

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 5, 2015

J. Dong
H. Wang
Huawei Technologies
May 4, 2015

Pseudowire Redundancy on S-PE
draft-ietf-pals-redundancy-spe-01

Abstract

This document describes Multi-Segment Pseudowire (MS-PW) protection scenarios in which the pseudowire redundancy is provided on the Switching-PE (S-PE). Operations of the S-PEs which provide PW redundancy are specified in this document. Signaling of the preferential forwarding status as defined in [RFC 6870] is reused. This document does not require any change to the T-PEs of MS-PW.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 <http://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 5, 2015.

Copyright Notice

Copyright (c) 2015 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

(<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Typical Scenarios of PW Redundancy on S-PE	2
2.1.	MS-PW Redundancy on S-PE	3
2.2.	MS-PW Redundancy on S-PE with S-PE Protection	3
3.	S-PE Operations	4
3.1.	Operations of Scenario 1	5
3.2.	Operations of Scenario 2	6
4.	VCCV Considerations	7
5.	IANA Considerations	7
6.	Security Considerations	7
7.	Acknowledgements	7
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	8
	Authors' Addresses	8

[1.](#) Introduction

[RFC6718] describes the framework and requirements for pseudowire (PW) redundancy, and [RFC6870] specifies Pseudowire (PW) redundancy mechanism for scenarios where a set of redundant PWs is configured between provider edge (PE) nodes in single-segment pseudowire (SS-PW) [RFC3985] applications, or between terminating provider edge (T-PE) nodes in multi-segment pseudowire (MS-PW) [RFC5659] applications.

In some MS-PW scenarios, there are benefits to provide PW redundancy on S-PEs, such as reducing the burden on the access T-PE nodes, and faster protection switching. This document describes some scenarios in which PW redundancy is provided on S-PEs, and specifies the operations of the S-PEs. Signaling of the preferential forwarding status as defined in [RFC6870] is reused. This document does not require any change to the T-PEs of MS-PW.

2. Typical Scenarios of PW Redundancy on S-PE

In some MS-PW deployment scenarios, there are benefits to provide PW redundancy on S-PEs. This section describes typical scenarios of PW redundancy on S-PE.

2.1. MS-PW Redundancy on S-PE

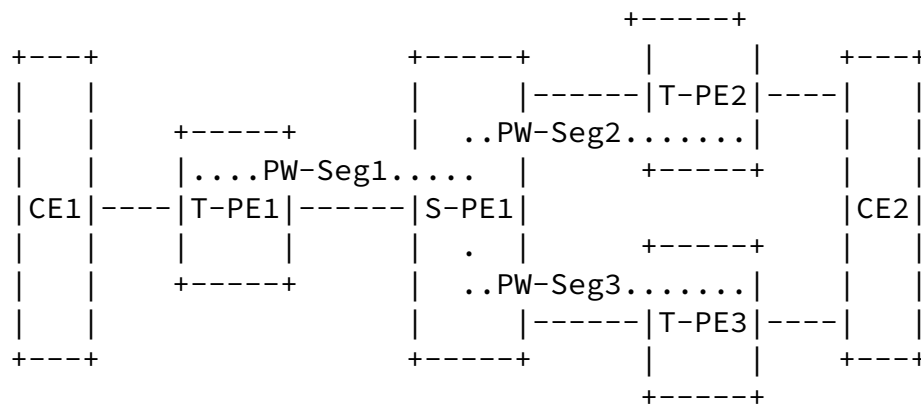


Figure 1. MS-PW Redundancy on S-PE

As illustrated in Figure 1, CE1 is connected to T-PE1 while CE2 is dual-homed to T-PE2 and T-PE3. T-PE1 is connected to S-PE1 only, and S-PE1 is connected to both T-PE2 and T-PE3. The MS-PW is switched on S-PE1, and PW-Seg2 and PW-Seg3 provides resiliency on S-PE1 for failure of T-PE2 or T-PE3 or the connected ACs. PW-Seg2 is selected as the primary PW segment, and PW-Seg3 is the secondary PW segment.

MS-PW redundancy on S-PE is beneficial for the scenario in Figure 1 since T-PE1 as an access node may not support PW redundancy. Besides, with PW redundancy on S-PE, the number of PW segments required between T-PE1 and S-PE1 is only half of the number of PW segments needed when end-to-end MS-PW redundancy is used. In addition, in this scenario PW redundancy on S-PE could provide faster protection switching, compared with end-to-end protection switching of MS-PW.

