

Pseudowire Redundancy on S-PE
draft-dong-pwe3-redundancy-spe-00

Abstract

This document describes Multi-Segment Pseudowire (MS-PW) protection scenarios in which the pseudowire redundancy is provided on the Switching-PE (S-PE). Signaling of preferential forwarding defined in [[I-D.ietf-pwe3-redundancy-bit](#)] is reused.

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 April 26, 2012.

Copyright Notice

Copyright (c) 2011 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 [3](#)
- [2.](#) PW Redundancy on S-PE [3](#)
- [3.](#) S-PE Operations [4](#)
- [4.](#) IANA Considerations [6](#)
- [5.](#) Security Considerations [6](#)
- [6.](#) Acknowledgements [6](#)
- [7.](#) References [6](#)
 - [7.1.](#) Normative References [6](#)
 - [7.2.](#) Informative References [6](#)
- Authors' Addresses [7](#)

1. Introduction

[I-D.ietf-pwe3-redundancy] and [I-D.ietf-pwe3-redundancy-bit] describe Pseudowire (PW) redundancy mechanism for scenarios where a set of redundant PWs terminate on either provider edge (PE) nodes in single-segment pseudowire (SS-PW) [RFC3985] applications, or on terminating provider edge (T-PE) nodes in multi-segment pseudowire (MS-PW) [RFC5659] applications. This document describes the scenarios where PW redundancy is provided on S-PEs of MS-PW. Signaling of preferential forwarding defined in [I-D.ietf-pwe3-redundancy-bit] is reused for these scenarios, and operations on S-PE are specified.

2. PW Redundancy on S-PE

In some MS-PW deployment scenarios, PW redundancy may need to be provided on S-PE. This section gives some examples of PW redundancy on S-PE.

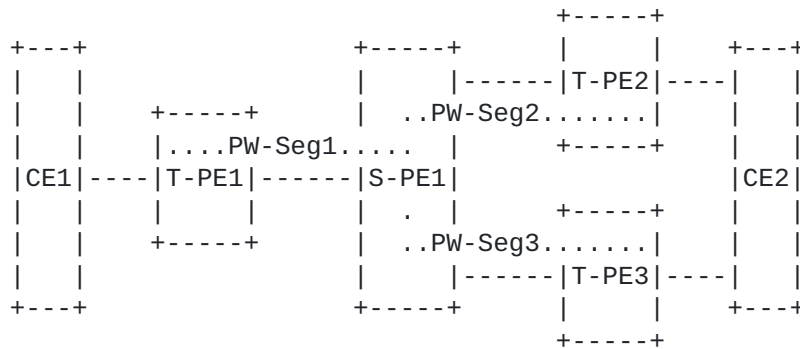


Figure 1. MS-PW Redundancy on S-PE

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 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 primary PW segment, and PW-Seg3 is secondary PW segment.

MS-PW redundancy on S-PE is beneficial for scenario in Figure 1 since on T-PE1 side it may be impossible to provide PW redundancy, especially when the PW-Seg1 between T-PE1 and S-PE1 is statically configured. For PW redundancy on S-PE, the number of PW segments needed between T-PE1 and S-PE1 is only half of the number of PW segments needed for End-to-End PW redundancy. Also PW redundancy on S-PE could provide faster protection switching than end-to-end protection switching of MS-PW.

