                Requirements for MPLS Transport Profile (MPLS-TP)
                           Shared Mesh Protection

Abstract

   This document presents the basic network objectives for the behavior
   of Shared Mesh Protection (SMP) that are not based on control-plane
   support. This document provides an expansion of the basic
   requirements presented in RFC 5654 ("Requirements of an MPLS
   Transport Profile") and RFC 6372 ("MPLS Transport Profile (MPLS-TP)
   Survivability Framework"). This document provides requirements for
   any mechanism that would be used to implement SMP for MPLS-TP data
   paths, in networks that delegate protection switch coordination to
   the data plane.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

    The MPLS Transport Profile (MPLS-TP) is described in [RFC5921].
    [RFC6372] provides a survivability framework for MPLS-TP and is the
    foundation for this document.

    Terminology for recovery of connectivity in networks is provided in
    [RFC4427] and includes the concept of surviving network faults
    (survivability) through the use of re-established connections
    (restoration) and switching of traffic to pre-established backup
    paths (protection). 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] and [RFC7271] define a set of triggers and
    coordination protocols 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 [RFC4427], applying 1+1 protection requires that
    resources are allocated for use by both the working and protection
    paths. Applying 1:1 protection requires that the same resources are
    allocated but allows the resources of the protection path to be
    utilized for preemptible 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. m:n protection
    architecture provides m protection paths to protect n working paths,
    where m or n can be 1.

    This document provides requirements for any mechanism that would be
    used to implement SMP for MPLS-TP data paths, in networks that
    delegate protection switch coordination to the data plane.

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 and Terminology

   This document uses the following acronyms:

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

   This document defines the following term:

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

3. Shared Mesh Protection Reference Model

   As described in [RFC6372], 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.

```
              A----B----C----D----E
               \                  /
                \                /
                 \              /
                  P-----Q-----R
                 /              \
                /                \
               /                  \
              V----W----X----Y----Z
```

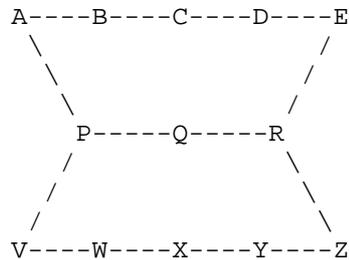   Figure 1: Basic SMP Architecture

3.1. Protection or Restoration

   [RFC6372], based upon the definitions in [RFC4427], differentiates
   between "protection" and "restoration", depending on the dynamism of
   the resource allocation. The same distinction is used in [RFC3945],
   [RFC4426], and [RFC4428].

   This document also uses the same distinction between protection and
   restoration as the distinction stated in [RFC6372].

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 on the behavior of SMP in networks that delegate
   protection switching for resiliency to the data plane. Definitions
   of logic and specific protocol messaging are out of scope for 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 from 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.

```
                A----B----C----D----E---N
                 \             /    /     \
                  \         M ---/--       \
                   \          /    \        \
                  P-----Q-----R-----S----T
                  /|       \     \     \     \
                 / F---G---H    J--K---L      \
                /                              \
              V------W-------X-------Y-------Z
```

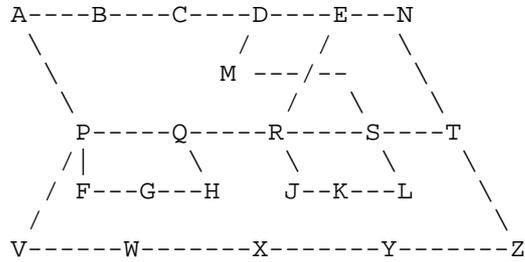   Figure 2: Example of a Larger SMP Architecture

   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>, and {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 are 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 contested,
priority can be applied to determine which parties utilize the
protection resources.

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 dedicated, the
protection path in the case of SMP consists of segments that are used
for 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
allocated but would not be utilized until requested and resolved in
relation to other members of the SMP Protection Group as part of a
protection switchover.