2.2. MS-PW Redundancy on S-PE with S-PE Protection

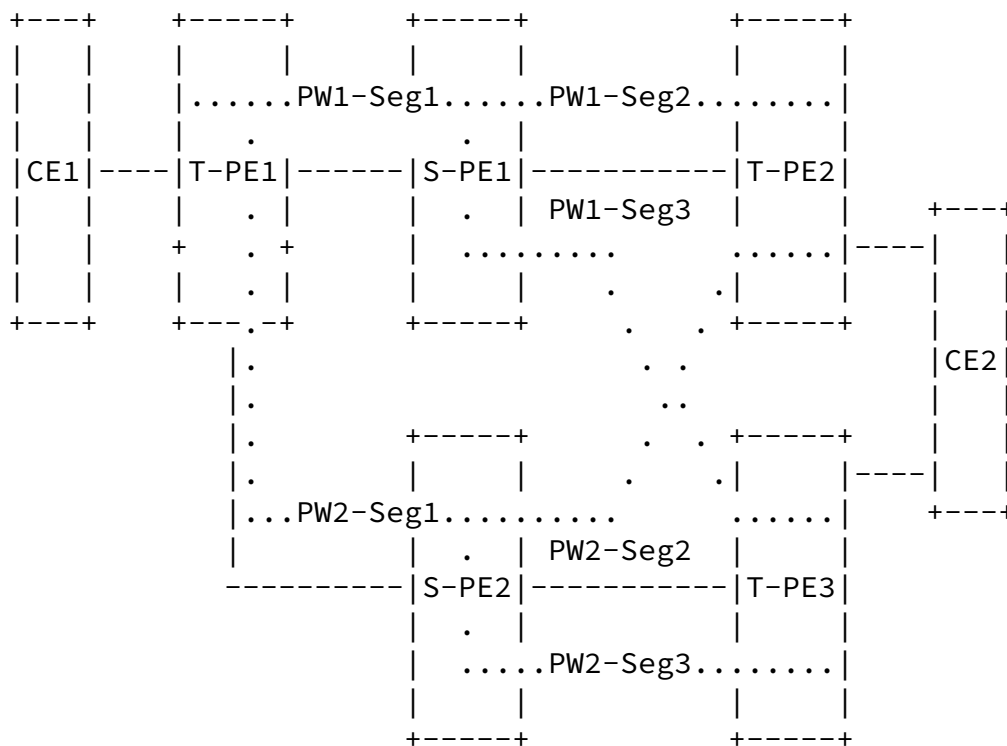


Figure 2. MS-PW Redundancy on S-PE with S-PE protection

As illustrated in Figure 2, CE1 is connected to T-PE1 while CE2 is dual-homed to T-PE2 and T-PE3. T-PE1 is connected to S-PE1 and S-PE2, and both S-PE1 and S-PE2 are connected to both T-PE2 and T-PE3. There are two MS-PWs which are switched at S-PE1 and S-PE2 respectively to provide S-PE node protection. For MS-PW1, S-PE1 provides resiliency using PW1-Seg2 and PW1-Seg3. For MS-PW2, S-PE2 provides resiliency using PW2-Seg2 and PW2-Seg3. MS-PW1 is the

primary PW and PW1-Seg2 between S-PE1 and T-PE2 is the primary PW segment. MS-PW2 is the secondary PW.

MS-PW redundancy on S-PE is beneficial for this scenario since it reduces the number of end-to-end MS-PWs required for both T-PE and S-PE protection. In addition, PW redundancy on S-PE could provide faster protection switching, compared with end-to-end protection switching of MS-PW.

3. S-PE Operations

For an S-PE which provides PW redundancy for MS-PW, it is important to advertise proper preferential forwarding status to the PW segments on both sides and perform protection switching according to the received status information. Note that when PW redundancy for MS-PW is provided on S-PE, the optional S-PE Bypass Mode as defined in [[RFC6478](#)] MUST NOT be used. This section specifies the operations of S-PEs on which PW redundancy is provisioned. This section does not make any change to the T-PEs of MS-PW.

The S-PEs connect to the neighboring T-PEs or other S-PEs on two sides with PW segments. For the S-PE which provides PW redundancy for an MS-PW, on one side there is a single PW segment, which is called the single-homed side, and on the other side there are multiple PW segments, which is called the multi-homed side. The scenario in which the S-PE has two multi-homed sides is out of scope.

In general, the S-PE MUST work as a Slave node for the single-homed side, and MUST work in Independent mode for the multi-homed side. The S-PE MUST pass the preferential forwarding status received from the single-homed side unchanged to the PW segments on the multi-homed side. The S-PE MUST advertise Standby status to the single-homed side if it receives Standby status from all the PW segments on the multi-homed side, and it MUST advertise Active status to the single-homed side if it receives Active status from any of the PW segments on the multi-homed side. For the single-homed side, the active PW segment is determined by the T-PE on this side, which works as the Master node. On the multi-homed side, the PW segment which has both local and remote Up/Down status and Preferential Forwarding status as Up and Active MUST be selected for traffic forwarding.

The Signaling of Preferential Forwarding bit as defined in [[RFC6870](#)]

and [\[RFC6478\]](#) is reused in these scenarios.

[3.1.](#) Operations of Scenario 1

For the scenario in Figure 1, assume the AC from CE2 to T-PE2 is active. In normal operation, S-PE1 would receive Active Preferential Forwarding status bit on the single-homed side from T-PE1, then it would advertise Active Preferential Forwarding status bit on both PW-Seg2 and PW-Seg3. T-PE2 and T-PE3 would advertise Active and Standby preferential status bit to S-PE1 respectively, reflecting the forwarding state of the two ACs connected to CE2. By matching the local and remote Up/Down status and Preferential Forwarding status, PW-Seg2 would be used for traffic forwarding.

