

Network Working Group
INTERNET-DRAFT
Intended status: Informational
Expires: December 12, 2014

Y. Weingarten

S. Aldrin
Huawei Technologies
P. Pan
Infinera
J. Ryoo
ETRI
G. Mirsky
Ericsson
June 10, 2014

Requirements for MPLS-TP Shared Mesh Protection
draft-ietf-mpls-smp-requirements-06.txt

Abstract

This document presents the basic network objectives for the behavior of shared mesh protection (SMP) which are not based on control plane support. This is an expansion of the basic requirements presented in [RFC 5654](#) "Requirements for the Transport Profile of MPLS" and [RFC 6372](#) "MPLS Transport Profile (MPLS-TP) Survivability Framework". This document is to be used as a basis for the definition of any mechanism that would be used to implement SMP for MPLS-TP data paths, in networks that delegate executive action for resiliency to the data plane.

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 December 12, 2014.

Copyright Notice

Copyright (c) 2014 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.	Terminology and Notation	3
2.1.	Acronyms	3
2.2.	Terms defined in this document	4
3.	Shared Mesh Protection Reference Model	4
3.1.	Protection or Restoration	4
3.2.	Scope of document	5
3.2.1.	Relationship to MPLS	5
4.	SMP Architecture	5
4.1.	Coordination of resources	7
4.2.	Control plane or data plane	8
5.	SMP Network Objectives	8
5.1.	Resource reservation and coordination	8
5.1.1.	Checking resource availability for multiple protection paths	9
5.2.	Multiple triggers	9
5.3.	Notification	10
5.4.	Revertive protection switching	10
5.5.	Protection switching time	11
5.6.	Timers	11
5.7.	Communicating information and channel	11
6.	Manageability Considerations	12
7.	Security Considerations	12
8.	IANA Considerations	12
9.	Acknowledgements	12
10.	Normative References	12
11.	Contributing Authors	13
12.	Authors' Addresses	13

1. Introduction

The MPLS Transport Profile (MPLS-TP) provides tools to construct and operate a set of connections between nodes in an MPLS network. The MPLS network is a mesh comprising nodes and the links between them. A connection may be between neighboring nodes (i.e. spanning a single physical link) or between non-adjacent nodes (spanning a path over multiple nodes). The connections in a network constitute the Label Switched Paths (LSP) that transport packets between the endpoints of these paths. The survivability of these connections, as described in [\[RFC6372\]](#), is a critical aspect for various service providers that are bound by Service Level Agreements (SLAs) with their customers.

MPLS provides control plane tools to support various survivability schemes, some of which are identified in [\[RFC4426\]](#). In addition, recent efforts in the IETF have started providing for data plane tools to address aspects of data protection. In particular, [\[RFC6378\]](#) defines a set of triggers and coordination protocol for 1:1 and 1+1 linear protection of point-to-point paths.

When considering a full-mesh network and the protection of different paths that traverse the mesh, it is possible to provide an acceptable level of protection while conserving the amount of protection resources needed to protect the different data paths. As pointed out in [\[RFC6372\]](#) and [\[RFC4428\]](#), applying 1+1 linear protection requires that resources are committed for use by both the working and protection paths. Applying 1:1 protection requires that the same resources are committed, but allows the resources of the protection path to be utilized for pre-emptible extra traffic. Extending this to 1:n or m:n protection allows the resources of the protection path to be shared in the protection of several working paths. However, 1:n or m:n protection architecture is limited by the restriction that all of the n+1 or m+n paths must have the same endpoints.

2. Terminology and Notation

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 [\[RFC2119\]](#). Although this document is not a protocol specification, the use of this language clarifies the instructions to protocol designers producing solutions that satisfy the requirements set out in this document.

The terminology used in this document is based on the terminology defined in the MPLS-TP Survivability Framework document [\[RFC6372\]](#) which in-turn is based on [\[RFC4427\]](#).

