally, when the remote ES fails and the PE receives the "mass withdrawal" message associated with the failed ES per [RFC7432], a PE device can quickly update its next-hop adjacency list (adjacency list) for all VPWS service tunnels sharing the ESI field and achieve fast convergence for multi-homing scenarios. Even if fast convergence were not needed, there would still be a need for signaling each AC failure (via its corresponding VPWS service tunnel) associated with the failed ES so that the adjacency list gets updated and the packets are sent to a backup PE (in case of Single-Active multi-homing) or to other PEs in the redundancy group (in case of All-Active multi-homing). In the absence of updating the adjacency list properly, the traffic for that VPWS service tunnel will be dropped by the egress PE with a failed ES/AC.

When a single VPWS service tunnel carries multiple ACs across various ESs (physical interfaces) without signaling the ACs via EVPN BGP to remote PE devices, those remote PE devices lack the information to associate the received ES with these ACs or with their local ACs. They also lack the association between the VPWS service tunnel (e.g., EVPN service label) and the far-end ACs. This means that, while the remote PEs can associate their local ACs with the VPWS service tunnel, they cannot make similar associations for the far-end ACs.

Consequently, in case of a connectivity failure to the ES, the remote PEs are unable to redirect traffic via another multi-homing PE to that ES. In other words, even if an ES failure is signaled via EVPN to the remote PE devices, they cannot effectively respond because they do not know the relationship between the remote ES, the remote ACs, and the VPWS service tunnel.

To address this issue when multiplexing a large number of ACs onto a single VPWS service tunnel, two mechanisms have been developed: one to support VPWS services between two single-homed endpoints and another one to support VPWS services where one of the endpoints is multi-homed.

For single-homed endpoints, it is acceptable not to signal each AC in BGP because, in the event of a connection failure to the ES, there is no alternative path to that endpoint. However, the implication of not signaling an AC failure is that the traffic destined for the failed AC is sent over the MPLS/IP core and then discarded at the destination PE, thereby potentially wasting network resources.

This waste of network resources during a connection failure may be transient, as it can be detected and prevented at the application layer in certain cases. Section 3.2 outlines a solution for such single-homing VPWS services.

For VPWS services where one of the endpoints is multi-homed, there are two options:

1. Signal each AC via BGP, allowing the adjacency list to be updated upon a failure affecting those ACs. This solution is described in Section 3.3 and is referred to as the "VLAN-signaled FXC service".
2. Bundle several ACs on an ES together per destination endpoint (e.g., ES, MAC-VRF, etc.) and associate such a bundle with a single VPWS service tunnel. This approach is similar to the VLAN bundle service interface described in [RFC8214]. This solution is described in Section 3.2.1.

3. Solution

This section specifies how to provide a new VPWS service between two PE devices where a large number of ACs (such as VLANs) that span across multiple ESs (physical interfaces) on each PE are multiplexed onto a single P2P EVPN service tunnel. Since the multiplexing involves several physical interfaces, there can be overlapping VLAN IDs (VIDs) across these interfaces. In such cases, the VIDs must be translated into unique VIDs to prevent collisions. Furthermore, if the number of VLANs being multiplexed onto a single VPWS service tunnel exceeds 4095, then a single tag to double tag translation must be performed. This translation of VIDs into unique VIDs (either single or double) results in a "normalized VID".

When a single normalized VID is used, the lower 12 bits of the Ethernet Tag ID field in EVPN routes MUST be set to that VID. When a double normalized VID is used, the lower 12 bits of the Ethernet Tag ID field MUST be set to the inner VID, while the higher 12 bits are set to the outer VID. 24-bit VPWS service instance identifiers [RFC8214] as well as 12-bit VPWS service instance identifiers representing normalized VIDs MUST be right-aligned.

