MPLS Working Group F. Zhang, Ed.
Internet-Draft WJ. He
Intended status: Standards Track ZTE
Expires: January 5, 2012 July 04, 2011

MPLS-TP Shared Mesh Protection
draft-zhang-mpls-tp-shared-mesh-protection-00

Abstract

The MPLS Transport Profile (MPLS-TP) requirements document [RFC5654], describes that MPLS-TP MUST support sharing of protection resources.

This document describes a shared mesh protection processing based on the existing MPLS-TP linear protection switching mechanisms.

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 January 5, 2012.

Table of Contents

1. Introduction 3
2. Conventions used in this document 3
2.1. Acronyms 3
3. Shared Mesh Protection Architecture 4
3.1. Planning 4
4. Changes to PSC 5
5. Basic Operation 6
5.1. Preemption 6
6. IANA Considerations 7
7. Security Considerations 7
8. Acknowledgement 7
9. References 8
9.1. Normative references 8
9.2. Informative References 8
Authors' Addresses 9

1. Introduction

For shared mesh protection, the protection resources are used to protect multiple LSPs which do not all share the same end points. In this way, mesh protection can substantially reduce the network resources that have to be reserved in oder to provide protection. The requirements have been described in [RFC5654] (Req. 66, 67, 68 and 69). Furthermore, the MPLS Transport Profile (MPLS-TP) Survivability Framework [I-D.ietf-mpls-tp-survive-fwk] outlined the operation. The shared mesh recovery schemes are also discussed in [RFC4428] and [G.smp]. In (1:1)^n protection decribed in [RFC4428], n working paths are protected by n dedicated protection paths while sharing the same protection bandwidth.

This document describes a shared mesh protection processing based on the concept of (1:1)^n recovery scheme defined in [RFC4428], and on the protecion mechanism being developed in [I-D.ietf-mpls-tp-linear-protection] [I-D.ezy-mpls-1ton-protection].

2. Conventions used in this document

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 [RFC2119].

2.1. Acronyms

This draft uses the following acronyms:

Ack	Acknowledge
DNR	Do not revert
FS	Forced Switch
LER	Label Edge Router

LO	Lockout of protection
LSP	Label Switched Path
MEP	Maintenance Entity Group End Point
MIP	Maintenance Entity Group Intermediate point
MPLS-TP	Transport Profile for MPLS
MS	Manual Switch
NR	No Request
P2P	Point-to-point
PSC	Protection State Coordination Protocol
SF	Signal Fail
SMPG	Shared Mesh Protection Group
WTR	Wait-to-Restore

3. Shared Mesh Protection Architecture

A simple case of shared mesh protection is illustrated in Figure 1. Consider three paths W1 [via nodes A, B, C], W3 [via nodes D, E, F] and W2 [via nodes X, Y, Z]. For these working paths do not share end points, they cannot make use of 1:N protection even though they also do not share any common points of failure. W1 may be protected by the path P1 [via nodes A, G, P, Q, H, C}, W3 may be protected by the path P3 [via nodes D, I, S, T, J, F], and W2 may be protected by the path P2 [via nodes X, V, P, Q, R, S, T, W, Z]. For all these cases, 1:1 protection may be used.

In the event that the failure only affect one of the working paths, the shared segment PQ or/and ST only carry traffic from the working path being affected. Thus, it is possible for the network resources on the segment PQ and ST to be shared by the two protection paths. In this way, shared mesh protection can substantially reduce the amount of network resources that have to be reserved to provide protection.

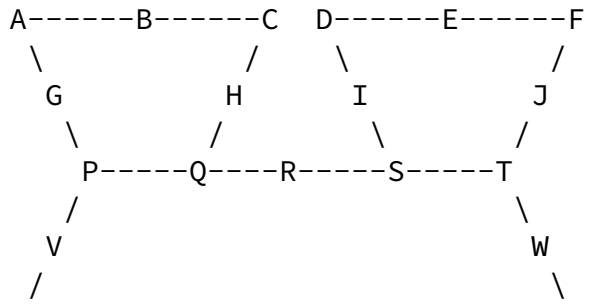

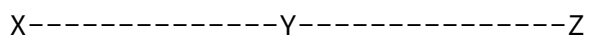

Figure 1: An example of shared mesh protection

3.1. Planning

As described in [I-D.ietf-mpls-tp-survive-fwk], the network becomes more and more complex and the number of LSPs increases, the potential for shared mesh protection also increase. However, this can quickly become unmanageable owing to the increased complexity. Therefore, shared mesh protection is normally pre-planned and configured by the operator, although an automated sytem cannot be ruled out. This will include but not limited to:

- Planning the shared mesh protection group (SMPG) which includes the protected paths and protecting paths. Different SPMGs do not share protection resources and are protected independently. This means that working paths which have the higher protection

switching priority are planned to be in different SMPGs, in such a way the higher priority paths will be protected mostly when one failure affects different SMPGs.

