Use Cases and Requirements for MPLS-TP multi-failure protection
draft-ietf-mpls-tp-mfp-use-case-and-requirements-03

Abstract

   For the Multiprotocol Label Switching Transport Profile (MPLS-TP)
   linear protection capable of 1+1 and 1:1 protection has already been
   defined. That linear protection mechanism has not been designed for
   handling multiple, simultaneously occuring failures, i.e. multiple
   failures that affect the working and the protection entity during the
   same time period. In these situations currently defined protection
   mechanisms would fail.

   This document introduces use cases and requirements for mechanisms
   that are capable of protecting against such failures. It does not
   specify a multi-failure protection mechanism itself.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

   Today's packet optical transport networks concentrate large volumes
   of traffic onto a relatively small number of nodes and links. As a
   result, the failure of a single network element can potentially
   interrupt a large amount of traffic. For this reason, ensuring
   survivability through careful network design and appropriate
   technical means is important.

   In MPLS-TP networks, a basic end-to-end linear protection
   survivability technique is available as specified in [RFC6378],
   [RFC7271] and [RFC7324]. That protocol however is limited to 1+1 and
   1:1 protection and not designed to handle multiple failures that
   affect both the working and protection entity at the same time.

There are various scenarios where multi-failure protection is an
important requirement for network survivability. E.g. for disaster
recovery, after catastrophic events such as earthquakes or tsunamis.
During the period after such events, network availability is crucial,
in particular for high-priority services such as emergency telephone
calls. Existing 1+1 or 1:n protection however is limited to cover
single failures which has proven as not sufficient during past
events.

Beyond the natural disaster use case above, multi-failure protection
is also beneficial in situations where the network is particularly
vulnerable, e.g., when a working entity or protection entity was
closed for maintenance or construction work. During this time, the
network service becomes vulnerable to single failures since one
entity is already down. If a failure occurs during this time, an
operator might not be able to meet service level agreements (SLA).
Thus, a technical means for multi-failure protection could take
pressure off network operations.

1.1. Document scope

This document describes use cases and requirements for m:1 and m:n
protection in MPLS-TP networks without the use of control plane
protocols. Existing solutions based on a control plane such as GMPLS
may be able to restore user traffic when multiple failures occur.
Some networks however do not use full control plane operation for
reasons such as service provider preferences, certain limitations or
the requirement for fast service restoration (faster than achievable
with control plane mechanisms). These networks are the focus of this
document which defines a set of requirements for m:1 and m:n
protection not based on control plane support. This document imposes
no formal time constraints on detection times.

1.2. Requirements 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].

1.3. Terminology

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].

In particular, the following protection schemes are defined in
[RFC4427] and used as terms in this document:

o 1+1 protection

o 1:n (n >= 1) protection

o m:n (m, n > 1, n >= m > 1) protection

o Further, the following additional terminology is from [RFC4427] is
   used:

o "broadcast bridge"

o "selector bridge"

o "working entity"

o "protection entity"

This document defines a new protection type:

o m:1 (m > 1) protection: A set of m protection entities protecting
   a single working entity

2. General m:n protection scenario

The general underlying assumption of this work is that an m:n
relationship between protection entity and working entity exists,
i.e. there is no artificial limitation on the ratio between
protection and working entities.

This general scenario is illustrated in Figure 1 which shows a
protection domain with n working entities and m protection entities
between Node A and Node Z.

At Node A, traffic is transported over its respective working entity
and may be simultaneously transported over one of its protection
entities (in case of a broadcast bridge), or it is transported over
its working entity and only in case of failure over one of the
protection entities (in case of a selector bridge). At Node Z, the
traffic is selected from either its working entity or one of the
protection entities. Note that any of the n working entities and m
protection entities should follow a disjoint path through the network
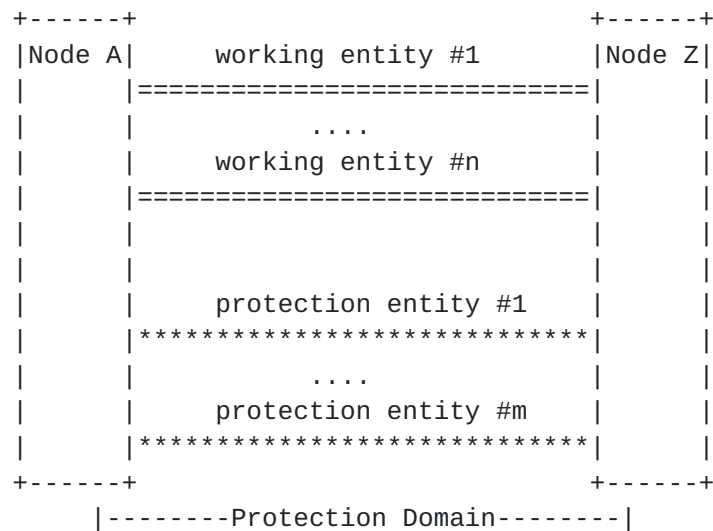from Node A to Node Z.

```
               +------+                          +------+
               |Node A|     working entity #1     |Node Z|
               |      |=============================|      |
               |      |          ....              |      |
               |      |     working entity #n       |      |
               |      |=============================|      |
               |      |                             |      |
               |      |                             |      |
               |      |     protection entity #1    |      |
               |      |*****************************|      |
               |      |          ....              |      |
               |      |     protection entity #m    |      |
               |      |*****************************|      |
               +------+                          +------+
                  |--------Protection Domain--------|
```