Since there is only a single EVPN-VPWS service tunnel associated with many normalized VIDs (either single or double) across multiple physical interfaces, an MPLS lookup at the disposition PE is no longer sufficient to forward the packet to the correct egress endpoint or interface. Therefore, in addition to an EVPN label lookup corresponding to the VPWS service tunnel, a VID lookup (either single or double) is also required. At the disposition PE, the EVPN label lookup identifies a VID-VRF, and the lookup of the normalized VIDs within that table identifies the appropriate egress endpoint or interface. The tag manipulation (translation from normalized VIDs to the local VID) SHOULD be performed either as part of the VID table lookup or at the egress interface itself.

Since the VID lookup (single or double) needs to be performed at the disposition PE, VID normalization MUST be completed prior to MPLS encapsulation on the ingress PE. This requires that both the imposition and disposition PE devices be capable of VLAN tag manipulation, such as rewriting (single or double), adding, or deleting (single or double) at their endpoints (e.g., their ESs, MAC-VRFs, IP-VRFs, etc.). Operators should be informed of potential trade-offs from a performance standpoint, compared to typical pseudowire processing.

3.1. VPWS Service Identifiers

In [RFC8214], a unique value identifying the service is signaled in the context of each PE's EVI. The 32-bit Ethernet Tag ID field MUST be set to this VPWS service instance identifier value. Translation at an Autonomous System Border Router (ASBR) is needed if re-advertising to another AS affects uniqueness.

For FXC, this same Ethernet Tag ID field value is an identifier that

may represent:

- * VLAN Bundle Service Interface: a unique value for a group of VLANs
- * VLAN-Aware Bundle Service Interface: a unique value for individual VLANs and is considered the same as the normalized VID

Both the VPWS service instance identifier and normalized VID are carried in the Ethernet Tag ID field of the Ethernet A-D per EVI route. For FXC, in the case of a 12-bit ID, the VPWS service instance identifier is the same as the single-tag normalized VID and will be the same on both VPWS service endpoints. Similarly in the case of a 24-bit ID, the VPWS service instance identifier is the same as the double-tag normalized VID.

3.2. Default Flexible Cross-Connect

In this mode of operation, many ACs across several ESs are multiplexed into a single EVPN-VPWS service tunnel represented by a single VPWS service ID. This is the default mode of operation for FXC, and the participating PEs do not need to signal the VLANs (normalized VIDs) in EVPN BGP.

Regarding the data plane aspects of this solution, the imposition PE performs VID normalization, and the disposition PE carries out VID lookup and translation. Both imposition and disposition PE devices MUST be aware of the VLANs through configuration. There should ideally be a single point-to-point (P2P) EVPN-VPWS service tunnel between a pair of PEs for a specific set of ACs.

As previously mentioned, because the EVPN-VPWS service tunnel is employed to multiplex ACs across various ESs or physical interfaces, the EVPN label alone is not sufficient for accurate forwarding of the received packets over the MPLS/IP network to egress interfaces. Therefore, normalized VID lookup is REQUIRED in the disposition direction to forward packets to their proper egress endpoints; the EVPN label lookup identifies a VID-VRF, and a subsequent normalized VID lookup in that table identifies the egress interface.

In this solution, for each PE, the single-homing ACs represented by their normalized VIDs are associated with a single VPWS service instance within a specific EVI. The generated EVPN route is an Ethernet A-D per-EVI route with an ESI of 0, the Ethernet Tag field is set to a VPWS service instance ID, and the MPLS label field is set to a dynamically generated EVPN service label representing the EVPN-VPWS service tunnel. This route is sent with a Route Target (RT) that represents the EVI, which can be auto-generated from the EVI according to Section 5.1.2.1 of [RFC8365]. Additionally, this route is sent with the EVPN Layer 2 Attributes Extended Community defined in Section 3.1 of [RFC8214] with two new flags (outlined in Section 4) that indicate: 1) the VPW service tunnel (set to default Flexible Cross-Connect) and 2) the normalized VID type (set to normalized single VID or double VID). The receiving PE uses these new flags for a consistency check and MAY generate an alarm if it detects inconsistencies, but it will not disrupt the VPWS service.

It should be noted that in this mode of operation, a single Ethernet A-D per-EVI route is transmitted upon the configuration of the first AC with the normalized VID. As additional ACs are configured and associated with this EVPN-VPWS service tunnel, the PE does not advertise any additional EVPN BGP routes and only locally associates these ACs with the pre-established VPWS service tunnel.

