TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2020

J. He
I. Busi
Huawei Technologies
J. Ryoo
B. Yoon
ETRI
P. Park
KT

March 8, 2020

GMPLS Signaling Extensions for Shared Mesh Protection
draft-ietf-teas-gmpls-signaling-smp-03

Abstract

ITU-T Recommendation G.808.3 [G808.3] defines the generic aspects of a Shared Mesh Protection (SMP) mechanism, where the difference between SMP and Shared Mesh Restoration (SMR) is also identified. ITU-T Recommendation G.873.3 [G873.3] defines the protection switching operation and associated protocol for SMP at the Optical Data Unit (ODU) layer. RFC 7412 [RFC7412] provides requirements for any mechanism that would be used to implement SMP in a Multi-Protocol Label Switching - Transport Profile (MPLS-TP) network.

This document updates RFC 4872 [RFC4872] to provide the extensions to the Generalized Multi-Protocol Label Switching (GMPLS) signaling to support the control of the shared mesh protection.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

RFC 4872 [RFC4872] defines extension of Resource Reservation Protocol - Traffic Engineering (RSVP-TE) to support Shared Mesh Restoration (SMR) mechanism. SMR can be seen as a particular case of pre-planned Label Switched Path (LSP) rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, and an explicit restoration signaling is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the provisioning phase.

ITU-T Recommendation G.808.3 [G808.3] defines the generic aspects of a shared mesh protection (SMP) mechanism. ITU-T Recommendation G.873.3 [G873.3] defines the protection switching operation and associated protocol for SMP at the Optical Data Unit (ODU) layer. RFC 7412 [RFC7412] provides requirements for any mechanism that would be used to implement SMP in a Multi-Protocol Label Switching - Transport Profile (MPLS-TP) network.

SMP differs from SMR in the activation/protection switching operation. The former activates a protecting LSP via the automatic protection switching (APS) protocol in the data plane when the working LSP fails, while the latter does it via the control plane signaling. It is therefore necessary to distinguish SMP from SMR during provisioning so that each node involved behaves appropriately in the recovery phase when activation of a protecting LSP is done.

This document updates [RFC4872] to provide the extensions to the Generalized Multi-Protocol Label Switching (GMPLS) signaling to support the control of the SMP mechanism. Only the generic aspects for signaling SMP are addressed by this document. The technology-specific aspects are expected to be addressed by other documents.

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.

In addition, the reader is assumed to be familiar with the terminology used in [RFC4872] and RFC 4426 [RFC4426].

3. SMP Definition

[G808.3] defines the generic aspects of a SMP mechanism. [G873.3] defines the protection switching operation and associated protocol for SMP at the ODU layer. [RFC7412] provides requirements for any mechanism that would be used to implement SMP in a MPLS-TP network.

The SMP mechanism is based on pre-computed protection transport entities that are pre-configured into the network elements. Pre-configuration here means pre-reserving resources for the protecting LSPs without activating a particular protecting LSP (e.g. in circuit networks, the cross-connects in the intermediate nodes of the protecting LSP are not pre-established). Pre-configuring but not activating the protecting LSP allows the common link and node resources in a protecting LSP to be shared by multiple working LSPs

that are physically (i.e., link, node, Shared Risk Link Group (SRLG), etc.) disjoint. Protecting LSPs are activated in response to failures of working LSPs or operator's commands by means of the APS protocol that operates in the data plane. The APS protocol messages are exchanged along the protecting LSP. SMP is always revertive.

SMP has a lot of similarity to SMR except that the activation in case of SMR is achieved by control plan signaling during the recovery operation while SMP is done by APS protocol in the data plane. SMP has advantages with regard to the recovery speed compared with SMR.

4. GMPLS Signaling Extension for SMP

Consider the following network topology:

```
                          A---B---C---D
                           \         /
                            E---F---G
                           /         \
                          H---I---J---K
```

Figure 1: An example of SMP topology

The working LSPs [A,B,C,D] and [H,I,J,K] could be protected by [A,E,F,G,D] and [H,E,F,G,K], respectively. Per RFC 3209 [RFC3209], in order to achieve resource sharing during the signaling of these protecting LSPs, they must have the same Tunnel Endpoint Address (as part of their SESSION object). However, these addresses are not the same in this example. Similar to SMR, a new LSP Protection Type of the secondary LSP is defined as "Shared Mesh Protection" (see Section 5.1) to allow resource sharing along nodes E, F, and G. In this case, the protecting LSPs are not merged (which is useful since the paths diverge at G), but the resources along E, F, G can be shared.