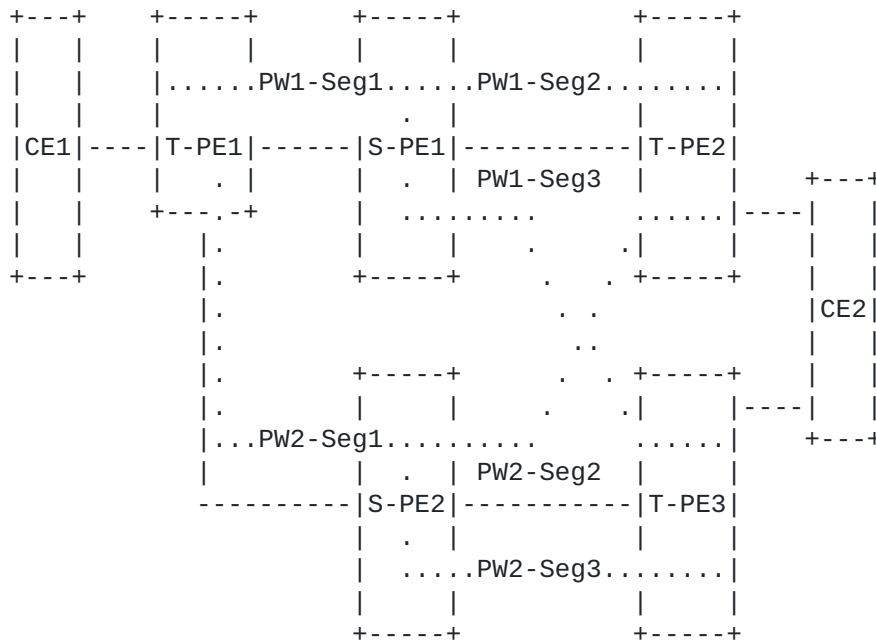


Figure 2. MS-PW Redundancy on S-PE with S-PE protection

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, both S-PE1 and S-PE2 are connected to 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, the S-PE1 provides resiliency using PW1-Seg2 and PW1-Seg3. For MS-PW2, the S-PE2 provides resiliency using PW2-Seg2 and PW2-Seg3. MS-PW1 is the primary PW and PW1-Seg2 is the primary PW segment.

MS-PW redundancy on S-PE is beneficial for scenario in Figure 2 since it reduces the number of end-to-end MS-PWs required for both T-PE and S-PE protection. Also redundancy on S-PE could provide faster protection switching than end-to-end protection switching of MS-PW.

3. S-PE Operations

When S-PE redundancy is provisioned, it is necessary that S-PE could perform protection switching according to the status change of PW segments and announce appropriate PW status to adjacent PEs. Signaling of preferential forwarding defined in [\[I-D.ietf-pwe3-redundancy-bit\]](#) is reused for these scenarios, and operation on S-PE is specified as below.

For scenario of Figure 1, assume the AC from CE2 to T-PE2 is active. if S-PE1 knows PW-Seg1 is in "PW forwarding" State, it would

advertise "Preferential Forwarding" status bit of "Active" on both PW-Seg2 and PW-Seg3. T-PE2 advertises the preferential status "Active" and T-PE3 advertises the preferential status "Standby", by matching the local and remote 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 would then advertise the preferential status "Active" to S-PE1, and T-PE2 would advertise the preferential status "Standby". S-PE1 would perform the switchover according to the updated local and remote preferential forwarding status, and select PW-Seg3 to forward traffic. If S-PE selects a new Active PW segment successfully, it SHOULD NOT advertise any change of the PW status to T-PE1. Hence T-PE1 would not be aware of the failure on the remote side.

For scenario of Figure 2, assume the AC from CE2 to T-PE2 is active. T-PE1 would advertise preferential status "Active" on PW1-Seg1 and "Standby" on PW2-Seg1. According to the received preferential status, S-PE1 SHOULD advertise preferential status "Active" on both PW1-Seg2 and PW1-Seg3, and S-PE2 SHOULD advertise preferential status "Standby" on both PW2-Seg2 and PW2-Seg3. T-PE2 advertises preferential status "Active" on both PW1-Seg2 and PW2-Seg3, and T-PE3 advertises preferential status "Standby" on both PW1-Seg3 and PW2-Seg3. By matching the local and remote preferential forwarding status, PW1-Seg2 would be used for traffic forwarding. Since S-PE1 connects to the primary PW segment PW1-Seg2, it would advertise preferential status "Active" to T-PE1. S-PE2 would advertise preferential status "Standby" to T-PE1 since it does not connect to the primary PW segment.

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 the preferential status "Active" to S-PE1, and T-PE2 would advertise the preferential status "Standby". S-PE1 would perform the switchover according to the updated local and remote preferential forwarding status, and select PW1-Seg3 to forward traffic. Since S-PE1 selects a new Active PW segment successfully, it SHOULD NOT advertise any change of the PW status to T-PE1, and T-PE would not be aware of the failure on the remote side.

When the S-PE1 fails, T-PE1 would advertise the preferential status "Active" to S-PE2, On receiving the change of preferential status, S-PE2 SHOULD advertise the preferential status "Active" on both PW2-Seg2 and PW2-Seg3. Then by matching the local and remote preferential forwarding status, PW2-Seg2 would be selected as primary PW segment, and traffic would be forwarded on MS-PW2.

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

This document has the same security properties as in the PWE3 control protocol [[RFC4447](#)] and [[I-D.ietf-pwe3-redundancy-bit](#)].

6. Acknowledgements

The authors would like to thank Mach Chen for his review and suggestions.

7. References

7.1. Normative References

- [I-D.ietf-pwe3-redundancy]
Muley, P., Aissaoui, M., and M. Bocci, "Pseudowire Redundancy", [draft-ietf-pwe3-redundancy-05](#) (work in progress), September 2011.
- [I-D.ietf-pwe3-redundancy-bit]
Muley, P. and M. Aissaoui, "Pseudowire Preferential Forwarding Status Bit", [draft-ietf-pwe3-redundancy-bit-05](#) (work in progress), September 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", [RFC 3985](#), March 2005.
- [RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", [RFC 5659](#), October 2009.

7.2. Informative References

- [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.

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