

Internet Engineering Task Force (IETF)
Request for Comments: 8234
Updates: 7271
Category: Standards Track
ISSN: 2070-1721

J. Ryoo
T. Cheung
ETRI
H. van Helvoort
Hai Gaoming BV
I. Busi
G. Wen
Huawei Technologies
August 2017

Updates to MPLS Transport Profile (MPLS-TP) Linear Protection in
Automatic Protection Switching (APS) Mode

Abstract

This document contains updates to MPLS Transport Profile (MPLS-TP) linear protection in Automatic Protection Switching (APS) mode defined in RFC 7271. The updates provide rules related to the initialization of the Protection State Coordination (PSC) Control Logic (in which the state machine resides) when operating in APS mode and clarify the operation related to state transition table lookup.

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 <http://www.rfc-editor.org/info/rfc8234>.

Copyright Notice

Copyright (c) 2017 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. Conventions Used in This Document	4
3. Abbreviations	4
4. Updates	4
4.1. Initialization Behavior	5
4.2. State Transition Modification	6
4.3. Operation Related to State Transition Table Lookup	6
5. Security Considerations	7
6. IANA Considerations	7
7. References	8
7.1. Normative References	8
7.2. Informative References	8
Acknowledgements	8
Authors' Addresses	9

1. Introduction

MPLS Transport Profile (MPLS-TP) linear protection in Automatic Protection Switching (APS) mode is defined in RFC 7271 [RFC7271]. It defines a set of alternate and additional mechanisms to perform some of the functions of linear protection described in RFC 6378 [RFC6378]. The actions performed at initialization of the Protection State Coordination (PSC) Control Logic are not described in either [RFC7271] or [RFC6378]. Although it is a common perception that the state machine starts at the Normal state, this is not explicitly specified in any of the documents and various questions have been raised by implementers and in discussions on the MPLS working group mailing list concerning the detailed actions that the PSC Control Logic should take.

The state machine described in [RFC7271] operates under the assumption that both end nodes of a linear protection domain start in the Normal state. In the case that one node reboots while the other node is still in operation, various scenarios may arise resulting in problematic situations. This document resolves all the problematic cases and minimizes traffic disruptions related to initialization, including both cold and warm reboots that require re-initialization of the PSC Control Logic.

This document contains updates to the MPLS-TP linear protection in APS mode defined in [RFC7271]. The updates provide rules related to initialization of the PSC Control Logic (in which the state machine resides) when operating in APS mode. The updates also include modifications to the state transition table defined in Section 11.2 of [RFC7271]. The changes in the state transition table have been examined to make sure that no new problems are introduced.

This document does not introduce backward compatibility issues with implementations of [RFC7271]. In case a node implementing this document restarts, the new state changes will not cause problems at the remote node implementing [RFC7271], and the two ends will converge to the same local and remote states. In case a node implementing [RFC7271] restarts, the two ends behave as they do today.

This document also provides some clarifications on the operation related to state transition table lookup.

The reader of this document is assumed to be familiar with [RFC7271].

2. Conventions Used in This Document

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. Abbreviations

This document uses the following abbreviations:

APS	Automatic Protection Switching
DNR	Do-not-Revert
E::R	Exercise state due to remote EXER message
EXER	Exercise
MS-P	Manual Switch to Protection path
MS-W	Manual Switch to Working path
MPLS-TP	MPLS Transport Profile
NR	No Request
PF:DW:R	Protecting Failure state due to remote SD-W message
PF:W:L	Protecting Failure state due to local SF-W
PF:W:R	Protecting Failure state due to remote SF-W message
PSC	Protection State Coordination
RR	Reverse Request
SA:MP:R	Switching Administrative state due to remote MS-P message
SA:MW:R	Switching Administrative state due to remote MS-W message
SD	Signal Degrade
SD-W	Signal Degrade on Working path
SF	Signal Fail
SF-P	Signal Fail on Protection path
SF-W	Signal Fail on Working path
UA:P:L	Unavailable state due to local SF-P
WTR	Wait-to-Restore

4. Updates

This section specifies the actions that will be performed at the initialization of the PSC Control Logic and the modifications of the state transition table defined in Section 11.2 of [RFC7271]. Some clarifications on the operation related to state transition table lookup are also provided.

4.1. Initialization Behavior

This section defines initialization behavior that is not described in [RFC7271].

When the PSC Control Logic is initialized, the following actions MUST be performed:

- o Stop the WTR timer if it is running.
- o Clear any operator command in the Local Request Logic.
- o If an SF-W or SF-P exists as the highest local request, the node being initialized starts at the PF:W:L or UA:P:L state, respectively.
- o If the node being initialized has no local request:
 - * If the node being initialized does not remember the active path or if the node being initialized remembers the working path as the active path, the node starts at the Normal state.
 - * Else (the node being initialized remembers the protection path as the active path), the node starts at the WTR state sending NR(0,1) or at the DNR state sending DNR(0,1) depending on the configuration that allows or prevents automatic reversion to the Normal state.
- o In case any local SD exists, the local SD MUST be considered as an input to the Local Request Logic only after the local node has received the first protocol message from the remote node and completed the processing (i.e., updated the PSC Control Logic and decided which action, if any, is to be sent to the PSC Message Generator).
- o If the local node receives an EXER message as the first protocol message after initialization and the remote EXER becomes the top-priority global request, the local node MUST set the position of the bridge and selector according to the Path value in the EXER message and transit to the E::R state.