On failure of the AC between CE2 and T-PE2, the forwarding state of AC on T-PE3 is changed to Active. T-PE3 then advertises Active Preferential Status to S-PE1, and T-PE2 would advertise a PW status Notification message to S-PE1, indicating that the AC between CE2 and T-PE2 is down. S-PE1 would perform the switchover according to the updated local and remote Preferential Forwarding status and the status of "Pseudowire forwarding", and select PW-Seg3 as the new PW Segment for traffic forwarding. Since S-PE1 still connects to an Active PW segment on the multi-homed side, it will not advertise any change of the PW status to T-PE1. If S-PE1 supports the SP-PE TLV processing as defined in [\[RFC6073\]](#), it SHOULD advertise the updated SP-PE TLVs by sending a Label Mapping message to T-PE1.

[3.2.](#) Operations of Scenario 2

For the scenario of Figure 2, assume the AC from CE2 to T-PE2 is active. T-PE1 works in Master mode and it would advertise Active and Standby Preferential Forwarding status bit respectively to S-PE1 and S-PE2 according to configuration. According to the received Preferential Forwarding status bit, S-PE1 would advertise Active Preferential Forwarding status bit to both T-PE2 and T-PE3, and S-PE2 would advertise Standby Preferential Forwarding status bit to both T-PE2 and T-PE3. T-PE2 would advertise Active Preferential Forwarding status bit to both S-PE1 and S-PE2, and T-PE3 would advertise Standby Preferential Forwarding status bit to both S-PE1 and S-PE2, reflecting the forwarding state of the two ACs connected to CE2. By matching the local and remote Up/Down Status and Preferential Forwarding status, PW1-Seg2 from S-PE1 to T-PE2 would be

used for traffic forwarding. Since S-PE1 connects to the Active PW segment on the multi-homed side, it would advertise Active Preferential Forwarding status bit to T-PE1, and S-PE2 would advertise Standby Preferential Forwarding status bit to T-PE1 since it does not have any Active PW segment on the multi-homed side.

On failure of the AC between CE2 and T-PE2, the forwarding state of AC on T-PE3 is changed to Active. T-PE3 would then advertise Active Preferential Forwarding status bit to both S-PE1 and S-PE2, and T-PE2 would advertise a PW status Notification message to both S-PE1 and S-PE2, indicating that the AC between CE2 and T-PE2 is down. S-PE1 would perform the switchover according to the updated local and remote Preferential Forwarding status and the status of "Pseudowire forwarding", and select PW1-Seg3 for traffic forwarding. Since S-PE1 still has an Active PW segment on the multi-homed side, it would not advertise any change of the PW status to T-PE1. If S-PE1 supports the SP-PE TLV processing as defined in [\[RFC6073\]](#), it SHOULD advertise the updated SP-PE TLVs by sending a Label Mapping message to T-PE1.

If S-PE1 fails, T-PE1 would notice this through some detection mechanism and then advertise the Active Preferential Forwarding status bit to S-PE2, and PW2-Seg1 would be selected by T-PE1 for traffic forwarding. On receipt of the newly changed Preferential Forwarding status, S-PE2 would advertise the Active Preferential Forwarding status to both T-PE2 and T-PE3. T-PE2 and T-PE3 would also notice the failure of S-PE1 by some detection mechanism. Then by matching the local and remote Up/Down and Preferential Forwarding status, PW2-Seg2 would be selected for traffic forwarding.

[4.](#) VCCV Considerations

PW VCCV [\[RFC5085\]](#) CC type 1 "PW ACH" can be used with S-PE redundancy mechanism. VCCV CC type 2 "Router Alert Label" is not supported for MS-PW as specified in [\[RFC6073\]](#). If VCCV CC type 3 "TTL Expiry" is to be used, the PW label TTL MUST be set to the appropriate value to reach the target PE. The hop count from one T-PE to the target PE can be obtained either via SP-PE TLVs, through MS-PW path trace or

based on management plane information.

5. IANA Considerations

This document makes no request of IANA.

6. Security Considerations

This document has the same security properties as in the PWE3 control protocol [[RFC4447](#)], [[RFC6870](#)] and [[RFC6478](#)].

7. Acknowledgements

The authors would like to thank Mach Chen, Lizhong Jin, Mustapha Aissaoui, Luca Martini, Matthew Bocci and Stewart Bryant for their valuable comments and discussions.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", [RFC 4447](#), April 2006.
- [RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", [RFC 6073](#), January 2011.
- [RFC6478] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", [RFC 6478](#), May 2012.
- [RFC6870] Muley, P. and M. Aissaoui, "Pseudowire Preferential Forwarding Status Bit", [RFC 6870](#), February 2013.

8.2. Informative References

- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", [RFC 3985](#), March 2005.
- [RFC5085] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", [RFC 5085](#), December 2007.
- [RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", [RFC 5659](#), October 2009.
- [RFC6718] Muley, P., Aissaoui, M., and M. Bocci, "Pseudowire Redundancy", [RFC 6718](#), August 2012.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Haibo Wang
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: rainsword.wang@huawei.com