2.1. Acronyms

This document uses the following acronyms:

LSP Label Switched Path
SLA Service Level Agreement
SMP Shared Mesh Protection
SRLG Shared Risk Link Group

2.2. Terms defined in this document

This document defines the following terms:

SMP Protection Group: the set of different protection paths that share a common segment.

3. Shared Mesh Protection Reference Model

As described in [[RFC6372](#)] Shared Mesh Protection (SMP) supports the sharing of protection resources, while providing protection for multiple working paths that need not have common endpoints and do not share common points of failure. Note that some protection resources may be shared, while some others may not be. An example of data paths that employ SMP is shown in Figure 1. It shows two working paths <ABCDE> and <VWXYZ> that are protected employing 1:1 linear protection by protection paths <APQRE> and <VPQRZ> respectively. The two protection paths that traverse segment <PQR> share the protection resources on this segment.

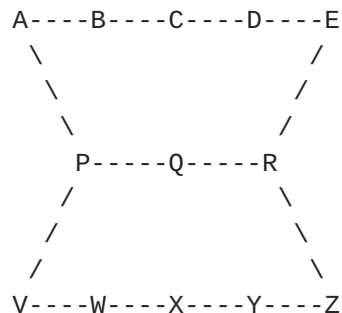


Figure 1: Basic SMP architecture

3.1. Protection or Restoration

[[RFC6372](#)], based upon the definitions in [[RFC4427](#)], differentiates between "protection" and "restoration" dependent upon the dynamism of the resource allocation. In SMP, the resources of the protection paths are planned at the time of path creation. However, the commitment of the resources, at least for the shared segments, will only be finalized when the protection path is actually activated.

Therefore, for the purists - regarding the terminology - SMP lies somewhere between protection and restoration.

3.2. Scope of document

[RFC5654] establishes that MPLS-TP SHOULD support shared protection (Requirement 68) and that MPLS-TP MUST support sharing of protection resources (Requirement 69). This document presents the network objectives and a framework for applying SMP within an MPLS network, without the use of control plane protocols. Although there are existing control plane solutions for SMP within MPLS, a data plane solution is required for networks that do not employ a full control plane operation for some reason (e.g. service provider preferences or limitations), or require service restoration faster than is achievable with control plane mechanisms.

The network objectives will also address possible additional restrictions of the behavior of SMP in networks that delegate executive action for resiliency to the data plane. Definition of logic and specific protocol messaging is out of scope of this document.

3.2.1. Relationship to MPLS

While some of the restrictions presented by this document originate from the properties of transport networks, nothing prevents the information presented here being applied to MPLS networks outside the scope of the Transport Profile of MPLS.

4. SMP Architecture

Figure 1 shows a very basic configuration of working and protection paths that may employ SMP. We may consider a slightly more complex configuration, such as the one in Figure 2 in order to illustrate characteristics of a mesh network that implements SMP.

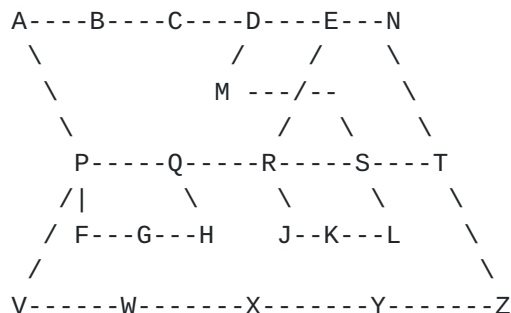


Figure 2: Larger sample SMP architecture

Consider the network presented in Figure 2. There are five working paths

- <ABCDE>
- <MDEN>
- <FGH>
- <JKL>
- <VWXYZ>

Each of these has a corresponding protection path

- <APQRE> (p1)
- <MSTN> (p2)
- <FPQH> (p3)
- <JRSL> (p4)
- <VPQRSTZ> (p5)

The following segments are shared by two or more of the protection paths - <PQ> is shared by p1, p3, and p5, <QR> is shared by p1 and p5, <RS> is shared by p4 and p5, and <ST> is shared by p2 and p5. In Figure 2, we have the following SMP Protection Groups - {p1, p3, p5} for <PQ>, {p1, p5} for <QR>, {p4, p5} for <RS>, {p2, p5} for <ST>.

