Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2013

Y. Weingarten

S. Aldrin
Huawei Technologies

P. Pan
Infinera

October 18, 2012

Requirements for MPLS Shared Mesh Protection
draft-weingarten-mpls-smp-requirements-01.txt

Abstract

This document presents the basic network objectives for the behavior of shared mesh protection (SMP) not based on control-plane support. This is an expansion of the basic requirements presented in the MPLS Transport Profile Requirements (RFC5654) and MPLS Transport Profile Survivability Framework (RFC6372) documents. This document should be used as a basis for the definition of the mechanism that would be used to implement SMP for MPLS-TP data paths, in networks that do not employ a control plane for their operation.

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 21, 2013.

Table of Contents

1. Introduction

MPLS transport networks can be characterized as being a network of connections between nodes within a mesh of nodes and the links between them. The connections, that may be between neighboring nodes, i.e. spanning a single physical link, or spanning a path of several nodes, 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 (SLA) with their customers.

MPLS provides control-plane tools to support various survivability schemes (Editor's note - add references). 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 p2p paths.

When considering a full-mesh network and the protection of different paths that criss-cross the mesh, it is possible to conserve 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 allocated and used by both the working and protection paths. Applying 1:1 protection requires that all of the resources are allocated, 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, there is a limitation in 1:n protection architectures - that all of the n+1 paths must have identical endpoints.

As described in [RFC6372] Shared Mesh Protection (SMP) supports a form of sharing protection resources, while providing protection for multiple data paths that may not have common endpoints and do not share common points of failure. It should be noted that some protection resources may not be shared by multiple protection paths, while other resources are shared. The basic configuration for data paths that employ SMP is shown in Figure 1. In this figure, we show two working paths [ABCDE] and [VWXYZ] that are protected employing 1:1 linear protection by protection paths [APQRE] and [VPQRZ] respectively. The segment [PQR] and all of its protection resources are shared by both of the protection paths.

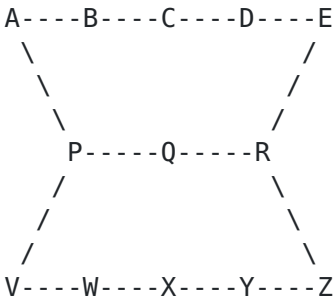

Figure 1: Basic SMP architecture

1.1. 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 reserved at the time of path creation. However, the full allocation 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.

1.2. 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. There are existing control-plane solutions for SMP within MPLS, however we address those networks that for some reason, e.g. service provider preferences or limitations, do not employ a full control plane operation, or require service restoration faster than achievable with control plane mechanisms.

The network objectives will also address possible additional restrictions of the behavior of SMP in statically configured operator networks. Definition of logic and specific protocol messaging is out of scope of this document.

1.2.1. Relationship to MPLS-TP

While some of the restrictions presented by this framework originate from the considerations of transport networks, there is no real constraint of the information presented here being applied to general MPLS networks, and not necessarily as part of the Transport Profile

of MPLS.

1.3. Contributing Authors

David Allan, Gregory Mirsky, Daniel King, Jeong-Dong Ryoo, Taesik Cheung

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].

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 draft uses the following acronyms:

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

3. SMP Architecture

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

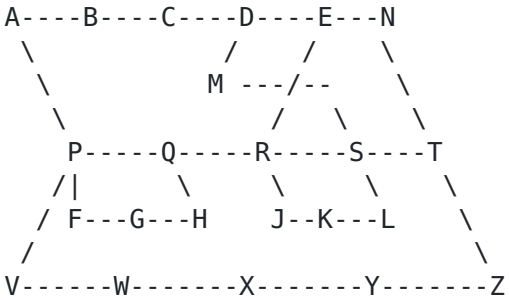

Figure 2: Larger sample SMP architecture

Consider the network presented in Figure 2. There are five working paths - [ABCDE], [MDEN], [FGH], [JKL], and [VWXYZ]. Each of these has a corresponding protection path - [APQRE] (p1), [MSTN] (p2), [FPQH] (p3), [JRSL] (p4), and [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 addition, 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 the main feature of the network that defines it as an SMP network is the fact that the segment [PQRST] is the union of all the shared segments of other protection paths (p1, p2, p3 and p4) while being a whole shared segment of one of the protection paths (p5).

In other words, the main feature of an SMP "protection domain" will be the segment that is the union of all the shared segments of the protection paths. We can further identify "protection group" as the different protection paths that share a common segment. For example, referring to Figure 2, we have the following protection groups - {p1, p3, p5} for [PQ], {p1, p5} for [QR], {p4, p5} for [RS], {p2, p5} for [ST].

Typical deployment of SMP would require various network planning activities. These would include:

- Identification of key services that require protection, and determining the number of working and protection paths.

- Reviewing network topology to determine which working or protection paths are required to be disjointed from each other, and exclude specified resources such as links, nodes, or shared risk link groups (SRLGs).

- Determining the size (bandwidth) of the shared resource

3.1. Coordination of resources

When a protection switch is triggered by any fault condition or operator command, the SMP network must perform two operations almost simultaneously - switch data traffic over to a protection path and verify that the shared resources are allocated for this protection path. The allocation of resources is dependent upon their availability at each of the shared segments.