When a failure, such as Signal Fail (SF) and Signal Degrade (SD), occurs on one of the working LSPs (say working LSP [A,B,C,D]), the end-node (say node A) that detects the failure initiates the protection switching operation. The end-node A will send a protection switching request APS message (for example SF) to its adjacent (downstream) intermediate node (say node E) to activate setting up the corresponding protecting LSP and will wait for a confirmation message from node E. If the protection resource is available, node E will send the confirmation APS message to the end-node A and forward the switching request APS message to its adjacent (downstream) node (say node F). When the confirmation APS message is received by node A, the cross-connection on node A is established. At this time the traffic is bridged to and selected from the

protecting LSP at node A. After forwarding the switching request APS message, node E will wait for a confirmation APS message from node F, which triggers node E to set up the cross-connection for the protecting LSP being activated. If the protection resource is not available (due to failure or being used by higher priority connections), the switching will not be successful; the intermediate node may send a message to notify the end node, or may keep trying until the resource is available, or the switching request is cancelled. If the resource is in use by a lower priority protecting LSP, the lower priority service will be removed and then the intermediate node will follow the procedure as described for the case when the protection resource is available for the higher priority protecting LSP.

The SMP preemption priority of a protecting LSP that the APS protocol uses to resolve the competition for shared resources among multiple protecting LSPs, is indicated in the TBD1 field of the PROTECTION object in the Path message of the protecting LSP. In SMP, the Setup and Holding priorities in the SESSION_ATTRIBUTE object can be used to configure or pre-configure a LSP, but is irrelevant to resolving the competition among multiple protecting LSPs, when controlled by the APS.

When an intermediate node on the protecting LSP receives the Path message, the preemption priority value in the TBD1 field MUST be stored for that protecting LSP. When resource competition among multiple protecting LSPs occurs, the APS protocol will use their priority values to resolve the competition.

[EDITOR'S NOTE: The TBD1 field for the preemption priority will be defined in the next version of this draft.]

In SMP, a preempted LSP SHOULD not be torn down. Once the working LSP and the protecting LSP are configured or pre-configured, the end node SHOULD keep refreshing both working and protecting LSPs regardless of failure or preempted situation.

When a lower priority protecting LSP is preempted, the intermediate node that performed preemption MAY send a Notify message with a new sub-code "Shared resources unavailable" under "Notify Error" code (see [RFC4872]) to the end nodes of that protecting LSP. Upon receipt of this Notify message, the end node MAY stop sending and selecting normal traffic to/from its protecting LSP and try switching the traffic to another protection LSP, if available.

When the shared resources become unavailable, the same Notify message MAY also be generated by the intermediate node to all the end nodes of the protecting LSPs that have lower preemption priorities than the

one that has occupied the shared resources. These end nodes, in case of a failure of the working LSP, MAY avoid trying to switch the traffic to these protection LSPs that have been configured to use the shared resources and try switching the traffic to other protection LSPs, if available.

When the shared resources become available, a Notify message with a new sub-code "Shared resources available" under "Notify Error" code MAY be generated by the intermediate node. The recipients of this Notify message are the end nodes of the lower priority protecting LSPs that have been preempted and/or all the end nodes of the protecting LSPs that have lower preemption priorities than the one that does not need the shared resources any more.

The following subsections detail how LSPs using SMP can be signaled in an interoperable fashion using GMPLS RSVP-TE extensions (see RFC 3473 [RFC3473]). This includes;

(1) the ability to identify a "secondary protecting LSP" (hereby called the "secondary LSP") used to recover another primary working LSP (hereby called the "protected LSP"),

(2) the ability to associate the secondary LSP with the protected LSP,

(3) the capability to include information about the resources used by the protected LSP while instantiating the secondary LSP,

(4) the capability to instantiate during the provisioning phase several secondary LSPs in an efficient manner, and

(5) the capability to support activation of a secondary LSP after failure occurrence via APS protocol in the data plane.

4.1. Identifiers

To simplify association operations, both LSPs (i.e., the protected and the secondary LSPs) belong to the same session. Thus, the SESSION object MUST be the same for both LSPs. The LSP ID, however, MUST be different to distinguish between the protected LSP carrying normal traffic and the secondary LSP.

A new LSP Protection Type "Shared Mesh Protection" is introduced to the LSP Flags of PROTECTION object (see [RFC4872]) to set up the two LSPs. This LSP Protection Type value is applicable only to bidirectional LSPs as required in [G808.3].

4.2. Signaling Primary LSPs

The PROTECTION object (see [RFC4872]) is included in the Path message during signaling of the primary working LSPs, with the LSP Protection Type value set to "Shared Mesh Protection".

Primary working LSPs are signaled by setting in the PROTECTION object the S bit to 0, the P bit to 0, the N bit to 1 and in the ASSOCIATION object, the Association ID to the associated secondary protecting LSP_ID.

Note: N bit is set to indicate that the protection switching signaling is done via data plane.

4.3. Signaling Secondary LSPs

The PROTECTION object (see [RFC4872]) is included in the Path message during signaling of the secondary protecting LSPs, with the LSP Protection Type value set to "Shared Mesh Protection".

Secondary protecting LSPs are signaled by setting in the PROTECTION object the S bit and the P bit to 1, the N bit to 1 and in the ASSOCIATION object, the Association ID to the associated primary working LSP_ID, which MUST be known before signaling of the secondary LSP. Moreover, the Path message used to instantiate the secondary LSP SHOULD include at least one PRIMARY_PATH_ROUTE object (see [RFC4872]) that further allows for recovery resource sharing at each intermediate node along the secondary path.