We assume that the available protection resources for these shared segments are not sufficient to support the complete traffic capacity of the respective working paths that may use the protection paths. We can further observe that with a method of coordinating sharing and preemption there is no co-routing constraints on shared components at the segment level.

The use of preemption in the network is typically a business or policy decision such that when protection resources are contended, priority can be applied to determine to which parties the protection resources are committed.

As opposed to the case of simple linear protection, where the relationship between the working and protection paths is defined and

the resources for the protection path are fully committed, the protection path in the case of SMP consists of segments that are dedicated to the protection of the related working path and also segments that are shared with other protection paths such that typically the protection resources are oversubscribed to support working paths that do not share common points of failure. What is required is a preemption mechanism to implement business priority when multiple failure scenarios occur. As such, the protection resources may be planned but would not be committed until requested and resolved in relation to other members of the SMP Protection Group as part of a protection switchover.

[RFC5712] defines two types of preemption that can be considered for how the resources of SMP Protection Groups, are shared. These are "soft preemption" whereby traffic of lower priority paths is degraded and "hard preemption" where traffic of lower priority paths is completely blocked. "Hard Preemption" requires the programming of selectors at the ingress of each shared segment to specify which backup path has the highest priority when committing protection resources, the others being preempted. When any protection mechanism whereby the protection end point may have a choice of protection paths (e.g. n:1 or m:n) is deployed the shared segment selectors require coordination with the protection end points as well.

Typical deployment of services that use SMP requires various network planning activities. These include:

- o Determining the number of working and protection paths required to achieve resiliency targets for the service.
- o Reviewing network topology to determine which working or protection paths are required to be disjoint from each other, and exclude specified resources such as links, nodes, or shared risk link groups (SRLGs).
- o Determining the size (bandwidth) of the shared resource

4.1. Coordination of resources

When a protection switch is triggered, the SMP network performs two operations simultaneously - switch data traffic over to a protection path and commit the associated resources. The commitment of resources is dependent upon their availability at each of the shared segments.

When the reserved resources of the shared segments are committed to a particular protection path, there may not be sufficient resources available for an additional protection path. This then implies that if another working path of the SMP domain triggers a protection

switch, the commitment of the resources may fail. In order to optimize the operation of the commitment and preparing for cases of multiple working paths failing, the commitment of the shared resources are be coordinated between the different working paths in the SMP network.

4.2. Control plane or data plane

As stated in both [[RFC6372](#)] and [[RFC4428](#)], full control of SMP including both configuration and the coordination of the protection switching is potentially very complex. Therefore, it is suggested that this be carried out under the control of a dynamic control plane similar to GMPLS [[RFC3945](#)]. Implementations for SMP with GMPLS exist and the general principles of its operation are well known, if not fully documented.

However, there are operators, in particular in the transport sector, that do not operate their MPLS-TP networks under the control of a control plane or for other reasons have delegated executive action for resilience to the data plane, and require the ability to utilize SMP protection. For such networks it is imperative that it be possible to perform all required coordination of selectors and end points for SMP via data plane operations.

5. SMP Network Objectives

5.1. Resource reservation and coordination

SMP is based on pre-configuration of the working paths and the corresponding protection paths. This configuration may be based on either a control protocol or static configuration by the management system. However, even when the configuration is performed by a control protocol, e.g. Generalized MPLS (GMPLS), the control protocol SHOULD NOT be used as the primary mechanism for detecting or reporting network failures, or for initiating or coordinating protection switch-over. That is, it SHOULD NOT be used as the primary resilience mechanism.

The protection relationship between the working and protection paths SHOULD be configured and the shared segments of the protection path MUST be identified prior to use of the protection paths. Relative priority for working paths to be used to resolve contention for protection path usage by multiple working paths MAY also be specified ahead of time.