3.2.1. Multi-homing

The default FXC mode can also be used for multi-homing. In this

mode, a group of normalized VIDs representing ACs on a single ES, all destined to a single endpoint, are multiplexed into a single EVPN-VPWS service tunnel, which is identified by a unique VPWS service ID. When employing the default FXC mode for multi-homing, rather than using a single EVPN-VPWS service tunnel, there may be multiple service tunnels per pair of PEs. Specifically, there is one tunnel for each group of VIDs per pair of PEs, and there can be many such groups between a pair of PEs, resulting in numerous EVPN service tunnels.

3.3. VLAN-Signaled FXC

In this mode of operation, similar to the default FXC mode described in Section 3.2, many normalized VIDs representing ACs across several ESs and interfaces are multiplexed into a single EVPN-VPWS service tunnel. However, this single tunnel is represented by multiple VPWS service IDs (one per normalized VID), and these normalized VIDs are signaled using EVPN BGP.

In this solution, on each PE, the multi-homing ACs represented by their normalized VIDs are configured with a single EVI. There is no need to configure a separate VPWS service instance ID in here, as it corresponds to the normalized VID. For each normalized VID on each ES, the PE generates an Ethernet A-D per-EVI route where the ESI field represents the ES ID, the Ethernet Tag field is set to the normalized VID, and the MPLS label field is set to a dynamically generated EVPN label representing the P2P EVPN service tunnel. This label is the same for all ACs multiplexed into a single EVPN-VPWS service tunnel. This route is sent with an RT representing the EVI. As before, this RT can be auto-generated from the EVI per Section 5.1.2.1 of [RFC8365]. Additionally, this route includes the EVPN Layer 2 Attributes Extended Community defined in Section 3.1 of [RFC8214] with two new flags (outlined in Section 4) that indicate: 1) this VPWS service tunnel for VLAN-signaled FXC and 2) the normalized VID type (single versus double). The receiving PE uses these new flags for a consistency check and may generate an alarm if it detects inconsistency, but it will not disrupt the VPWS service.

It should be noted that in this mode of operation, the PE sends a single Ethernet A-D per-EVI route for each AC that is configured. Each normalized VID that is configured per ES results in generation of an Ethernet A-D per EVI.

This mode of operation enabled automatic cross-checking of normalized VIDs used for Ethernet Virtual Private Line (EVPL) services because these VIDs are signaled in EVPN BGP. For instance, if the same normalized VID is configured on three PE devices (instead of two) for the same EVI, then when a PE receives the second remote Ethernet A-D per EVI route, it generates an error message unless the two Ethernet A-D per EVI routes include the same ESI. Such cross-checking is not feasible in the default FXC mode because the normalized VIDs are not signaled.

3.3.1. Local Switching

When cross-connection occurs between two ACs belonging to two multi-homed ESs on the same set of multi-homing PEs, the forwarding between the two ACs must be performed locally during normal operation (e.g., in absence of a local link failure). This means that traffic between the two ACs MUST be locally switched within the PE.

In terms of control plane processing, this means that when the receiving PE processes an Ethernet A-D per-EVI route whose ESI is a local ESI, the PE does not modify its forwarding state based on the received route. This approach ensures that local switching takes precedence over forwarding via the MPLS/IP network. This method of

prioritizing locally switched traffic aligns with the baseline EVPN principles described in [RFC7432], where locally switched preference is specified for MAC/IP routes.

In such scenarios, the Ethernet A-D per-EVI route should be advertised with the MPLS label either associated with the destination AC or with the destination ES in order to avoid any ambiguity in forwarding. In other words, the MPLS label cannot represent the same VID-VRF outlined in Section 3.3, as the same normalized VID can be reachable via two ESs. In the case of using an MPLS label per destination AC, this approach can also be applied to VLAN-based VPWS or VLAN bundle VPWS services as per [RFC8214].