- Configuring the mumbers of the SMPG. The working paths are disjoint as far as possible in one SMPG, so they will not be subject to common failures. Furthermore, each of the working paths may be assigned a relative priority that could be used to decide which working path would be protected in cases of conflict. The relative priority is recommend to be reflected by the entity number of the working path, which is compatible with 1:N linear protection [I-D.ezy-mpls-1ton-protection]. When equal priority requests occur simultaneously, the conflict is resolved in favour of the request with the lowest entity number.

4. Changes to PSC

Protection State Coordination Protocol (PSC) defined in [I-D.ezy-mpls-1ton-protection] is a multi-phased protocol, the end-points perform any protection switching with waiting for acknowledgement from the far end Label Edge Router (LER). The

protocol messages are transmitted using the G-ACh.

In order to support shared mesh protection, there is a need to make changes to the format of the PSC message. In particular there is the need to carry TTL TLV [I-D.ietf-mpls-lsp-ping-ttl-tlv] as one optional TLV to indicate which node to receive the return PSC message. Due to this change, the value of the Ver field for PSC messages for a shared mesh protection domain MUST be set to 3.

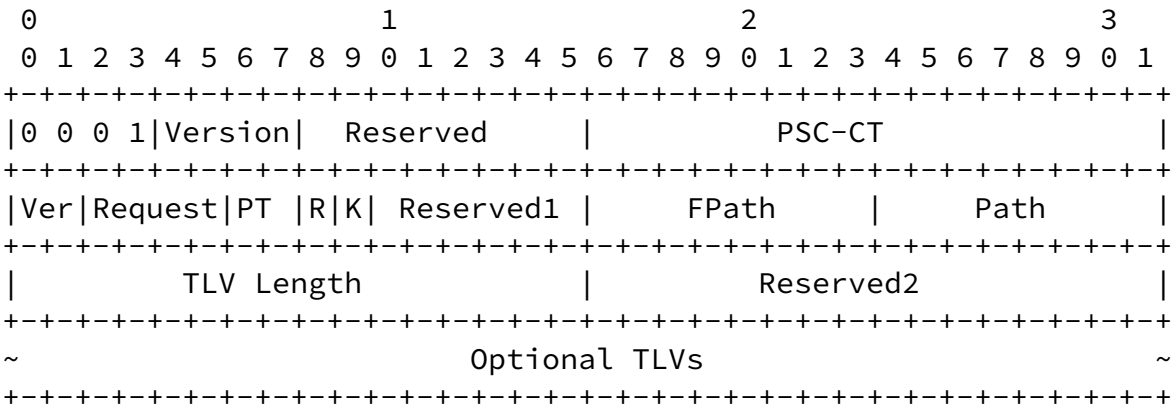

Figure 2: Format of shared mesh protection PSC packet with a G-ACh header

5. Basic Operation

This section illustrates the basic operation of the shared mesh protection for the topology shown in Figure 1 based on the PSC.

If a unidirectional failure occurs on the W2 in the direction from node Z to node X at time zero, the shared mesh protection will operate as follows:

a. Node X detects the signal failure (SF), sends SF(2,0) to node P with MPLS label TTL to control the hops between the node X and node P. TTL TLV may be inserted in the PSC message as one optional TLV to indicate that node P SHOULD use this value to return the PSC message.
b. Node P compares the protection switching priority of W2 and W1.

For W1 is in normal state, node P sends SF(2,0) to node P with MPLS label TTL to control the hops between the node P and node Q, similarly, TTL TLV may be inserted in the PSC message as one optional TLV to indicate that node S SHOULD use this value to return the PSC message. The same processing is done on node S and node T.

c. When Node Z receives the PSC message, it bridges and selects traffic from P2. Then sends NR(0,2)Ack to node T, TTL TLV may be inserted in the PSC message for the same reason.
d. Node T, S, Q, P relay the message until it arrives at node X.
e. Node X bridges and selects traffic from P2, then sends SF(2,2). When node Z receives this message, it responses with NR(0,2).