[RFC6372] defines two types of preemption that can be considered for
how the resources of SMP Protection Groups are shared: "soft
preemption", where traffic of lower-priority paths is degraded; and
"hard preemption", where traffic of lower-priority paths is
completely blocked. The traffic of lower-priority paths in this
document can be viewed as the extra traffic being preempted, as
described in [RFC6372]. "Hard preemption" requires the programming
of selectors at the ingress of each shared segment to specify the
priorities of backup paths, so that traffic of lower-priority paths
can be preempted. When any protection mechanism where the protection
endpoint may have a choice of protection paths (e.g., m:n or m:1) is
deployed, the shared segment selectors require coordination with the
protection endpoints as well.

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

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
   excluding 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 -- switching data traffic over to a protection path and
   coordinating the utilization of the associated shared resources.
   Both operations should occur at the same time, or as close together
   as possible, to provide fast protection. The resource utilization
   coordination is dependent upon their availability at each of the
   shared segments.

   When the reserved resources of the shared segments are utilized by 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 resource utilization coordination may fail. The
   different working paths in the SMP network are involved in the
   resource utilization coordination, which is a part of a whole SMP
   protection switching coordination.

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
   based on Generalized MPLS (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
endpoints 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., GMPLS, the control protocol SHALL NOT be used
as the primary mechanism for detecting or reporting network failures,
or for initiating or coordinating protection switchover. That is, it
SHALL 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 --
switch data traffic over to a protection path, and coordinate the
utilization of the associated shared resources. To provide fast
protection, both operations MUST occur at the same time or as close
to the same time as possible.

In the case of multiple working paths failing, the shared resource
utilization coordination SHALL be 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 endpoint identifies a
protection switching trigger and has more than one potential action
(e.g., m: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 have already been
utilized for the protection of, for example, one or more higher-
priority working paths.

5.2. Multiple Triggers

   If more than one working path is triggering a protection switch such
   that a protection segment is oversubscribed, there are two different
   actions that the SMP network can choose -- soft preemption and hard
   preemption [RFC6372].

5.2.1. Soft Preemption

   For networks that support multiplexing packets over the shared
   segments, the requirement is as follows:

   o All of the protection paths MAY be allowed to share the resources
      of the shared segments.

5.2.2. Hard Preemption

   There are networks that require the exclusive use of the protection
   resources when a protection segment is oversubscribed. Traffic of
   lower-priority paths is completely blocked. These include networks
   that support the requirements in [RFC5654], and in particular support
   Requirement 58. For such networks, the following requirements apply:

   1. 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.

   2. Resources of the shared segments SHALL be utilized by the
      protection path according to the highest priority amongst those
      requesting use of the resources.

   3. If multiple protection paths of equal priority are requesting the
      shared resources, the resources SHALL be utilized on a first come
      first served basis. Traffic of the protection paths that request
      the shared resources late SHALL be preempted. In order to cover
      the situation where the first come first served principle cannot
      resolve the contention among multiple equal-priority requests,
      i.e., when the requests occur simultaneously, tie-breaking rules
      SHALL be defined in the scope of an SMP domain.

   4. If a higher-priority path requires the protection resources that
      are being utilized by a lower-priority path, the resources SHALL
      be utilized by the higher-priority path. Traffic with the lower
      priority SHALL be preempted.

   5. Once resources of shared segments have been successfully utilized
      by a protection path, the traffic on that protection path SHALL
      NOT be interrupted by any protection traffic whose priority is
      equal to or lower than the protecting path currently in use.

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

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 coordination of the shared resource utilization. If the
   necessary shared resources are unavailable, the endpoints of the
   requesting working path SHALL be notified of protection switchover
   failure, and switchover will not be completed.

   Similarly, if preemption is supported and the resources currently
   utilized by 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.
   As described in [RFC6372], the event of preemption may be detected by
   Operations, Administration, and Maintenance (OAM) and reported as a
   fault or a degradation of traffic delivery.