3.4. Service Instantiation

The V field defined in Section 4 is OPTIONAL. However, if transmitted, its value may indicate an error condition that could lead to operational issues. In such cases, merely notifying the operator of an error is insufficient; the VPWS service tunnel must not be established.

If both endpoints of a VPWS tunnel are signaling a matching normalized VID in the control plane, but one is operating in single-tag mode and the other in double-tag mode, then the signaling of the V-bit facilitates the detection and prevention of this tunnel's instantiation.

If single VID normalization is signaled in the Ethernet Tag ID field (12 bits), yet the data plane is operating based on double tags, the VID normalization applies only to the outer tag. Conversely, if double VID normalization is signaled in the Ethernet Tag ID field (24 bits), VID normalization applies to both the inner and outer tags.

4. BGP Extensions

This document uses the EVPN Layer 2 Attributes Extended Community as defined in [RFC8214] with two additional flags incorporated into this Extended Community (EC) as detailed below. This EC is sent with Ethernet A-D per-EVI route per Section 3 and SHOULD be sent for both Single-Active and All-Active redundancy modes.

```

+-----+
| Type (0x06) / Sub-type (0x04) (2 octets) |
+-----+
| Control Flags (2 octets)                  |
+-----+
| L2 MTU (2 octets)                        |
+-----+
| Reserved (2 octets)                      |
+-----+

          1 1 1 1 1 1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
| MBZ          | V | M | | C | P | B |      (MBZ = MUST Be Zero)
+---+---+---+---+---+---+---+---+---+---+

```

The following bits in the Control Flags are defined; the remaining bits MUST be set to zero when sending and MUST be ignored when receiving this community.

```

+=====+
| Name   | Meaning                                     |
+=====+
| B,P,C | per definition in [RFC8214]                |
+-----+

```


M	00 mode of operation as defined in [RFC8214]
	01 VLAN-Signaled FXC
	10 Default FXC
V	00 operating per [RFC8214]
	01 single-VID normalization
	10 double-VID normalization

Table 1

The M and V fields are OPTIONAL. The M field is ignored at reception for forwarding purposes and is used for error notifications.

5. Failure Scenarios

The two following examples analyze the failure scenarios.

The first scenario is a default Flexible Cross-Connect with a multi-homing solution, and it is depicted in Figure 1. In this case, VID normalization is performed, and a single Ethernet A-D per-EVI route is sent for the bundle of ACs on an ES. That is, PE1 will advertise two Ethernet A-D per-EVI routes: The first one will identify the ACs on port p1's ES, and the second one will identify the AC2 in port p2's ES. Similarly, PE2 will advertise two Ethernet A-D per-EVI routes.