When a protection switch is triggered by any fault condition or operator command, the SMP network MUST perform two operations simultaneously - switch data traffic over to a protection path and

commit the associated resources.

In the case of multiple working paths failing, the commitment of the shared resources SHALL be coordinated between the different working paths in the SMP network.

5.1.1. Checking resource availability for multiple protection paths

In a hard-preemption scenario, when an end point identifies a protection switching trigger and has more than one potential action (e.g. n:1 protection) it MUST verify that the necessary protection resources are available on the selected protection path. The resources may not be available because they already have been committed to the protection of, for example, a higher priority working path.

5.2. Multiple triggers

If more than one working path is triggering a protection switch such that a protection segment is oversubscribed, there are different possible actions that the SMP network may apply. The basic MPLS action MAY allow all of the protection paths to share the resources of the shared segments (soft-preemption), for networks that support multiplexing packets over the shared segments.

There are networks that require the exclusive use of the protection resources (i.e. hard preemption). These include networks that support the requirements in [[RFC5654](#)], and in particular support requirement 58. For such networks, the following requirements apply:

- o Relative priority MAY be assigned to each of the working paths of an SMP domain. If the priority is not assigned, the working paths are assumed to have equal priority.
- o Resources of the shared segments SHALL be committed to the protection path according to the highest priority amongst those requesting use of the resources.
- o If multiple protection paths of equal priority are requesting allocation of the shared resources, the resources SHOULD be committed on a first come first served basis. Tie-breaking rules SHALL be defined in scope of an SMP domain.
- o If the protection resources are committed to a protection path, whose working path has a lower priority, resources required for the higher priority path SHALL be committed to the higher priority path. Traffic with lower priority MAY use available resources or

MAY be interrupted.

- o When triggered, protection switching action SHOULD be initiated immediately to minimize service interruption time.
- o If the protection resources are already committed to a higher priority protection path the protection switching SHALL NOT be performed.
- o Once resources of shared segments have been successfully committed to a protection path, the traffic on that protection path SHALL NOT be interrupted by any protection traffic whose priority is equal or lower than the protecting path currently in-use.
- o During preemption, shared segment resources MAY be used by both existing traffic (that is being preempted) and higher priority traffic.
- o During preemption, if there is an over-subscription of resources protected traffic SHOULD be treated as defined in [[RFC5712](#)] or [[RFC3209](#)].

5.3. Notification

When a working path endpoint has a protection switch triggered, it SHOULD attempt to switch the traffic to the protection path and request the commitment of protection resources. If the necessary shared resources are unavailable to be committed to the protection path, the endpoints of the requesting working path SHALL be notified of protection switchover failure, and switchover MAY not be completed.

Similarly, if preemption is supported and the resources currently committed for a particular working path are being preempted then the endpoints of the affected working path whose traffic is being preempted SHALL be notified that the resources are being preempted.

5.4. Revertive protection switching

When the working path detects that the condition that triggered the protection switch has cleared, it is possible to either revert to using the working path resources or continue to utilize the protection resources. Continuing the use of protection resources allows the operator to delay the disruption of service caused by the switchover until periods of lighter traffic. The switchover would need to be performed via an explicit operator command unless the protection resources are preempted by a higher priority fault. Hence, both automatic and manual revertive behaviors MUST be supported for

hard-preemption in an SMP domain. Normally the network should revert to use of the working path resources in order to clear the protection resources for protection of other path triggers. However, the protocol MUST support non-revertive configurations.

5.5. Protection switching time

Protection switching time refers to the transfer time (T_t) defined in [G.808.1] and recovery switching time defined in [RFC4427], and is defined as the interval after a switching trigger is identified until the traffic begins to be transmitted on the protection path. This time does not include the time needed to initiate the protection switching process after a failure occurred, and the time needed to complete preemption of existing traffic on the shared segments as described in Section 4.2. The former, which is known as detection and correlation time in [RFC4427], is related to the OAM or management process, but the latter is related to the actions within an SMP domain. Support for a protection switching time of 50ms is dependent upon the initial switchover to the protection path, but the preemption time SHOULD also be taken into account to minimize total service interruption time.