5.4. Reversion

   When the condition that triggered the protection switch is 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 (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 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 time needed to initiate the protection
switching process, which is known as detection time or correlation
time in [RFC4427], is related to the OAM or management process, but
the time needed to complete preemption is related to the actions
within an SMP domain. Support for a protection switching time of
50 ms 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.

When triggered, protection switching action SHOULD be initiated
immediately to minimize 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 as described in [RFC6372].

In order to prevent an unstable recovering working path from invoking
intermittent switching operations, SMP SHOULD employ a
Wait-To-Restore timer during any reversion switching, as described in
[RFC6372].

5.7. Communication Channel and Fate-Sharing

SMP SHOULD provide a communication channel, along the protection
path, between the endpoints of the protection path, to support fast
protection switching.

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.

Bidirectional protection switching SHOULD be supported in SMP.

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].

   Security considerations for any proposed solution should consider
   exhaustion of resources related to preemption, especially by a
   malicious actor as a threat vector against which the resources should
   be protected. Protections should also be considered to prevent a
   malicious actor from attempting to create an alternate path on which
   to force traffic from a sensor/device, thereby enabling pervasive
   monitoring [RFC7258].

8. Normative References

   [RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
               Requirement Levels", BCP 14, RFC 2119, March 1997,
               <http://www.rfc-editor.org/info/rfc2119>.

   [RFC3945]   Mannie, E., Ed., "Generalized Multi-Protocol Label
               Switching (GMPLS) Architecture", RFC 3945, October 2004,
               <http://www.rfc-editor.org/info/rfc3945>.

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

   [RFC4427]   Mannie, E., Ed., and D. Papadimitriou, Ed., "Recovery
               (Protection and Restoration) Terminology for Generalized
               Multi-Protocol Label Switching (GMPLS)", RFC 4427,
               March 2006, <http://www.rfc-editor.org/info/rfc4427>.

   [RFC4428]   Papadimitriou, D., Ed., and E. Mannie, Ed., "Analysis of
               Generalized Multi-Protocol Label Switching (GMPLS)-based
               Recovery Mechanisms (including Protection and
               Restoration)", RFC 4428, March 2006,
               <http://www.rfc-editor.org/info/rfc4428>.

   [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed.,
              "MPLS Generic Associated Channel", RFC 5586, June 2009,
              <http://www.rfc-editor.org/info/rfc5586>.

   [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed.,
              Sprecher, N., and S. Ueno, "Requirements of an MPLS
              Transport Profile", RFC 5654, September 2009,
              <http://www.rfc-editor.org/info/rfc5654>.

   [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,
              <http://www.rfc-editor.org/info/rfc5921>.

   [RFC5951] Lam, K., Mansfield, S., and E. Gray, "Network Management
              Requirements for MPLS-based Transport Networks", RFC 5951,
              September 2010, <http://www.rfc-editor.org/info/rfc5951>.

   [RFC6372] Sprecher, N., Ed., and A. Farrel, Ed., "MPLS Transport
              Profile (MPLS-TP) Survivability Framework", RFC 6372,
              September 2011, <http://www.rfc-editor.org/info/rfc6372>.

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

   [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an
              Attack", BCP 188, RFC 7258, May 2014,
              <http://www.rfc-editor.org/info/rfc7258>.

   [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, June 2014,
              <http://www.rfc-editor.org/info/rfc7271>.

   [G.7710]   International Telecommunication Union, "Common equipment
              management function requirements", ITU-T Recommendation
              G.7710/Y.1701, February 2012.

   [G.808.1] International Telecommunication Union, "Generic Protection
              Switching - Linear trail and subnetwork protection", ITU-T
              Recommendation G.808.1, May 2014.

Contributors

   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

Authors' Addresses

   Yaacov Weingarten
   34 Hagefen St.
   Karnei Shomron, 4485500
   Israel

   EMail: wyaacov@gmail.com

   Sam Aldrin
   Huawei Technologies
   2330 Central Expressway
   Santa Clara, CA 95050
   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