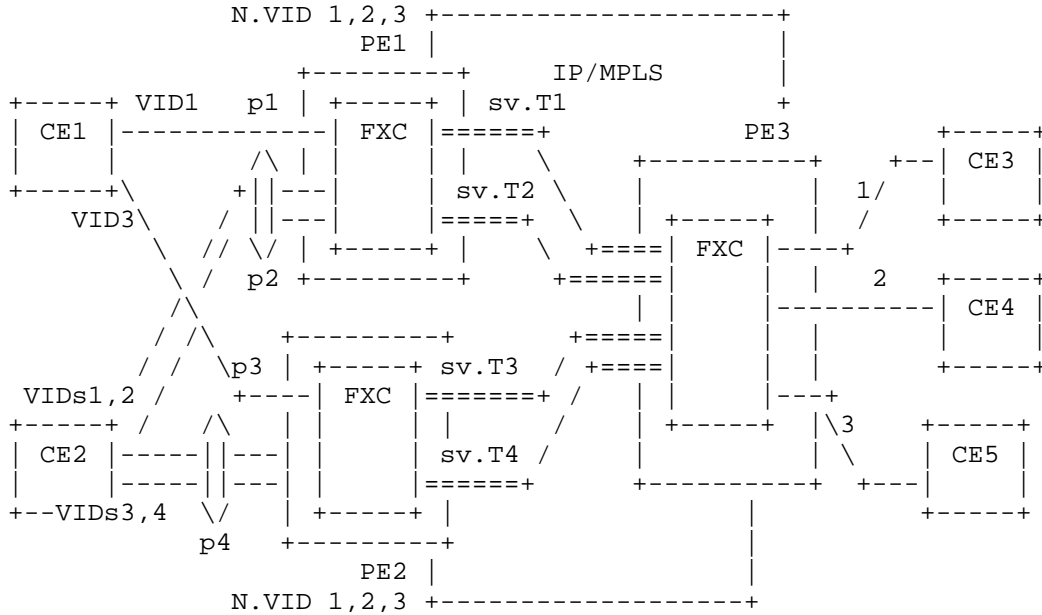


Figure 1: Default Flexible Cross-Connect

The second scenario, depicted in Figure 2, illustrates the VLAN-signaled FXC mode with multi-homing. In this example:

- * CE1 is connected to PE1 and PE2 via (port,VID)=(p1,1) and (p3,3), respectively. CE1's VIDs are normalized to value 1 on both PEs, and CE1 is cross-connected to CE3's VID 1 at the remote end.
- * CE2 is connected to PE1 and PE2 via ports p2 and p4, respectively:
 - The combinations (p2,1) and (p4,3) identify the ACs used to cross-connect CE2 to CE4's VID 2 and are normalized to value 2.

- The combinations (p2,2) and (p4,4) identify the ACs used to cross-connect CE2 to CE5's VID 3 and are normalized to value 3.

In this scenario, PE1 and PE2 advertise an Ethernet A-D per EVI route for each normalized VID (values 1, 2, and 3). However, only two VPWS Service Tunnels are required: 1) VPWS Service Tunnel 1 (sv.T1) between PE1's FXC and PE3's FXC and 2) VPWS Service Tunnel 2 (sv.T2) between PE2's FXC and PE3's FXC.

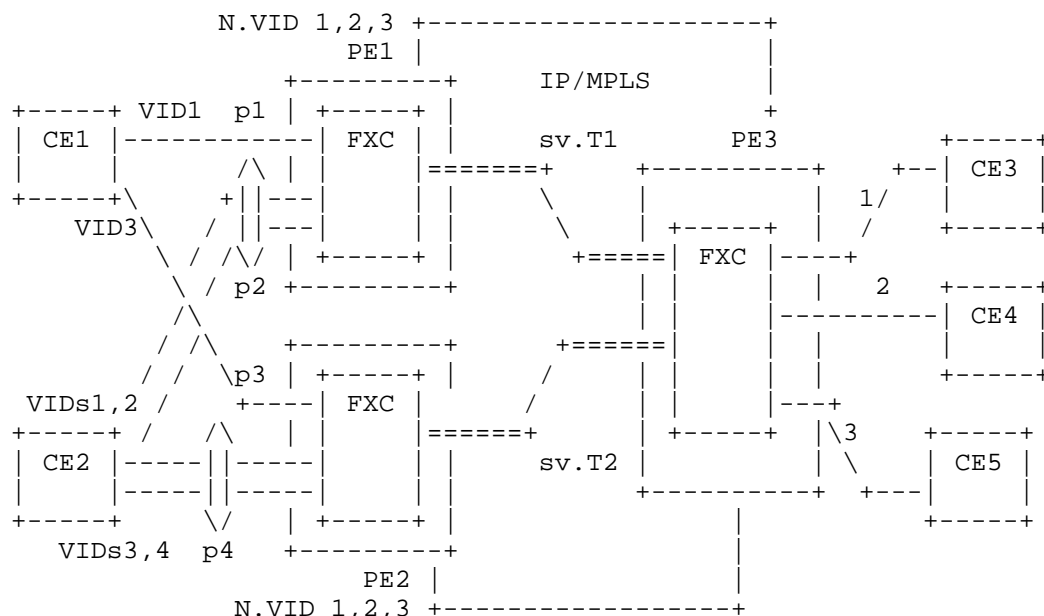


Figure 2: VLAN-Signaled FXC

5.1. EVPN-VPWS Service Failure

The failure detection of an EVPN-VPWS service can be performed via OAM mechanisms such as Bidirectional Forwarding Detection (BFD) for the pseudowire Virtual Circuit Connectivity Verification (VCCV) [RFC5885], and upon such failure detection, the switch over procedure to the backup PE is the same as the one described above.