Figure 1: m:n protection domain

3. Use cases

3.1. m:1 (m > 1) protection

With MPLS-TP linear protection such as 1+1/1:1 protection, when a
single failure is detected on the working entity, the service can be
restored using the protection entity. However, during the time the
protection is active the traffic is unprotected until the working
entity is restored.

m:1 protection can increase service availability and reduce
operational pressure since multiple protection entities are
available. For any m > 1, m - 1 protection entities may fail and the
service still would have a protection entity available.

There are different ways to provision these alternative protection
entities which are outlined in the following sub-sections.

3.1.1. Pre-configuration

The relationship between the working entity and the protection
entities is part of the system configuration and needs to be
configured before the working entity is being used. The same applies
to additional protection entities.

Unprotected traffic can be transported over the m protection entities
as long as these entities do not carry protected traffic.

3.1.2. On-demand configuration

The protection relationship between a working entity and a protection
entity is configured while the system is in operation.

Additional protection entities are configured by either a control
plane protocol or static configuration using a management system
directly after failure detection and/or notification of either the
working entity or the protection entities. In case a management
system is used, there is no need for a standardized solution.

3.2. m:n (m, n > 1, n >= m > 1) protection

Because m:1 protecion introduces additional protection entities
compared to 1:1 protection, an additional cost has to be paid. In
order to reduce the cost of these additional protection entities, in
the m:n scenario, m dedicated protection transport entities are
sharing protection resources for n working transport entities.

The bandwidth of each protection entity should be allocated in such a
way that it may be possible to protect any of the n working entities
in case at least one of the m protection entities is available. When
a working entity is determined to be impaired, its traffic first must
be assigned to an available protection transport entity followed by a
transition from the working to the assigned protection entity at both
Node A and Node Z of the protected domain. It is noted that when
more than m working entities are impaired, only m working entities
can be protected.

4. Requirements

Recovery requirements are defined in section 2.5 of RFC 5654
[RFC5654]. More specifically, RFC 5654 outlines protection
requirements in subsections 2.5.1.1. and 2.5.1.2. These however are
limited to cover single failure cases and not multiple,
simultaneously occuring failures. This section extends the list of
requirements to support multiple failures scenarios.

R1. MPLS-TP SHOULD support m:1 (m > 1) protection.

1. An m:1 protection mechanism MUST protect against multiple
    failures that are detected on both the working entity and one or
    more protection entities.

2. Pre-configuration of protection entities SHOULD be supported.

3. On-demand protection entity configuration MAY be supported.

4. On-demand protection resource activation MAY be supported.

5. A priority scheme MUST be provided, since a protection entity has
    to be chosen out of two or more protection entities.

R2. MPLS-TP SHOULD support m:n (m, n > 1, n >= m > 1) protection.

1. An m:n protection mechanism MUST protect against multiple
    failures that are simultaneously detected on both a working
    entity and a protection entity or multiple working entities.

2. A priority scheme MUST be provided, since protection resources
    are shared by multiple working entities dynamically.

If a solution is designed based on an existing mechanism such as PSC,
then this solution MUST be backward compatible and not break such
mechanisms.

5. 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]. The requirements described in
this document are extensions to the requirements presented in
[RFC5654] and does not introduce any new security risks.

6. IANA Considerations

This document makes no request of IANA.

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

7. Normative References

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

[RFC4427] Mannie, E. and D. Papadimitriou, "Recovery (Protection and
           Restoration) Terminology for Generalized Multi-Protocol
           Label Switching (GMPLS)", RFC 4427, March 2006.

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

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

   [RFC5921]   Bocci, M., Bryant, S., Frost, D., Levrau, L., and L.
               Berger, "A Framework for MPLS in Transport Networks", RFC
               5921, July 2010.

   [RFC6372]   Sprecher, N. and A. Farrel, "MPLS Transport Profile (MPLS-
               TP) Survivability Framework", RFC 6372, September 2011.

   [RFC6378]   Weingarten, Y., Bryant, S., Osborne, E., Sprecher, N., and
               A. Fulignoli, "MPLS Transport Profile (MPLS-TP) Linear
               Protection", RFC 6378, October 2011.

   [RFC7271]   Ryoo, J., Gray, E., 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.

   [RFC7324]   Osborne, E., "Updates to MPLS Transport Profile Linear
               Protection", RFC 7324, July 2014.

Authors' Addresses

   Zhenlong Cui
   NEC

   Email: c-sai@bx.jp.nec.com

   Rolf Winter
   NEC

   Email: Rolf.Winter@neclab.eu

   Himanshu Shah
   Ciena

   Email: hshah@ciena.com

   Sam Aldrin
   Huawei Technologies

   Email: aldrin.ietf@gmail.com

Masahiro Daikoku
KDDI

Email: ms-daikoku@kddi.com