5.1. Preemption

If a unidirectional failure occurs on the W1 in the direction from node C to node A at time one, the shared mesh protection will operate as follows:

a. Node A detects the SF on W1, sends SF(1,0) to node P with MPLS label TTL to control the hops between the node A and node P. TTL TLV may be inserted in the PSC message as one optional TLV to indicate that node P SHOULD use this value to return the PSC message.
b. Node P compares the protection switching priority of W2 and W1. For W1's entity number is samller than W2's, it has the higher priority under the SF events on both working paths. So node P sends SF(1,0) on the path P1, and sends LO(2,2) on the path P2.
c. Node Q relays the message SF(1,0) on the path P1 to node A, and sends LO(2,2) to node S when it receives LO(2,2) from the previous hop on the path P2.

d. When Node C receives the SF(1,0) message, it bridges and selects traffic from P1 and sends NR(0,1)Ack to node Q, TTL TLV may be inserted in the PSC message for the same reason. While node Z receives the LO(2,2) message, it bridges and selects traffic from W2, and responses with NR(0,0). It should be noted that this may cause loss of user data since W1 is still in a failure condition.
e. Node A bridges and selects traffic from P1, then sends SF(1,1) when it receives NR(0,1)Ack. Node C will response with NR(0,1)

while receives this message. According to the received NR(0,0), node X will bridge and select traffic from W2, and response with SF(2,0), then node P relays with LO(2,0) towards node Z.

If a unidirectional failure occurs on the working path W3 (not on the W1) in the direction from node F to node P at time one. Although the resource preemption will fail, the basic PSC processing will be similarly.

6. IANA Considerations

TBD.

7. Security Considerations

The generic security considerations for the data-plane of MPLS-TP are described in the security framework document [I-D.ietf-mpls-tp-security-framework] together with the required mechanisms needed to address them. The security considerations for the generic associated control channel are described in [RFC5586]. The security considerations for protection and recovery aspects of MPLS-TP are addressed in [I-D.ietf-mpls-tp-survive-fwk].

The extensions to the protocol described in this document are extensions to the protocol defined in [I-D.ietf-mpls-tp-linear-protection] [I-D.ezy-mpls-1ton-protection] and does not introduce any new security risks.

8. Acknowledgement

The authors would like to thank all members of the teams (the Joint Working Team, the MPLS Interoperability Design Team in IETF and the T-MPLS Ad Hoc Group in ITU-T) involved in the definition and specification of MPLS Transport Profile.

9. References

9.1. Normative references

[I-D.ietf-mpls-tp-linear-protection]
Bryant, S., Osborne, E., Sprecher, N., Fulignoli, A., and Y. Weingarten, "MPLS-TP Linear Protection", draft-ietf-mpls-tp-linear-protection-07 (work in progress), June 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

9.2. Informative References

[I-D.ezy-mpls-1ton-protection]
Osborne, E., Zhang, F., and Y. Weingarten, "MPLS-TP 1toN Protection", draft-ezy-mpls-1ton-protection-00 (work in progress), June 2011.

[I-D.ietf-mpls-lsp-ping-ttl-tlv]
Boutros, S., Sivabalan, S., Swallow, G., Saxena, S., and V. Manral, "Definition of Time-to-Live TLV for LSP-Ping Mechanisms", draft-ietf-mpls-lsp-ping-ttl-tlv-00 (work in progress), June 2011.

[I-D.ietf-mpls-tp-security-framework]
Fang, L., Niven-Jenkins, B., and S. Mansfield, "MPLS-TP Security Framework", draft-ietf-mpls-tp-security-framework-01 (work in progress), May 2011.

[I-D.ietf-mpls-tp-survive-fwk]
Sprecher, N. and A. Farrel, "Multiprotocol Label Switching Transport Profile Survivability Framework", draft-ietf-mpls-tp-survive-fwk-06 (work in progress), June 2010.

[RFC4428] Papadimitriou, D. and E. Mannie, "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", RFC 4428, March 2006.

[RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", RFC 5586, June 2009.

Authors' Addresses

Fei Zhang (editor)
ZTE

Email: zhang.fei3@zte.com.cn

Wenjuan He
ZTE

Email: hewenjuan@zte.com.cn