When the reserved resources of the shared segments are allocated for a particular protection path, there may not be sufficient resources

available for an additional protection path. This then implies that if an additional working path triggers a protection switch, the allocation of the resources may fail and MUST be treated as described below in Section 4.3. In order to optimize the operation of the allocation and preparing for cases of multiple working path failures, the allocation of the shared resources SHALL be coordinated between the different working paths in the SMP network.

4. SMP Network Objectives

4.1. Configuration and resource reservation

SMP is a survivability mechanism that is based on pre-configuration of the network working paths and the corresponding protection paths. This configuration may be based on either a control protocol or static configuration by the management system. 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.

As opposed to the case of simple linear protection, where the relationship between the working and protection paths is defined, the resources for the protection path may be fully committed for the unshared portions of the protection path. 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. On the shared segments, the protection resources may be reserved but would not be allocated until requested as part of a protection switch.

4.1.1. Checking resource availability

When a working path identifies a protection switching trigger it MUST verify that the necessary protection resources are available on the protection path. The resources may not be available because they have been allocated to the protection of a higher priority working path, as described above.

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]. In fact, implementations for SMP with GMPLS exist and the general principles of its operation are well known, if not fully documented.

There are, however, operators, in particular in the transport sector, that do not operate their MPLS networks under the control of a control plane and require the ability of performing SMP protection while utilizing data-plane tools for coordination of the protection switching. This requirement is emphasized in different areas of [RFC5654] for MPLS-TP environments. Therefore, it is imperative that it be possible to perform all of the coordination needed for SMP via data plane operations.

4.3. Multiple triggers

If more than one working path is triggering a protection switch 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, for those networks that support multiplexing packets over the shared segments. For those networks, in particular for networks that support the requirements in [RFC5654] [and in particular support for requirement 58], that require the exclusive use of the protection resources, the following behavior SHOULD be supported:

- Relative priority MAY be assigned to each of the working paths that share a common protection segment

- Resources of the shared segments SHALL be allocated to the protection path according to the highest priority amongst those requesting use of the resources.

- If multiple protection paths of equal priority are requesting allocation of the shared resources, the resources SHOULD be allocated on a first come first served basis. Tie-breaking rules SHALL be defined by the SMP process.

- If the protection resources are currently in use by a protection path, whose working path has a lower priority, resources SHALL be allocated to the path with higher priority. Traffic with lower priority MAY use available resources or MAY be interrupted.

- When trigerred, protection switching action SHOULD be initiated immediately to minimize service interruption time. If the protection resources are already allocated to a higher priority protection path the protection switching MAY not be performed.

- Once a protection path occupies the resource of a shared segments successfully, 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.

- During preemption, shared segment resources MAY be used by both existing traffic (that is being preempted) and higher priority traffic for a short period.

- During preemption, if there is an oversubscription of resources protected traffic SHOULD be treated as defined in [RFC5712] or [RFC3209]

4.4. Notification

When a working path identifies a trigger for implementing a switchover to the protection path, it SHOULD attempt to switchover the traffic to the protection path and requesting the allocation of the resources for this protected traffic. If the necessary shared resources are in use by a protection path of higher priority or are unavailable to be allocated to the protection path, a notification SHALL be sent to both endpoints of the requesting working path and the switchover MAY not be completed.

Similarly, if preemption is supported and as a result of the allocation of resources to a different working path that triggered a protection switch, the resources currently allocated for a particular working path are being preempted then a notification SHALL be sent to the endpoints of the working path whose traffic is being preempted indicating that the resources are being preempted.

4.5. Protection switching time

Protection switching time refers to the transfer time (Tt) 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 is exclusive of 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 4.3. 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 SMP process. 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.

4.6. Timers

In order to prevent multiple switching actions for a single switching trigger, SMP SHOULD be controlled a hold-off timer that would allow lower level mechanisms to complete their switching actions before

invoking SMP protection actions.

In addition, 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. Managability Considerations

To be added in future version.

6. Security Considerations

To be added in future version.

7. IANA Considerations

This document makes no request of IANA.

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

8. Acknowledgements

TBD

9. Normative References

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

[RFC5654] Niven-Jenkins, B., Nadeau, T., and C. Pignataro, "Requirements for the Transport Profile of MPLS", RFC 5654, Sept 2009.

[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.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, Oct 2004.

[G.808.1] ITU, "Generic Protection Switching - Linear trail and subnetwork protection", ITU-T G.808.1, Feb 2010.

[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.

[RFC5712] Meyer, M. and JP. Vasseur, "Recovery (Protection and Restoration) Terminology for GMPLS", RFC 5712, January 2010.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., and V. Srinivasan, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

Authors' Addresses

Yaacov Weingarten
34 Hagefen St.
Karnei Shomron, 4485500
Israel

Phone:
Email: wyaacov@gmail.com

Sam Aldrin
Huawei Technologies
2330 Central Express Way
Santa Clara, CA 95951
United States

Email: aldrin.ietf@gmail.com

Ping Pan
Infinera

Email: ppan@infinera.com