In the case of no local request, remembering the active path minimizes traffic switchovers when the remote node is still in operation. This approach does not cause a problem even if the remembered active path is no longer valid due to any local input that occurred at the remote node while the initializing node was out of operation.

Note that in some restart scenarios (e.g., cold rebooting), no valid SF/SD indications may be present at the input of the Local Request Logic. In this case, the PSC Control Logic restarts as if no local requests are present. If a valid SF/SD indication is detected later, the PSC Control Logic is notified and state change is triggered.

4.2. State Transition Modification

In addition to the initialization behavior described in Section 4.1, four cells of the remote state transition table need to be changed to make two end nodes converge after initialization. State transition by remote message as defined in Section 11.2 of [RFC7271] is modified as follows (only modified cells are shown):

	MS-W	MS-P	WTR	EXER	RR	DNR	NR
N			(13)			DNR	
PF:W:R						DNR	
PF:DW:R						DNR	

The changes in two rows of remote protecting failure states lead to the replacement of note (10) with DNR; therefore, note (10) is no longer needed. The resultant three rows read:

	MS-W	MS-P	WTR	EXER	RR	DNR	NR
N	SA:MW:R	SA:MP:R	(13)	E::R	i	DNR	i
PF:W:R	SA:MW:R	SA:MP:R	(9)	E::R	i	DNR	(11)
PF:DW:R	SA:MW:R	SA:MP:R	(9)	E::R	i	DNR	(11)

In the tables above, the letters 'i' and 'N' stand for "ignore" and "Normal state", respectively. Other abbreviations can be found in Section 3.

4.3. Operation Related to State Transition Table Lookup

In addition to the rules related to the state transition table lookup listed in Section 11 of [RFC7271], the following rule is also applied to the operation related to the state transition table lookup:

- o When the local SF-P is cleared and the priorities of the local and remote requests are re-evaluated, the last received remote message may no longer be valid due to the previous failure of the protection path. Therefore, the last received message MUST be treated as if it were NR and only the local request shall be evaluated.

The last paragraph in Section 11 of [RFC7271] is modified as follows:

Old text:

In the state transition tables below, the letter 'i' stands for "ignore" and is an indication to remain in the current state and continue transmitting the current PSC message.

New text:

In the state transition tables below, the letter 'i' is the "ignore" flag; if it is set, it means that the top-priority global request is ignored.

If re-evaluation is triggered, the ignore flag is checked. If it is set, the state machine will transit to the supposed state, which can be Normal or DNR as indicated in the footnotes to the state transition table in Section 11.1 of [RFC7271]. If the ignore flag is not set, the state machine will transit to the state indicated in the cell of the state transition table.

If re-evaluation is not triggered, the ignore flag is checked. If it is set, the state machine will remain in the current state, and the current PSC message continues to be transmitted. If the ignore flag is not set, the state machine will transit to the state indicated in the cell of the state transition table.

5. Security Considerations

No specific security issue is raised in addition to those ones already documented in [RFC7271]. Note that tightening the description of the initializing behavior may help to protect networks from restart attacks.

6. IANA Considerations

This document does not require any IANA actions.

7. References

7.1. 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>>.
- [RFC7271] Ryoo, J., Ed., Gray, E., Ed., van Helvoort, H., D'Alessandro, A., Cheung, T., and E. Osborne, "MPLS Transport Profile (MPLS-TP) Linear Protection to Match the Operational Expectations of Synchronous Digital Hierarchy, Optical Transport Network, and Ethernet Transport Network Operators", RFC 7271, DOI 10.17487/RFC7271, June 2014, <<https://www.rfc-editor.org/info/rfc7271>>.
- [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>>.

7.2. Informative References

- [RFC6378] Weingarten, Y., Ed., Bryant, S., Osborne, E., Sprecher, N., and A. Fulignoli, Ed., "MPLS Transport Profile (MPLS-TP) Linear Protection", RFC 6378, DOI 10.17487/RFC6378, October 2011, <<https://www.rfc-editor.org/info/rfc6378>>.

Acknowledgements

The authors would like to thank Joaquim Serra for raising the issue related to initialization of the PSC Control Logic at the very beginning. The authors would also like to thank Adrian Farrel and Loa Andersson for their valuable comments and suggestions on this document.

Authors' Addresses

Jeong-dong Ryoo
ETRI

Email: ryoo@etri.re.kr

Taesik Cheung
ETRI

Email: cts@etri.re.kr

Huub van Helvoort
Hai Gaoming BV

Email: huubatwork@gmail.com

Italo Busi
Huawei Technologies

Email: Italo.Busi@huawei.com

Guangjuan Wen
Huawei Technologies

Email: wenguangjuan@huawei.com