5.6. Timers

In order to prevent multiple switching actions for a single switching trigger, when there are multiple layers of networks, SMP SHOULD be controlled by a hold-off timer that would allow lower layer mechanisms to complete their switching actions before invoking SMP protection actions.

In order to prevent an unstable recovering working path from invoking intermittent switching operation, SMP SHOULD employ a wait-to-restore timer during any reversion switching.

5.7. Communicating information and channel

SMP in hard-preemption mode SHOULD include support for communicating information to coordinate the use of the shared protection resources among multiple working paths. The message encoding and communication channel between the nodes of the shared protection resource and the endpoints of the protection path are out of the scope of this document.

SMP in hard-preemption mode SHOULD provide a communication channel, along the protection path, between the endpoints of the protection path to support fast protection switching.

6. Manageability Considerations

The network management architecture and requirements for MPLS-TP are specified in [[RFC5951](#)]. They derive from the generic specifications described in ITU-T G.7710/Y.1701 [G.7710] for transport technologies. This document does not introduce any new manageability requirements beyond those covered in those documents.

7. Security Considerations

General security considerations for MPLS-TP are covered in [[RFC5921](#)]. The security considerations for the generic associated control channel are described in [[RFC5586](#)]. This document introduces no new security considerations beyond those covered in those documents.

8. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

9. Acknowledgements

This document is the outcome of discussions on Shared Mesh Protection for MPLS-TP. The authors would like to thank all contributors to these discussions, and especially Eric Osborne for facilitating them.

We would also like to thank Matt Hartley for working on the English review.

10. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., and V. Srinivasan, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [RFC 3945](#), Oct 2004.
- [RFC4426] Lang, J., Rajagopalan, B., and Papadimitriou, D.E. "GMPLS Recovery Functional Specification", [RFC 4426](#), March 2006.
- [RFC4427] Mannie, E. and D. Papadimitriou, "Recovery (Protection and Restoration) Terminology for GMPLS", [RFC 4427](#), March 2006.

- [RFC4428] Mannie, E. and D. Papadimitriou, "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", [RFC 4428](#), March 2006.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), June 2009.
- [RFC5654] Niven-Jenkins, B., Nadeau, T., and C. Pignataro, "Requirements for the Transport Profile of MPLS", [RFC 5654](#), Sept 2009.
- [RFC5712] Meyer, M. and JP. Vasseur, "MPLS Traffic Engineering Soft Preemption", [RFC 5712](#), January 2010.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), July 2010.
- [RFC5951] Lam, K., Mansfield, S., and E. Gray, "Network Management Requirements for MPLS-based Transport Networks", [RFC 5951](#), September 2010.
- [RFC6372] Sprecher, N. and A. Farrel, "MPLS-TP Survivability Framework", [RFC 6372](#), Sept 2011.
- [RFC6378] Sprecher, N., Bryant, S., Osborne, E., Fulignoli, A., and Y. Weingarten, "MPLS-TP Linear Protection", [RFC 6378](#), Nov 2011.
- [G.808.1] ITU, "Generic Protection Switching - Linear trail and subnetwork protection", ITU-T G.808.1, May 2014.

11. Contributing Authors

David Allan
Ericsson
Email: david.i.allan@ericsson.com

Daniel King
Old Dog Consulting
Email: daniel@olddog.co.uk

Taesik Cheung
ETRI
Email: cts@etri.re.kr

12. Authors' Addresses

Yaacov Weingarten
34 Hagefen St.
Karnei Shomron, 4485500
Israel

Email: wyaacov@gmail.com

Sam Aldrin
Huawei Technologies
2330 Central Express Way
Santa Clara, CA 95051
United States

Email: aldrin.ietf@gmail.com

Ping Pan
Infinera

Email: ppan@infinera.com

Jeong-dong Ryoo
ETRI
218 Gajeongno
Yuseong, Daejeon 305-700
South Korea

Email: ryoo@etri.re.kr

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