With this setting, the resources for the secondary LSP SHOULD be pre-reserved, but not committed at the data plane level, meaning that the internals of the switch need not be established until explicit action is taken to activate this LSP. Activation of a secondary LSP and protection switching to the activated protecting LSP is done using APS protocol in the data plane.

After protection switching completes the protecting LSP SHOULD be signaled with the S bit set to 0 and O bit set to 1 in the PROTECTION object. At this point, the link and node resources must be allocated for this LSP that becomes a primary LSP (ready to carry normal traffic). The formerly working LSP MAY be signaled with the A bit set in the ADMIN_STATUS object (see [RFC3473]).

Support for extra traffic in SMP is for further study. Therefore, mechanisms to setup LSPs for extra traffic are also for further study.

4.4. SMP APS Configuration

SMP relies on APS protocol messages being exchanged between the nodes along the path to activate a SMP protecting LSP.

In order to allow exchange of APS protocol messages, an APS channel has to be configured between adjacent nodes along the path of the SMP protecting LSP. This should be done before any SMP protecting LSP has been setup by other means than GMPL signaling which are therefore outside the scope of this document.

Depending on the APS protocol message format, the APS protocol may use different identifiers than GMPLS signaling to identify the SMP protecting LSP.

Since APS protocol is for further study in [G808.3], it can be assumed that APS message format and identifiers are technology-specific and/or vendor-specific. Therefore, additional requirements for APS configuration are outside the scope of this document.

[EDITOR'S NOTE: Three options for APS configuration: a) out-of-scope, b) define a new object whose content is vendor-specific, c) define a new object with a TLV structure. The above paragraph (option a) can be confirmed or modified depending on a reply LS from the ITU-T SG15.

5. Updates to PROTECTION Object

GMPLS extension requirements for SMP introduce several updates to the Protection Object (see [RFC4872]).

5.1. New Protection Type

A new LSP protection type "Shared Mesh Protection" is added in the protection object. This LSP Protection Type value is applicable to only bidirectional LSPs.

LSP (Protection Type) Flags:

0x11: Shared Mesh Protection

5.2. Other Updates

N bit and O bit in the Protection object as defined in [RFC4872] are also updated to include applicability to SMP.

Notification (N): 1 bit

When set to 1, this bit indicates that the control plane message exchange is only used for notification during protection switching. When set to 0 (default), it indicates that the control plane message exchanges are used for protection-switching purposes. The N bit is only applicable when the LSP Protection Type Flag is set to either 0x04 (1:N Protection with Extra-Traffic), or 0x08 (1+1 Unidirectional Protection), or 0x10 (1+1 Bidirectional Protection). In SMP, N bit MUST be set to 1. The N bit MUST be set to 0 in any other case.

Operational (O): 1 bit

When set to 1, this bit indicates that the protecting LSP is carrying the normal traffic after protection switching. The O bit is only applicable when the P bit is set to 1, and the LSP Protection Type Flag is set to either 0x04 (1:N Protection with Extra-Traffic), or 0x08 (1+1 Unidirectional Protection), or 0x10 (1+1 Bidirectional Protection), or 0x11 (Shared Mesh Protection). The O bit MUST be set to 0 in any other case.

6. IANA Considerations

IANA actions required by this document will be described later.

7. Security Considerations

No further security considerations than [RFC4872].

8. Contributor

The following person contributed significantly to the content of this document and should be considered as a co-author.

Yuji Tochio
Fujitsu
Email: tochio@fujitsu.com

9. References

9.1. Normative References

[G808.3] International Telecommunication Union, "Generic protection switching - Shared mesh protection", ITU-T Recommendation G.08.3, October 2012.

[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>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <https://www.rfc-editor.org/info/rfc3473>.

[RFC4426] Lang, J., Ed., Rajagopalan, B., Ed., and D. Papadimitriou, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Recovery Functional Specification", RFC 4426, DOI 10.17487/RFC4426, March 2006, <https://www.rfc-editor.org/info/rfc4426>.

[RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, DOI 10.17487/RFC4872, May 2007, <https://www.rfc-editor.org/info/rfc4872>.

[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>.

9.2. Informative References

[G873.3] International Telecommunication Union, "Optical transport network - Shared mesh protection", ITU-T Recommendation G.873.3, September 2017.

[RFC7412] Weingarten, Y., Aldrin, S., Pan, P., Ryoo, J., and G. Mirsky, "Requirements for MPLS Transport Profile (MPLS-TP) Shared Mesh Protection", RFC 7412, DOI 10.17487/RFC7412, December 2014, <https://www.rfc-editor.org/info/rfc7412>.

Authors' Addresses

Jia He
Huawei Technologies
F3-1B, R&D Center, Huawei Industrial Base, Bantian, Longgang District
Shenzhen
China

Email: hejia@huawei.com

Italo Busi
Huawei Technologies

Email: italo.busi@huawei.com

Jeong-dong Ryoo
ETRI
218 Gajeongno
Yuseong-gu, Daejeon 34129
South Korea

Phone: +82-42-860-5384
Email: ryoo@etri.re.kr

Bin Yeong Yoon
ETRI

Email: byyun@etri.re.kr

Peter Park
KT

Email: peter.park@kt.com