5.2. Attachment Circuit Failure

In the event of an AC failure, the VLAN-Signaled and default FXC modes exhibit distinct behaviors:

- * Default FXC (Figure 1): In the default mode, a VLAN or AC failure is not signaled. Consequently, in case of an AC failure, such as VID1 on CE2, there is nothing to prevent PE3 from directing traffic from CE4 to PE1, leading to a potential packet loss at the egress PE with a failed AC. Application layer OAM may be utilized if per-VLAN fault propagation is necessary in this scenario.
- * VLAN-Signaled FXC (Figure 2): In the case of a VLAN or AC failure, such as VID1 on CE2, this triggers the withdrawal of the Ethernet A-D per-EVI route for the corresponding Normalized VID, specifically Ethernet-Tag 2. Upon receiving the route withdrawal, PE3 will remove PE1 from its outgoing adjacency list for traffic originating from CE4.

5.3. PE Port Failure

In the event of a PE port failure, the failure will be signaled, and the other PE will assume forwarding in both scenarios:

- * Default FXC (Figure 1): In the case of a port failure, such as p2, the route for Service Tunnel 2 (sv.T2) will be withdrawn. Upon receiving the fault notification, PE3 will remove PE1 from its adjacency list for traffic originating from CE4 and CE5.
- * VLAN-Signaled FXC (Figure 2): A port failure, such as p2, triggers the withdrawal of the Ethernet A-D per EVI routes for normalized VIDs 2 and 3, along with the withdrawal of the Ethernet A-D per-ES route for p2's ES. Upon receiving the fault notification, PE3 will remove PE1 from its adjacency list for the traffic originating from CE4 and CE5.

5.4. PE Node Failure

In the case of PE node failure, the operation is similar to the steps described above, albeit that EVPN route withdrawals are performed by the route reflector instead of the PE.

6. Security Considerations

Since this document describes a muxing capability that leverages EVPN-VPWS signaling, no additional functionality beyond the muxing service is added, and thus no additional security considerations are needed beyond what is already specified in [RFC8214].

7. IANA Considerations

This document has allocated bits 8-11 in the "EVPN Layer 2 Attributes Control Flags" registry with names M and V:

M Signaling mode of operation (bits 10-11)
 V VLAN-ID normalization (bits 8-9)

8. References

8.1. Normative References

- [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>>.
- [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>>.

8.2. Informative References

- [BGP-PIC] Bashandy, A., Ed., Filsfils, C., and P. Mohapatra, "BGP Prefix Independent Convergence", Work in Progress, Internet-Draft, draft-ietf-rtgwg-bgp-pic-21, 7 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-bgp-pic-21>>.
- [RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", RFC 5659,

DOI 10.17487/RFC5659, October 2009,
<<https://www.rfc-editor.org/info/rfc5659>>.

[RFC5885] Nadeau, T., Ed. and C. Pignataro, Ed., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, DOI 10.17487/RFC5885, June 2010,
<<https://www.rfc-editor.org/info/rfc5885>>.

[RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018,
<<https://www.rfc-editor.org/info/rfc8365>>.

Contributors

In addition to the authors listed on the front page, the following coauthors have also contributed substantially to this document:

Wen Lin
Juniper Networks
Email: wlin@juniper.net

Luc Andre Burdet
Cisco
Email: lburdet@cisco.com

Authors' Addresses

Ali Sajassi (editor)
Cisco Systems
Email: sajassi@cisco.com

Patrice Brissette
Cisco Systems
Email: pbrisset@cisco.com

James Uttaro
Individual
Email: juttaro@ieee.org

John Drake
Independent
Email: je_drake@yahoo.com

Sami Boutros
Ciena
Email: sboutros@ciena.com

Jorge Rabadan
Nokia
Email: jorge.rabadan@nokia.com